Content Server Troubleshooting Guide 10g Release 3 (10.1.3.3.0)

March 2007



Troubleshooting Guide, 10g Release 3 (10.1.3.3.0) Copyright © 2007, Oracle. All rights reserved.

Contributing Authors: Deanna Burke, Sandra Christiansen, Eric Raney, Bruce Silver

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software--Restricted Rights (June 1987). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of thirdparty products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Table of Contents



Chapter 1: Introduction

Overview	
About This Guide	
Product Documentation	

Chapter 2: Troubleshooting Overview

Overview	.2-1
Symptoms, Problems, and Solutions	.2-1
General Problem-Solving Model	.2-2
Preparing for Issues	.2-3

Chapter 3: Finding Error and Status Information

Overview
Log Files
Log File Characteristics
Accessing the Log Files
Using Content Server Logs
Using Console Output Logs
Setting Up Console Output Capturing
Opening Console Output Logs
Using Archiver Logs
Inbound Refinery Logs
Configuration Information
System Audit Information
Web Server Filter Configuration
Tracing
Server-Wide Tracing
Applet-Specific Tracing

Environment Packager	3-21
Content Server Analyzer	3-23
Content Server Analyzer: Configuration Tab	
Content Server Analyzer: Progress Tab	3-26
Using Content Server Analyzer	3-27
Accessing the Content Server Analyzer	
Specifying a Custom Analyzer Log Directory	3-27
Invoking the Analysis Process	3-28
Analyzing the Content Server Database	3-28
Analyzing the Content Server Search Index	3-29
Analyzing the Content Server File System	3-30
Viewing the Analysis Progress and Results	3-30
Generating a Status Report	3-31
Cancelling the Status Report	3-31
Configuration Entries	3-32
Stack Traces	3-33
Chapter 4: Web Server Issues	

Overview
Permission Issues
Login Errors After Manual Reinstallation
Difficulties Accessing a Specific Document
File Not Found Error on Download
Content Server or Web Server Won't Start4-4
Web Server Runs But Content Server Won't Start
Cannot Start the Website or Content Server
Admin Server Page will not Open4-8
Port Configuration Issues
Changing the Port Setting for IIS
iPlanet-Specific Issues
iPlanet Freezes and Content Server is Unavailable
Miscellaneous Web Server Issues
Setting Up Web Server Logging4-12
Portal Page Buttons/Links not Working
Turning Off Logging When Running Content Server As Windows Service 4-14

Chapter 5: Database Issues

Overview	
Content Server Terminates or Won't Start	
Permission Problem with IDC Content Service	
Database Password Changed	
Changed/Forgot Content Server 'sa' Password	
Content Server Can't Connect to Database	
General Connectivity Problem Analysis	
ODBC / System DSN Connection Information	
	5-10
Pointing to Wrong Database	
Content Server Doesn't Connect to Manually Configured SQL Server Database	
New User Passwords Are Not Recognized	
Database Password Is Not Encrypted	5-15
Database Connection Terminates or Fails	5-17
Database Connection Repeatedly Fails	
Database Connection Intermittently Fails	5-18
Database Runs out of Connections	5-19
Need to Increase Database Connection Limits	5-19
Unable to Execute Service and Function	
Unable to Load User Information	5-21
Database Sizing Issues	5-22
Estimating Sizes for Database Files	
Estimating the Open File Limit Size (UNIX)	
Exceeded the Open File Limit (UNIX)	
Distributed Installation Issues	
Moving or Using an Alternate Database	
Linking to the Database Through a Firewall	
SQL Server-Specific Issues	
Setting Up Low Resource Alerts and MS Paging	
Oracle-Specific Issues	
Check-In Problems/Errors	
Increasing Tablespace.	
Checking Oracle Driver Information.	
DB2 Deadlocks or Limeouts Occur	

Miscellaneous Issues	5-32
New Check-Ins Compete for IDs	5-32
Changing User Permissions	5-33
Changing Host Name / Instance Name	5-33
Turning On System Database Tracing	5-34

Chapter 6: Archiving Issues

0-1
6-1
6-2
6-3
6-3
6-4
6-5
6-6
6-8
6-9
6-11
6-12
6-12
6-13
6-13
6-13 6-14
6-13 6-14 6-15
6-13 6-14 6-15 6-16
6-13 6-14 6-15 6-16 6-17
6-13 6-14 6-15 6-16 6-17 6-18
6-13 6-14 6-15 6-16 6-17 6-18 6-19
6-13 6-14 6-15 6-16 6-17 6-18 6-19 6-19
6-13 6-14 6-15 6-16 6-17 6-18 6-19 6-20
6-13 6-14 6-15 6-16 6-17 6-18 6-19 6-19 6-20 6-20
6-13 6-14 6-15 6-16 6-17 6-18 6-19 6-19 6-20 6-20 6-22
6-13 6-14 6-15 6-16 6-17 6-18 6-19 6-19 6-20 6-20 6-22 6-22
6-13 6-14 6-15 6-16 6-17 6-18 6-19 6-20 6-20 6-22 6-22 6-22 6-23
6-13 6-14 6-15 6-16 6-17 6-18 6-19 6-19 6-20 6-20 6-22 6-22 6-23 6-23

Chapter 7: Workflow Issues

Overview	.7-1
Workflow Item Stuck in EDIT or GENWWW Status	.7-1
Workflow Item Stuck in REVIEW Status	.7-3
Workflow Item Entered in Wrong Workflow	.7-4

Chapter 8: International Issues

Overview
Display Issues
Characters Are Displayed on Screen as Question Marks
Characters Are Displayed on Screen as Boxes
Characters Are Displayed on Screen as Garbage
Some Multi-Byte Characters Are Displayed Slightly Differently Than Expected8-4
Text in Forms With Re-used Data Is Displayed in &#xxxx; Form8-5
Characters Are Not Displayed Entirely Correctly in Internet Explorer 5.x8- 6
Some Bookmarked URLs Don't Work in Internet Explorer8-6
PDF Files Cannot Be Viewed in Internet Explorer
Web-Viewable Files Cannot Be Viewed
Some Thumbnails Are Not Displayed8-8
Sort Order is Not Always Consistent
Login Issues
Users With Accented User Name And/Or Password Cannot Log In8-10 Some Users Cannot Log In After Content Server
File Encoding Is Changed8-10

Appendix A: Understanding the Database

Overview A-1
Data Dictionary Terminology A-1
Data Dictionary A-3
Alias A-4
AliasUser
ArchiveChangedRows
ArchiveHistory
Collaborations
Config A-10
Counters A-11
DatedCaches A-12

DeletedRows A-13
DocFormats A-14
DocMeta A-15
DocMetaDefinition A-16
DocTypes A-18
DocumentAccounts
DocumentHistory A-20
Documents
ExtensionFormatMap A-24
HtmlConversions
HtmlConversionSums A-26
OptionsList
ProblemReports A-28
ProjectDocuments
RegisteredProjects A-32
Revisions
RoleDefinition A-38
SecurityGroups A-39
Subscription A-40
Users A-42
UserSecurityAttributes A-45
WorkflowAliases A-46
WorkflowCriteria A-47
WorkflowDocAttributes A-48
WorkflowDocuments A-49
WorkflowHistory
Workflows A-54
WorkflowStates A-56
WorkflowSteps A-57
sysconstraints A-58
syssegments A-59

Index

INTRODUCTION

OVERVIEW

This section covers the following topics:

- ✤ About This Guide (page 1-1)
- Product Documentation (page 1-2)

ABOUT THIS GUIDE

This *Troubleshooting Guide* is intended for people who are responsible for maintaining, administering, or supporting Content Server environments.

Chapters 2 and 3 of this guide provide general information about troubleshooting a Content Server environment and how to diagnose issues:

- Troubleshooting Overview (chapter 2)
- Finding Error and Status Information (chapter 3)

Chapters 4 to 8 provide more in-depth information about troubleshooting in specific areas:

- Web Server Issues (chapter 4)
- Database Issues (chapter 5)
- Archiving Issues (chapter 6)

- ✤ Workflow Issues (chapter 7)
- International Issues (chapter 8)

Generally, the chapters are organized using general problem topics. Each general problem topic is further divided into specific related issues.

Symbols

The following symbols are used throughout this document:

Symbols	Description
?	This is a note. It is used to bring special attention to information.
٢	This is a technical tip. It is used to identify information that can be used to make your tasks easier.
	This is an important notice. It is used to identify a required step or required information.
	This is a caution. It is used to identify information that might cause loss of data or serious system problems.

PRODUCT DOCUMENTATION

The Content Server software is shipped with a comprehensive set of electronic documentation, which is available from two main locations:

- Before installation of the Content Server software, it is available on the software distribution media (typically a separate CD labeled 'documentation').
- After installation of the Content Server software, it is also available on the computer the software was installed on, in the following directory: <Install_Dir>\weblayout\help\documentation

(where *<Install_Dir>* is Content Server's installation directory).

The easiest way to access the documentation set is through a special navigation menu, which can be called up by starting the file *menu.pdf* in the documentation directory. If you installed the Content Server software under Microsoft Windows, there is a shortcut to this file in Start—Programs—Content Server—*[Instance_Name]*—Utilities— Documentation.

In addition to these documents, there are also the release notes, which list new and enhanced features of each new software release, and also provide special, up-to-theminute considerations for installing and using the software. The release notes are important documents. Always make sure you read them before installing or updating Oracle software.



Note: To view and print the documentation files, you need the Adobe Acrobat Reader. This is a free utility, which can be downloaded from the Internet at http://www.adobe.com/products/acrobat/readstep2.html.



Note: The documentation root directory contains a file called *tips_en.pdf*, which provides useful information about using the PDF documentation set. This file can also be accessed from the navigation menu.

Chapter

2

TROUBLESHOOTING OVERVIEW

OVERVIEW

Dependency on content management systems has grown tremendously over the past number of years. In today's global economy, an organization's success is highly dependent on the availability and reliability of its content management system. Content server troubleshooting has therefore become a crucial element to many organizations.

In addition to increasingly relying on content management systems, organizations are also moving toward increasingly complex environments, involving multiple content servers, database types, web server types, search engines, and often interconnection to various third-party products. These variables add to the complexity and importance of content server reliability.

This section covers the following topics:

- Symptoms, Problems, and Solutions (page 2-1)
- General Problem-Solving Model (page 2-2)
- Preparing for Issues (page 2-3)

SYMPTOMS, PROBLEMS, AND SOLUTIONS

Problems with content management systems are characterized by specific symptoms. These symptoms may be general (for example, no content can be checked in), or more specific (for example, internal PowerPoint links not being converted to hyperlinks in PDF). You can trace symptoms to one or more problems or causes by using specific troubleshooting tools and techniques. After being identified, problems can be resolved by implementing a solution consisting of a series of actions.

This guide describes how to define symptoms, identify problems, and implement solutions in general environments. You should always apply the specific context in which you are troubleshooting to determine how to detect symptoms and diagnose problems for your specific environment.

GENERAL PROBLEM-SOLVING MODEL

When you are troubleshooting a content server, a systematic approach works best. An unsystematic approach to troubleshooting can result in wasting valuable time and resources, and can sometimes make symptoms even worse. Define the specific symptoms, identify all potential problems that could be causing the symptoms, and then systematically eliminate each potential problem (from the most likely to the least likely) until the symptoms disappear.

This process flow is not a rigid outline for troubleshooting a content management system. Rather, it is a foundation from which you can build a problem-solving process to suit your particular environment.

The following steps detail the problem-solving process:

1. Create a clear problem statement. You should define the problem in terms of a set of symptoms and potential causes.

To properly analyze the problem, identify the general symptoms and then determine what types of problems (causes) could result in these symptoms.

2. Gather the facts to isolate possible causes.

Ask questions to affected users, network administrators, managers, and other key people. Collect information from sources such as system, server, or applet error messages, logs, or software release notes.

3. Consider possible problems based on the facts that you gathered. Using the facts, you can eliminate some of the potential problems from your list.

Depending on the data, for example, you might be able to eliminate the operating system as a problem so that you can focus on web server problems. At every opportunity, try to narrow the number of potential problems so that you can create an efficient plan of action.

4. Create an action plan based on the remaining potential problems. Begin with the most likely problem, and devise a plan in which only one variable is manipulated.

Changing only one variable at a time enables you to reproduce a given solution to a specific problem. If you alter more than one variable simultaneously, you might solve the problem, but identifying the specific change that eliminated the symptom becomes far more difficult and will not help you solve the same problem if it occurs in the future.

- 5. Implement the action plan, performing each step carefully while testing to see whether the symptom disappears.
- 6. Whenever you change a variable, be sure to gather results. Generally, you should use the same method of gathering facts that you used in step 2 (that is, working with the key people affected, in conjunction with utilizing your diagnostic tools).
- 7. Analyze the results to determine whether the problem has been resolved. If it has, then the process is complete.
- 8. If the problem has not been resolved, you must create an action plan based on the next most likely problem on your list. Return to step 4, change one variable at a time, and repeat the process until the problem is solved.

Note: If you exhaust all the general causes and actions—either those identified in this guide or ones that you have identified for your environment—contact your Oracle technical support representative.

PREPARING FOR ISSUES

It is always easier to recover from an issue if you are prepared ahead of time. Possibly the most important requirement in any content server environment is to have current and accurate information about that content server instance available to the network support personnel at all times. Only with complete information can intelligent decisions be made about network change, and only with complete information can troubleshooting be done as quickly and efficiently as possible.

During the process of troubleshooting, content server is expected to exhibit abnormal behavior. Therefore, it is always a good practice to set up a development system for troubleshooting to minimize any business impact. Always document any changes being made so that it is easier to back out if troubleshooting has failed to identify the problem within the maintenance window.

To determine whether you are prepared for a content management problem, answer the following questions:

- Do you have an accurate metadata and security model of your content server?
- Does your organization or department have an up-to-date map that outlines the physical location of all content server instances on the network and how they are connected, as well as a logical map of network addresses, IP numbers, etc.?
- Do you know how content is being routed?
- For routed content, do you have correct, up-to-date archiving and provider configuration?
- Do you know all the points of contact with third-party products, including any connections to the Internet?
- For each external network connection, do you know what routing protocol is being used?
- Has your organization documented normal content server behavior and performance at different times of the day so that you can compare the current problems with a baseline reference?

If the answer to all of these questions is 'yes', you will be able to recover from problems more quickly and more easily than if you are not prepared. Finally, for every problem solved, document the problems with solutions provided. This way, you will create a problem/answer database that others in your organization can refer to in case similar problems occur later. This will invariably reduce the time to troubleshoot your content server and, consequently, minimize your business impact.

Chapter

FINDING ERROR AND STATUS INFORMATION

OVERVIEW

Effective troubleshooting relies on the availability of useful, detailed information. The Content Server products provide various sources of information that may be helpful in the troubleshooting process.

This section covers the following topics:

- ✤ Log Files (page 3-2)
- Configuration Information (page 3-7)
- System Audit Information (page 3-9)
- Web Server Filter Configuration (page 3-15)
- Tracing (page 3-18)
- Environment Packager (page 3-21)
- Content Server Analyzer (page 3-23)
- Using Content Server Analyzer (page 3-27)
- Configuration Entries (page 3-32)
- Stack Traces (page 3-33)

LOG FILES

Content Server stores status information and errors in log files. Log files are used to register system events, together with their date and time of occurrence. They can be valuable tools for troubleshooting, especially if verbose logging is turned on. Not only do logs indicate that specific events occurred, they also provide important clues about a chain of events that led to an error or problem.



Note: Verbose logging can quickly increase the size of a log file and possibly cause the content server to slow down. It is recommended that verbose logging is only used when troubleshooting a specific issue.

This section covers the following topics:

- Log File Characteristics (page 3-2)
- Accessing the Log Files (page 3-3)
- Using Content Server Logs (page 3-4)
- Using Archiver Logs (page 3-6)
- Inbound Refinery Logs (page 3-7)

Log File Characteristics

The log files associated with Content Server have the following characteristics:

- They are created only once each day at the time the first status, error, or fatal error occurs.
- ✤ No empty log files are generated.

Each log file contains the following columns:

- Type—Specifies the kind of incident that prompted the log entry: Information, Error, or Fatal.
- Time—Lists the date and time the log entry occurred.
- Description—Describes the incident that occurred.

The log files are standard HTML pages and are maintained for each content server instance. Logs are kept in revolving file name format for a maximum of 30 files. When the 31st file is created, the oldest one is deleted. Therefore, log file names in Content Server bear no relation to the date they were generated. To find a certain day in the log file, view

the index file in a browser and select that day's link. The file name is displayed in the browser's status bar (if it is enabled).



Tech Tip: Bookmark your log file pages. This will help you to troubleshoot problems, even if the content server is unavailable. Also, know where your configuration files are so you can find them if the content server is unavailable.

Accessing the Log Files

The log files of a content server are normally accessed from the Log Files folder under the Administration tray.



Note: You must be logged into the content server as an administrator to be able to view the log files.

Figure 3-1 Content Server access link location for log files



If, for whatever reason, you cannot view the log files from the Administration tray, you can also access them on the file system of the content server computer. The log files are located in the following locations:

Log Files	Found in:
Content Server	<install_dir>/<instance_dir>/weblayout/groups/secure/logs</instance_dir></install_dir>
Console Output Logs	<install_dir>/<instance_dir>/bin/<classname>.log</classname></instance_dir></install_dir>
Refinery	<install_dir>/<instance_dir>/weblayout/groups/secure/logs/refinery</instance_dir></install_dir>
Archiver	<install_dir>/<instance_dir>/weblayout/groups/secure/logs/archiver</instance_dir></install_dir>

Using Content Server Logs

The content server logs are listed by date and time. One file is generated for each day. Entries are added to the file throughout the day as events occur.

The following types of server log entries are generated:

- Info—Displays basic status information. For example, status information is logged if the server is ready and waiting.
- Error—Displays errors that occur but do not stop the software from functioning.
 For example, an error is logged if a user requests secure information that they are not allowed to access.
- Fatal—Displays errors that stop the software from functioning. For example, a fatal error is logged if the content server cannot access the database.

Туре	Time	Description
Info	4/15/04 3:13 AM	IdcAdmin: Starting the service 'IDC Content Admin Service'.
Info	4/15/04 3:14 AM	IdcAdmin: Admin Server Version 7-stable (040408) ready and waiting for connection on port 4440.
Info	4/15/04 3:14 AM	Starting the service 'Idc Content Service idcm1'.
Info	4/15/04 3:14 AM	Server version 7-stable (040408) ready and waiting for connection on port 4444.
Error	4/15/04 8:53 AM	Published schema directory could not be swapped into its proper location. Unable to create result set for query 'SELECT child dFormat, child dConversion, child dDescription FROM DocFormats child ORDER BY child dFormat ASC'. Network error- Connection reset by peer: socket write error

Figure 3-2Example of Content Server log file page

Opening a Content Server Log

To open a server log, complete the following steps:

- 1. Make sure that you are logged into Content Server as an administrator.
- 2. Click the **Content Server Logs** link, found on the Administration page or in the Administration tray's **Log Files** folder.

3. Select the link that corresponds to the date and the time of the log that you want to view.

3

Tech Tip: The error "*Error with path to collection. Directory 'verity.1' does not exist.*" might appear in the content server or Verity logs every few minutes or irregularly on instances where the Verity search engine add-on module has been installed. The system may otherwise seem to be functioning correctly. To solve this problem, locate and delete the file <*Install_Dir*>/*search/activeindex.hda*. The system will recreate it automatically and the error should go away.

Using Console Output Logs

When the Content Server is running as a Windows service, the console output logs are created automatically when the Content Server is launched and is properly configured [refer to Setting Up Console Output Capturing (page 3-5)]. In the event of a server crash, this feature enables the capture of output from the Java Virtual Machine (VM). This includes logging output from any enabled tracing facilities (such as script errors and Verity output by default) as well as stack dumps resulting from the VM crash (data sent to System.out and System.err).

Enabling this feature is useful to:

- ✤ Capture all data output by the VM
- Capture data specifically related to server crashes

Setting Up Console Output Capturing

To set up the configuration parameter to capture console output:

- 1. In a text editor, open the config.cfg file: <*install dir*>/config/config.cfg
- 2. Locate the #General Option Variables section and add the following line:

UseRedirectedOutput=true

- 3. Save and close the file.
- 4. Restart the Content Server.

Opening Console Output Logs

To open a server log:

1. Locate the following file on the computer where Content Server is installed:

<install dir>/bin/<classname>.log

2. Open the applicable classname log file in a standard text editor.

Using Archiver Logs

Archiver logs show information about imports, exports, and replications. The Archiver logs are listed by date and time. They are generated once a day when the first Archiver information status, fatal error, or error occurs.

Figure 3-3 Example of Archiver log file page

		Archiver Log File Created: 4/22/04 11:25 AM
Туре	Time	Description
Info	4/22/04 11:25 AM	Log organization created by application.
Info	4/22/04 11:25 AM	Event generated by user 'sysadmin' at host 'ref2'. Added archive 'all' to collection 'idem1'.
Info	4/22/04 11:25 AM	Exporting 'all' in 'idem1': Started.
Info	4/22/04 11:25 AM	Exporting 'all' in 'idcm1': Finished. Successfully exported 4 revisions.

The following types of archiver log entries are generated:

- Info—Displays basic status information. For example, status information is logged when an export and an import starts and finishes.
- Error—Displays user/administration errors that occur but do not stop the software from functioning. For example, an error is logged if there is no file information for a content item that you are trying to export.
- Fatal—Displays errors that stop the software from functioning. For example, a fatal error is logged if the content server cannot access the database. Check the connection string, user name, and password.

Opening an Archiver Log

To open an Archiver log, complete the following steps:

1. Make sure that you are logged into Content Server as an administrator.

- 2. Click the **Archiver Logs** link, found on the Administration page or in the Administration tray's **Log Files** folder.
- 3. Select the link that corresponds to the date and the time of the log.

A table showing the type, date and time, and description of each action is displayed. It also includes the name of the content server instance that created the archive.

Inbound Refinery Logs

With the release of Inbound Refinery version 10gR3, all Refinery logging is accessed through the Inbound Refinery interface. For more information see the *Refinery Products System Administration Guide*.

CONFIGURATION INFORMATION

Content Server provides a page which displays configuration information for a content server instance, which may be useful while troubleshooting a problem or working with Oracle's support organization. To access this page, click the **Configuration for** *[Instance]* link in the content server Administration tray.

The following configuration information is presented:

- Server information, such as name, description, and host filter.
- Installation directories, such as the locations of the core content server software, the native file repository ('Vault'), and the web-viewable file repository ('Web Layout').
- * Internet properties, such as the mail server and HTTP server names.
- Database properties, such as the JDBC driver name and the JDBC connection string.
- Version and build information, such as the Content Server version and Java version.
- License properties, such as the serial number and feature code.
- Server options, which lists the current value of a number of server-specific options.
- Content Security options, which specify what users can do with content.
- ★ Java properties, such as the JVM vendor and version.
- Server components, which lists all components that are currently enabled for the content server instance.



Note: Some of these options are specified during the software installation, while others are set using the System Properties utility.

-System Configuration		
Server Name: idcm_80 Version: 8.0-dev ajk ((Build:7.2.0.177)	D_current 061029T161757)	Server Configurations
Install Directory: C:/stellent/idcm_80_cu	rrent/	Hide Details
Key Name	V	alue
Install Directory	C:/stellent/idcm_80)_current/
Shared Library and Resources Directory	C:/stellent/idcm_80)_current/shared/
Server Data State Directory	C:/stellent/idcm_80)_current/data/
Weblayout Directory	C:/stellent/idcm_80)_current/weblayout/
Vault Directory	C:/stellent/idcm_80)_current/vault/
Database Type: Micros Database Version: 08. HTTP Server Address: Mail Server: mail.stello Search Engine Name: DATABASE.METADATA Index Engine Name: DATABASE.METADATA Active Index: verity.1	oft SQL Server 00.0760 BSILVERNOTE ent.com	Database Connection Details Internet Configurations
Features And Compo Number of Installed Feat Number of Enabled Comp Number of Disabled Com	nents ures: 8 oonents: O ponents: 2	Feature Details Disabled Component Details
- Options And Others - Auto Number Prefix: Use Accounts: FALSE NtlmSecuirty Enabled:	FALSE	Server Options
Allow get copy for user v FALSE Allow only original contri FALSE	vith read privilege: butor to check out:	Content Security Details
Java Version: 1.5.0_0	7	Java Properties

SYSTEM AUDIT INFORMATION

Content Server provides a page which displays system audit information for a content server instance, which may be useful while troubleshooting a problem or tweaking a server's performance. To access this page, click the **System Audit Information** link, found on the Administration page or in the Administration tray.



		Actions	Select an action	
General Inform	nation	Actions.	Select an action	
Server has not h	nad too many request	thread	5.	
Total JVM Memor Total JVM Availal	y: 64MB ble Memory: 29MB		Memory De	tail
Total Threads: 17	7		Thread De	tail
Total Active Data Total Audit Messa	base Connections: O ages: O		Database Conne De	tio: tail
Tracing Section	ns Information ⁽ⁱ⁾ —			
Verbose	🗹 Full Verbose Tracing	ı		
Active Sections	system, idocscript, requestaudit, r	esourceloa	de	
	💌 🗖 Sav	e		
	Update Reset			
Cache Informa	ition			
Permanently loa Temporary cach No temporary it	aded 196 pages and 1 le capped at 10 millio ems loaded.	3 resou n doubl	rce files. e-byte character	s.
Total O distinct so Total O rows stor	earch queries being exe red in cache, 0% full.	cuted	Search Cache De	tail
Total 17 items st 12,414 bytes use permitted. 0% use	ored in schema cache ed out of 10,485,760 sed.		Schema Cache De	tail

The System Audit Information page has three sections:

- System Audit General Information (page 3-10)
- System Audit Tracing Sections Information (page 3-11)
- System Audit Cache Information (page 3-12)

System Audit General Information

Total JVM Memory: 62MB Total JVM Available Memory: 28MB	Hide Details
Key Name	Value
Free memory	29 MB
Total memory	62 MB
Maximum memory	381 MB
Available processors	2
Free memory before garbage collection	27 MB
Total memory before garbage collection	62 MB
Finalization time	245439 ns.
Time for garbage collection	102 ms.
Free memory after garbage collection	29 MB
Total memory after garbage collection	62 MB
Total Threads: 16	Thread Details
Total Active Database Connections: O Total Audit Messages: O	Hide Detail:
Number of Read Actions: 1054 Number of Write Actions: 2 Waiting to get a connection: 0 Waiting to perform database action: 0	
No active connections	
No audit messages	

The General Information section of the System Audit Information page provides the following:

- Information regarding whether the system is receiving too many requests. If it is receiving too many requests, an e-mail is sent to the system administrator regarding load performance.
- Information about the memory cache for the system, and is useful in troubleshooting any "out of memory" errors you may receive. This is important when running the content server with many users and a large quantity of data.
- Information about which Java threads are currently running. This is useful in determining the cause of an error.
- Information about database activity.
- Listing of any audit messages.

System Audit Tracing Sections Information



Tracing in Content Server can be activated on a section-by-section basis. Tracing for active sections is displayed on the Server Output Page (page 3-13). Section tracing is useful for determining which section of the server is causing trouble, or when you want to view the details of specific sections. Sections can be added by appending extra sections to create a comma separated list. A listing of the sections available for tracing, with brief descriptions, is available by clicking the i (i) next to the *Tracing Sections Information* heading. The wildcard character * is supported so that *schema** will trace all sections that begin with the prefix *schema*.

Some tracing sections also support verbose output. Enable **Full Verbose Tracing** if you wish to see in-depth tracing for any active section that supports it. See Tracing (page 3-18) for more information.



Important: Any options set here will be lost when the content server is restarted unless you enable **Save** and click **Update**.

System Audit Cache Information

-Cache Informatio	n				
Permanently loade Temporary cache c No temporary item	d 196 page apped at 11 s loaded.	s and 13 re 0 million do	source file uble-byte	es. charact	ers.
Total O distinct sear Total O rows stored	ch queries bo in cache, 0%	eing execute 6 full.	d		Hide Details
Total 0 distinct search q Total 0 items stored in c Total 0 rows stored in ca Total 0 hits for 0 rows. Total 0 misses for 0 row Cache is searching acro There are currently 0 ac Maximum age of a cach Time between cache cle No most recent item exi No least recent item exi	ueries being e cache. ache, 0% full. ss. ss 0 providers tive searches e item: 240 m anup attempts sts. sts.	xecuted inutes. :: 120 seconds	;.		
Total 15 items store 11,286 bytes used o used.	d in schema out of 10,48!	cache 5,760 permit	tted. 0%		Hide Details
Total 15 items stored in Least recent item was u Most recent item was us 11,286 bytes used out o	cache. sed at 9/13/06 ed at 9/13/06 f 10,485,760 ;	5 3:47 PM. 3:47 PM. permitted. 0%	used.		
Buffer Cache Sumr	nary				Hide Details
IdcStringBuilder avera IdcStringBuilder capac	ige capacity: city capacity:	6,985 bytes 201,180 per	million		
Pool Name	Buffer Memory	Total Memory	Total Buffers	Reused Buffers	Outstanding Buffers
IdcStringBuilder	67584	67772	6	207	0
FileUtils	16640	16812	4	673	0
IdcCharArrayWriter	1654784	1655444	24	153	0
ParseOutput	2365440	2365772	26	3924	2

Content Server caches various items for quick access. The Cache Information section displays current information of three main caches:

- Searches—This pertains to the number of searches currently being run, how many executed searches are currently in cache, and when the cache is emptied. These details are useful when troubleshooting any search related issues.
- Schema—This lists details of any schema items currently in cache.
- Buffer—This displays information about Java objects in cache and how much memory each object is using, which is reflected in the memory information under the System Audit General Information (page 3-10) section. This information can be useful in pinpointing which object may be responsible for any memory leaks or other memory issues.

Server Output Page

Console output from the Content Server: Refresh Clear
requestaudit 09.14 16:00:00.494 Audit Request Monitor ****End Audit Report*****
requestaudit 09.14 16:00:00.494 Audit Request Monitor Request Audit
Report over the last 3600 Seconds****
requestaudit 09.14 16:00:00.494 Audit Request Monitor -Num Requests 10
Errors 0 Reqs/sec. 0.003 Avg. Latency (secs)0.056 Max Thread Count 1
requestaudit 09.14 16:00:00.494 Audit Request Monitor 1 Service
GET_SYSTEM_AUDIT_INFO Total Elapsed Time (secs) 0.297 Num requests 2 Num
errors 0 Avg. Latency (secs) 0.148
requestaudit 09.14 16:00:00.494 Audit Request Monitor 2 Service
GET_SEARCH_RESULTS Total Elapsed Time (secs) 0.172 Num requests 2 Num
errors 0 Avg. Latency (secs) 0.086
requestaudit 09.14 16:00:00.494 Audit Request Monitor 3 Service
GET_DOC_PAGE Total Elapsed Time (secs) 0.062 Num requests 1 Num errors 0
Avg. Latency (secs) 0.062
requestaudit 09.14 16:00:00.494 Audit Request Monitor 4 Service
GENERATE_GUIDS Total Elapsed Time (secs) 0.024 Num requests 5 Num errors
0 Avg. Latency (secs) 0.005
requestaudit 09.14 16:00:00.494 Audit Request Monitor ****End Audit
Report****
requestaudit 09.14 16:02:00.014 Audit Request Monitor Request Audit
Report over the last 120 Seconds****
requestaudit 09.14 16:02:00.014 Audit Request Monitor No requests
occured during this period
requestaudit 09.14 16:02:00.014 Audit Request Monitor ****End Audit
Report *****

The server output page displays the console output of Content Server. This is the same information that is located in the *<install_dir>/<instance_dir>/bin/<classname>.*log file. It includes information pertaining to all the sections selected for audit tracing in the System Audit Tracing Sections Information section (page 3-11) of the System Audit Information page (page 3-9). The Server Output page is accessed by clicking **View Server Output** on the System Audit Information page.

Localization Audit Page

s and the second	ystem Aud Stop au	lit Information
	Clear a	uditing
Key Name	Message	Stack Trace
en.UCF Test	missing	Show
en.Alpha Fields	missing	<pre>java.lang.Exception at intradoc.common.LocaleResources. getStringInternal(Unknown Source) at intradoc.common.LocaleResources. appendString(Unknown Source) at intradoc.common.LocaleResources. getString(Unknown Source) at intradoc.server.script.PageMerg- erScriptExtensions.evaluateFunction(Unknown Source) at intradoc.common.DynamicHt- mlMerger.computeFunction(Unknown Source) at intradoc.common.DynamicHt- mlMerger.evaluateGrammarElement(Unknown Source) at intradoc.common.DynamicHt- mlMerger.evaluateGrammarElement(Unknown Source) at intradoc.common.DynamicHt- mlMerger.substituteVariable(Unknown Source) at intradoc.common.DynamicHt- mlMerger.substituteVariable(Unknown Source) at intradoc.common.DynamicHtml. substituteVariable(Unknown Source) at intradoc.common.DynamicHtml. substituteVariable(Unknown Source) at intradoc.common.DynamicHtml. outputHmlEx(Unknown Source) at intradoc.common.DynamicHtml. outputHmlEx(Unknown Source) at intradoc.common.DynamicHtml. outputHtmlEx(Unknown Source) at intradoc.common.DynamicHtml. outputHtmlEx(Unknown Source) at intradoc.common.DynamicHtml. outputHtmlEx(Unknown Source) at intradoc.common.DynamicHtml. outputHtmlEx(Unknown Source) at intradoc.common.DynamicHtml. outputHtmlEx(Unknown Source) at intradoc.common.DynamicHtml.</pre>
		Hide
en.wwCpdTracingSection	missing	Show
en.bas2	missing	Show
en.bas1	missing	Show
en.wwCpdManageBaskets	missing	Show
en.Alpha and Beta Fields	missing	Show

The Localization Audit page is accessed by clicking **Localization Auditing** on the System Audit Information page (page 3-9). It displays information regarding the availability of localized variables for the Content Server user interface, and is useful in determining if any custom metadata field labels or other customized Content Server text requires localization. Clicking Show in the Stack Trace column displays the generated Java exceptions. Localization auditing is not persistent and must be started and stopped when using it.

WEB SERVER FILTER CONFIGURATION

Content Server provides a useful interface to configure and troubleshoot the web server filter. Configuring the web filter does two things:

- It generates a log file that provides information about all the activity on the web server associated with the content server's ISAPI filter.
- It maintains filter integrity and credential authentication between Windows, Active Directory Server, LDAP, or another external authentication protocol.

To access the information and activate the filter, web server logging must be set up. Do this using the **Filter Administration** link in the Administration tray of the content server instance. After the web server filter is configured, diagnostic information is generated and stored in a log file specific to the web server being used. Typically, this file provides the most valuable diagnostic information if there are problems accessing the content server using a browser.

General Options		
Cache Timeout	2	
This value specifies the number of minutes the web server filter will cache user data.		
Default Authentication		Basic -
This value specifies the default authentication method to apply to users who have never visited the Content Server before. The two valid choices are 'NTLM' and 'Basic'. 'NTLM' will use the Microsoft Login method to login users while 'Basic' will attempt to log users into the Content Server.		
		FALSE -
Disable GZIP Compression		1
Disable GZIP Compression By default, the content server compresses the HTML response pages for performance reasons. You may wish to disable it if CGI_RECEIVE_DUMP or CGI_DEBUG is enabled. Logging Options		1
Disable GZIP Compression By default, the content server compresses the HTML response pages for performance reasons. You may wish to disable it if CGI_RECEIVE_DUMP or CGI_DEBUG is enabled. Logging Options CGI_DEBUG		[
Disable GZIP Compression By default, the content server compresses the HTML response pages for performance reasons. You may wish to disable it if CGI_RECEIVE_DUMP or CGI_DEBUG is enabled. CGI_DEBUG Log summary of data and headers sent between the web server filter and the Content Server.		[
Disable GZIP Compression By default, the content server compresses the HTML response pages for performance reasons. You may wish to disable it if CGI_RECEIVE_DUMP or CGI_DEBUG is enabled. Logging Options CGI_DEBUG Log summary of data and headers sent between the web server filter and the Content Server. CGI_SEND_DUMP	_	[
Disable GZIP Compression By default, the content server compresses the HTML response pages for performance reasons. You may wish to disable it if CGI_RECEIVE_DUMP or CGI_DEBUG is enabled. CGI_DEBUG Log summary of data and headers sent between the web server filter and the Content Server. CGI_SEND_DUMP Log data and headers sent from the web server filter to the Content Server.		[
Disable GZIP Compression By default, the content server compresses the HTML response pages for performance reasons. You may wish to disable it if CGI_RECEIVE_DUMP or CGI_DEBUG is enabled. CGI_DEBUG Log summary of data and headers sent between the web server filter and the Content Server. CGI_SEND_DUMP Log data and headers sent from the web server filter to the Content Server. CGI_RECEIVE_DUMP		Г. 202 _ Г Г
Disable GZIP Compression By default, the content server compresses the HTML response pages for performance reasons. You may wish to disable it if CGI_RECEIVE_DUMP or CGI_DEBUG is enabled. CGI_DEBUG Log summary of data and headers sent between the web server filter and the Content Server. CGI_SEND_DUMP Log data and headers sent from the web server filter to the Content Server. CGI_RECEIVE_DUMP Log data and headers sent from the Content Server to the web server filter.		[
Disable GZIP Compression By default, the content server compresses the HTML response pages for performance reasons. You may wish to disable it if CGI_RECEIVE_DUMP or CGI_DEBUG is enabled. CGI_DEBUG Log summary of data and headers sent between the web server filter and the Content Server. CGI_SEND_DUMP Log data and headers sent from the web server filter to the Content Server. CGI_RECEIVE_DUMP Log data and headers sent from the Content Server to the web server filter. FILTER_DEBUG		
Disable GZIP Compression By default, the content server compresses the HTML response pages for performance reasons. You may wish to disable it if CGI_RECEIVE_DUMP or CGI_DEBUG is enabled. CGI_DEBUG Log summary of data and headers sent between the web server filter and the Content Server. CGI_SEND_DUMP Log data and headers sent from the web server filter to the Content Server. CGI_RECEIVE_DUMP Log data and headers sent from the Content Server to the web server filter. FILTER_DEBUG Log events inside the web server filter.		
Disable GZIP Compression By default, the content server compresses the HTML response pages for performance reasons. You may wish to disable it if CGI_RECEIVE_DUMP or CGI_DEBUG is enabled. CGI_DEBUG Log summary of data and headers sent between the web server filter and the Content Server. CGI_SEND_DUMP Log data and headers sent from the web server filter to the Content Server. CGI_RECEIVE_DUMP Log data and headers sent from the Content Server to the web server filter. FILTER_DEBUG Log events inside the web server filter. PLUGIN_DEBUG		

Figure 3-6 Web server filter configuration

The web server filter is configured in two sections:

- ✤ General Options
- ✤ Logging Options

General Options

This section allow you to configure and maintain filter integrity and includes the following options:

***** Cache Timeout:

When activated, this filter option changes the amount of time the web server holds user credentials.

Default Authentication:

This option specifies the default authentication method to apply to users who have never previously logged into the content server.

***** Disable GZIP Compression:

This option specifies whether the HTML response pages of the content server are compressed or not.

Logging Options

This section allows you to set the options for a log file that contains information about the communication between the web server filter and the content server. The configuration options include:

CGI_DEBUG:

This option helps to determine password and user authentication problems.

CGI_SEND_DUMP:

This option logs detailed communications from the web server filter to the Oracle content server.

CGI_RECEIVE_DUMP:

This option logs detailed communications from the Oracle content server to the web server filter.

FILTER_DEBUG:

This option creates information about user credentials and logs detailed information about events inside the filter.

*** PLUGIN_DEBUG:**

This option creates information about events that happen inside the plug-in filters.

For more detailed information about setting up web server logging, refer to Setting Up Web Server Logging (page 4-12). More detailed information about the web server filter applet and setting the web server filter configurations is included in the *Content Server Infrastructure, Integration, and Security Guide*.

TRACING

You can activate content server tracing to display detailed system information that may be very useful for troubleshooting and optimizing system performance. There are two options:

- ✤ Server-wide tracing
- ✤ Applet-specific tracing

Server-Wide Tracing

Server-wide tracing is used to view activities throughout the system.

There are two ways to activate server-wide tracing. To activate tracing from the Administration page, complete the following steps:

- 1. Make sure that you are logged into Content Server as an administrator.
- 2. Click on the System Audit Information link in the Administration tray.
- 3. Enable **Full Verbose Tracing** if you want to see in-depth tracing for any active section that supports it.
- 4. Specify the traces you wish to activate.
- 5. Click Update.
- 6. Click View Server Output.



Tech Tip: Tracing options are lost on system restart. To ensure your settings are retained after restarting the content server, enable Save before clicking Update.

To activate tracing from an applet, follow these steps:

- 1. Start an administrative applet.
- 2. Select **Options** and then **Tracing**.
- 3. Select Server tracing.
- 4. Select the tracings to activate or **all** and click **OK**.

The following tracing options are available:

 applet—This trace contains result sets from initialized applets, such as the Configuration Manager or User Admin.

- archiver—This trace provides information about archiving activities, including the reading and writing of archiver data files and the time the activities were initiated and finished.
- archiverlocks—This trace provides information about the locks put on files during archiving activities, including time initiated.
- chunkedrequest—This trace displays the messages and headers that are created when large requests are 'chunked' into smaller requests.
- docprofile—This trace displays the computation of content profiles, specifically the evaluation of the rules that determine which fields are labels, hidden, and so on.
- encoding—This trace provides information about encoding transformations that have occurred and the activities where encoding occurred.
- filelock—This trace displays information about short-term system locks put on directories (during activities like archiving, for example) with a focus on collisions that occur and timeouts.
- filelonglock—This trace displays information about the creation, removal, and maintenance of long term locks imposed by the system.
- **filequeue**—This trace displays information about accesses to a file queue.
- indexer—This trace displays information about index functions that occur when the database is updated, including the steps taken to update the index and the time elapsed for each step.
- indexermonitor—This trace provides a brief summary of automatic index activities, including time started and ended.
- indexerprocess—This trace displays information about a manually launched index process and indicates if the process terminated properly.
- ◆ localization—This trace displays information about localization usage and activities.
- * mail—This trace provides a description of mail sent by the content server.
- pagecreation—This trace displays information about the creation of displayed pages, including the server thread and the time taken to generate the page.
- requestaudit—This trace provides summary reports on service requests, including the elapsed time for the requests and the number of requests made.
- scheduledevents—This trace provides a list of hourly or daily background scheduled events.

- schema—This trace provides information about schema publishing (tables and views published as .js files) and caching (tables cached into content server memory).
- searchquery—This trace displays information about recent searches, including the fields used to search on and the order of sorting for results.
- socketrequests—This trace displays the date, time, and thread number of socket requests as well as the actions during the request.
- system—This trace displays internal system messages, such as system socket requests and responses.
- systemdatabase—This trace provides information about database activities, including queries executed, index updates, threads used, and time initiated.
- transfermonitor—This trace displays information about the archiver and the batch file transfer activities.
- userstorage—This trace describes the access of external user repositories, including what actions were taken during access.
- workflow—This trace displays a list of metadata on content items going through workflow, including document title and revision number.

Note: To facilitate international support, most tracing messages are in English and do not have translations.

Applet-Specific Tracing

For applet-specific tracing, the output goes to the browser Java console. To perform tracing by applet, complete the following steps:

- 1. Start the administration applet that you want to trace.
- 2. Select **Options** and then **Tracing**.
- 3. Make your selections, and click **OK**. The output is directed to the browser Java console.
| C Applet Tracing C Server Tracing
Is Full Verbose Tracing
all
all
applet
archiver
carchiverlocks
chunkedrequest
filelock
filelock |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Is Full Verbose Tracing all applet applet archiver chunkedrequest filelock filelock filelonglock |
| □ all ☑ applet ☑ archiver ☑ archiverlocks □ chunkedrequest ☑ filelock □ filelonglock |
| I applet
I archiver
I archiverlocks
I chunkedrequest
I filelock
I filelock
I filelock |
| I▼ archiver
I▼ archiverlocks
I□ chunkedrequest
I▼ filelock
I□ filelonglock |
| I⊽ archiverlocks
I⊂ chunkedrequest
I⊽ filelock
I⊂ filelonglock |
| ☐ chunkedrequest
☞ filelock
☐ filelonglock |
| I⊽ filelock
I⊓ filelonglock |
| Filelonglock |
| |
| 🔽 filequeue |
| indexermonitor — |
| Iocalization |
| 🔽 mail |
| pagecreation |
| scheduledevents |
| 🔽 schema 💆 |
| Add Section |
| OK Cancel |

D* 37 Applate ogific tr

ENVIRONMENT PACKAGER

The Environment Packager is a diagnostic tool for you and Stellent Support. It creates a zip file of the desired state directories, log files, and other component and resource directories.

To create an environment zip file, complete the following steps:

- 1. Make sure that you are logged into Content Server as an administrator.
- 2. From the Administration tray, click on the Environment Packager link. The Environment Packager page is displayed.



Environment Pac	kager
Check the boxes be packaged. This will a and your operating	low to select which parts of the environment to be also gather information about your Java environment system.
	Z Unix 'etc' Log Directory
1	Server Data State Directory
E.	Z Search Engine State Directory
E.	Content Server Logs
E.	Z Archiver Logs
E.	Inbound Refinery Logs
R	Z Verity Logs
_	Custom Components Directory
	🗆 All Files 🛛 🗹 Digests Only
_	Classes Directory
	🗆 All Files 🛛 🗹 Digests Only
	Weblayout Common Directory for Applets
	🗆 All Files 🛛 🗹 Digests Only
_	Shared Library and Resources Directory Image: All Files Image: Digests Only
_	Start Packaging Reset

- 3. Select which parts of the environment should be packaged.
- 4. When you are ready to create the environment zip file, click on Start Packaging.

A message is displayed while the zip file is being built, with a link to the zip file. The packaging process may take several minutes. The zip file link will not be available until the process has finished.



Note: The packaged zip is named server_environment_*.zip. While content server builds the packaged zip file, it will be located in *<install_dir>*/vault/~temp. Once the build of the zip file is complete, it is moved to *<install_dir>*/weblayout/groups/secure/logs/env.

CONTENT SERVER ANALYZER

Content Server Analyzer is a utility that enables you to confirm the integrity of the content server repository components, including the file system, database, and search index. It can also assist system administrators in repairing some problems that are detected in the repository components.

The Content Server Analyzer utility enables system administrators to do any of the following:

- Confirm the accuracy of synchronization between three important content server database tables (Revisions, Documents, and DocMeta).
- Confirm that the dRevClassID and dDocName fields are consistent across all revisions of content items.
- Determine if the file system (native and web-viewable file repositories) contains any duplicate or missing files.
- Ensure the accuracy of synchronization between the search index and the file system.
- Ensure the accuracy of synchronization between the search index and the Revisions database table.
- ✤ Ensure that the file system contains all necessary files.
- Remove duplicate files from the content server repository either permanently or provisionally by moving them into the logs/directory.
- Produce a general report on the state of content items in the content server.

The method to start the Content Server Analyzer depends on the operating system:

- Windows: Choose Start—Programs—Stellent Content Server—<*Instance_Name*>— Utilities—Content Server Analyzer.
- UNIX: Change to the <*Install_Dir*>/bin directory and run the Content Server Analyzer program.

🐉 Content Server Analyzer	r for idcm_80_current 📃 🗆 🖾	
Options Help		
Content Server Analyzer		
Check database Check RevClassIDs Clean database Check search index Clean search index	Analyze the content server database for errors. The analysis checks that the data is consistent between the DocMeta, Revisions and Documents tables. As an optional check, the dRevClassID and dDocName is checked to make sure they are consistent across revisions of content items. Analyze the content server search index. Checks for double indexed documents and any documents which may be missing from the search collection.	
Check file system	Analyze the weblayout and vault directories to make sure all the necessary files are present. The file names any renditions (thumbnails, etc.) are calculated from the document's info in	
Generate report	Generates a report on the state of the content items in the content server.	
Range Analyzer log dir	(Format: 'dID:dID' or '(date):(date)')	
Ready	Start Analysis Cancel	

Figure 3-9 Content Server Analyzer

Content Server Analyzer: Configuration Tab

The Configuration tab of the Content Server Analyzer Configuration table is used to configure analysis options and specify a customized logging directory structure:

Option	Description
Check database	Performs all checks on the database, ensures the integrity of the database columns, and confirms the consistency of data between the DocMeta, Revision, and Documents tables.

Option	Description
Check RevClassIDs	Ensures the accurate data synchronization between the dRevClassID and dDocName tables.
Clean database	Removes inconsistent rows from the database. Extra entries in the DocMeta table are deleted, inadequately defined entries in the Document table are deleted, and entries without a corresponding reference in the Revisions table are deleted.
Check search index	Analyzes the search index to ensure its integrity and checks for duplicate data records for indexed documents and any documents that might be missing from the search collection.
Clean search index	Re-indexes the search index and replaces missing data records of any omitted documents.
Check file system	Analyzes the file system (weblayout and vault file repositories) to ensure all necessary files are present.
Delete	Permanently deletes extra files that were found during the file system analysis.
Safe delete	Creates a safe delete directory in the logs/ directory and copies the extra files that were found during the file system analysis into this directory.
Check for extra files	Identifies any possible extra files that might be in the file system.
Generate report	Uses the console window to report statistics about the content items in the repository. It includes information pertaining to the status, release and processing states of content items in the file system and provides prior and current totals. Progress and error messages are also logged to the console window.
Range	Specifies the first and last of the criteria analyze.
Log Directory	The default directory used by Content Server Analyzer is <i><install_dir>/bin/logs/</install_dir></i> . Optionally, you can also enter a custom directory name. If the Safe delete option is selected, the files are moved to this directory.

Content Server Analyzer: Progress Tab

The Progress tab of the Content Server Analyzer displays the progress of the analysis processes and all generated information. To access this tab, click the tab on the Content Server Analyzer Application, or click **Start Analysis** on the Content Server Analyzer Application.

Feature	Description
Task progress bar	Displays the combined progress of the specific analysis tasks selected on the Configuration tab.
Overall progress bar	Displays the overall progress of the analysis process.
Console area	Area that displays the information collected and summarized during the analysis processes. Displays applicable information for each selected option. Also displays progress and error messages generated during the analysis processes.

USING CONTENT SERVER ANALYZER

This section describes the following Content Server Analyzer tasks:

- ✤ Accessing the Content Server Analyzer (page 3-27)
- Specifying a Custom Analyzer Log Directory (page 3-27)
- Invoking the Analysis Process (page 3-28)
- Analyzing the Content Server Database (page 3-28)
- Analyzing the Content Server Search Index (page 3-29)
- Analyzing the Content Server File System (page 3-30)
- Viewing the Analysis Progress and Results (page 3-30)
- Generating a Status Report (page 3-31)
- Cancelling the Status Report (page 3-31)

Accessing the Content Server Analyzer

To display the Content Server Analyzer, use one of the following methods:

- Windows: Select Start—Programs—Stellent Content Server— Master on server—Utilities—Content Server Analyzer.
- Unix: Change to the <install_dir>/bin directory and run the Content Server Analyzer program.

The Content Server Analyzer Application is displayed.

Specifying a Custom Analyzer Log Directory

The logs/ directory is the default logging directory for the Content Server Analyzer. Analysis output files are written to this directory and extra files detected during a file system analysis process can be transferred here as well. Optionally, the default logs/ directory name and path can be changed as desired.

To customize the Analyzer log directory name and path:

- 1. On the Content Server Analyzer: Configuration Tab (page 3-24), place the cursor in the Analyzer log dir field.
- 2. Enter the desired directory path.

During the next analysis process, the Content Server Analyzer automatically creates the specified directory or directories in the *<install_dir>/*bin/ directory hierarchy.

Invoking the Analysis Process

To invoke the analysis process:

- 1. On the Content Server Analyzer: Configuration Tab (page 3-24), select and activate the desired options (checking the corresponding check boxes).
- 2. Click Start Analysis.



Note: If this is the very first time the Content Server Analyzer has been run, the output files in the logs/ directory are automatically created. On subsequent analysis processes, a confirmation message is displayed asking to overwrite the existing log file.

3. Click Yes to overwrite the existing log file.

The Content Server Analyzer: Progress Tab (page 3-26) is displayed automatically.



Note: If you click **No**, the analysis process is terminated and you are prompted to manually remove files from the logs/ directory before running the Content Server Analyzer again.

A completion message is displayed when all of the selected analysis processes are finalized.

4. Click OK.

The results are displayed in the console area on the Progress tab.

Analyzing the Content Server Database

These options are used to check the integrity of the database columns. The available options enable users to examine the three tables that are used to store content item revision information (DocMeta, Documents, and Revisions). The DocMeta file is examined for extra entries that are not found in the Revisions table. Similarly, the Documents table is examined to verify that there are sufficient entries to correspond to the entries in the Revisions table.





Note: The *Check RevClassIDs* and *Clean database* options are activated and selectable only when the *Check database* option is selected.

To analyze the Content Server database:

- 1. On the Content Server Analyzer: Configuration Tab (page 3-24), select the applicable options.
- 2. Click Start Analysis.

The results are displayed in the console area on the Content Server Analyzer: Progress Tab (page 3-26). See Invoking the Analysis Process (page 3-28) for information about the analysis procedure.

Analyzing the Content Server Search Index

These options are used to check the entries in the Revisions table to ensure that all of the documents that belong in the index are properly listed. Additionally, a check can be performed to ensure that there are no duplicate entries in the search index.

```
Check search index
```

🗌 🖂 csiDCAnalyzeCleanindex



Note: The *csIDCAnalyzeCleanIndex* option is activated and selectable only when the *Check search index* option is selected.

To analyze the content server search index:

- 1. On the Content Server Analyzer: Configuration Tab (page 3-24), select the applicable options.
- 2. Click the **Start Analysis** button (refer to Invoking the Analysis Process (page 3-28) for information about the analysis procedure).

The results are displayed in the console area on the Content Server Analyzer: Progress Tab (page 3-26).

Analyzing the Content Server File System

These options check the integrity of the file system (weblayout and vault file repositories). Using the information in the database, these options ensure that every file in the Revisions table contains accurate entries corresponding to the items in the file system. A check can also be completed to locate any extra files in the vault and weblayout file repositories.

Check file system	
🗖 Delete	
🗌 Safe delete	
Check for extra files	S



Note: The *Delete*, *Safe delete*, and *Check for extra files* options are activated and selectable only when the *Check file system* option is selected.

To analyze the content server file system (vault and weblayout file repositories):

- 1. On the Content Server Analyzer: Configuration Tab (page 3-24), select the applicable options.
- 2. Click Start Analysis.

The results are displayed in the console area on the Content Server Analyzer: Progress Tab (page 3-26). See Invoking the Analysis Process (page 3-28) for information about the analysis procedure.

Viewing the Analysis Progress and Results

The Content Server Analyzer: Progress Tab (page 3-26) is displayed automatically when the **Start Analysis** button is clicked. The progress bars show when the Content Server Analyzer has completed processing the selected analysis options. The following image shows a partially finished analysis:

When the analysis process is complete, the results are displayed in the console area of the Progress tab. The results depend on what analysis options were selected. The following image of the console area shows the results from selecting database, search index, and file system options:



Note: The Generate report option was not selected for this example (refer to Generating a Status Report (page 3-31) for an example of the generated status report).





Generating a Status Report

The status report generated by the Content Server Analyzer provides statistics about the content items in the repository. The status report output is displayed in the console area of the Progress tab.

To generate a status report:

- 1. On the Content Server Analyzer: Configuration Tab (page 3-24), select the Generate report check box.
- 2. Click Start Analysis.

When the analysis process is complete, the status report information is displayed immediately following the standard analysis results in the console area of the Content Server Analyzer: Progress Tab (page 3-26).

Cancelling the Status Report

The report generation feature can be suppressed after the analysis process has already started. To cancel the content item status report during the analysis process:

 During the analysis process, click Cancel on the Content Server Analyzer Application.

You are prompted about cancelling after the current task is finished.

2. Click Yes to suppress the status report.

The status report is not included with the analysis results that are displayed in the console area of the Progress tab.

CONFIGURATION ENTRIES

The content server also provides debugging configuration variables that, when appropriately set, contribute applicable diagnostic information. Generally, these configuration variables are manually entered and set in the content server's configuration file (*<Install_Dir>/config/config.cfg*), although some are set using system utilities or administration applets (for example, the web server filter configuration settings).

Some of the more standard configuration variables include:

DebugMode

This variable defines whether the content server should run in debug mode. It is set during installation and when the content server is updated.

DebugStdConversion

This variable defines whether the standard conversion process should be debugged. It is set during installation and when the content server is updated.

SearchDebugLevel

This variable defines how detailed the indexing messages are in the indexer log files. The debug level selected determines the degree of detailed information to be collected from each file accessed or each activity performed. This setting can be updated from Repository Manager.

✤ ScriptDebugTrace

This variable enables a trace of all specified resource includes and displays the results of the trace on the page.

setResourceInclude

This variable enables a resource include to be assigned dynamically to a constructed script.

ScriptErrorTrace

This variable enables a trace of script errors. If used as a parameter to a service call, script error information can be added to the bottom of the displayed page.

✤ UseRedirectedOutput

In the event of a server crash, this variable enables the capture of output from the Java

Virtual Machine (VM). This includes logging output from any enabled tracing facilities as well as stack dumps resulting from the VM crash.

 \bigcirc

Note: For further details refer to the Oracle Idoc Script Reference Guide.

STACK TRACES

The stack trace allows you to see what threads are currently running in the Content Server. It is a useful troubleshooting tool that provides information about the threads and allows you to monitor the Content Server's processing.

To initiate the current stack trace for the Content Server:

UNIX

- 1. Navigate to the *<install_dir>/*etc directory.
- 2. Open the *pid* file to obtain the process id.
- 3. Run the following command:

kill -3 <pid>

Windows

- 1. Open a command console window.
- 2. Launch the Content Server application:

<install_dir>\bin\idcserver.exe

3. Press Ctrl+Break in the same command console used to start Content Server.



Note: You must launch Content Server from a command console window to use this Windows troubleshooting method. If Content Server is running as a service, you will need to use third-party tools to initiate a stack trace.

Finding Error and Status Information

C h a p t e r

WEB SERVER ISSUES

OVERVIEW

This section provides solutions to several common web server issues. Please attempt the solutions recommended in this guide before contacting Support.

- Permission Issues (page 4-1)
- Content Server or Web Server Won't Start (page 4-4)
- Port Configuration Issues (page 4-8)
- ✤ iPlanet-Specific Issues (page 4-10)
- Miscellaneous Web Server Issues (page 4-12)

PERMISSION ISSUES

This section covers the following topics:

- Login Errors After Manual Reinstallation (page 4-2)
- Difficulties Accessing a Specific Document (page 4-3)

Login Errors After Manual Reinstallation

Symptom

It was necessary to manually configure the web server. Afterwards, however, errors are encountered when users try to log into the content server.

Problem

It is difficult to log into the content server and when searching for specific documents, an error message is displayed stating that the document does not exist.

Recommendation

The following recommended solutions are examples using the IIS web server. Problems logging into the content server could result from improper administrative settings for IIS. To check IIS settings, complete the following steps:

- 1. Open the IIS web server and double-click the server name.
- 2. Expand the Default Web Site directory.
- 3. Right-click the idcplg directory and select Properties.
- 4. Ensure the following settings are enabled on the Virtual Directory tab:
 - Read
 - Log visits
 - Index this resource
 - Execute Permissions: Scripts and Executables
- 5. Click OK.
- 6. Right-click the Administration Web Site directory and select Properties.
- 7. Ensure the following settings are enabled on the Home Directory tab:
 - Read
 - Log visits
 - Index this resource
 - Execute Permissions: Scripts only
- 8. Click OK.
- 9. Restart the IIS service and the content server.

Difficulties Accessing a Specific Document

If the virtual directory name and the HTTP relative web root (HttpRelativeWebRoot value) are not identical, the content server will not be able to locate a requested document. Instead, an error message is displayed stating that no such document exists in the specified directory. To check these values, complete the following steps:

- 1. Determine the name of the virtual directory (the installation directory for the content server).
- 2. Open the following file in a text editor:

<instance_dir>/config/config.cfg

- 3. Locate the **#Internet Variables** section and find the **HttpRelativeWebRoot** configuration setting.
- 4. Compare the name of the virtual directory with the HttpRelativeWebRoot value. If they are not identical, modify the HttpRelativeWebRoot value.

For example, if the HttpRelativeWebRoot value is 'Stellent' and the virtual directory name is 'stellent,' the HttpRelativeWebRoot value must be modified to match the virtual directory name.

- 5. Save the changes and close the config.cfg file.
- 6. Restart the content server.

FILE NOT FOUND ERROR ON DOWNLOAD

Windows 2003 comes with IIS 6.0. Citing security reasons, IIS 6.0 will not serve a file if its extension is not registered as a MIME type for the web site. MIME (Multipurpose Internet Mail Extensions) is a standard used to communicate the type of data being delivered. When an extension is not registered, IIS returns a 404 (File Not Found) status code to the client browser. In IIS logs, it will record the sc-status as a 404 and the sc-substatus as 3 when this happens.

To add a MIME registration, use the steps below.

- 1. Open the IIS 6.0 Manager from Administrative Tools.
- 2. Open the Properties dialog for the server level object. The location and text for this object can be different depending on how the tool was started. In any view, it is the level above the "Web Sites" folder.

3. Click the MIME types button. On the MIME type window, you can view, add, edit, remove file extension registrations for the entire server. The MIME type to specify depends on the application. Typical settings for some common missing registrations are as follows:

Extension	МІМЕ Туре
	application/octet-stream
.log	text/plain
.dwg	drawing/x-dwg



You can identify MIME types by checking the list of registered MIME types at http://www.iana.org/assignments/media-types/index.html. Other sites that list MIME types include http://filext.com/ and http://www.webmaster-toolkit.com/mime-types.shtml.



Note: The ability to add the extension of "." (dot) to serve files without an extension was added in a recent update of Windows 2003. Using the no extension registration is preferred over the wildcard extension (*). The wildcard extension would serve all unknown extensions, and thus remove the security benefits of this IIS 6 feature.

- 4. Restart IIS.
- 5. Test the configuration. When the MIME type has been configured, you will be able to view or download the files.



Note: If you prefer to change the configuration for a specific web site, use the MIME types dialog on the HTTP Headers tab of the properties dialog for that site. At site level, the server wide registrations do not display, making this harder to use.

CONTENT SERVER OR WEB SERVER WON'T START

This section covers the following topics:

- Web Server Runs But Content Server Won't Start (page 4-5)
- Cannot Start the Website or Content Server (page 4-6)

Web Server Runs But Content Server Won't Start

Symptom

The web server is running, but the content server won't start.

Problem

Attempting to start the content server from a command window results in the following error:

Could not listen on port 80. Address in use.

Recommendation

This error message indicates that another service is already using port 80. You need to determine which ports the web server and the content serve are configured to use. The following recommended solution is an example using the IIS web server.

To determine the port IIS is configured to run on, complete the following steps:

- 1. Open the IIS web server and double-click the server name.
- 2. Right-click on the website and select Properties.

The TCP Port setting is listed on the Web Site tab.

3. Click OK.

If IIS has been configured to run on port 80, then another service is also trying to run on it as well. This could be the reason that the content server will not start. There may be a port conflict between the IIS and the content server if both have been configured to run on port 80. By default, the content server should be configured to run on port 4444 and IIS should be configured to run on port 80.

To determine what port the content server is configured to run on, complete the following steps:

1. Open the following file:

<*CS_Instance_Dir>/install/<date>-<install_type>.txt* for example, 2-5-2004-new.txt.

 Locate the following configuration setting: <Instance_Name>/IntradocServerPort=<Port_Number> The IntradocServerPort configuration setting defines the port that the web server filter or any other application must use to talk to the content server. If both the content server and IIS are configured to use the same port, there will be a sharing violation.

In this situation, you can try to edit the configuration of the content server using Admin Server. To do this, you will need to stop IIS so that the content server will start.

To change the port setting of the content server, complete the following steps:

- 1. Log into the content server as an administrator.
- 2. Go to the Administration page and click the Admin Server link.
- 3. Click the Edit Server link.
- 4. Select the content server to be edited from the list.
- 5. Click Edit.
- 6. Enter the correct server port number.
- 7. Click Finish.
- 8. Restart the content server.

Although the content server configuration settings can be edited using Admin Server, typically changing any settings other than the description or allowed actions is unusual. It might actually be easier and more reliable to simply uninstall and then reinstall the content server to ensure that it is configured to use the correct port.

Cannot Start the Website or Content Server

Symptom

There are problems with starting the default website and the content server.

Problem

When the web server is stopped, the master content server instance starts without problems. However, when the web server is running and an attempt is made to start the content server from the services window, the following error message is issued:

Could not start the IDC Content Service <master_instance> service on Local Computer. Error 1064: An exception occurred in the service when handling the control request. Also, if an attempt is made to start the default website when the web server is stopped and the master content server instance is running, the following message is issued:

Address already in use.

Recommendation

Error 1064 is usually issued when the content server cannot connect to the database. Generally, the connectivity problem happens when the database and the content server reside on the same machine. The content server tries to connect to the database before the database has completed its initialization process. Additional information about problems associated with the content server database, refer to Chapter 5 (*Database Issues*).

Error 1064 can also be issued if the user name and password are incorrect. On the Windows platform, this can be tested through the ODBC connection.

The selected web server used with the content server is handling all the web services that are running on the machine. For this reason, if the web server is not started, the default website cannot be accessed.

To make sure that there are no other connectivity problems when the web server is running, try to start the content server using a command line:

- 1. Open a command line window.
- 2. Give the content server start command.

Microsoft Windows:

- a. Go to *<CS* Instance Dir*>/bin* directory.
- Run the *IdcServer.exe* application: IdcServer.exe -start

UNIX:

- a. Go to *<CS* Instance Dir*>/bin* directory.
- Run the *IdcServer.exe* application: idcserver_start
- 3. This should provide a descriptive error message.

Admin Server Page will not Open

Symptom

Visiting the Admin Server page gives the following error:

Unable to retrieve root IdcAdmin page. Unable to execute service method 'loadAllIdcServerData'.

A Java Stack Dump follows this error.

Problem

This problem usually occurs because a content server system configuration file has become corrupt.

Recommendation

- 1. Navigate to the *<install_dir>/<instance_dir>/admin/data/servers* directory.
- 2. Select the servers.hda file, back it up, then delete it.
- 3. Still in the *<install_dir>/<instance_dir>/*admin/data/servers directory, select the directory named after the instance, back it up, and then delete it.
- 4. Restart the content server.
- 5. Go to the Admin page and manually re-add the content server.

PORT CONFIGURATION ISSUES

This section covers the following topics:

Changing the Port Setting for IIS (page 4-8)

Changing the Port Setting for IIS

Question

How do I change the port for my web server from 80 (the default) to another number?

Recommendation

It is possible to change the port setting for the web server from the default value. After determining the current web server port setting, an alternate port for web server to run on must be defined and the HTTP server address of the content server must be revised to use the new port setting.

The following recommended solution is an example using the IIS web server.

Defining an alternate port setting for IIS

- 1. If applicable, open the IIS web server and double-click the server name.
- 2. Right-click on the website and select **Properties**.
- 3. Select the **Web Site** tab, locate the **TCP Port** field and change the default value (80) to the desired setting.
- 4. Click OK.

Revising the HTTP server address of the content server

- 1. Start the System Properties utility by choosing Start—Programs—Content Server— *<Instance Name>*—Utilities—System Properties.
- 2. Select the Internet tab.
- 3. Modify the HTTP Server Address field as follows:
 - a. Add a colon (:) after the server name.
 - b. Enter the new port number after the colon.
 For example, if the server name is 'docserver' and the new port setting is 82, then the HTTP Server Address field entry is 'docserver:82'.
 - c. Click OK.
- d. Using the Systems Properties utility automatically updates most of the HTTP server address settings when the content server is restarted. However, the *server.hda* file must be manually updated.
 - a. Open the following file in a text editor: <CS_Inst_Dir>/admin/data/servers/<server_name>/server.hda
 - b. Locate the following line: HttpServerAddress=<server_name>
 - c. Add a colon (:) after the server name.
 - d. Enter the new port number after the colon.

- 4. If there are any outgoing providers or preview providers configured on this computer, ensure that the HTTP server address specified in them is also correct. To check the providers currently configured, complete the following steps:
 - a. Log into the Content Server as an administrator.
 - b. Go to the Administration page and click on Providers link.
 - c. Click the Info link to open the Provider Information page.
 - d. Locate the HTTP Server Address entry and check the setting.
 - e. If necessary, click on the **Edit** button to update the HTTP server address; enter a colon (:) after the server name followed by the new port number.
- 5. Restart the content server.
- 6. After restarting the content server, it may be necessary to rebuild your portal page to reflect the changes. To edit the portal page, use the Web Layout Editor (Options—Update Portal).

IPLANET-SPECIFIC ISSUES

This section covers the following topics:

iPlanet Freezes and Content Server is Unavailable (page 4-10)

iPlanet Freezes and Content Server is Unavailable

Symptom

After submitting a query to the content server, iPlanet occasionally locks up. Although the world icon is turning, either no results are returned from the query or the results cannot be viewed. After closing the browser and reconnecting to cmsprod_source, the login box does not display.

Problem

When iPlanet has locked tight, multiple warning messages are recorded in the iPlanet logs that indicate the host is attempting to find specific class files, but is unsuccessful. After the iPlanet process is terminated to delete the tasks and restarted, it works fine again.

Occasionally, when the system crashes, neither the content server nor the admin server are accessible. Trying to connect to the content server URL produces the following error:

Content Server is unavailable, click here for administration server.

Clicking the link to the Admin Server does not work, but iPlanet is still available. Although the content server and admin server are unavailable, it is still possible to connect to iPlanet.

Recommendation

If this problem occurs only occasionally, rather than shutting iPlanet down and restarting it, try starting the content server and Admin Server through the Start menu commands. However, if the occurrences are significantly more frequent (for example, two or three times in one or two hours), it is probably necessary to increase the hard file descriptor limit. It is possible that a process is not closing the file descriptors properly, which, in turn, would cause iPlanet to hang.

Each time iPlanet starts up, it issues the following warning message:

[<date:time>] warning: The server configuration may require more file descriptors than the operating system provides. If you encounter PR_PROC_DESC_TABLE_FULL_ERROR errors, you may wish to increase the operating system hard file descriptor limit from 1024 to 5763 (see your operating system documentation) or decrease one or more of the following settings: MaxFiles in nsfc.conf and ConnQueueSize, RqThrottle, and MaxKeepAliveConnections in magnus.conf.

If it is issued, the PR_PROC_DESC_TABLE_FULL_ERROR error message is logged into the iPlanet logs. Depending on the particular operating system currently in use, the default value for the file descriptor limit varies. There are two types of file descriptor limits: a hard (system-wide) limit and a soft limit. The hard limit can only be changed at the root but the soft limit can be changed by any user. The only restriction is that the soft limit must never be set higher than the corresponding hard limit.

There are several ways to increase the default value for the file descriptor limit:

The ulimit(1) command

Use the ulimit(1) command from your shell to increase the soft file descriptor limit from its default value to a specified value less than the hard file descriptor limit.

The setrlimit(2) call

Use the setrlimit(2) call to change both the soft and hard file descriptor limits. You will need root access privileges to be able to change the hard file descriptor limit.

✤ The /etc/system file

Modify the */etc/system* file and include the following line to increase the hard file descriptor limit (for example, to 128 — 0x80 in hexadecimal): set rlim fd cur=0x80

After changing the */etc/system* file, reboot the system to ensure the changes are activated.

If you chose to set the hard file descriptor limit to 5763 (0x1683 in hexadecimal) as stipulated in the iPlanet startup warning message (included above), modify the */etc/system* file and include the following line:

set rlim_fd_cur=0x1683

MISCELLANEOUS WEB SERVER ISSUES

This section covers the following topics:

- Setting Up Web Server Logging (page 4-12)
- Portal Page Buttons/Links not Working (page 4-13)
- Turning Off Logging When Running Content Server As Windows Service (page 4-14)

Setting Up Web Server Logging

Question

What content server log files are available that provide information concerning problems connecting to the content server through a browser?

Recommendation

The content server log files provide basic diagnostic information for the content server and each log file contains information for one day. These log files are named IdcLognn (where nn is a sequential, two-digit identification number) and are written to the 'logs' directory (<<u>CS_Instance_Dir</u>>/weblayout/groups/secure/logs).

However, typically if you are having problems accessing the content server using the browser (and it seems to be a universal problem), the most valuable diagnostic information can be acquired from the web server. To access this information, you will need to set up web server logging for the content server:

1. Open the following file in a text editor:

<CS_Instance_Dir>/config/config.cfg

2. Locate the **#General Option Variables** section in this file and add the following lines:

CGI_DEBUG=true CGI_RECEIVE_DUMP=true CGI_SEND_DUMP=true FILTER_DEBUG=true

These configuration settings add a log file named *idc_cgi_isapi-<server_name>.dll* located in the *<Install_Dir>/idcplg* directory. This file provides information about all of the activity on the web server associated with the content server's ISAPI filter.

- 3. Save the changes and close the *config.cfg* file.
- 4. Restart the content server.

Portal Page Buttons/Links not Working

Symptom

After completing the content server installation process, the initial portal page does not function as expected.

Problem

There are problems using the portal page (i.e., content server home page). When I open the page, none of the images display and none of the buttons/links work.

Recommendation

The problems you have experienced are not actually errors. After the content server installation is complete, the portal page that you are trying to view directly is actually the basic framework of the content server. At this point, the building process is not complete. Therefore, what you see is normal and does not indicate any error.

Turning Off Logging When Running Content Server As Windows Service

Question

How do I turn off logging to the *[Stellent_Dir]\bin\IdcServerNT.log* file when running the content server as a service?

Recommendation

If you are running the content server as a Windows service, the standard output and error messages are written to the file [Stellent_Dir]\bin\IdcServerNT.log. This file contains the output from any tracing options that you turn on.

To turn off logging to the *IdcServerNT.log* file, you need to include the following configuration variable in your *[Stellent_Dir]\bin\intradoc.cfg* file: UseRedirectedOutput=false

You will need to restart the content server service after setting this configuration variable.

Chapter

5

DATABASE ISSUES

OVERVIEW

This section provides solutions to several common database issues. Please attempt the solutions recommended in this guide before Support.

- Content Server Terminates or Won't Start (page 5-2)
- Content Server Can't Connect to Database (page 5-6)
- Database Connection Terminates or Fails (page 5-17)
- Database Runs out of Connections (page 5-19)
- Database Sizing Issues (page 5-22)
- Distributed Installation Issues (page 5-25)
- ✤ SQL Server-Specific Issues (page 5-26)
- Oracle-Specific Issues (page 5-27)
- DB2-Specific Issues (page 5-29)
- Miscellaneous Issues (page 5-32)

CONTENT SERVER TERMINATES OR WON'T START

This section covers the following topics:

- Permission Problem with IDC Content Service (page 5-2)
- Database Password Changed (page 5-3)
- Changed/Forgot Content Server 'sa' Password (page 5-4)

Permission Problem with IDC Content Service

Symptom

The content server will not start and users cannot access the system.

Problem

When trying to start the content server, one of the following error messages (or similar) is issued:

Could not start the "IDC Content Service <instance_name>" service on Local Computer, error 1053. The service "IDC Content Service <instance_name>" is failing to start.

Recommendation

If an incorrect user permission setting has been configured during the database installation, the IDC Content Service will not start. Consequently, the content server will also be unable to start, because the IDC Content Service launches the content server, which must be running for users to access the system. During the database installation process, the wrong user type may have been specified for the service startup. If so, this is a permission problem.

For example, if, during a database installation (such as SQL Server), the service was installed as a network user, then the service is set up to start as a "LocalSystem" and a network user must physically start the service. In this case, the setting must be changed to a Network Admin account. This should resolve the problem and ensure that the IDC Content Service starts and properly launches the content server.

Database Password Changed

Symptom

The content server does not start.

Problem

The database password was changed on the database side and now the content server will not start.

Recommendation

Generally, during the database installation process, you are prompted to enter a database password. When the content server is installed, the database password information must be provided to set up proper connection data (to ensure proper communication between the database and the content server). Because of this, changes in one location are generally required in other locations.

On the content server side, the database password is used as the JDBC user password (set using the System Properties utility) and the JdbcPassword configuration setting (stored in <*CS_Instance_Dir*>/config/config.cfg). When the database password is changed, it must be properly updated in two other locations:

- on the database side, in the ODBC Data Source Administrator
- on the content server side, using the System Properties utility.

The ODBC system data source stores information about how the connection is made to the designated data provider. The password used for the database must be updated in the ODBC DSN information to ensure that the database can verify the authenticity of the user logging in. If this process fails, the process to start the content server will also fail. Make the changes in the order specified below.

Editing the Database Password in the ODBC Data Source Administrator

To update the database password for the database side, complete the following steps:

- Open Control Panel, and then choose Administrative Tools—Data Source (ODBC). The ODBC Data Source Administrator is started.
- 2. Open the System DSN tab.

- Select the applicable database ODBC system data source and click Configure. The DSN Configuration wizard opens.
- 4. If necessary, enter the applicable **data source** and **database** information and click **Next**.
- 5. Enter the new password and click Next.
- 6. Complete the remaining wizard dialogs as required.
- Note: On UNIX platforms, DSNs are set up and stored in the environment file named odbc.ini located in the <CS_Install_Dir>/_platform/odbc/ directory. The odbc.ini file is used by ODBC applications to resolve DSN names. It must be updated for each DSN that is used and is referenced by an environment variable named PDBCINI.

Editing the Database Password using the Systems Properties Utility

To update the database password on the content server side, complete the following steps:

- 1. Start the System Properties utility.
- 2. Open the **Database** tab.
- 3. Change the JDBC User Password field value.
- 4. Click OK.
- 5. Restart the content server.

Changed/Forgot Content Server 'sa' Password

Symptom

The original 'sa' password has been changed or forgotten. The content server will not start and the database cannot be accessed.

Problem

If I changed the 'sa' password on the content server side, what do I need to do to get the content server started? Alternatively, if the 'sa' password has been forgotten, how can it be cleared out and a new one created?

Recommendation

The 'sa' password is actually the default value for the JDBC user name (set using the System Properties utility) and the JdbcUser configuration setting (stored in <*CS_Instance_Dir>/config/config.cfg*). The acronym 'sa' stands for system administrator. The password associated with the 'sa' user is the JDBC user password (set using the System Properties utility) and the JdbcPassword configuration setting (stored in <*CS_Instance_Dir>/config/config.cfg*). The password is also referred to as the database password.

Correcting a Changed 'sa' Password

Presuming that the database password was changed correctly on the database side (using the applicable database administrative tool) and the ODBC DSN properly changed, all that would need to be done on the content server side is to update the JDBC user password using the System Properties utility. For further details refer to Editing the Database Password using the Systems Properties Utility (page 5-4).

Be sure to test the ODBC connections to verify that it is working properly and check the JDBC user password on the Database tab in the System Properties utility. If these are both properly set up and the content server still won't start, there might be a problem with the JdbcPassword setting in *CS_Instance_Dir>/config/config.cfg*. When the database password is changed using the System Properties utility, the password encryption function is automatically activated and the database password is stored in encrypted form in the *config.cfg* file. For further details refer to Database Password Is Not Encrypted (page 5-15).

If the database password configuration setting (JdbcPassword) in *config.cfg* is manually changed to the new password, the content server will not start. In this situation, it will be necessary to disable the database password encryption function. To do this, follow the procedure in Disabling the Encryption Function (page 5-16).

Changing a Forgotten 'sa' Password

The database 'sa' password can be changed, but this must be performed properly on both the database side and the content server side. To change the password on the database side, use the database administrative tool (for example, the SQL Server Enterprise Manager). Run the database administrative tool as the local admin to make the change in the database.

Creating a new password will have additional implications on the system. First, you will need to update the ODBC DSN information. Second, the content server database password fields and configuration settings must be updated using the System Properties utility. Both of these procedures are detailed in Database Password Changed (page 5-3).

CONTENT SERVER CAN'T CONNECT TO DATABASE

This section covers the following topics:

- General Connectivity Problem Analysis (page 5-6)
- Wrong Port Configured (page 5-10)
- Pointing to Wrong Database (page 5-10)
- Content Server Doesn't Connect to Manually Configured SQL Server Database (page 5-12)
- New User Passwords Are Not Recognized (page 5-14)
- Database Password Is Not Encrypted (page 5-15)

General Connectivity Problem Analysis

When the content server is not connecting to the database, various error messages are issued that indicate connectivity problems. Normally, this situation can be quickly and easily resolved by checking:

- ODBC / System DSN Connection Information (page 5-6)
- JDBC Connection Information (page 5-9)

ODBC / System DSN Connection Information

An ODBC system data source stores information about how to connect to the designated data provider (in this case, the content server database). The ODBC system data source is frequently referred to as the system DSN (data source name). System DSN information includes some of the data derived from the content server's Java Database Connectivity (JDBC) connection information. Specifically, the system DSN includes the following:

✤ Data Source Name

This value is the database instance name from the JDBC connection string (set on the Database tab of the System Properties utility) and the JdbcConnectionString configuration setting (stored in *<CS_Instance_Dir>/config/config.cfg*).

Login ID

This value is the JDBC user name (set on the Database tab of the System Properties utility) and the JdbcUser configuration setting (stored in <*CS_Instance_Dir>/config/config.cfg*).

Password

This value is the JDBC user password (set on the Database tab of the System Properties utility) and the JdbcPassword configuration setting (stored in <*CS_Instance_Dir*>/config/config.cfg).

The system DSN information for the database must correspond exactly with the content server's JDBC information in all locations. If this information is inaccurate or inconsistent, the content server will not be able to connect with the database. If there is a connectivity problem between the content server and the database, it is recommended that you try the following corrective procedures regarding the ODBC connection information:

- Checking/Testing the Database Provider (page 5-7)
- Verifying the System DSN Listing (page 5-8)

Checking/Testing the Database Provider

Checking and testing the database provider can ensure that the ODBC information is correct and that there is proper connectivity between the content server and the database. To check or test the database provider, complete the following steps:

- 1. Log into the content server as an administrator.
- 2. Go to the Administration page and click on the **Providers** link.

The Providers page is displayed.

3. Click on the Info link of the appropriate database provider.

The Database Provider Information page opens. Verify the JDBC information.

4. Return to the Providers page and click on the **Test** link of the appropriate database provider.

Verifying the System DSN Listing

Verifying that a system DSN record exists with the correct information for the database ensures that there is proper ODBC connectivity between the content server and the database. To verify the system DSN listing, complete the following steps:

- Open Control Panel, and then choose Administrative Tools—Data Source (ODBC). The ODBC Data Source Administrator is started.
- 2. Open the System DSN tab.
- 3. Select the applicable database ODBC system data source and click **Configure**.

The DSN Configuration wizard opens.

- 4. Verify that the correct data source (database) name and server are listed.
- 5. Click Next.
- 6. Verify that the correct login ID and password are listed.
- 7. Click Next.
- 8. Verify that the correct database is selected.
- 9. Click **Cancel** to exit the wizard and avoid creating a new and potentially unnecessary ODBC data source.
- 10. Click OK.



Note: If the proper ODBC system data source is not listed on the System DSN tab, one needs to be created (beyond the scope of this topic). If the proper ODBC system data source exists but any of the fields are incorrect, these need to be updated as appropriate. For example, refer to Editing the Database Password in the ODBC Data Source Administrator (page 5-3).



Note: On UNIX platforms, DSNs are set up and stored in the environment file named *odbc.ini* located in the *<CS_Install_Dir>/_platform/odbc/* directory. The *odbc.ini* file is used by ODBC applications to resolve DSN names. It must be updated for each DSN that is used and is referenced by an environment variable named PDBCINI.
JDBC Connection Information

The JDBC connection information is generally created automatically during the content server installation process if there are no problems that make it impossible for the system to connect to the database. These settings must be updated manually if any database parameters were changed (for example, a changing password) or if the system could not find the database during installation.

The content server's JDBC information for the database must correspond exactly with the system DSN information in all locations. If this information is inaccurate or inconsistent, the content server will not be able to connect with the database. If there is a connectivity problem between the content server and the database, it is recommended that you check the JDBC field settings using the System Properties utility:

- 1. Start the System Properties utility.
- 2. Open the **Database** tab.
- 3. Verify that the following JDBC fields are identical to their corresponding fields in the applicable ODBC system data source (listed in parenthesis):
 - The database instance name (ODBC data source name)
 - The JDBC user name (database login ID)
 - The JDBC user password (database password)
- 4. Click OK.



Note: If any of the JDBC field values are incorrect, these need to be updated as appropriate. For example, refer to Editing the Database Password using the Systems Properties Utility (page 5-4).

Editing JDBC Information in the System Properties Utility

To edit one or more JDBC connection information fields:

- 1. Start the System Properties utility.
- 2. Open the **Database** tab.
- 3. Change the JDBC field value that is incorrect.

These field values correspond to the JDBC configuration settings in *<CS_Instance_Dir>/config/config.cfg*.

- 4. Click OK.
- 5. Restart the content server.

Wrong Port Configured

Symptom

The content server starts, but does not connect to the database.

Problem

After I start the content server, the following error messages are issued:

Unable to execute service WORK_IN_PROGRESS and function createResultSetSQL. Unable to instantiate the system database. Unable to create database connection for the database <database_name> with connection string <JdbcConnectionString>. Please make sure that the connection string, user and password are correct.

Recommendation

These error messages usually indicate that there is a database connectivity problem. Check the ODBC connection information to ensure that it is correct. For additional information about the ODBC connection and how to verify its accuracy, refer to ODBC / System DSN Connection Information (page 5-6). If the ODBC connection information is configured correctly, then there must be another problem that prevents the content server from connecting to the database.

Perhaps the wrong port was specified during the content server installation process. The server port is used by content server's outgoing providers and/or web servers to communicate with other system entities as well as for data transfer. Depending upon the specific web server being used to connect to the database, the port that is set up for the web server may be inaccurate.

The assigned port number must be correct to ensure proper connectivity and communication between the content server and the system database. Additionally, the listener process in the content server database must be started and listening on the correct port. Check to see what port was configured for your database connection and ensure that the listener process is started and that it is listening on the configured port.

Pointing to Wrong Database

Symptom

The content server starts, but does not connect to the database.

Problem

After I start the content server, the following error message is issued:

The Network Adapter could not establish the connection. Unable to instantiate the system database. Unable to create database connection for the database '<database_name>' with connection string '<jdbc_connection_string>'.

Recommendation

This error message usually means that either the wrong port is configured for the database connection or that the database listener process has not started. Check to see what port was configured for your database connection and ensure that the listener process is started and that it is listening on the configured port. (Your system administrator can help determine and verify the listener process information.)

Checking the Content Server Port Configuration Setting

To check which port number was configured to for data transfer between the Content Server and database, complete the following steps:

- 1. Log into the content server as an administrator.
- 2. Go to the Administration page and click on the Admin Server link.

The Administration for Servers page is displayed.

- 3. Click the Edit Server link.
- 4. Select the appropriate content server and click Edit.

The configured port number is listed in the Server Port field.

5. Click **Cancel** to exit the page.

Checking the Database Connectivity Information

If the database port is configured correctly and the listener process is working, then there must be another problem that prevents the Content Server from connecting to the database. Try pinging the machine running the database. If the machine cannot be pinged, the Content Server cannot connect to the database on that machine. If the database connection information is wrong or points to a machine where the database does not reside, then the Content Server will try to connect with the wrong machine and will be unable to find the system database.

Make sure that the ODBC and JDBC connection information is correct:

- 1. Check the database provider to verify that the ODBC connection is valid and working [refer to Checking/Testing the Database Provider (page 5-7)
- 2. Verify that the ODBC data Source name information is correct [refer to Verifying the System DSN Listing (page 5-8)].
- 3. Ensure that the Jdbc Connection String configuration setting is correct [refer to JDBC Connection Information (page 5-9)].

Content Server Doesn't Connect to Manually Configured SQL Server Database

Symptom

The content server does not connect to the manually configured SQL Server database.

Problem

Content Server was installed without specifying the specific database and has been manually configured. How can the content server be set up to connect to the newly installed SQL Server database?

Recommendation 1

Check Installation Scripts

During the content server installation process, you are prompted to specify the database that will be used with the content server. In addition to a list of supported databases, an option is provided that causes the content server to run from a manually configured database. If this option is selected during installation, then the content server will not be able to connect to the database until the database is properly configured manually. Otherwise, if the database had been specified during the content server installation process, the installer would automatically configure the database correctly.

If you chose to configure your database manually, then the content server database tables need to be set up by loading and executing the proper installation scripts that are supplied with the Content Server software. The scripts are located at:

<CS_Instance_Dir>/database/<Database_Type>

Currently, there are script directories for various databases including Microsoft SQL Server, Oracle, Sybase, Informix, and DB2. These scripts must be executed in the correct order; otherwise the content server cannot connect to the database. Because each database type has slightly different scripts in their respective directory, obtain the proper sequence information from the Content Server installation guides or the Oracle Support Hotline.

Additional information about establishing the proper connections between the content server and the manually configured database as well as how to run the database table setup scripts, refer to the Content Server installation guides. These documents provide comprehensive information about manually setting up and connecting to specific databases and manually configuring the appropriate web servers.

Recommendation 2

Check config.cfg

The database connection information in *CS_Instance_Dir>/config/config.cfg* must be correct to connect content server to the database. Because the database has been manually configured, you cannot use the System Properties utility to enter all of the necessary database configuration settings. Instead, this information must be manually entered into the *config.cfg* file as follows:

1. Open the following file in a text editor:

<CS_Instance_Dir>/config/config.cfg

- 2. Locate or create the **#Database Variables** section.
- 3. Enter the required database configuration settings or make any changes as necessary.

The variables in the sample set of database configuration settings provided below will be the same for your database. However, the setting values, will be unique for every database. The sample is provided only as a guide.

```
#Database Variables
IsJdbc=true
JdbcDriver=sun.jdbc.odbc.JdbcOdbcDriver
JdbcConnectionString=JDBC:ODBC:SQLFor<Instance_Name>
JdbcUser=sa
JdbcPassword=stellent
DatabasePreserveCase=false
```

- 4. Save the changes and close the *config.cfg* file.
- 5. Restart the content server.

In addition to the proper database configuration settings, you should also ensure that the ODBC and JDBC connection information is also correct. For details refer to

ODBC / System DSN Connection Information (page 5-6) and JDBC Connection Information (page 5-9).

 \bigcirc

Note: Ensure that no changes have been made to the master database because it is uses as a reference for each new database created later. Subsequently, the changes in the master database are automatically inherited by new databases. Therefore, if the master database has been irreparably altered, it may probably be necessary to rebuild the database, running the SQL scripts in the correct order, and ensure the master database is restored back to the normal default.

Additional information about establishing the proper connections between the content server and the manually configured database as well as how to run the database table setup scripts, refer to the Content Server installation guides. These documents provide comprehensive information about manually setting up and connecting to specific databases and manually configuring the appropriate web servers.

New User Passwords Are Not Recognized

Symptom

The content server does not immediately recognize new user passwords.

Problem

The passwords of some existing users were changed on the database side. Now, the content server does not recognize the changes until it is stopped and restarted.

Recommendation

The content server only dynamically recognizes password changes that are made via the User Admin applet on the Administration page. For user password changes, the content server does not interact directly with the database. Therefore, to avoid having to stop and restart the content server, use the User Admin applet to make the changes rather than making them through the database administrative tool. This should trigger an update of the user cache. Alternatively, a content server's IdcCommand file can be created and executed that would batchload a list of user information.

Database Password Is Not Encrypted

Symptom

The database connection password is not stored in encrypted form.

Recommendations

When the content server database is initially set up for new versions of Content Server older than 7.5.1 or updated versions of content server, the password configuration setting (*JdbcPassword*) in *CS_Instance_Dir>/config/config.cfg* is not listed in encrypted form. However, if this configuration setting is edited using the System Properties utility, the new password will be stored in the *config.cfg* file in encrypted form. This is because the default password encoding configuration setting (*JdbcPasswordEncoding*) is set to use encryption by default after the password is edited in the System Properties utility.



Note: Beginning with version 7.5.1, new installations of Content Server encrypt the database password upon installation. Update installations do not change the status of the database password.

Password Configuration Settings for Encryption Function

The JdbcPasswordEncoding parameter has two settings:

Intradoc

This is the default setting. However, this setting is only activated the first time the password is edited using the System Properties utility. If the *JdbcPasswordEncoding* configuration setting in *<CS_Instance_Dir>/config/config.cfg* is set to this value, then the *JdbcPassword* setting will be encrypted (or more accurately, hashed). The encrypted password is only usable by the content server. For example: JdbcPasswordEncoding=Intradoc

JdbcPassword=FUx7gf8h1CpNmMuh1NAxx0hVtVhTYi7daUe+pIeQ+OU=

ClearText

This is the setting to override the default setting. This setting must be used if the encryption mode has been activated and you do not want the password to be encrypted. Manually changing the encrypted *JdbcPassword* configuration setting in <*CS_Instance_Dir*>/config/config.cfg will prevent the content server from starting. For details refer to Changed/Forgot Content Server 'sa' Password (page 5-4). If the *JdbcPasswordEncoding* configuration setting in <*CS_Instance_Dir*>/config/config.cfg is set to this value, then the *JdbcPassword* setting will be in clear text form. For example: JdbcPasswordEncoding=ClearText JdbcPassword=stellent

Activating the Encryption Function

To activate the default password configuration setting (to enable encryption), complete the following steps:

- 1. Start the System Properties utility.
- 2. Open the **Database** tab.
- 3. Change the JDBC User Password field value.

This field value corresponds to the *JdbcPassword* configuration setting in <*CS_Instance_Dir*>/config/config.cfg.

- 4. Click OK.
- 5. Restart the content server.

Disabling the Encryption Function

If the password is edited and is currently stored in encrypted form in <*CS_Instance_Dir*>/*config/config.cfg*, it should not be manually changed as this will prevent the content server from starting. Instead, if you want to disable password encryption, the default password encoding configuration setting (*JdbcPasswordEncoding*) needs to be changed in <*CS_Instance_Dir*>/*config/config.cfg*.

Conversely, if the password has not been edited but you want to allow the option to change the password without activating the encryption function, then the default password encoding configuration setting (*JdbcPasswordEncoding*) needs to be added to <*CS_Instance_Dir*>/config/config.cfg.

To disable the encryption default mode, complete the following steps:

1. Open the following file in a text editor:

<CS_Instance_Dir>/config/config.cfg

- 2. Locate or create the **#Database Variables** section.
- Find the configuration setting *JdbcPasswordEncoding* and change the value (or add the configuration setting) as follows: JdbcPasswordEncoding=ClearText

Alternatively, if the *JdbcPasswordEncoding* line is already in the *config.cfg* file, it can be commented out to stop the encryption function.

- Optional: If the password is encrypted, edit the *JdbcPassword* configuration setting by entering the new password (make sure it is identical to the password on the Database tab of the System Properties utility).
 Alternative, after saving and closing the *config.cfg* file, you can use the System Properties utility to change the password.
- 5. Save the changes and close the *config.cfg* file.
- 6. Restart the content server.

DATABASE CONNECTION TERMINATES OR FAILS

This section covers the following topics:

- Database Connection Repeatedly Fails (page 5-17)
- Database Connection Intermittently Fails (page 5-18)

Database Connection Repeatedly Fails

Symptom

The connection to the database repeatedly fails.

Problem

The database password was changed on the database side and now the connection fails.

Recommendation

If the database password is changed on the database side without properly updating the necessary ODBC and JDBC connection information, database connectivity will terminate or fail.

For more detailed information on properly setting up the required connectivity information on the database side, refer to Database Password Changed (page 5-3). For more detailed information on properly setting up the necessary connectivity information on the Content Server side, refer to ODBC / System DSN Connection Information (page 5-6) and JDBC Connection Information (page 5-9).

Database Connection Intermittently Fails

Symptom

The connection to the database repeatedly and abruptly stops functioning.

Problem

When the database connection fails, the error messages reference ODBC problems along with indexer type errors. The content server log file contains multiple "Indexing aborted" error messages that occur repetitively every five minutes. In addition to the index-related errors, the content server suddenly is unable to connect to the database. After the database connection stops, it is possible to reboot and restart the content server.

Recommendation

Although the index-related error messages seem to indicate an indexing problem, this is actually a connectivity problem. When the database connection fails, the content server cannot 'write' to the database. The automatic update cycle in the content server system is configured to try every five minutes to reconnect to the database, which causes the errors to be logged consistently at five minute intervals. However, the indexer errors are a result of the content server's inability to connect to the database.

Since the content server is able to restart and function properly, the ODBC and JDBC connection configurations are obviously correct. This eliminates the possible problem with the indexer or the content server. It may be helpful to ensure that the database has enough space to grow with the demands of writing data to the database. Otherwise, there must be a different source for the problems.

Since the problem is consistently intermittent and seems to be repetitive, the interruption pattern may indicate that a scheduled process (such as an automatic database backup) is interfering with the connection between the content server and the database. This is particularly plausible if the error logs show that the interruption consistently occurs at a specific time. If so, the scheduled incremental or full backups could interfere with the database connectivity and would subsequently stop the SQL service.

It might help to look at the specific network interruptions listed in the operating system event viewer at the time(s) that the content server stopped. If the connection problems consistently occur simultaneously with the scheduled backups, you may have to consider scheduling the incremental or full backups at a time when the content server is unlikely to write to the database.

Additionally, the connection and current server activity can be checked using the SQL Server Enterprise Manager: <Server_Group>—Management—Current Activity— Process Info. The process connections directed to the content server database should show a "Sleeping" status, which indicates the process is waiting for a lock or user input.

DATABASE RUNS OUT OF CONNECTIONS

This section covers the following topics:

- Need to Increase Database Connection Limits (page 5-19)
- Unable to Execute Service and Function (page 5-20)
- Unable to Load User Information (page 5-21)

Need to Increase Database Connection Limits

Question

The system currently uses 100 RDBMS connections and needs to be increased. What information is available about the maximum limit for connections and potential performance issues related to the number of connections?

Recommendation

Due to the tremendous number of variables that influence system performance (such as the particular user operations being executed), specific limits about database connections and statistics pertaining to subsequent performance are not available. However, as a general guideline, when adding more connections, performance cannot be expected to improve much beyond the existing number of CPUs on the servers.

Unable to Execute Service and Function

Symptom

The database is inaccessible and/or the available connections to the database are completely expended. Additionally, the following error messages (or others that are similar) may be displayed:

```
Unable to execute service <service_name> and function <function_name>. (System
Error: There are no connections available from pool for provider
    'provider_name>'.)
```

Problem

The system has run out of connections to the database and the content server cannot access the database.

Recommendation

This error message indicates that the activity on the server exceeds the default number of connections in the database pool. The number of database connections can be increased by editing the *NumConnections* variable value in *CS_Instance_Dir*/*config/config.cfg*. Alternatively, the number of connections parameter can be increased by editing the database provider information, because the default number may be insufficient. Editing the NumConnections configuration setting can be accomplished in one of two ways:

- Changing Connections in the 'config.cfg' File (page 5-20)
- Changing Connections on the Providers Page (page 5-21)

Changing Connections in the 'config.cfg' File

1. Open the following file in a text editor:

<CS_Instance_Dir>/config/config.cfg

 Locate or create the #Additional Variables section, and make sure this section contains the following line. NumConnections=10

Numeonneccions-10

Due to the vast differences in system implementations, a standard recommended value is not available. It is, however, recommended to start with a value of perhaps 10 or 20. Monitor the effects and, if necessary, increase this value in increments of 5 followed by testing until the desired results are achieved.

- 3. Save the changes and close the *config.cfg* file.
- 4. Restart the content server.

Changing Connections on the Providers Page

- 1. Log into the content server as an administrator.
- Go to the Administration page and click on the Providers link. The Providers page is displayed.
- 3. Click on the **Info** link of the appropriate database provider. The Database Provider Information page is displayed.
- 4. Click the **Edit** button to modify the parameter values.
- 5. Enter the desired number in the Number of Connections field.
- 6. Click the **Update** button to save the changes.
- 7. Restart the content server.

Unable to Load User Information

Symptom

The database is inaccessible and/or the available connections to the database are completely expended. Additionally, the following error messages (or others that are similar) may be displayed:

Event generated by user `<user>' at host `<IP_address>'. Unable to load user information. (System Error: There are no connections available from pool for provider `<provider_name>'.)

Problem

The system has run out of connections to the database and content server cannot access the database.

Recommendation

This error message indicates that the activity on the server exceeds the default number of connections in the database pool. The number of database connections can be increased by editing the *NumConnections* variable value in *<CS_Instance_Dir>/config/config.cfg*. Alternatively, the number of connections parameter can be increased by editing the

database provider information, because the default number may be insufficient. Editing the NumConnections configuration setting can be accomplished in one of two ways:

- Changing Connections in the 'config.cfg' File (page 5-20)
- Changing Connections on the Providers Page (page 5-21)

DATABASE SIZING ISSUES

This section covers the following topics:

- Estimating Sizes for Database Files (page 5-22)
- Estimating the Open File Limit Size (UNIX) (page 5-23)
- Exceeded the Open File Limit (UNIX) (page 5-23)

Estimating Sizes for Database Files

Question

What formula can be used to estimate the table sizes for database files? Are estimated sizes based on the amount of metadata to be collected multiplied by the number of documents to be maintained?

Recommendation

Due to the numerous types of databases supported by Oracle in addition to the dependency on the type of metadata used, a standard formula for estimating database table sizes is not available. However, it is possible to collect the type and sizes of the columns in each Oracle database table (the Table Properties window lists the data type and size of each table column name).

With that information plus what you add in custom metadata, you can calculate the size of the table rows. Estimate the number of rows you will have in each table and process these numbers through an estimation formula provided by your database vendor and you can probably calculate a database size that will meet your needs.

Estimating the Open File Limit Size (UNIX)

Question

What factors must be considered to properly estimate the open file limit size that needs to be set to accommodate the content server processes? For example, each site is expected to have approximately 50,000 pages and we are using a UNIX operating system which has a default limit of 1024 open files per process.

Recommendation

The content server does not have any specific requirements regarding the UNIX file size limit. If you anticipate that the content server will need to access large files, your system administrator may need to increase the server's file size limit. Generally, most UNIX systems can dynamically adjust the kernel to manage increasing file sizes. However, some UNIX machines require manual kernel adjustments to accommodate various file sizes of specific implementations (for example, Solaris).

If your sites exceed the default open file size limit, an applicable error message will probably be written to the content server logs. In that case, you will need to increase the open file size limit. For this procedure, refer to Exceeded the Open File Limit (UNIX) (page 5-23).

Exceeded the Open File Limit (UNIX)

Symptom

The following related error messages are recurrently reported in the content server logs:

```
Socket accept failed: 4444 Too many open files.
Too many open files: file permissions deny server access.
```

Problem

After I increased the number of open files per process to 4096, I checked the plimit using the process id of the content server, but shows 1024 rather than 4096. I discovered that the 1024 limit is hard-coded. How can I override this number?

Recommendation

The reported errors occur when the number of open file descriptors for the content server process is inadequate and the limit needs to be increased. One of the default values in the

UNIX startup scripts is the open file descriptor configuration variable and its default limit size (INTRADOCMS_FDLIMIT=1024). This variable sets the number of open file descriptors that the content server will be allowed to use.

The default file descriptor limit can be changed, but care should be taken not to set the file descriptor limit higher than the system hard limit. The default file descriptor limit can be verified and, if necessary, changed.

Checking File Descriptor Limits

To check the maximum soft and hard file descriptor limits in a UNIX operating system, complete the following steps:

- 1. Login as root.
- 2. Switch to ksh.
- To obtain the soft (current) limit, enter: ulimit -n
- 4. To obtain the hard (max) limit, enter: ulimit -nH
- 5. To change the soft limit (for that shell only), enter: ulimit -n 512

Overriding Startup Values

The default number of open file descriptors available to the content server is 1024. If you want to increase this number (for example, to 4096), complete the following steps:

- 1. Create a file called *config* in the *CS_Instance_Dir>/etc/* directory.
- 2. In the *config* file, add the following line: INTRADOCMS_FDLIMIT=4096
- 3. Make sure that the kernel is configured to allow the number of open file descriptors per process that you set in the *config* file (check the maximum soft and hard file descriptor limits).



Note: If you set the file descriptor limit in the *config* file higher than the system hard limit, a "bad ulimit" error message will be issued.

DISTRIBUTED INSTALLATION ISSUES

This section covers the following topics:

- Moving or Using an Alternate Database (page 5-25)
- Linking to the Database Through a Firewall (page 5-26)

Moving or Using an Alternate Database

Question

What is involved in either moving an existing database to another server or setting up and connecting to a new database on another server?

Recommendation

The general procedures for either case involve the following general requirements after the database is installed or moved to the other server:

- The content server database tables must be set up by loading and executing the proper install scripts. For details refer to Content Server Doesn't Connect to Manually Configured SQL Server Database (page 5-12).
- 2. The content server must properly connect to the database, which requires manually entering the database connection variables into the content server's *config.cfg* file. For details refer to page 5-13.
- Make sure that the ODBC and JDBC connection information is correct. For details refer to ODBC / System DSN Connection Information (page 5-6) and JDBC Connection Information (page 5-9).

Linking to the Database Through a Firewall

Question

What is involved in linking to a database from a content server on the other side of a firewall?

Recommendation

To connect to a database that is part of a distributed system architecture and a firewall separates the database from the content server, a correct ODBC connection must be established, accurate DSN information added to applicable files and configuration settings, and all JDBC information updated. For details refer to ODBC / System DSN Connection Information (page 5-6) and JDBC Connection Information (page 5-9).

SQL SERVER-SPECIFIC ISSUES

This section covers the following topics:

Setting Up Low Resource Alerts and MS Paging (page 5-26)

Setting Up Low Resource Alerts and MS Paging

Question

Are there any negative effects on Content Server if low resource alerts are set up in SQL Server in addition to setting up the Microsoft paging service?

Recommendation

Neither of these services should have a detrimental impact on Content Server, because the content server's database in SQL Server is not integrated with SQL control functions. The content server's database only requires a properly configured ODBC connection to ensure unobstructed data handling.

ORACLE-SPECIFIC ISSUES

This section covers the following topics:

- Check-In Problems/Errors (page 5-27)
- Increasing Tablespace (page 5-27)
- Checking Oracle Driver Information (page 5-28)
- Using Debug Builds for Oracle Drivers (page 5-29)

Check-In Problems/Errors

Symptom

There are some problems when documents are checked into the content server.

Problem

When the Repository Manager is launched, the following error message is issued: ORA - 3232 Unable to allocate an extent of 32 blocks.

Recommendation

The Max Extents parameter must be increased for the Oracle database tablespace and tables including:

- Docmeta
- Revisions
- Documents

Increasing Tablespace

Symptom

Error messages are issued that state the database size has reached its maximum limit.

Problem

The tablespace in the database has reached its maximum capacity and needs to be increased. Currently, there are 20,000 documents checked in and an additional 10,000 must be added. How much should the database be increased?

Recommendation

Due to the tremendous number of variables that influence database tablespace (such as the number of documents, number of metadata fields, number of users, amount of activity, etc.), specific size recommendations are not available. However, the current database size and space available is listed in the database properties window. If your database reaches its maximum size limitation, it is possible to increase the size as necessary. (For example, a database set to 3 MB can be reconfigured to 20 MB.)

Checking Oracle Driver Information

Question

I am getting some really odd Oracle behavior. Should I check the driver? Which Oracle driver should I use?

Recommendation

Always check the driver. The Oracle drivers shipped with Content Server may have newer versions out that fix various bugs. However, which driver to use can be quite confusing. Here is a simple algorithm to determine which version to use:

- 1. If you are using Oracle 8i, always use *classes12.zip*.
- 2. If you are using Java JDK 1.2 or 1.3, always use *classes12.zip*.
- 3. In all other cases (i.e. with JDK 1.4 and/or Oracle 9i), use ojdbc14.jar.



Note: The default class file used with Content Server 7.*x* is *ojdbc14.jar*.

To determine the exact version of the Oracle driver being used by the content server, complete the following steps:

1. Open the *classes12.zip* or *ojdbc14.jar* file (for example, using WinZip) in the following location:

<CS_Instance_Dir>/shared/classes

2. Look for the file *Manifest.mf* (it may be located in a *Meta-inf* subdirectory). This is a plain-text file that includes version and timestamp information about the driver file, for example:

```
Specification-Version: "Oracle JDBC Driver version - 9.0.2.0.0"
Specification-Vendor: "Oracle Corporation" .
Implementation-Title: "ojdbc14.jar"
Implementation-Version: "Oracle JDBC Driver version - 9.0.2.0.0"
Implementation-Vendor: "Oracle Corporation"
Implementation-Time: "Wed Feb 19 15:32:24 2003"
```

Using Debug Builds for Oracle Drivers

Question

On the Oracle install media, there are drivers with _g in the file names in the following directory:

/packages/allplatform/oracle_jdbc/10.1.0.2.0

What are these and should they be used with Oracle 10g?

Recommendation

The *_g designates a debug build for the Oracle driver. You could use this to debug issues with the JDBC layer.

DB2-SPECIFIC ISSUES

This section covers the following topics:

- DB2 Deadlocks or Timeouts Occur (page 5-29)
- Troubleshooting DB2 JDBC Issues (page 5-31)

DB2 Deadlocks or Timeouts Occur

Symptom

Deadlock or timeout errors are reported during publication of a content item, and the following (or similar) errors are issued:

Unable to execute service CHECKIN_PUBLISH and function docRefinery.

```
(System Error: Unable to execute query 'UrevisionExtension(UPDATE RevisionS
SET dWebExtension = 'hcsp', dStatus = 'DONE', dProcessingState =
'Y',dIndexerState = ' ' WHERE dID=999999)'. [IBM][CLI Driver][DB2/SUN]
SQL0911N The current transaction has been rolled back because of a deadlock or
timeout. Reason code "2". SQLSTATE=40001)
Error detected uploading [stellent3^ADVWEPG_S2_010894]. Response: Unable to
execute service CHECKIN PUBLISH and function docRefinery.
```

Problem

DB2 has four transaction isolation levels that can be specified:

- Repeatable Read (RR)
- Read Stability (RS)
- ✤ Cursor stability (CS)
- ♦ Uncommitted Read (UR—allow "dirty" read).

Repeatable Read (RR) uses "Next Key Locking," which is a mechanism that basically places locks not only on the current row, but also the row containing the next key in an index, in order to prevent phantom reads (which means the result set could grow in two consecutive reads because a second thread has changed a different row to match the query criteria). This implicit locking could cause two seemingly unrelated transactions to clash and form a deadlock.

Recommendation

The potential deadlock and timeout problems can be resolved by turning on DB2's DB2_RR_TO_RS environment variable: db2set DB2_RR_TO_RS=on



Important: Make sure that you use 'on' as the value, not 'YES' (as reported incorrectly by some IBM DB2 administration guides).

This will do all of the following:

- It sets a default isolation level for user tables to RS isolation level.
- It reduces locking of the next key for inserted or changed rows.
- It reduces deadlocks when using reverse scan indexes.

Troubleshooting DB2 JDBC Issues

Question

How do I troubleshoot issues with DB2 Java Database Connectivity (JDBC)?

Recommendation

It is recommended that you use the following path to troubleshoot DB2 JDBC issues:

1. Make sure that the *db2java.zip* file was correctly copied from:

Windows: <Drive>:/SQLLIB/java UNIX: <DB2_Root>/java

to <*CS_Instance_Dir*>/shared/classes.

- 2. Make sure that the above file is in the Classpath specified in <<u>CS_Instance_Dir</u>>/bin/intradoc.cfg, for example: CLASSPATH=\$COMPUTEDCLASSPATH;C:/stellent/shared/os/win32/classes/db2java.zip
- 3. Make sure that the "DB2 JDBC Applet Server" service is running:

Windows: In Task Manager, check that *db2jds.exe* is running. **UNIX:** Run the "ps -ef |grep db2jstrt" command.

4. Make sure that the "Database Variables" section in

<CS_Instance_Dir>/config/config.cfg is completed correctly, for example: #Database Variables IsJdbc=true JdbcDriver=COM.ibm.db2.jdbc.net.DB2Driver JdbcConnectionString=jdbc:db2://myserver.company.com:6789/MYDB JdbcUser=administrator JdbcPassword=password DatabasePreserveCase=1

where:

- "myserver.company.com" is the DNS or IP of the server running the DB2 Applet Server.
- "6789" is the TCP port that the DB2 Applet Server is listening on.
- "MYDB" is the name of the DB2 database running under the main DB2 instance.

If you are still unable to connect to the DB2 database, try the steps below to ensure you can properly connect to the DB2 JDBC Applet Server (and thus Content Server as well):

1. Windows: Stop the "DB2 JDBC Applet Server" service.

- 2. Windows: Stop the "DB2 JDBC Applet Server Control Center" service.
- 3. Start the JDBC Applet Server:

Windows: C:\Program Files\SQLLIB\bin\db2jd 6789 UNIX: <db2root>/bin/db2jstrt 6789

where "6789" is the TCP port that the DB2 Applet Server is listening on. If you start the Applet Server without specifying the port, it will use the default port of 6789.

4. **Windows:** Start the JDBC Applet Server Control Center from the Windows command prompt:

C:\Program Files\SQLLIB\bin\db2cc 6789

where "6789" is the TCP port that the DB2 Applet Server is listening on (as specified earlier).



Note: If you get a login dialog box, log in with the same user ID and password that you are using in the content server JDBC connection string.

5. If you are able to log in and the IBM DB2 Command Center applet shows up, try starting the content server from the command-line console:

Windows: <CS_Instance_Dir>\bin\IdcServer -console -debug UNIX: <CS_Root>/etc/idcserver_ctrl start

MISCELLANEOUS ISSUES

This section covers the following topics:

- New Check-Ins Compete for IDs (page 5-32)
- Changing User Permissions (page 5-33)
- Changing Host Name / Instance Name (page 5-33)
- Turning On System Database Tracing (page 5-34)

New Check-Ins Compete for IDs

Symptom

There are counter discrepancies between three database files: Counters, Documents, and Revisions. Corresponding fields in these files that should match show a one-digit inconsistency.

Problem

When new documents from the imported collection are checked in, they compete for the available IDs that results in an imbalance. For example:

DocID (Counters table) = 2287 dDocID (Documents table) max = 2642 RevID (Counters table) = 876 dID (Revisions table) max = 875 RevClassID (Counters table) = 1008 dRevClassID (Revisions table) max = 1007

Recommendation

The database table field "counters" must be reset so that the new document check-ins do not compete for the IDs during the collection import process. The three fields that are of particular importance are the DocID, RevID, and RevClassID. Reset each of these counters to 10000 and the conflicting ID problem should be resolved.

Changing User Permissions

Question

Is it possible to change content server user permissions (either downgrade or upgrade) without adversely effecting the current content server database(s)?

Recommendation

The content server user must have the database role of 'db_owner'. As long as the assigned user server roles permit this, there should not be any problems.

Changing Host Name / Instance Name

Question

What problems might be associated with changing the host name of the SQL database? How can the database host name be changed?

Recommendation

Until the host name information is changed in the system, the content server will continue to reference the old host name or instance name. Therefore, the primary issues involved with changing the host name or instance name of the SQL database concern editing the actual names as well as ensuring the ODBC and JDBC connection information is accurate.

Changing the Host Name

To change the host name, complete the following steps:

- 1. Start the System Properties utility.
- 2. Open the **Internet** tab.
- 3. Edit the host name in the HTTP Server Address field.
- 4. Click OK.
- 5. Restart the Content Server.

Changing the Instance Name

To change the instance name, complete the following steps:

- 1. Open the following file in a text editor: <*CS Instance Dir*>/config/config.cfg
- 2. Locate or create the IDC_Name configuration setting and edit its value.
- 3. Save the changes and close the *config.cfg* file.
- 4. Restart the content server.

Make sure that the ODBC and JDBC connection information is correct. Refer to ODBC / System DSN Connection Information (page 5-6) and JDBC Connection Information (page 5-9).

Turning On System Database Tracing

Question

My Content Server log file contains errors related to the database. How do I obtain more detailed information in order to troubleshoot the database error?

Recommendation

When performing services related to the database, errors received are entered into the Content Server log. Situations which may cause a database error are unique to the database configuration and customization of the content server. To troubleshooting a database error you can enable the 'systemdatabase' trace to get additional information.

Enabling System Database Tracing

To enable system database tracing, complete the following steps:

- 1. Launch a browser and log into the content server with administrator rights.
- 2. Go to the Administration page of the content server instance, and click on the System Audit Information link.
- 3. Scroll down to the Edit Active Output Console Tracing section.
- 4. Select 'systemdatabase' from the drop-down list. If desired, select the Full Verbose Tracing check box for extra detailed log information.
- 5. Click Update.

The content server output log will now contain detailed log entries related to the system database.

6. Click on the View Server Output link and examine the console output log for additional troubleshooting information related to the database.



Note: Enabling trace options will generate additional log entries. When you have completed your troubleshooting, make sure that you disable your trace to avoid large log files.

Chapter 60

ARCHIVING ISSUES

OVERVIEW

This section provides solutions to several common archiving issues. Please attempt the solutions recommended in this guide before Support.

- Importing Issues (page 6-1)
- Exporting Issues (page 6-11)
- Transfer Issues (page 6-14)
- ✤ WebDAV Issues (page 6-19)
- Replication Issues (page 6-20)
- Oracle-Specific Issues (page 6-22)
- Miscellaneous Issues (page 6-23)

IMPORTING ISSUES

This section covers the following topics:

- File Extension Errors on Import Machine (page 6-2)
- Selecting Specific Batch Files for Import (page 6-3)
- Import Maps Don't Work After Archive Import (page 6-3)
- Identifying Imported Content Items From Archive (page 6-4)
- Duplicate Content Items in Content Servers (page 6-5)

- Importing Archived Content to Proxied Server Fails (page 6-6)
- No Importing Errors But Documents Are Missing (page 6-6)
- Errors About Invalid Choice List Values (page 6-7)
- Import Fails Due to Missing Required Field (page 6-8)
- Changed Metadata Field Makes the Archiver Freeze During an Import (page 6-9)

File Extension Errors on Import Machine

Symptom

I am receiving errors on the importing machine indicating that there are transfer and file extension problems with the documents.

Problem

The following errors were issued to the Archiver log:

```
Error: Event generated by user <user_name> at host <host_name>. File I/O
error. Saving to file collection.hda. Write error.
Error: Import error for archive <archive_name> in collection
<collection_name>: Content item <item_name> was not successfully checked
in. The primary and alternate files must have different extensions.
```

Recommendation

The I/O error on the export side probably corrupted the batch file and is, in turn, causing the file extension error on the import side. Possible solutions include:

- Open the batch file in a text editor and check for invalid data. Try deleting the exported *collection.hda* file and manually re-run the export/import function.
- On the exporting server, open the applicable *collection.hda* file and look for the lines associated with the content items that caused the file extension error. Some of the revisions of these content items may have the native file in the vault location listed in the alternate file location. There might also be a format entry for the alternate file. Delete these lines and re-import the files.
- Add an alternate extensions configuration setting to the content server's configuration config.cfg file (<CS_Instance_Dir>/config.cfg) on the importing server:
 - 1. Open the *config.cfg* file in a text editor:

<CS_Instance_Dir>/config/config.cfg

- 2. Locate the General Option Variables section
- 3. Enter the following configuration setting: AllowSamePrimaryAlternateExtensions=true

This configuration setting allows checked in content items to use identical document extensions for both the alternate and primary files.

4. Save and close the *config.cfg* file.



Note: Although it probably is not necessary to add this configuration setting to the content server *config.cfg* file on the exporting server, it may be worthwhile to do so for general preventative measures.

5. Restart the content server.

Selecting Specific Batch Files for Import

Question

How can I select and re-run specific batch files from the General tab of the Archiver utility without deleting the remaining files that are required for backup purposes?

Recommendation

The most efficient method would be to create a new collection, copy the desired archives to the new collection, and run the import from there.

Import Maps Don't Work After Archive Import

Symptom

I configured a value map to change metadata values during the import on an archive collection. But after the transfer, the import maps don't work.

Problem

The metadata values didn't reflect the configured metadata value changes.

Recommendation

To ensure that metadata value changes are retained when the files are exported into an archive and then later imported from that archive, the value maps must be configured on

both sides of the transfer process. This means that the same value map must be configured on both the source (exporting) server as well as the target (importing) server.

Identifying Imported Content Items From Archive

Question

Due to a system crash, I need to import content from the old archive into a new archive without changing the content information (metadata) of the documents. How can I preface each content item using a letter or number to indicate that all the documents with this designation are new imports (but actually originated from the old archive)?

Recommendation

The archived documents can be re-imported and appropriately marked to distinguish them from other imported content items by applying an import map using the Content ID metadata field. An import map allows you to configure how values are copied from one metadata field to another during import. To set up the import mapm, complete the following steps:

1. On the Import Maps tab of the Archiver utility, click Edit in the Field Maps section.

The Edit Value Maps screen is displayed.

- 2. Select the All check box (leave the Input Value field blank).
- 3. Select Content ID from the Field drop-down list.
- 4. Enter X<\$dDocName\$> in the Output Value field.

Where 'X' is the letter or number used to distinguish the re-imported content items and 'dDocName' is the database table field value for the document Content ID.

5. Click OK.

After you re-import the archive, the letter or number used for 'X' should be added to the content ID of each content item. Be sure to configure the same value map on both the source (exporting) server and the target (importing) server. This ensures that the metadata value changes are retained when the files are imported from the archive.

Duplicate Content Items in Content Servers

Symptom

When I try to check in or import a content item, the following error message is issued: Content item already exists.

Recommendation

This error is issued when archiving is done between contribution servers that are using the same autonumbering scheme for content IDs. For example:

- Content ID 003' is checked into content server A and later archived to content server B. If a file is checked into content server B and the next autogenerated number happens to be 003, the error occurs.
- 'Content ID 005' is checked into both content server A and content server B. If this same content item is archived from content server A to content server B, the error occurs.

Possible solutions include:

- Set up an import value map that will add a prefix to the content ID of the imported files. For details refer to Identifying Imported Content Items From Archive (page 6-4).
- In each content server, use the System Properties utility to set up an automatic numbering prefix for checked-in content items:
 - 1. Start the System Properties utility.
 - 2. Open the Options tab.
 - 3. Select the Automatically assign a Content ID on check in check box.
 - 4. Enter the desired prefix in the Auto Name Prefix field.
 - 5. Click OK.
 - 6. Restart the content server.

Importing Archived Content to Proxied Server Fails

Symptom

I am trying to import content from an exported archive to my proxied content server, but the import fails.

Recommendation

For more information about Archiver problems, open and view the Archiver logs (accessible from the content server's Administration page). These logs provide the type of message along with more descriptive information about the logged messages.

For example, if the Archiver log indicates that an import problem involves a metadata field option value that is unavailable, information about configured option lists for metadata fields can be found on the Information Fields tab of the Configuration Manager utility (accessible from the Administration page).

Using this information, compare the option list for the problem metadata field on both the exporting and importing servers. If there are any differences, corrections in one of the servers will make both option lists identical. This would resolve the unavailable option discrepancy.

No Importing Errors But Documents Are Missing

Symptom

When I run the import function, no errors are issued, but not all of the documents are being imported.

Problem

I exported 428 documents from the development server along with the configuration information (the metadata fields). Then, I transferred the archive to the main production server and ran the import. No errors were issued, so I thought everything had gone well. Unfortunately, when I searched the documents, I discovered that only 198 of the original 428 were actually imported.

Recommendation

Suggestions to resolve this problem include:

Make sure that all Microsoft Word documents are included in the search index.

Particular versions of the search component do not include Microsoft Word documents with embedded links in the search index. Thus, these files will not be found in search queries.

You can remove all embedded hyperlinks from the affected documents or add the following configuration setting to *CS_Instance_Dir>/config/config.cfg*: CheckMkvdkDocCount=true

This configuration setting ensures that all Word files are included in the search index. However, only the metadata is included, not the full text.

Try exporting the original set of documents and ensure that the source files are deleted. Then re-import the archive that was just exported.

Errors About Invalid Choice List Values

Symptom

My imports are failing.

Problem

The system issues error messages indicating that there are invalid choice list values. I am currently using an option list in the Dependent Choice List applet to configure and control the values.

Recommendation

Apparently a specific metadata taxonomy has been established for your option lists such that there are probably fields that are dependent on each other. In this case, certain values in option lists are available based on what values have been selected in a previous option list. Unfortunately, when using the Archiver, the dependencies in your option lists are obviously conflicting with the content server's capacity to work with custom metadata fields.

A workaround for the conflict involves using the content server's Configuration Manager utility rather than the Dependent Choice List applet. This necessitates that you enter the metadata fields and corresponding option list values on the Information Fields tab of Configuration Manager:

- 1. Log into the content server as an administrator.
- 2. Go to the Administration page and click on the **Configuration Manager** link.

The Configuration Manager utility is started.

- 3. Open the Information Fields tab.
- 4. Click the **Add** button and enter one of your metadata field names in the Add Custom Info Field dialog.
- 5. Click OK.

The Add Custom Info Field window is displayed.

- 6. Complete the fields as appropriate.
- 7. In the Option List Type field, choose the Select List Not Validated option. This option ensures that content whose specified value does not match one currently entered in the Use Option List are nevertheless checked in with the specified value. The Use Option List field lists the name for the list of values a user may choose from for the specified field.
- 8. Click OK.
- 9. Click the Update Database Design button.
- 10. Click the Rebuild Search Index button.

Use this method for the duration of your import process.

Import Fails Due to Missing Required Field

Symptom

I used the Archiver to export documents. Now, I'm trying to import them and the process fails.

Problem

When I try to import the previously exported documents, the content server issues an error indicating that the 'Company' metadata file is required.
Recommendation

You will need to use the content server's Configuration Manager utility to edit the 'Company' field and make it a non-required field:

- 1. Log into the content server as an administrator.
- 2. Go to the Administration page and click on the **Configuration Manager** link.

The Configuration Manager utility is started.

- 3. Open the Information Fields tab.
- 4. Select the Company metadata field from the Field Info list.
- 5. Click Edit.

The Edit Custom Info Field window is displayed.

- 6. Deselect the Require Value check box.
- 7. Click OK.
- 8. Click the Update Database Design button.
- 9. Click the Rebuild Search Index button.

You should now be able to successfully re-import the archive.

Changed Metadata Field Makes the Archiver Freeze During an Import

Symptom

Some of our product names have changed and we need to update one of the metadata fields in the affected documents. After exporting all the documents with the old product name metadata field, I then attempt to import the documents using the new product name metadata field. But, every time I try this, the Archiver processes only a portion of the total archiving task and then stops.

Problem

Once the Archiver freezes, I am unable to navigate the content server user interface and I must shut down all of the open browsers. Also, during the next five minutes after shutting down the browsers, I have no connectivity to the content server at all. After this five-minute interval, I can access the content server again.

In addition to this freezing problem, the following error message is issued: Stream error (299) - SKIPPING

Recommendation

One or more processes seem to be interrupting the import. Some possible problem solutions could be any of the following:

- Checking the Metadata Field Properties (page 6-10)
- Checking the Indexing Automatic Update Cycle (page 6-11)

Checking the Metadata Field Properties

The product name metadata field may not have been properly updated in Configuration Manager. Depending on the type of metadata field that the 'product name' is, changing the value could be the reason for the lock-up problem. Is the product name metadata field a (long) text field only or also an option list? If it is an option list, make sure that the new name value is a selection on the corresponding list.

- 1. Log into the content server as an administrator.
- 2. Go to the Administration page and click on the Configuration Manager link.

The Configuration Manager utility is started.

- 3. Open the Information Fields tab.
- 4. Select the product name metadata field from the Field Info list.
- 5. Click Edit.
- 6. The Edit Custom Info Field window is displayed.
- 7. If the **Field Type** value is Text or Long Text *AND* the **Enable Option List** check box is disabled, click **OK** or **Cancel** (this should not cause the lock-up problem).

Otherwise,

If the **Enable Option List** check box is selected, then make sure that the new product name metadata field value is included as a selection on the corresponding list:

- a. Locate the Use Option List field and click Edit.
- b. Enter the new product name metadata field value in the Option List dialog.
- c. Click OK.
- 8. Click **OK** again (on the Edit Custom Info Field window).

- 9. Click the Update Database Design button.
- 10. Click the Rebuild Search Index button.

Checking the Indexing Automatic Update Cycle

The lock-up problem may be due to the indexer's automatic update cycle. The error message indicates that the indexer is failing because it loses connectivity. Every five minutes, the indexer executes an automatic update cycle and could somehow be grabbing the index file and locking it. If so, it might be useful to disable the indexer's automatic update cycle while you run the import.

- 1. Log into the content server as an administrator.
- 2. Go to the Administration page and click on the **Repository Manager** link.

The Repository Manager utility is started.

- 3. Open the **Indexer** tab.
- 4. Click the **Configure** button in the Automatic Update Cycle section.

The Automatic Update Cycle dialog displays.

- 5. Deselect the Indexer Auto Updates check box.
- 6. Click OK.



Note: Be sure to reactivate the automatic update cycle after completing the import. Otherwise, the server will no longer automatically update the index database, which could adversely impact future search results.

EXPORTING ISSUES

This section covers the following topics:

- Total Export Possible with Blank Export Query (page 6-12)
- New Check-Ins and Batch File Transfers (page 6-12)
- Exporting User Attributes (page 6-13)
- Folder Archive Export Doesn't Work If Collections Table Has Many Records (page 6-13)

Total Export Possible with Blank Export Query

Question

If I do not create an export query to define the content items to export, will the entire contents of my content server be exported?

Recommendation

Yes, test exports have confirmed that leaving the Export Query section blank (not defining an export query) will ensure that the content server contents are completely exported.

New Check-Ins and Batch File Transfers

Question

If I check some documents into the content server after I have initiated a large export (but before it completes), will these documents be included in the export? Or, does the Archiver read the timestamp information and determine that the new files are more recent than those originally allocated for the export and not include them? Also, what happens to the archive export if the connection between the servers is interrupted or lost during the export?

Recommendation

When the export is initiated, Archiver runs a query on the system to build a list of the documents that are to be exported. This information is cached and used to build the export archive. Therefore, any new documents that are checked in during the export process will not be included even if they match the export query definition.

If the connection between servers is disrupted, the export process on the source server continues but the transfer to the target server stops. The source server accumulates a number of batch files. While waiting to transfer these files, the source server continues to ping the target server for a connection at regular interval. When the connection is reestablished, the accumulated batch files are transferred to the target server.

If you have used an automated (replicated) transfer, the batch files and their associated content files are removed from the source content server. If you have used a manual transfer, the batch files and their associated content files remain in the source content server.

Exporting User Attributes

Question

How can I export users in an archive?

Recommendation

You can export a *users.hda* file, which contains the user attributes from the Users database table, as follows:

- 1. Log into the content server as an administrator.
- 2. Go to the Administration page and click on the Archiver link.

The Archiver utility is started.

- 3. Open the Export Data tab.
- 4. Click Edit in the Additional Data section.

The Edit Additional Data dialog is displayed.

- 5. Select the **Export User Configuration Information** check box.
- 6. Click OK.

Folder Archive Export Doesn't Work If Collections Table Has Many Records

Symptom

I use the folder archive export feature to move my website hierarchy created by Site Studio. Initially, I can export folders by using the Virtual Folder Administration Configuration page without any problem. However, as my website grows, this function does not work anymore. The following errors are issued during the export procedure:

```
Error <timestamp> Event generated by user '<user>'
at host '<host_name>'. Referred to by
http://<host>/intradoc-cgi/nph-idc_cgi.exe?IdcService=
COLLECTION_GET_ADMIN_CONFIG. Unable to retrieve content. Unable to execute
service method 'loadCollectionArchive'. (System Error: Unknown error.)
```

Error <timestamp> IdcAdmin: Event generated by user '<user>' at host '<host>'. Unable to obtain the console output. Unable to execute the service 'GET_SERVER_OUTPUT' on Content Server 'contribution'. Unable to receive request. Response from host has been interrupted. Read timed out.

There is also an out-of-memory error in the content server output console:

```
<timestamp> SystemDatabase#Thread-13: SELECT * FROM Collections, ColMeta
WHERE Collections.dCollectionID=ColMeta.dCollectionID AND
dParentCollectionID=564
java.lang.OutOfMemoryError
Reporting error on thread Thread-13 occurring at <timestamp>.
java.lang.OutOfMemoryError
java.lang.OutOfMemoryError
```

Problem

Depending on the size of the folder hierarchy that is being exported as an archive file, the default heap size value for the Java Virtual Machine (JVM) may not be adequate.

Recommendation

In the content server's <*CS_Instance_Dir*>/bin/intradoc.cfg file, comment out the JvmCommandLine setting and increase heap size to 512 M:

```
#JvmCommandLine=/home/contribution/shared/os/<os>/j2sdk1.4.1/bin/java -cp
$CLASSPATH $STARTUPCLASS
JAVA OPTIONS= -Xmx512m
```

After restarting the content server, the archive export function should work correctly again.

TRANSFER ISSUES

This section covers the following topics:

- Transfer Stopped When Target Locked Up (page 6-15)
- Aborting / Deleting a Running Transfer (page 6-16)
- Verifying the Integrity of Transferred Files (page 6-17)
- Transfer Process Is Not Working (page 6-18)

Transfer Stopped When Target Locked Up

Symptom

The automated transfer function stopped when the target server locked up.

Problem

After restarting the target server, the log file listed an error message stating that there was a problem with a security group and that this prevented the import on the target server.

Recommendation

In this case, obviously the security group problem on the target server must be corrected before the transfer can proceed. Two additional procedures to perform that can help include:

- Verifying and Testing the Outgoing Provider (page 6-15)
- Restarting the Content Server (page 6-15)

Verifying and Testing the Outgoing Provider

Verifying and testing the outgoing provider ensures that it is set up and working properly:

- 1. Log into the *source* content server as an administrator.
- 2. Go to the Administration page and click on the **Providers** link.

The Providers page is displayed.

3. Click the Info link of the appropriate outgoing provider.

The Outgoing Provider Information page is displayed.

- 4. Verify the information.
- 5. Return to the Providers page and click the **Test** link corresponding to the outgoing provider.

Restarting the Content Server

In some cases, after problems have been corrected on either the source or the target server, the source server may stop transferring or possibly the automation function no longer works. In either case, restarting the content server should resolve the problem.

Aborting / Deleting a Running Transfer

Question

I accidentally started transferring an excessively large file to the production content server. What is the most efficient way to stop the transfer process while it is running?

Recommendation

There are several methods to abort or delete a transfer, including:

- Disabling the Outgoing Provider (page 6-16)
- Deleting a Transfer from the Transfer To Tab (page 6-16)
- Deleting an Automated Transfer from the Automation for <Instance> Screen (page 6-17)

Disabling the Outgoing Provider

The fastest method to abort a running transfer is to disable the source server's outgoing provider:

- 1. Log into the source content server as an administrator.
- 2. Go to the Administration page and click on the **Providers** link.

The Providers page is displayed.

3. Click the **Info** link of the appropriate outgoing provider.

The Outgoing Provider Information page is displayed.

4. Click the **Disable** button.

Deleting a Transfer from the Transfer To Tab

To delete a transfer from the Transfer To tab, complete the following steps:

- 1. Log into the source content server as an administrator.
- 2. Go to the Administration page and click on the Archiver link.

The Archiver utility is started.

- 3. Select Options—Open Archive Collection.
- 4. Select the applicable collection from the list.

- 5. Click Open.
- 6. On the Archiver window, select the source archive in the Current Archives list.
- 7. Open the **Transfer To** tab.
- 8. Click **Remove** in the Transfer Destination section.
- 9. You are prompted to confirm the action.
- 10. Click Yes.

Deleting an Automated Transfer from the Automation for <Instance> Screen

To delete an automated transfer from the Automation for *<Instance>* screen, complete the following steps:

- 1. Log into the source content server as an administrator.
- 2. Go to the Administration page and click on the Archiver link.

The Archiver utility is started.

3. Select Options—View Automation For <Instance>.

The Automation For <Instance> window is displayed.

- 4. Open the **Transfers** tab.
- 5. Select the automated transfer to delete.
- 6. Click Remove.

The automated transfer is removed from the list.

Verifying the Integrity of Transferred Files

Question

What is the best approach to verify the integrity of the files that have been transferred between two servers? Obviously, the documents in the target content server instance should be identical to those in the source instance. I need to ensure that all documents were in fact transferred and if some were not transferred, I must determine which ones failed to transfer.

Recommendation

To ensure that the transferred documents are identical to those on the source server, two items can easily be checked.

The Revisions table:

Specifically, match the contents of the dDocName and dRevLabel columns on both instances and verify the accuracy or discrepancies between them.

The file system:

Check the native file repository (*CS_Instance_Dir>/vault/<content_type>*) and web-viewable file repository

(*<CS_Instance_Dir>/weblayout/groups/public/documents/<content_type>*) on each server and verify the accuracy or discrepancies between them.

Transfer Process Is Not Working

Symptom

The transfer process is not setting up properly.

Recommendation

If the transfer process is not functioning correctly, check the outgoing provider on the source server and ensure that the information is correct. In particular, make sure that the server host name is correct and matches the HTTP server address.

To verify the server host name on the source server, complete the following steps:

- 1. Start the System Properties utility.
- 2. Open the Internet tab.
- 3. Note the HTTP server address setting.
- 4. Go to the Administration page and click on the **Providers** link.

The Providers page is displayed.

5. Click the Info link of the appropriate outgoing provider.

The Outgoing Provider Information page is displayed.

6. Check the server host name and make sure it corresponds exactly to the HTTP server address setting in System Properties.

- 7. If the server host name setting is different than the HTTP server address, click the **Edit** button.
- 8. Modify the Server Host Name setting as necessary.
- 9. Click Update.
- 10. Restart the content server.

WEBDAV ISSUES

This section covers the following topics:

Archiver Error With WebDAV and Content Server (page 6-19)

Archiver Error With WebDAV and Content Server

Symptom

I am using both WebDav and the content server to check in documents. I would like to be able to view all of the checked-in documents through the WebDAV interface but I can see only the documents that were physically checked in through WebDAV. Currently, I am using the Archiver utility to import the files and have configured a value map using the xCollectionID field. This field is currently set to zero or null and I'm trying to update the value to 182.

Problem

For each file that I try to import, the following error is logged in the Archiver log: Unable to load collection mappings, too many items.

Recommendation

The error message indicates that you are exceeding the maximum number of items per folder. There is a configurable limit to how many items a folder can hold. To check and update your current limit settings, complete the following steps:

- 1. Log into the source content server as an administrator.
- 2. Go to the Administration page and click on the Folder Configuration link.

The Virtual Folder Administration Configuration page is displayed.

- 3. Check the following two limit settings:
 - Maximum Folders Per Virtual Folder
 - Maximum Content Per Virtual Folder
- 4. Increase both of these limit settings to accommodate your import requirements.

or,

Implement an 'infinite' limit setting by removing the limit setting.

- 5. Click Update.
- 6. Restart the content server.



Note: You should be aware, however, that the system performance will decrease with the increased number of items that you have in a folder.

REPLICATION ISSUES

This section covers the following topics:

Stopping the Automatic Import Function (page 6-20)

Stopping the Automatic Import Function

Question

How can I stop the automatic import function?

Recommendation

When content meets the specified criteria, the automatic importer is, by default, configured to automatically perform an import every five minutes. However, there are two ways to disable the automatic import function:

- Unregistering an Importer from the Replication Tab (page 6-21)
- Deleting a Registered Importer from the Automation for <Instance> Screen (page 6-21)

Unregistering an Importer from the Replication Tab

To unregister an importer from the Replication tab, complete the following steps:

- 1. Log into the source content server as an administrator.
- 2. Go to the Administration page and click on the **Archiver** link.

The Archiver utility is started.

- 3. Select the archive in the Current Archives list.
- 4. Open the **Replication** tab.
- 5. Click Unregister.

The automatic import function is disabled from the selected archive.

Deleting a Registered Importer from the Automation for </br><Instance> Screen

To delete a registered importer from the Automation for *<Instance>* screen, complete the following steps:

- 1. Log into the source content server as an administrator.
- 2. Go to the Administration page and click on the Archiver link.

The Archiver utility is started.

3. Select Options—View Automation For < Instance>.

The Automation For Instance screen is displayed.

- 4. Open the **Importers** tab.
- 5. Select the registered importer to delete.
- 6. Click Remove.

The registered importer is removed from the list.

ORACLE-SPECIFIC ISSUES

This section covers the following topics:

✤ Allotted Tablespace Exceeded (page 6-22)

Allotted Tablespace Exceeded

Symptom

I can't transfer files. Every time I try to transfer files, I get 'max extents' error messages.

Problem

The following error messages (or similar) are issued:

```
IdcApp: Unable to execute query '<query_name>'. Error: ORA-01631: max #
extents (50) reached in table <table_name>.
ORA-01631 max # extents (<text_string>) reached in table <table_name>.
```

Recommendation

When the content server creates its database tablespace, it only allocates 50 extents. As the database grows and is re-indexed, it uses more space (extents). Eventually, the 50 extents limit is exceeded. At some point in the transfer, one of your files tried to extend past the 'max extents' limit. In this case, try implementing one or more of the following solutions:

- Look for weblayout queries that are excessively large, eliminate them, and retry your transfer.
- Perhaps an Oracle content server user doesn't have the right permission grants (resource and connect) to the content server schema. That user must have the temporary tablespace and default tablespace set to the content server defaults.
- If the system 'max extents' limit is less than the system maximum, you will need to increase the number of extents that are available. Refer to your Oracle documentation or ask your database administrator for the appropriate Oracle SQL command to increase the tablespace extents.
- You can optionally choose to re-create the database using larger initial, next or percent to grow parameters for the tablespaces. In this case, it is advisable to set the initial extents and next extents to 1Mb. Set the percent to grow parameter (PCTINCREASE) to 0% to allow the tables to automatically grow on an as-needed basis.

MISCELLANEOUS ISSUES

This section covers the following topics:

- Archiving Doesn't Work With Shared File System (page 6-23)
- Archiving Doesn't Work Over Outgoing Provider (page 6-23)

Archiving Doesn't Work With Shared File System

Symptom

I am trying to transfer between two content server s with access to a shared file system but it isn't working.

Recommendation

When transferring between content servers on a shared file system, the mapped or mounted drive must be available to both content servers. This means that the computers must be on and logged in as a user who has system access to both content servers. Make sure that all of the following conditions are met:

- ✤ Both computers are turned on.
- Both computers are logged in as a user with system access to both content server file systems.
- The shared drive has been properly mapped or mounted so the content server can "see" it. Having network access to the computer is not sufficient.

Archiving Doesn't Work Over Outgoing Provider

Symptom

I am trying to transfer between two content server s over an outgoing provider but it isn't working.

Recommendation

The content server that has an outgoing provider set up is considered the 'local' server, and the target content server for the outgoing provider is considered to 'proxied' server. Files are always transferred in the direction of the outgoing provider, from the local (source) instance to the proxied (target) instance.

It is possible that when the outgoing provider was added and defined for the source content server, the 'Proxied' check box was selected. However, because the relative web root is the same for both content servers, the outgoing provider is confused. The 'Proxied' check box should be selected only if the target content server was installed as an actual proxy of the local (master) content server. This server option should not be selected if the relative web root is the same for both content server.

WORKFLOW ISSUES

OVERVIEW

This section provides solutions to several common workflow issues. Please attempt the solutions recommended in this guide before Support.

- ♦ Workflow Item Stuck in EDIT or GENWWW Status (page 7-1)
- Workflow Item Stuck in REVIEW Status (page 7-3)
- ♦ Workflow Item Entered in Wrong Workflow (page 7-4)

WORKFLOW ITEM STUCK IN EDIT OR GENWWW STATUS

Symptom

A content item in a workflow is in EDIT or GENWWW status, and no reviewers were notified by e-mail that action is required.

Problem

Inbound Refinery failed to convert the file properly.

Recommendation

- View the content item's status in Repository Manager or the Content Items for Workflow page (accessed from Content Manager—Active Workflows— *<Workflow_Name>*). For more details on the conversion failure, display the Content Information page.
- 2. Determine why the conversion failed:
 - If the problem was with Inbound Refinery or a conversion add-on product:
 - a. Correct the problem so that the file will convert properly.
 - b. Resubmit the file for conversion from Repository Manager or Content Information page. The content item will continue through the workflow.
 - If the problem was with a file in a criteria workflow and the content item is the only item in the workflow:
 - a. Save a copy of the native file.
 - b. Disable the workflow to release the content item.
 - c. Delete the released revision using Repository Manager or the Content Information page.
 - d. Correct the problem so that the file will convert properly.
 - e. Check in the content item again so that it goes through the entire workflow process.
 - If the problem was with a file in a criteria workflow and there are multiple content items in the workflow:
 - a. Save a copy of the native file.
 - b. Delete the "stuck" revision using Repository Manager. (You could disable the workflow, but all content items in the workflow would be released.)
 - c. Correct the problem so that the file will convert properly.
 - d. Check in the content item again so that it goes through the entire workflow process.
 - If the problem was with a file in a basic workflow and the content item is the only item in the workflow:
 - a. Save a copy of the native file.
 - b. Cancel the workflow to delete the revision from the system.
 - c. Correct the problem so that the file will convert properly.
 - d. Contribute the content item again to the Basic workflow so that it goes through the entire workflow process.

- If the problem was with a file in a basic workflow and there are multiple content items in the workflow:
 - a. Save a copy of the native file.
 - b. Delete the "stuck" revision using Repository Manager. (You could cancel the workflow, but all content items in the workflow would be deleted from the system.)
 - c. Correct the problem so that the file will convert properly.
 - d. Contribute the content item again to the Basic workflow so that it goes through the entire workflow process.

WORKFLOW ITEM STUCK IN REVIEW STATUS

Symptom

A content item in a workflow is in REVIEW status, and the minimum number of reviewers have approved it, but the revision is not moving to the next step.

Problem

The content item does not meet the exit criteria of the workflow.

Recommendation

- 1. Check the exit condition definition in the workflow step to see why the document does not meet the exit criteria.
- 2. Determine if the file is finished interacting with an external process. Resolve any problems with the external process to see if the exit condition can be met.
- 3. If there is an error in the exit condition and/or external processing, you can check out the document and check it back in with the correct metadata to meet the exit condition.

WORKFLOW ITEM ENTERED IN WRONG WORKFLOW

Symptom

A content item entered the wrong criteria workflow.

Problem

Two or more criteria workflows have the same criteria defined. The content item matched this criteria, so it entered the first matching workflow in the workflow list.

Recommendation

Redefine the criteria so that each workflow has a unique combination of the security group and metadata field value. If necessary, use jumps and sub-workflows to define additional criteria within the main workflow.

C h a p t e r

INTERNATIONAL ISSUES

OVERVIEW

This section provides solutions to several common internationalization issues. Please attempt the solutions recommended in this guide before Support.

This section covers the following topics:

- Display Issues (page 8-2)
- ✤ Login Issues (page 8-9)



Important: A very useful document in this respect is *Using Content Server in International Environments*. This document is part of the Oracle Content Server product documentation set, which can be accessed through the online Help. It is also available as a PDF file called *internat_environments_XXen.pdf* (where 'XX' is the version number), which can be found in the *CS_Instance_Dir>/weblayout/help/documentation/ integrator* directory or on the Content Server distribution CD labeled 'Documentation' (in the */integrator* directory).

DISPLAY ISSUES

This section covers the following topics:

- Characters Are Displayed on Screen as Question Marks (page 8-2)
- Characters Are Displayed on Screen as Boxes (page 8-3)
- Characters Are Displayed on Screen as Garbage (page 8-4)
- Some Multi-Byte Characters Are Displayed Slightly Differently Than Expected (page 8-4)
- Text in Forms With Re-used Data Is Displayed in &#xxxx; Form (page 8-5)
- Characters Are Not Displayed Entirely Correctly in Internet Explorer 5.x (page 8-6)
- Some Bookmarked URLs Don't Work in Internet Explorer (page 8-6)
- PDF Files Cannot Be Viewed in Internet Explorer (page 8-7)
- Web-Viewable Files Cannot Be Viewed (page 8-8)
- Some Thumbnails Are Not Displayed (page 8-8)
- Sort Order is Not Always Consistent (page 8-9)

Characters Are Displayed on Screen as Question Marks

Symptom

Characters are displayed on screen as question marks.

Problem

If characters are displayed on screen as question marks, this is typically caused by either of two things:

The most common cause is an incorrect language encoding of the operating system. For example, if you are running English Windows without the Japanese language pack and you are using a Japanese content server, you will see question marks on screen rather than Japanese characters. This is because Windows cannot properly handle the Japanese characters without the appropriate language pack and, as a result, does not display the characters correctly. Another, less common cause is corrupted data in the database. For example, this will occur if you specify a multi-byte content title (for example, Japanese or Korean) on an English content server and the database has not been set up to use Unicode. The multi-byte title is then written to the database in non-Unicode form and cannot be displayed correctly.

Recommendation

To prevent characters from being displayed on screen as question marks, always make sure that the language encoding of the operating system is compatible with the language(s) you want to use. Also, if you anticipate languages of different file encodings to be used in the content server environment (for example, a combination of European and Asian languages), it is wise to set up the database to use Unicode encoding.



Note: Refer to the *Using Content Server in International Environments* guide for further details on language encoding of the operating system and database encoding.

Characters Are Displayed on Screen as Boxes

Symptom

Characters are displayed on screen as boxes.

Problem

If characters are displayed on screen as boxes, this is typically caused by either of two things:

- The fonts required to display the characters correctly are incompatible, corrupted, or missing. If a content item uses very specific, non-standard fonts, there may be font compatibility issues which prevent the text from being displayed correctly.
- The web browser does not have the required language support enabled. For example, a Japanese web page will only display correctly in a browser if the browser has been set up to support that language (even if the computer has all required Japanese fonts).

Recommendation

To prevent characters from being displayed on screen as boxes, always make sure that all required fonts are available on the computer (especially if your content uses special fonts). Also, make sure that your browser is set up to support all required languages or set up the web server to use UTF-8 encoding.

 \bigcirc

Note: Refer to the *Using Content Server in International Environments* guide for further details on language support in browsers.

Characters Are Displayed on Screen as Garbage

Symptom

Characters are displayed on screen as garbage.

Problem

If characters are displayed on screen as garbage, this is typically caused by a mismatch of encodings. For characters to be displayed correctly, the language encoding of the operating system, the file encoding of the content server, and the language encoding of the database all need to be compatible. If they are not, characters are not handled correctly and and are then displayed on screen as garbage.

Recommendation

To prevent characters from being displayed on screen as garbage, always make sure that all encodings used in the content management system are compatible with each other. For further details refer to the encoding compatibility matrix in the *Using Content Server in International Environments* guide.

Some Multi-Byte Characters Are Displayed Slightly Differently Than Expected

Symptom

Some characters are displayed slightly differently than expected. For example, an em-dash (--) may be displayed as a somewhat shorter dash.

Problem

In setups with mixed encoding (for example, UTF-8 encoding for the content server and native encoding for the database), the data needs to be converted back and forth between one encoding and the other. A "round trip" from native to UTF-8 and back again for certain characters in multi-byte languages does not return to the exact same character.

This can cause problems anywhere data is encoded either as Unicode or UTF-8. In particular, the content server uses Java's Unicode strings, so this problem will always occur if the content being submitted is in its original native encoding.

Fortunately, the characters involved are usually fairly obscure and the changes are typically fairly small (such as changing a long em-dash to a slightly shorter dash). End users do need to be aware that there are a few characters that do not work very well when submitted as part of a browser post or when used in file names where you use WebDAV to check in content. It can also affect certain types of searches.

Recommendation

You can prevent problems like this by using only one type of encoding throughout the content management system (for example, UTF-8 and Unicode).

Text in Forms With Re-used Data Is Displayed in &#xxxx; Form

Symptom

In Internet Explorer, text in form fields that contain "re-used" data—for example, prefilled data fields on the "check-in similar" form or document update form—are displayed in Unicode form (&#xxxx;).

Problem

If you type a character into an Internet Explorer form and the current encoding of that page does not support that character, Internet Explorer does not stop you from actually entering the character. Instead, when you perform the post, it will encode the character in its Unicode encoding (&#xxxx;). This works fine as long as you view the same content inside an HTML page for a browser. However, problems arise if you try to use the content in a JavaScript construction, try to repurpose the values into other formats (such as text files), try to search for it, or try to make sure there is enough room in a database field to hold it.

Recommendation

You can prevent problems like this by using Unicode encoding in the database.

Characters Are Not Displayed Entirely Correctly in Internet Explorer 5.x

Symptom

Certain characters are not displayed entirely correctly in Internet Explorer 5.x.

Problem

If you use the ISO-8859-1 encoding in Internet Explorer, then IE 5.x has a slightly smaller set of supported characters than ISO-8859-1 in IE 6.0. If you submit content using IE 6.0 form posts, then you may run into some issues viewing or updating the content in IE 5.x.

Recommendation

To make sure all characters are displayed correctly, it is recommended that you upgrade from Internet Explorer 5.x to 6.0.

Some Bookmarked URLs Don't Work in Internet Explorer

Symptom

Some bookmarked URLs in Internet Explorer do not work, even though the links themselves are valid.

Problem

If your content server has been configured to use UTF-8 file encoding and you have a URL that has UTF-8 multi-byte sequences in it that have not been URL-escaped (%xx escape sequence), then the URL will work in Internet Explorer if clicked inside UTF-8 pages, but it will not work when bookmarked. Please note that this is not a problem with Mozilla-based browsers.

Recommendation

If you set the XmlEncodingMode=full parameter in *CS_Instance_Dir*/*config/ config.cfg*, it will reduce the consequences of this problem because many more parts of the content server CGI URL parameters will be fully 7-bit encoded. However, some content server URL constructions will not URL-escape the dDocName parameters (document name).

If you wish to avoid this problem with a UTF-8 content server altogether, then you will need to either do some customization of the content server pages or make sure your all dDocName values are clean 7-bit.

PDF Files Cannot Be Viewed in Internet Explorer

Symptom

End users cannot view PDF files in their web browser. Instead, a blank page is shown and Acrobat Reader does not start at all.

Problem

The most obvious problem is, of course, that Adobe Acrobat or Acrobat Reader is not installed on the end user computer. If that is not the issue, there may be an encoding issue. If you serve up a UTF-8 encoded page with a UTF-8 encoded URL but do not encode all non 7-bit ASCII values into %xx sequences in your URLs, then Acrobat or Acrobat Reader will send the URL in the encoding of the user computer. This may result in an encoding mismatch, which prevents the PDF file from being displayed in the browser.

Recommendation

If you set the XmlEncodingMode=full parameter in *CS_Instance_Dir/config/ config.cfg*, then the content server will always fully encode the static web URLs (except for URLs generated by Dynamic Converter to images and multiple pages).



Note: Please note that this is no longer an issue in Oracle Content Server 7.1.1, which fully encodes the URLs to PDF documents if the page is in UTF-8 and the browser is Internet Explorer.

Web-Viewable Files Cannot Be Viewed

Symptom

End users cannot retrieve web-viewable renditions of content that is checked into the content server. Typically, a blank page is displayed and no error messages are reported.

Problem

IIS 5.0, which is part of Microsoft Windows 2000 Server, cannot deliver UTF-8 URLs that have characters in them that are not supported by the native desktop encoding. This means that if you use UTF-8 as the file encoding for a content server (FileEncoding=UTF8 in *CS_Instance_Dir/config.cfg*) on an English system, you cannot use Japanese, Korean, or Chinese characters for your security groups, content types, or document titles (since these are part of a content item's URL). Please note that this is not an issue with IIS 6.0 (part of Microsoft Windows Server 2003) or Apache 2.0.

Recommendation

The easiest way to prevent this problem is to not use multi-byte characters in security groups, content types, or document titles. Alternatively, you could also use IIS 6.0 or Apache 2.0 as the web server.

Some Thumbnails Are Not Displayed

Symptom

Some thumbnails are not displayed in the content server user interface.

Problem

Some of the conversion solutions (such as thumbnails) do not support UTF-8 encoded paths and will fail for web-viewable paths that have non 7-bit byte sequences in them (i.e., characters other than basic ASCII). This means that if a security group, content types, or document title (all are part of a content item's URL) contains, for example, Japanese characters, no thumbnails will be created for the content item.

Recommendation

If you are using Inbound Refinery with UTF-8 as the content server file encoding, then you will need to enable the Copy Conversion Copy option (on Inbound Refinery's Local Configuration—General tab). This will allow Inbound Refinery to do the conversion locally and then copy the result to the web-viewable path.

Alternatively, you can also make sure that no non 7-bit byte sequences (i.e., characters other than basic ASCII) are used in security groups, content types, or document titles.

Sort Order is Not Always Consistent

Symptom

The sort order used on screen (for example, in search results pages) is not always consistent.

Problem

Even in English there are several different sort orders to choose from. If you are piecing together sorted data from many iterative requests to a database and your sorting choice is not in alignment with the database, problems may arise. This issue is a prevalent problem with the Select User applet.

Recommendation

The only fix is to keep educating the Select User applet about all the different sort orders used by databases and make sure that the applet configuration parameters agree with the current database sort order. For example, SQL Server's default sort order is a Dictionary sort order and Oracle's is a raw ASCII value sort order.

LOGIN ISSUES

This section covers the following topics:

- Users With Accented User Name And/Or Password Cannot Log In (page 8-10)
- Some Users Cannot Log In After Content Server File Encoding Is Changed (page 8-10)

Users With Accented User Name And/Or Password Cannot Log In

Symptom

End users that have special characters in their user name and/or password (for example, 'kmüller' or 'mdupré') cannot log into the content server.

Problem

If you are using the HTTP Basic protocol to log in and if the native encoding of the user computer does not match that of the content server, then non 7-bit ASCII characters (i.e., extended ASCII or UTF-8) cannot be used in user names and passwords. This is because the HTTP protocol does not have a mechanism for specifying the encoding used in user names and passwords (or in any other HTTP request header for that matter). If you use NTLM for authenticating Internet Explorer to IIS, then the actual authentication is done in wide characters, but when you ask for the REMOTE_USER variable when inside IIS, the encoding will be in the default encoding of the computer that is running IIS. This is a limitation in the original design of the web communication protocols.

Recommendation

If you want to use special characters in user names and passwords, you could use the CookieLoginPlugin component and do logins using cookies. This should eliminate any encoding issues with logins.

Some Users Cannot Log In After Content Server File Encoding Is Changed

Symptom

Some end users cannot log into the content server after the content server file encoding is changed. Their passwords are rejected.

Problem

If you change the file encoding of the content server after entering new users into the system, then any user name that includes non 7-bit ASCII characters (i.e., characters other than basic ASCII) will have their password invalidated. An administrator will have to

re-enter the passwords. This is because the SHA1 hashing algorithm used by Oracle Content Server is based on a byte sequence produced using the current encoding of the content server.

Recommendation

To avoid this problem, users should preferably only use basic ASCII characters in their passwords. To correct the problem once it has occurred, an administrator needs to reset the user passwords.



UNDERSTANDING THE DATABASE

OVERVIEW

Content Server solutions rely on database tables as pillars of the architecture. Various databases are supported in the Oracle environment.



Caution: Constraints are not in the database, so updates to the database are not safe.

This section covers these topics:

- Data Dictionary Terminology (page A-1)
- Data Dictionary (page A-3)

DATA DICTIONARY TERMINOLOGY

The following terminology is important with regard to the data dictionary:

Table Name

The name of the collective rows and columns used to store data in the database. Types of tables documented in this data dictionary are public tables and views.

Column Name

The representation of the modeled object. For example, many column names have corresponding metadata values in Oracle Content Server.

Primary Key

Unique identifier for the rows in a table. All of the columns should contain data specific to the primary key.

Type Name

The data type:

- char: A character data type that is limited to 1 character and is sorted by the code page used to store non-Unicode character data: Latin1.
- varchar: A character data type that holds variable-length non-Unicode data and is sorted by the code page used to store non-Unicode character data: Latin1.
- int: A numeric data type that holds whole numbers. For Oracle tables, the length requirement for int is 4, and the range is -2^{31} to $+2^{31}$.
- **smallint**: A numeric data type that holds whole numbers. For Oracle tables, the length requirement for smallint is 2, and the range is -2^{15} to $+2^{15}$.
- tinyint: A numeric data type that holds whole numbers. For Oracle tables, the length requirement for tinyint is 1, and the range is 0 to 255.

Length

The maximum number of characters needed to represent data.



Note: The database display values for *int* and *date* fields are not used for Content Server display.

Nullable

Identifies whether the column permits null data (i.e., data with no explicitly assigned value).

Constraints

These are notes identifying constraints for specific columns and/or tables. All tables with primary keys have primary key constraints.

DATA DICTIONARY

This section describes the following database tables. All are public unless specified otherwise:

- ✤ Alias (page A-4)
- AliasUser (page A-5)
- ArchiveChangedRows (page A-6)
- ArchiveHistory (page A-7)
- Collaborations (page A-9)
- Config (page A-10)
- Counters (page A-11)
- DatedCaches (page A-12)
- DeletedRows (page A-13)
- DocFormats (page A-14)
- DocMeta (page A-15)
- DocMetaDefinition (page A-16)
- DocTypes (page A-18)
- DocumentAccounts (page A-19)
- DocumentHistory (page A-20)
- Documents (page A-22)
- ExtensionFormatMap (page A-24)
- HtmlConversions (page A-25)
- HtmlConversionSums (page A-26)
- OptionsList (page A-27)
- ProblemReports (page A-28)
- ProjectDocuments (page A-30)
- RegisteredProjects (page A-32)
- Revisions (page A-33)
- RoleDefinition (page A-38)

- SecurityGroups (page A-39)
- ✤ Subscription (page A-40)
- ✤ Users (page A-42)
- UserSecurityAttributes (page A-45)
- WorkflowAliases (page A-46)
- WorkflowCriteria (page A-47)
- WorkflowDocAttributes (page A-48)
- WorkflowDocuments (page A-49)
- WorkflowHistory (page A-51)
- ✤ Workflows (page A-54)
- WorkflowStates (page A-56)
- WorkflowSteps (page A-57)
- sysconstraints (page A-58)
- syssegments (page A-59)

Alias

This table defines an alias name and a description of the alias. Users are placed in the alias through the AliasUser table.

Alias				
Column Name	Primary Key	Type Name	Length	Nullable
dAlias	Х	varchar	30	
	Column Description: Unique identifier for the alias.			
dAliasDescription		varchar	50	Х
	Column Description: Description of the alias conter or usage.			
Constraints: primary key				
AliasUser

This table assigns a user to an alias. Aliases are defined in the Alias table.

AliasUser				
Column Name	Primary Key	Type Name	Length	Nullable
dAlias	Х	varchar	30	
	Column Description: Unique identifier for the alias.			
dUserName	Х	varchar	50	
	Column Descrip also potentially b	tion: User inside o be an alias name as	f alias (us s well).	ser could
Constraints: primary key				

ArchiveChangedRows

ArchiveChangedRows				
Column Name	Primary Key	Type Name	Length	Nullable
dRowID	Х	varchar	16	
	Column Descrip	tion: Unique ident	ifier for tl	ne row.
dTable		varchar	30	
	Column Descrip	tion: The table tha	t has beer	h changed.
dPrimaryKeys		varchar	255	
	Column Description: The primary key is a unique key for each table. Each primary key is constrained so that no two values are equal.			
dPKColumns		varchar	150	
	Column Description: The table columns that comprise the primary key for each table.			
dPKTypes		varchar	50	Х
	Column Descrip	tion: The types of	primary k	keys.
dChangeDate		datetime	8	
	Column Descript modified.	tion: Date that the	table was	edited or
dActionDate		datetime	8	
	Column Description: The date the action occurred.			
dSourceID		varchar	50	x
	Column Description: The content ID.			
Constraints: primary key				

ArchiveHistory

This table provides a history of archived assets within the system.

Archive History				
Column Name	Primary Key	Type Name	Length	Nullable
dActionMillis	Х	biginit	8	
	Column Descrip seconds, millised	tion: Unique ident cond, and date.	ifier base	d on the
dActionDate	Х	datetime	8	
	Column Descrip	tion: Date, accurat	te to the n	ninute.
dID	Х	biginit	8	
	Column Description: Repository ID used throughout the system to uniquely identify the document.			
dDocName		varchar	30	Х
	Column Description: Unique name of a document. Revisions of a document have the same name.			ument. ne.
dDocType		varchar	30	Х
	Column Descript type from the Do	tion: Unique ident ocTypes table.	ifier of a o	document
dDocTitle		varchar	80	Х
	Column Descrip	tion: Descriptive t	itle of the	revision.
dRevClassID		bigint	8	Х
	Column Description: Unique identifier of the document up to the revision. Revisions of a documen have the same dRevClassID.			

Archive History				
dRevisionID		bigint	8	Х
	Column Descript increments for ea internal use only.	tion: Revision num ach revision of a d	iber. Start ocument.	s at 1 and For
dRevLabel		varchar	10	Х
	Column Descript duplicate and hav This is the extern	tion: Revision labe ve its own increme nal revision identif	el. Can be enting alg ier.	a orithm.
dSecurityGroup		varchar	30	Х
	Column Description: The security group classification assigned to the document revision.			
dArchiveName		varchar	100	Х
	Column Descript includes the colle	tion: Archive locat ection name and an	tion ident	ifier that ne.
dBatchFile		varchar	50	Х
	Column Descript includes the expo	tion: Batch file loc ort directory and b	ation iden atch file r	tifier that name.
dDocAccount		varchar	30	Х
	Column Description: Account that is assigned to the revision. used to control security for the revision.			ed to the sion.
Constraints: primary key				

Collaborations

Collaborations				
Column Name	Primary Key	Type Name	Length	Nullable
dClbraName	Х	varchar	50	
	Column Descrip	tion: The name of	the collab	ooration.
dClbraDescription	varchar		80	Х
	Column Description	tion: The descripti	on of the	
dClbraType		varchar	30	Х
	Column Description: The type of collaboration, similar to dUserType or dDocType.			
dClbraCreateDate		datetime	8	Х
	Column Description: The origination date of the collaboration.			f the
dClbraCreatedBy		varchar	50	Х
	Column Descrip originated the co	tion: The name of llaboration.	the user v	vho
dClbraChangeDate		datetime	8	Х
	Column Description: The date the collaboration was edited or modified.			
dClbraChangedBy		varchar	50	Х
	Column Description: The name of the user who edite or modified the collaboration.			ho edited
Constraints: primary key				

Config

Config				
Column Name	Primary Key	Type Name	Length	Nullable
dSection	Х	varchar	50	
	Column Descript	tion: Category for	the config	g entry.
dName	Х	varchar	50	
	Column Description: Name of this config entry.			
dVersion	Х	varchar	50	
	Column Description: Content Server version that applies to this config entry.			
dValue		varchar	200	Х
	Column Description: Counter that tracks use iterations of this config entry.			2
Constraints: primary key				

Counters

Counters				
Column Name	Primary Key	Type Name	Length	Nullable
dCounterName	Х	varchar	50	
	Column Description: Name of the counter.			
dNextIndex		bigint	8	
	Column Descript the counter.	tion: Next availabl	e index ni	umber for
Constraints: primary key				

DatedCaches

DatedCaches				
Column Name	Primary Key	Type Name	Length	Nullable
dDocName	Х	varchar	30	
	Column Descrip Revisions of a de	tion: Unique name ocument have the	e of a docu same nam	ument. ie.
dCacheAction	Х	varchar	10	
	Column Descrip cache.	tion: The actions p	berformed	on the
dName		varchar	50	Х
	Column Descrip	tion: The name of	the cache	-
dSecurityGroup		varchar	30	Х
	Column Description: The security group classification assigned to the document revision.			
dDocAccount		varchar	30	Х
	Column Descrip displayed in an c	tion: Predefined a option list for users	ccount to 5.	be
dLastUsedDate		datetime	8	Х
	Column Descrip used.	tion: The last date	the cache	was
dCacheHash		varchar	30	Х
	Column Descrip cache.	tion: The hashing	used to er	ncrypt the
dDocType		varchar	30	Х
	Column Description: Unique identifier of a document type from the DocTypes table.			
Constraints: primary key	L			

DeletedRows

DeletedRows				
Column Name	Primary Key	Type Name	Length	Nullable
dRowID	Х	varchar	16	
	Column Descript	tion: Unique ident	ifier for tl	ne row.
dTable		varchar	30	
	Column Descript	tion: The table tha	t has beer	n deleted.
dPrimaryKeys		varchar	255	
	Column Description: The primary key is a unique key for each table. Each primary key is constrained so that no two values are equal.			nique key ed so that
dPKColumns		varchar	150	
	Column Description: The table columns that comprise the primary key for each table.			
dPKTypes		varchar	50	Х
	Column Descript	tion: The types of	primary k	xeys.
dDeleteDate		datetime	8	
	Column Descript	tion: Date that the	table was	deleted.
dActionDate		datetime	8	
	Column Descript	tion: The date the	action oc	curred.
dSourceID		varchar	50	Х
	Column Description: The content ID.			
Constraints: primary key				

DocFormats

This table lists HTML-defined document formats and how they are to be handled by the Content Server. The dFormat column in this table is equivalent to the dFormat column in the Documents table.

DocFormats				
Column Name	Primary Key	Type Name	Length	Nullable
dFormat	Х	varchar	30	
	Column Description: A document format that Content Server understands.			
dConversion		varchar	30	
	Column Description: The conversion rule used to convert the document format.			sed to
dDescription		varchar	80	Х
	Column Descrip	tion: A descriptior	of the fo	rmat.
Constraints: primary key				

DocMeta

This table lists the custom meta-data data fields. Each column is a custom meta-data field. These may be user-defined or created by installing a Content Server component.

DocMeta				
Column Name	Primary Key	Type Name	Length	Nullable
dID	Х	bigint	8	
	Column Description: Repository ID used throughout the system to uniquely identify a document revision.			roughout revision.
xComments		varchar	255	Х
	Column Descrip document revisio	tion: General com	ments fiel	d for
Constraints: primary key				

DocMetaDefinition

Each entry in this table defines the Content Server properties for a column in the DocMeta table.

DocMetaDefinition				
Column Name	Primary Key	Type Name	Length	Nullable
dName	Х	varchar	30	
	Column Descript	tion: Name of field	d.	
dType		varchar	50	Х
	Column Description: Type of field (Text, BigText, Int, Date, or Memo).			gText,
dIsRequired ¹		bit	1	
	Column Description: Whether or not the field is required.			
dIsEnabled ²		bit	1	
	Column Description: Whether or not the field displays on browser screens.			d displays
dIsSearchable ³		bit	1	
	Column Descript be exported to th	tion: Whether or n e search engine.	ot the fiel	ld should
dCaption		varchar	50	Х
	Column Description: Text to use when externally referring to the field.			
dIsOptionList		bit	1	
	Column Description: Indicates whether the field is a restricted list of options.			field is a

DocMetaDefinition				
dOptionListKey		varchar	50	Х
	Column Descript key is used to fin	tion: If the field is ad options in the O	an option ptionsLis	list, this t table.
dDefaultValue		varchar	50	Х
	Column Descript	tion: Default value	e for the fi	eld.
dOrder		smallint	2	
	Column Description: Numerical order in which options should be presented.			
dOptionListType ⁴		varchar	8	Х
	Column Descript strict choice, chu combo = edit and	tion: Type of option nval = strict choic d choice, multi = n	on list (cho e unvalid nulti selec	bice = ated, et).
Constraints: primary key 1. default constraint (0) 2. default constraint (1) 3. default constraint (1) 4. default constraint (0)				

DocTypes

This table lists the document types that are used as meta-data for document revisions.

DocTypes					
Column Name	Primary Key	Type Name	Length	Nullable	
dDocType	Х	varchar	30		
	Column Description: Name of the document type.				
dDescription		varchar	80	Х	
	Column Description: Document type description.				
dGif		varchar	50	Х	
	Column Description: File name of the gif image for the document type (gifs are located in <intradocwebdir>/images/docgifs).</intradocwebdir>				
Constraints: primary key					

DocumentAccounts

This table defines account names. Accounts are assigned to documents in the Revisions table. Users gain rights to accounts through the UserSecurityAttributes table when the value of dAttributeType is set to Account.

DocTypes				
Column Name	Primary Key	Type Name	Length	Nullable
dDocAccount	Х	varchar	30	
	Column Description: Predefined account to be displayed in an option list for users.			be
Constraints: primary key				

DocumentHistory

This table defines a history of actions performed on a document.

DocumentHistory				
Column Name	Primary Key	Type Name	Length	Nullable
dActionMillis	Х	bigint	8	
	Column Descrip seconds, millised	tion: Unique ident cond, and date.	ifier base	d on the
dActionDate	Х	datetime	8	
	Column Descrip	tion: Date, accurat	te to the n	ninute.
dID		bigint	8	Х
	Column Description: Repository ID used throughout the system to uniquely identify a document.			
dRevClassID		bigint	8	
	Column Description: Unique identifier of the document up to the revision. Revisions of a document have the same dRevClassID.			
dUser		varchar	50	Х
	Column Descrip	tion: User that per	formed ac	ction.
dDocName		varchar	30	Х
	Column Descrip Revisions of a de	tion: Unique name ocument have the	e of a docu same nam	ument. ne.
dAction		varchar	50	Х
	Column Description: The action performed. Values include: Checkin, Check out, New checkin, Delete document, Delete revision, Undo checkout, Update, Resubmit.			Values Delete Update,

DocumentHistory				
dSecurityGroup		varchar	30	Х
	Column Description: The security group classificat assigned to the document revision.			sification
Constraints: primary key				

Documents

This table defines an asset (content) within the Content Server.

DeletedRows				
Column Name	Primary Key	Type Name	Length	Nullable
dDocID	Х	bigint	8	
	Column Descript file rendition. Fo dDodIDs (primat	tion: Unique ident r each dID, there o ry, alternate, and v	ifier of do can be up vebviewa	ocument to three ble).
dID		bigint	8	
	Column Descript the system to unit	tion: Repository II iquely identify a d	D used the ocument.	roughout
dIsPrimary		bit	1	
	Column Description: Indicates whether or not the file is the primary file for the revision.			
dIsWebFormat		bit	1	
	Column Descript displaying the do	tion: Indicates the ocument revision of	format to on the web	use for D.
dLocation		varchar	50	Х
	Column Descript drive. This is not key to be used in	tion: Repository lo a physical file add lookup).	ocation on dress but a	the hard an alias (a
dOriginalName		varchar	80	Х
	Column Description: Original file name.			
dFormat		varchar	30	Х
	Column Description: Mime type of the document (such as application/ms-word, text/html, etc.).			iment .).

DeletedRows				
dExtension		invarchart	8	Х
	Column Description: The file extension.			
dFileSize		bigint	8	Х
	Column Description: Size of the file in bytes.			3.
Constraints: primary key				

ExtensionFormatMap

This table maps a document extension (such as .doc for MS Word) to a DocFormat. The dFormat column in this table is equivalent to the dFormat column in the DocFormats table.

ExtensionFormatMap				
Column Name	Primary Key	Type Name	Length	Nullable
dExtension	Х	varchar	8	
	Column Description: The file extension.			
dFormat		varchar	30	
	Column Descrip	tion: The file form	at type.	
Constraints: primary key	·			

HtmlConversions

HtmlConversions				
Column Name	Primary Key	Type Name	Length	Nullable
dConversionKey	Х	varchar	50	
	Column Descript	tion: ????? Key us	ed to initi	ate the
dOutputFile		varchar	150	Х
	Column Description: The mapped file type that is used to convert the document format. Or ??? The path to the mapped output file.			
dConversionDate		datetime	8	Х
	Column Descript converted.	tion: The date the	document	twas
dVaultTimeStamp		varchar	50	Х
	Column Description: The date and time the converted document was written to the vault.			
dDependencyKey		varchar	100	Х
	Column Description: Key that determines the necessary parameter(s) for document conversion.			
Constraints: primary key				

HtmlConversionSums

HtmlConversions				
Column Name	Primary Key	Type Name	Length	Nullable
dTotalCached	Х	varchar	25	
	Column Description: The total number of documents available from cached storage.			ocuments
dTotalFiles	Х	varchar	25	
	Column Description: The total number of documents that have been converted.			
dLastAccessed	Х	datetime	8	
	Column Description: The last date the converted document was accessed.			
dLastKey	Х	varchar	100	
	Column Description: ?????????			
Constraints: primary key				

OptionsList

The rows in this table provide the information for the list of options contained in the DocMetaDefinition table. The dKey column in this table is equivalent to the dOptionListKey column in the DocMetaDefinition table.

OptionsList				
Column Name	Primary Key	Type Name	Length	Nullable
dKey	X	varchar	50	
	Column Description: Key used to search for a list of selection options.			a list of
dOption		varchar	50	
	Column Description: Display order for option list.			on list.
dOrder		bigint	8	
	Column Description: The field holding the option for a particular key. The options are found by retrieving all rows with the same dKey.			option for strieving
Constraints: primary key				

ProblemReports

This table contains reports of problems from Content Publisher projects. Refer to the RegisteredProjects table for publisher project definitions.

ProblemReports				
Column Name	Primary Key	Type Name	Length	Nullable
dPrID	Х	bigint	8	
	Column Descrip report.	tion: Unique ident	ifier for a	problem
dPrCaption		varchar	80	Х
	Column Descrip	tion: Friendly ider	ntifier.	
dPrAuthor		varchar	50	Х
	Column Descrip	tion: Author of the	e problem	report.
dPrState		varchar	20	Х
	Column Descrip report.	tion: Current statu	s of the p	roblem
dPrSeverity		varchar	20	Х
	Column Descrip problem report.	tion: Assigned sev	verity of th	ne
dPrCreateDate		datetime	8	Х
	Column Description: Origination date of the problem report.			
dPrChangeDate		datetime	8	X
	Column Description: Date the problem report was edited or modified.			rt was

ProblemReports				
dProjectID		varchar	30	Х
	Column Description: Unique identifier for the project.			ne project.
dDocName		varchar	30	Х
	Column Description: The document associated with this problem report.			
dDocTitle		varchar	80	Х
	Column Descript report.	tion: Descriptive ti	itle of this	problem
Constraints: primary key	<u>.</u>			

ProjectDocuments

This table defines Content Server assets that are being used by a registered Content Publisher project. Refer to the RegisteredProjects table for project definitions.

ProjectDocuments	etDocuments			
Column Name	Primary Key	Type Name	Length	Nullable
dDocName	Х	varchar	30	
	Column Descrip Revisions of a do	tion: Unique name ocument have the	e of a docu same nam	ument. ie.
dProjectID		varchar	30	Х
	Column Descrip	tion: The unique p	roject ide	ntifier.
dSourceInstanceName		varchar	50	Х
	Column Description: The Content Server instance of the source document.			
dSourceDocID		bigint	8	Х
	Column Description: The repository revision ID for the source document.			n ID for
dSourceDocName		varchar	30	Х
	Column Descript the revision.	tion: The source d	ocument	name for
dSourcePending		varchar	1	Х
	Column Description: Indicates if the source document is the latest revision. For example, the document could be in a contributor workflow.			
dPrjTopParent		varchar	30	Х
	Column Descrip	tion: The project's	top parer	nt.

ProjectDocuments				
dPrjImmediateParent		varchar	30	Х
	Column Descript project.	ion: The immedia	te parents	of the
dPrjMiddleParents		varchar	255	Х
	Column Description: The middle parent of the project.			
dPrjAgentID		varchar	30	Х
	Column Description: The publisher agent that performed the document action.			
dPrjAction		varchar	30	Х
	Column Description: The action, usually DELETED.			
Constraints: primary key				

RegisteredProjects

This table contains Content Publisher projects that have been stored in the Content Server.

RegisteredProjects				
Column Name	Primary Key	Type Name	Length	Nullable
dProjectID	Х	varchar	30	
	Column Descrip	tion: The unique p	roject ide	ntifier.
dPrjDescription		varchar	80	Х
	Column Description: The project description that is displayed on the user interface.			
dPrjSourcePath		varchar	255	Х
	Column Description: The project source path.			h.
dPrjUrlPath		varchar	255	Х
	Column Descrip	tion: The URL to t	the projec	t.
dPrjFunctions		varchar	80	Х
	Column Description: Functions to which this project is registered such as preview and staging workflow.			
Constraints: primary key				

Revisions

This table contains the meta-data for each revision of a document, as defined by the dID column.

Revisions				
Column Name	Primary Key	Type Name	Length	Nullable
dID	Х	bigint	8	
	Column Description the system to unit	tion: Repository II iquely identify the	D used the documen	oughout t.
dDocName		varchar	30	
	Column Description: Unique name of a document. Revisions of a document have the same name.			ument. 1e.
dDocType		varchar	30	
	Column Description: Unique identifier of a document type from the DocTypes table.			
dDocTitle		varchar	255	
	Column Descrip	tion: Descriptive t	itle of the	revision.
dDocAuthor		varchar	50	Х
	Column Descript revision.	tion: The most rec	ent author	r of the
dRevClassID		bigint	8	
	Column Description: The unique identifier of the document up to this revision. Revisions of a docume have the same dRevClassID.			of the document

Revisions				
dRevisionID		bigint	8	
	Column Descript at one and incren document. For in	tion: The revision nents for each revi nternal use only.	number tl ision of a	nat starts
dRevLabel		varchar	10	Х
	Column Description: The revision label. This is the external revision identifier. It can be a duplicate and have its own algorithm for incrementing.			
dIsCheckedOut ¹		bit	1	
	Column Description: A flag that indicates whether the document is checked out. Only the most recent revision can have this flag set.			
dCheckoutUser		varchar	50	Х
	Column Description: The user currently editing the document. Only the most recent revision can have a value assigned.			
dSecurityGroup ²		varchar	30	
	Column Descript assigned to the de	tion: The security gout the security gout the security gout the security of th	group clas	sification
dCreateDate		datetime	8	Х
	Column Description: The creation date. The date the revision entered the system, not necessarily the date the document was released.			
dInDate		datetime	8	
	Column Description: The suggested release date. The date the revision should enter the system for display on the web.			date. The display

Revisions				
dOutDate		datetime	8	Х
	Column Descript The standard imp as well as the cur	tion: The date the plementation expir rrent revision.	revision e res older r	expires. revisions
dStatus ³		varchar	20	
	Column Description: The current status. Values include: EDIT (Contribution), REVIEW (Review), PENDING (Waiting others to finish), DONE (Ready for release), RELEASED (Released), GENWWW (Waiting for conversion), DELETED (Deleted, waiting replication logic), and EXPIRED (Expired).			
dReleaseState		char	1	Х
	Column Description: The transitioning release state. Values include: N (not currently releasable), R (processing for release), Y (released), O (old release superseded by later version), E (workflow), U (released and updated), and I (released and being indexed).			
dFlag1		char	1	Х
	Column Description: A custom flag that can be used to mark a revision as it is being processed from one state to another. Provided to external applications.			
dWebExtension		varchar	8	Х
	Column Description: The extension used by the search engine.			

Revisions				
dProcessingState ⁴		char	1	Х
	Column Description: The conversion and indexing state. Values include: C (converting), F (failed conversion), Y (converted), M (Not full text indexed), P (conversion failed, web viewable created by copying native).			
dMessage		varchar	255	Х
	Column Description: A message describing the results of document conversion or indexing. Contains the most recent message.			he results ns the
dDocAccount		varchar	30	Х
	Column Description: A predefined account to be displayed in an option list for the user.			
dReleaseDate		datetime	8	Х
	Column Descript dStatus went to F	tion: The date that RELEASED.	the value	of
dRendition1		char	1	Х
	Column Description: An optional extra rendition entry. Currently, the value is either empty or null to indicate no extra rendition, or T to indicate that a thumbnail was produced.			
dRendition2		char	1	Х
	Column Description: A second optional extra rendition. For future use.			
dIndexerState		char	1	Х
	Column Descript indexing state.	tion: Indicates the	document	's current

Revisions				
dPublishType		char	1	Х
	Column Description: Indicates the type of published document this is. Values include: P (page), H (home or entry point), N (navigation), G (gallery), C (contributor object), etc.			
dPublishState		char	1	Х
	Column Description: Values include: P (publish), S (staging), and W (workflow).			lish), S
dWorkflowState		varchar	1	Х
	Column Description: The current state of the workflow.			
dRevRank		biginit	8	
	Column Descrip this indicates the the tip revision. zero for dRevRa	tion: For any revis number of revision The tip revision al nk.	tion of a d ons it is do ways has	ocument, own from a value of
Constraints: primary key				
^{1.} default constraint (0)				
² default constraint (Public) ³ default constraint (GENWWW)				
4. default constraint (C)	,			

RoleDefinition

This table defines Roles. Each Role must contain a unique role name (dRoleName) per Security Group (dGroupName). This table also contains admin privilege information. Document privileges are contained in the UserSecurityAttributes table where the value of dAttributeType is role.

RoleDefinition				
Column Name	Primary Key	Type Name	Length	Nullable
dRoleName	Х	varchar	30	
	Column Description: Name of the role.			
dGroupName	Х	varchar	30	
	Column Description: The group name and unique identifier.			nique
dPrivilege ¹		bigint	8	
	Column Description: The privilege level allowed to the group by members belonging to the role.			
Constraints: primary key ^{1.} default constraint (0)				

SecurityGroups

This table defines a security group using dGroupName (PK) and provides a description of the group. Permissions are handled in the UserSecurityAttributes table where the value of dAttributeType is Group.

SecurityGroups				
Column Name	Primary Key	Type Name	Length	Nullable
dGroupName	Х	varchar	30	
	Column Description: The group name and unique identifier.			nique
dDescription		varchar	80	Х
	Column Description: Security group description.			
Constraints: primary key				

Subscription

Subscription					
Column Name	Primary Key	Type Name	Length	Nullable	
dSubscriptionAlias	Х	varchar	50		
	Column Descript group of users.	tion: The user nam	ne or alias	for the	
dSubscriptionAliasType	Х	varchar	30		
	Column Description: Indicates whether alias is for a single user (user) or is an alias for a group of users (alias).				
dSubscriptionEmail		varchar	80	Х	
	Column Description: The email address if the email address or addresses can not be determined from the dSubscriptionAlias.				
dSubscriptionID	Х	varchar	255		
	Column Descript document or the documents fall u	tion: Either the uni filter data used to nder the subscripti	ique ident determine ion.	ifier for a e which	
dSubscriptionType	Х	varchar	30		
	Column Description: The rule to use for interpreting dSubscriptionID. The standard value is docName that indicates the subscription is for a single document.				
dSubscriptionCreateDate		datetime	8		
	Column Description: The date the subscription was first created.			on was	
Subscription					
--------------------------	--------------------------------------------------------------------------	----------	---	---	--
dSubscriptionNotifyDate		datetime	8	Х	
	Column Description: The last date a user was notified of a subscription.				
dSubscriptionUsedDate		datetime	8	Х	
	Column Description: The last date a user used the notification.				
Constraints: primary key					

Users

This table contains the information about the users that are defined in the Content Server.

Users					
Column Name	Primary Key	Type Name	Length	Nullable	
dName	Х	varchar	50		
	Column Descripthe system.	tion: Uniquely ide	ntifies the	e user in	
dFullName		varchar	50	Х	
	Column Description: The user's full name.				
dEmail		varchar	80	Х	
	Column Description: The user's email address.				
dPasswordEncoding	varchar 8 X				
	Column Description: The rule used to encrypt the user's password. Values include: OPENTEXT or blank if the password is in plain text and SHA1-C SHA1 is used in hashing the password.				
dPassword		varchar	50	Х	
	Column Description: The password that is encoded using the rule indicated by dPasswordEncoding.				
dUserType		varchar	30	Х	
	Column Description: Type of user; similar to dDocType.				

Users					
dUserAuthType		varchar	30	Х	
	Column Description: Determines what is exposed to the ns/isapi filters and where and how credentials are to be retrieved. Values include: GLOBAL, LOCAL, or EXTERNAL.				
dUserOrgPath		varchar	255	Х	
	Column Description: Used for searching and easier access. May be mapped to accounts.				
dUserSourceOrgPath		varchar	50	Х	
	Column Description: The external user source. Could refer to LDAP or the name of the provider.				
dUserSourceFlags		bigint	8	Х	
	Column Description: The flags that indicate which external information is provided by this instance.				
dUserArriveDate		datetime	8	Х	
	Column Description: Used for user list replication. Indicates that the current system has changed the user information and it can/should be replicated out.				
dUserChangeDate		datetime	8	Х	
	Column Description: Used for user list replication. Indicates that a foreign system has changed the user information and the current instance should update its information.				
dUserLocale		varchar	30	Х	
	Column Descrip locale.	tion: Used to mon	itor the us	er's	

03013				
dUserTimeZone		varchar	30	Х
	Column Description: Used to maintai zone information.		intain the u	iser's tim
Constraints: primary key	·			

UserSecurityAttributes

This table contains the user-based document security information for roles, groups, and accounts.

UserSecurityAttributes					
Column Name	Primary Key	Type Name	Length	Nullable	
dUserName	Х	varchar	50		
	Column Descript	tion: The user's na	ime.		
dAttributeName	Х	varchar	30		
	Column Description: Name of the attribute. For example, the name of the role if it is a role attribute or the account name it is an account attribute.				
dAttributeType	Х	varchar	8		
	Column Description: The type of attribute. Values include: role, account, or defacct (default account).				
dAttributePrivilege		bigint	8	Х	
	Column Description: The access privilege level. It is a combination of bit flags and values include: 1 (read), 2 (write), 4 (delete), and 8 (admin).				
dAttributeState		varchar	8	Х	
	Column Description: For future use.				
Constraints: primary key					

WorkflowAliases

This table defines a workflow step that includes an alias as defined in the Alias table.

WorkflowAliases					
Column Name	Primary Key	Type Name	Length	Nullable	
dWfStepID	X	bigint	8		
	Column Description: An incrementing counter used to uniquely identify a step.				
dWfID		bigint	8	Х	
	Column Description: The unique counter of the workflow containing the step.				
dAlias	Х	varchar	50		
	Column Description: Unique identifier for an alia				
dAliasType		varchar	20		
	Column Description: The type of alias (single user or alias for a group of users).				
Constraints: primary key					

WorkflowCriteria

	This table d	lefines the	criteria t	for entering	a step	in a	workflow.
--	--------------	-------------	------------	--------------	--------	------	-----------

WorkflowCriteria					
Column Name	Primary Key	Type Name	Length	Nullable	
dWfID	Х	bigint	8		
	Column Description: The unique counter ID for the workflow.				
dWfCriteriaName	Х	varchar	30		
	Column Description: The attribute having the criteria applied.				
dWfCriteriaOperator	ator varchar 8 X Column Description: The operator being applied against the criteria value. The only allowed entry matches.				
dWfCriteriaValue		varchar	30	Х	
	Column Description: The value being matched. Database wildcards are allowed.				
Constraints: primary key					

WorkflowDocAttributes

This table defines additional meta-data attributes for assets within a workflow.

WorkflowDocAttributes					
Column Name	Primary Key	Type Name	Length	Nullable	
dWfID	Х	bigint	8		
	Column Description: The unique counter ID for the workflow.				
dDocName	X varchar 30				
	Column Description: Unique name of a document in workflow. Revisions of a document have the same name.				
dWfAttribute	varchar 50 X				
	Column Description: The associated value of the document. For example, contributor might mean the original contributor. For future use.				
dWfAttributeType	Х	varchar	8		
	Column Description: Determines how dWfAttribute is to be interpreted. For future use.				
Constraints: primary key					

WorkflowDocuments

This table defines all documents that are part of a workflow. It includes the step in the workflow and the current state of the workflow (such as INPROCESS or INIT -- not active).

WorkflowDocuments					
Column Name	Primary Key	Type Name	Length	Nullable	
dWfID		bigint	8	Х	
	Column Description: The unique counter ID for the workflow that contains the step.				
dDocName	Х	varchar	30		
	Column Description: Unique name of a document. Revisions of a document have the same name.				
dWfDocState		varchar	20	Х	
	Column Description: The current global workflow state of the document. Values include: INIT (workflow not active) and INPROCESS (workflow active).				
dWfComputed		varchar	20	Х	
	Column Description: Derived information to minimize the complexity of queries and coding logic For future use.				
dWfCurrentStepID		bigint	8	Х	
	Column Description: The unique counter ID step that the workflow in. Refer to dWfStepI WorkflowSteps table.				
dWfDirectory		varchar	100	Х	
	Column Description: The location of the workflo document's companion file.				

WorkflowDocuments				
dClbraName		varchar	50	Х
	Column Descrip	tion: The name of	the collab	ooration.
Constraints: primary key	·			

WorkflowHistory

This table contains the history of the actions performed on an asset within a workflow.

WorkflowHistory	cflowHistory				
Column Name	Primary Key	Type Name	Length	Nullable	
dActionMillis	Х	bigint	8		
	Column Descrip seconds, millised	tion: Unique ident cond, and date.	ifier base	d on the	
dActionDate	Х	datetime	8		
	Column Description: date, accurate to the minute.				
dAction		varchar	30		
	Column Description: The workflow action. Values include: Checkin, Approve, Reject, Start, and Cancel.				
dWfName		varchar	30		
	Column Description: The user-assigned identifier for the workflow.				
dWfStepName		varchar	30	Х	
	Column Description: The user-assigned name of the workflow step.				
dUser		varchar	50		
	Column Description: The user who performed the action.				
dID		bigint	8	Х	
	Column Description: Repository ID used throughout the system to uniquely identify the document.				

WorkflowHistory				
dDocName		varchar	30	Х
	Column Description: Unique name of a document. Revisions of a document have the same name.			
dDocType		varchar	30	Х
	Column Description: Unique identifier of a document type from the DocTypes table.			
dDocTitle		varchar	80	Х
	Column Description: Descriptive title of the revision.			
dDocAuthor		varchar	50	Х
	Column Description: The most recent editor of the document.			
dRevClassID		bigint	8	Х
	Column Description: Unique identifier of the document up to the revision. Revisions of a document have the same dRevClassID.			
dRevisionID		bigint	8	Х
	Column Description: Revision number. Starts at 1 and increments for each revision of a document. For internal use only.			
dRevLabel		varchar	10	Х
	Column Description: Revision label. Can be a duplicate and have its own incrementing algorithm. This is the external revision identifier.			
dSecurityGroup		varchar	30	
	Column Description: The security group classification assigned to the document revision.			

WorkflowHistory				
dCreateDate		datetime	8	Х
	Column Description: The creation date. The date the revision entered the system, not necessarily the date the document was released.			
dInDate		datetime	8	Х
	Column Description: The suggested release date. The date the revision should enter the system for display on the web.			
dOutDate		datetime	8	Х
	Column Description: The date the revision expires. The standard implementation expires older revisions as well as the current revision.			
dDocAccount		varchar	30	Х
	Column Description: Account that is assigned to the revision. used to control security for the revision.			
dClbraName		varchar	50	Х
Column Description: The name of the col				ooration.
Constraints: primary key				

Workflows

This table defines a workflow in the Content Server.

Workflows				
Column Name	Primary Key	Type Name	Length	Nullable
dWfName		varchar	30	Х
	Column Description: The user-assigned identifier for the workflow.			
dWfID	Х	bigint	8	
	Column Description: The unique counter of the workflow containing the step.			
dWfDescription		varchar	80	Х
	Column Description: The workflow description.			tion.
dCompletionDate		datetime	8	Х
	Column Description: The date the workflow last reached completion and all the documents were released to the web. Used only for a compound workflow where multiple documents are released the web simultaneously.			v last vere ind eased to
dSecurityGroup		varchar	30	Х
	Column Description: The security group classification assigned to the document revision.			
dWfStatus		varchar	20	Х
	Column Description: The state of the workflow. Values include: INIT (not started or inactive) and INPROCESS (active).			

Workflows					
dWfType		varchar	20	Х	
	Column Description: The type of workflow. Values include: Criteria (documents enter workflow when satisfying criteria - docs enter and exit singly) and Basic (many docs released together with controlled entry of docs into workflow).				
dProjectID	varchar 30 X				
	Column Description: If set, indicates that this workflow is a staging workflow.				
dIsCollaboration		bit	1		
	Column Description: Indicates if this workflow is part of a collaboration.				
Constraints: primary key					

WorkflowStates

This table defines the users that have approved specific documents within a workflow.

WorkflowStates				
Column Name	Primary Key	Type Name	Length	Nullable
dID	Х	int	4	
	Column Description: Repository ID used throughout the system to uniquely identify the document.			
dDocName		varchar	30	Х
	Column Description: Unique name of a document. Revisions of a document have the same name.			
dWfID	Х	bigint	8	
	tion: The unique c	ounter ID	of the	
dUserName	Х	varchar	50	
	Column Description: The user that has approved the document. This table is used to track which users have approved which documents.			
dWfEntryTs		datetime	8	Х
	Column Description: The entry timestamp for a user into a particular step in the workflow.			
Constraints: primary key				

WorkflowSteps

This table defines the steps in a workflow. Workflows are defined in the Workflows table.

WorkflowSteps				
Column Name	Primary Key	Type Name	Length	Nullable
dWfStepName		varchar	30	Х
	Column Description: The name of the step assigned by the user.			
dWfStepID	Х	bigint	8	
	Column Description: An incrementing counter used to uniquely identify the step.			
dWfID		bigint	8	Х
	Column Description: The unique counter ID for the workflow that contains the step.			
dWfStepDescription		varchar	80	Х
	Column Description: Description of the step.			
dWfStepType		varchar	20	Х
	Column Description: The type of step. Values include: AutoContribution, Contribution, Reviewer/Contribution, and Reviewer.			
dWfStepIsAll		bit	1	
	Column Description: Requires all user to approve.			
dWfStepWeight		bigint	8	Х
	Column Description: The number of users needed to approve if the value of dWfStepIsAll is set to false.			
Constraints: primary key				

sysconstraints

sysconstraints				
Column Name	Primary Key	Type Name	Length	Nullable
constid		int	4	Х
id		int	4	Х
colid		smallint	2	Х
spare1		tinyint	1	Х
status		int	4	Х
actions		int	4	Х
error		int	4	Х
Constraints: View (not Table)				

syssegments

WorkflowStates					
Column Name	Primary Key	Type Name	Length	Nullable	
segment		int	4		
name		varchar	10		
status		int	4		
Constraints: View (not table)					

Appendix BBB

THIRD PARTY LICENSES

OVERVIEW

This appendix includes a description of the Third Party Licenses for all the third party products included with this product.

- Apache Software License (page B-1)
- ✤ W3C[®] Software Notice and License (page B-2)
- Zlib License (page B-3)
- General BSD License (page B-4)
- General MIT License (page B-5)
- Unicode License (page B-5)
- Miscellaneous Attributions (page B-7)

APACHE SOFTWARE LICENSE

- * Copyright 1999-2004 The Apache Software Foundation.
- * Licensed under the Apache License, Version 2.0 (the "License");
- * you may not use this file except in compliance with the License.
- * You may obtain a copy of the License at
- * http://www.apache.org/licenses/LICENSE-2.0
- *

- * Unless required by applicable law or agreed to in writing, software
- * distributed under the License is distributed on an "AS IS" BASIS,
- * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
- * See the License for the specific language governing permissions and
- * limitations under the License.

W3C® SOFTWARE NOTICE AND LICENSE

- * Copyright © 1994-2000 World Wide Web Consortium, * (Massachusetts Institute of Technology, Institut National de * Recherche en Informatique et en Automatique, Keio University). * All Rights Reserved. http://www.w3.org/Consortium/Legal/ * This W3C work (including software, documents, or other related items) is * being provided by the copyright holders under the following license. By * obtaining, using and/or copying this work, you (the licensee) agree that * you have read, understood, and will comply with the following terms and * conditions: * Permission to use, copy, modify, and distribute this software and its * documentation, with or without modification, for any purpose and without * fee or royalty is hereby granted, provided that you include the following * on ALL copies of the software and documentation or portions thereof, * including modifications, that you make: 1. The full text of this NOTICE in a location viewable to users of the redistributed or derivative work. 2. Any pre-existing intellectual property disclaimers, notices, or terms and conditions. If none exist, a short notice of the following form (hypertext is preferred, text is permitted) should be used within the
- * body of any redistributed or derivative code: "Copyright ©
- * [\$date-of-software] World Wide Web Consortium, (Massachusetts

* *

*

*

*

- * Institute of Technology, Institut National de Recherche en
- * Informatique et en Automatique, Keio University). All Rights
- * Reserved. http://www.w3.org/Consortium/Legal/"
- * 3. Notice of any changes or modifications to the W3C files, including the
 * date changes were made. (We recommend you provide URIs to the location
 * from which the code is derived.)
- * THIS SOFTWARE AND DOCUMENTATION IS PROVIDED "AS IS," AND COPYRIGHT HOLDERS * MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT * NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR * PURPOSE OR THAT THE USE OF THE SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE * ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS. *
- * COPYRIGHT HOLDERS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL OR
- * CONSEQUENTIAL DAMAGES ARISING OUT OF ANY USE OF THE SOFTWARE OR
- * DOCUMENTATION.
- *
- * The name and trademarks of copyright holders may NOT be used in advertising
- \ast or publicity pertaining to the software without specific, written prior
- * permission. Title to copyright in this software and any associated
- * documentation will at all times remain with copyright holders.

*

ZLIB LICENSE

 * zlib.h -- interface of the 'zlib' general purpose compression library version 1.2.3, July 18th, 2005

Copyright (C) 1995-2005 Jean-loup Gailly and Mark Adler This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

- The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
- 2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
- 3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly jloup@gzip.org

Mark Adler madler@alumni.caltech.edu

GENERAL BSD LICENSE

Copyright (c) 1998, Regents of the University of California

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

"Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

"Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

"Neither the name of the <ORGANIZATION> nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

GENERAL MIT LICENSE

Copyright (c) 1998, Regents of the Massachusetts Institute of Technology

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

UNICODE LICENSE

UNICODE, INC. LICENSE AGREEMENT - DATA FILES AND SOFTWARE

Unicode Data Files include all data files under the directories http://www.unicode.org/Public/, http://www.unicode.org/reports/, and http://www.unicode.org/cldr/data/ . Unicode Software includes any source code published in the Unicode Standard or under the directories http://www.unicode.org/Public/, http://www.unicode.org/reports/, and http://www.unicode.org/cldr/data/. NOTICE TO USER: Carefully read the following legal agreement. BY DOWNLOADING, INSTALLING, COPYING OR OTHERWISE USING UNICODE INC.'S DATA FILES ("DATA FILES"), AND/OR SOFTWARE ("SOFTWARE"), YOU UNEQUIVOCALLY ACCEPT, AND AGREE TO BE BOUND BY, ALL OF THE TERMS AND CONDITIONS OF THIS AGREEMENT. IF YOU DO NOT AGREE, DO NOT DOWNLOAD, INSTALL, COPY, DISTRIBUTE OR USE THE DATA FILES OR SOFTWARE.

COPYRIGHT AND PERMISSION NOTICE

Copyright © 1991-2006 Unicode, Inc. All rights reserved. Distributed under the Terms of Use in http://www.unicode.org/copyright.html.

Permission is hereby granted, free of charge, to any person obtaining a copy of the Unicode data files and any associated documentation (the "Data Files") or Unicode software and any associated documentation (the "Software") to deal in the Data Files or Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Data Files or Software, and to permit persons to whom the Data Files or Software are furnished to do so, provided that (a) the above copyright notice(s) and this permission notice appear with all copies of the Data Files or Software, (b) both the above copyright notice(s) and this permission notice appear in associated documentation, and (c) there is clear notice in each modified Data File or in the Software as well as in the documentation associated with the Data File(s) or Software that the data or software has been modified.

THE DATA FILES AND SOFTWARE ARE PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THE DATA FILES OR SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in these Data Files or Software without prior written authorization of the copyright holder.

_____Unicode and the Unicode logo are trademarks of Unicode, Inc., and may be registered in some jurisdictions. All other trademarks and registered trademarks mentioned herein are the property of their respective owners

MISCELLANEOUS ATTRIBUTIONS

Adobe, Acrobat, and the Acrobat Logo are registered trademarks of Adobe Systems Incorporated. FAST Instream is a trademark of Fast Search and Transfer ASA. HP-UX is a registered trademark of Hewlett-Packard Company. IBM, Informix, and DB2 are registered trademarks of IBM Corporation. Jaws PDF Library is a registered trademark of Global Graphics Software Ltd. Kofax is a registered trademark, and Ascent and Ascent Capture are trademarks of Kofax Image Products. Linux is a registered trademark of Linus Torvalds. Mac is a registered trademark, and Safari is a trademark of Apple Computer, Inc. Microsoft, Windows, and Internet Explorer are registered trademarks of Microsoft Corporation. MrSID is property of LizardTech, Inc. It is protected by U.S. Patent No. 5,710,835. Foreign Patents Pending. Oracle is a registered trademark of Oracle Corporation. Portions Copyright © 1994-1997 LEAD Technologies, Inc. All rights reserved. Portions Copyright © 1990-1998 Handmade Software, Inc. All rights reserved. Portions Copyright © 1988, 1997 Aladdin Enterprises. All rights reserved. Portions Copyright © 1997 Soft Horizons. All rights reserved. Portions Copyright © 1995-1999 LizardTech, Inc. All rights reserved. Red Hat is a registered trademark of Red Hat, Inc. Sun is a registered trademark, and Sun ONE, Solaris, iPlanet and Java are trademarks of Sun Microsystems, Inc. Sybase is a registered trademark of Sybase, Inc. UNIX is a registered trademark of The Open Group. Verity is a registered trademark of Autonomy Corporation plc

Index

#

&#xxxx; character format, 8-5

Α

accented characters in content types, 8-8, 8-8 in document titles, 8-6, 8-8, 8-8 in password, 8-10, 8-10 in security groups, 8-8, 8-8 in URLs, 8-6, 8-7, 8-8 in user names, 8-10 accessing Content Server Analyzer, 3-27 accessing log files, 3-3 Acrobat Reader, 1-3 Alias (database table), A-4 AliasUser (database table), A-5 analysis Content Server Analyzer database, 3-28 Content Server Analyzer file system, 3-30 Content Server Analyzer search index, 3-29 starting Content Server Analyzer, 3-28 Analyzer, 3-23, 3-24 applet-specific tracing, 3-20 ArchiveHistory (database table), A-7 Archiver log files, 3-3, 3-6 troubleshooting, 6-1 archiving issues, 6-1 exporting issues, 6-11 importing issues, 6-1

B

bookmarked URLs, 8-6 boxes, characters display as --, 8-3

С

canceling Content Server Analyzer status report, 3-31 character display boxes, 8-3 garbage, 8-4 question marks, 8-2 Collaboration (database table), A-9 Config (database table), A-10 configuration information, 3-7 configuration variables, 3-32 DebugMode, 3-32 DebugStdConversion, 3-32 ScriptDebugTrace, 3-32 ScriptErrorTract, 3-32 SearchDebugLevel, 3-32 setResourceInclude, 3-32 UseRedirectedOutput, 3-32 XmlEncodingMode, 8-6, 8-7 Content Server log files, 3-3, 3-4 Content Server Analyzer accessing, 3-27 canceling status report, 3-31 console display area, 3-30 database analysis, 3-28 file system analysis, 3-30 Progress tab, 3-26 results, 3-30 search index analysis, 3-29 specifying log directory, 3-27 starting analysis, 3-28 status report, 3-31 Content Server analyzer, 3-23, 3-24 content types accented characters in --, 8-8, 8-8 CookieLoginPlugin component, 8-10 Counters (database table), A-11

D

data dictionary, A-1, A-3 database Content Server Analyzer analysis, 3-28 database issues, 5-1 database structure, A-1 database tables, A-3 Alias, A-4 AliasUser, A-5 ArchiveHistory, A-7 Collaboration, A-9 Config, A-10 Counters, A-11 DatedCaches, A-12 DocFormats, A-14 DocMeta, A-15 DocMetaDefinition, A-16 DocTypes, A-18 DocumentAccounts, A-19 DocumentHistory, A-20 Documents, A-22 ExtensionFormatMap, A-24 HtmlConversions, A-25 HtmlConversionSums, A-26 OptionsList, A-27 ProblemReports, A-28 ProjectDocuments, A-30 RegisteredProjects, A-32 Revisions, A-33 RoleDefinition, A-38 SecurityGroups, A-39 Subscription, A-40 sysconstraints, A-58 syssegments, A-59 Users, A-42 UserSecurityAttributes, A-45 WorkflowAliases, A-46 WorkflowCriteria, A-47 WorkflowDocAttributes, A-48 WorkflowDocuments, A-49 WorkflowHistory, A-51 Workflows, A-54 WorkflowState, A-56 WorkflowSteps, A-57 DatedCaches (database table), A-12 debugging, 3-18, 3-32 DebugMode (configuration variable), 3-32 DebugStdConversion (configuration variable), 3-32 directory specifying Content Server Analyzer log, 3-27 display issues, 8-2 DocFormats (database table), A-14 DocMeta (database table), A-15

DocMetaDefinition (database table), A-16 DocTypes (database table), A-18 document titles accented characters in --, 8-6, 8-8, 8-8 DocumentAccounts (database table), A-19 DocumentHistory (database table), A-20 Documents (database table), A-22

Ε

Edit status workflow items stuck in --, 7-1 environment packager, 3-21 exporting issues (Archiver), 6-11 ExtensionFormatMap (database table), A-24

F

file system Content Server Analyzer analysis, 3-30 filter administration, 3-15

G

gargabe, characters display as --, 8-4 generating Content Server Analyzer status report, 3-31 GenWWW status workflow items stuck in --, 7-1

Η

HtmlConversions (database table), A-25 HtmlConversionSums (database table), A-26

importing issues (Archiver), 6-1 Inbound Refinery log files, 3-3, 3-7 international issues, 8-1 display issues, 8-2 login issues, 8-9 iPlanet-specific issues, 4-10

L

log files, 3-2 accessing --, 3-3 Archiver, 3-3, 3-6 characteristics, 3-2 columns in --, 3-2 Content Server, 3-3, 3-4 database tracing, 5-34 Inbound Refinery, 3-3, 3-7 turning off logging, 4-14 logging options, 3-17 login issues, 8-9 logs specifying directory for Content Server Analyzer, 3-27

Μ

multi-byte characters, 8-4

0

opening log files Archiver, 3-6 Content Server, 3-4 OptionsList (database table), A-27

Ρ

PDF files viewing in browser, 8-7 port configuration issues, 4-8 preparing for issues, 2-3 problem-solving model, 2-2 ProblemReports (database table), A-28 product documentation, 1-2 ProjectDocuments (database table), A-30 purpose of Troubleshooting Guide, 1-1

Q

question marks, characters display as --, 8-2

R

RegisteredProjects (database table), A-32 release notes, 1-3 Review status workflow items stuck in --, 7-3 Revisions (database table), A-33 RoleDefinition (database table), A-38

S

ScriptDebugTrace (configuration variable), 3-32 ScriptErrorTrace (configuration variable), 3-32 search index Content Server Analyzer analysis, 3-29 SearchDebugLevel (configuration variable), 3-32 security groups accented characters in --, 8-8, 8-8 SecurityGroups (database table), A-39 server start issues, 4-4 server-wide tracing, 3-18 setResourceInclude (configuration variable), 3-32 setting up web server logging, 4-12 sort order, 8-9 specifying Content Server Analyzer log directory, 3-27 start issues, 4-4 starting Content Server Analyzer analysis, 3-28 status report Content Server Analyzer, 3-31 Subscription (database table), A-40 symptoms, 2-1 sysconstraints (database table), A-58 syssegments (database table), A-59 system audit information, 3-9

Т

tables, see 'database tables' thumbnails, 8-8 tracing, 3-18 applet-specific, 3-20 database. 5-34 server-wide, 3-18 troubleshooting archiving issues, 6-1 configuration information, 3-7 configuration variables, 3-32 Content Server analyzer, 3-23, 3-24 database issues, 5-1 environment packager, 3-21 international issues, 8-1 iPlanet-specific issues, 4-10 log files, 3-2 port configuration issues, 4-8 preparing for issues, 2-3 problem-solving model, 2-2 server start problems, 4-4 symptoms, problems, and solutions, 2-1 system audit information, 3-9 tracing, 3-18

Index

web server filter configuration, 3-15 web server issues, 4-1 web server logging, 4-12 workflow issues, 7-1 Troubleshooting Guide purpose of --, 1-1 turning off logging, 4-14

U

URLs bookmarked -- don't work, 8-6 UseRedirectedOutput (configuration variable), 3-32 Users (database table), A-42 UserSecurityAttributes (database table), A-45

W

web browser viewing PDF in --, 8-7 viewing web-viewable files in --, 8-8 web server filter configuration, 3-15 web server issues

permissions, 4-1 server start problems, 4-4 web server logging, 4-12 web-viewable files viewing in browser, 8-8 workflow issues, 7-1 WorkflowAliases (database table), A-46 WorkflowCriteria (database table), A-47 WorkflowDocAttributes (database table), A-48 WorkflowDocuments (database table), A-49 WorkflowHistory (database table), A-51 workflows items in wrong workflow, 7-4 items stuck in Edit status, 7-1 items stuck in GenWWW status, 7-1 items stuck in Review status, 7-3 Workflows (database table), A-54 WorkflowState (database table), A-56 WorkflowSteps (database table), A-57

Χ

XmlEncodingMode (configuration variable), 8-6, 8-7