

Oracle® Universal Content Management
Cluster Configuration Guide for UNIX
10g Release 3 (10.1.3.3.1)

May 2010

Cluster Configuration Guide for UNIX, 10g Release 3 (10.1.3.3.1)

Copyright © 2010, Oracle. All rights reserved.

Contributing Authors: Jean Wilson, Bruce Silver

Contributors: Alec Kloss, Sam White, Bob Wicklund, Don Kries, Aaron Briel

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 third-party 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	1-1
Product Overview	1-1
About This Guide	1-2
Intended Audience	1-2
Symbols and Conventions	1-2
Chapter Overviews	1-3

Chapter 2: Pre-Implementation Tasks and Considerations

Overview	2-1
Requirements	2-1
Recommendations and Prerequisites	2-2
Hardware Overview	2-2
Configuring Automation Activity	2-3

Chapter 3: Implementing a Cluster Configuration

Overview	3-1
Setting Up the Install Node Server	3-2
Installing a Cluster	3-3
Using the Installer	3-3
Configuring Additional Nodes	3-4
Reconfiguring an Updated Cluster	3-5
Remote Administration	3-5
Configuring the Cluster	3-6
Backing Up Clustered Content Servers	3-6
Cluster Installation Details	3-7

Chapter 4: Content Server Configurations

Overview	4-1
Clusters and Oracle Products	4-1
CoreWebDAV and Folders Component	4-2
Content Publisher	4-2
Inbound Refinery	4-2
Dynamic Converter	4-3
Clusters and HTTP Servers	4-3

Appendix A: Configuration Settings

Overview	A-1
Introduction	A-1
Local Directory Variables	A-2
Automation Variables	A-3
Locking and Tracing Variables	A-4
Local Directory Variables for the Stub_Dir	A-6

Appendix B: Third Party Licenses

Overview	B-1
Apache Software License	B-1
W3C® Software Notice and License	B-2
Zlib License	B-3
General BSD License	B-4
General MIT License	B-5
Unicode License	B-5
Miscellaneous Attributions	B-7

INTRODUCTION

OVERVIEW

This section covers the following topics:

- ❖ [Product Overview](#) (page 1-1)
- ❖ [About This Guide](#) (page 1-2)



Important: These clustering guidelines are intended only as a reference and not as a complete example of how you should establish your cluster configuration. You are encouraged to contact Consulting Services for assistance as you design and implement your cluster.

PRODUCT OVERVIEW

A server cluster is a group of independent servers managed as a single system that can be used as a multipurpose platform for database management, file and intranet data sharing, and general business applications. Content Server can be a stand-alone system, or it can be modified and integrated in unlimited configurations to alleviate problems of maintenance, load balancing, and scalability. In a cluster configuration where multiple servers share a Content Server instance, all of the servers also use a common content server file system, database, and index collection. A load balancer can be used to provide high availability for contributors.



Important: Not all clustering configurations have been tested, and this guide does not specifically cover all possible clustering environments. In all cases, we recommend that you test your clustering configuration in a development environment before deploying it to production.

ABOUT THIS GUIDE

This configuration guide provides all information required to install and configure the server software for a cluster configuration where the machines are running the UNIX operating system (Sun Solaris, HP-UX, IBM AIX, or Red Hat Linux Suse). The information contained in this document is subject to change as the product technology evolves and as hardware, operating systems, and third-party software are created and modified.




Intended Audience

This configuration guide is intended for individuals who are responsible for designing, implementing, and/or maintaining systems that combine the Content Server product with multiple servers in a clustered environment. In addition, this manual will be useful for database administrators who maintain the integrity of the document information stored in the Content Servers. This document assumes that you are familiar with Oracle products and the architecture of the Content Server.

Symbols and Conventions

The following symbols and conventions are used throughout this manual:

Symbols	Description
	This is a note. It is used to bring special attention to information.

Symbols	Description
	This is a tech 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.

Chapter Overviews

The following topics are discussed in this guide:

- ❖ [Chapter 2 \(*Pre-Implementation Tasks and Considerations*\)](#): This chapter provides some cluster configuration background information. It also includes some tasks and considerations that need to be taken into account before installing the software onto and setting up the install node and additional nodes in the cluster configuration.
- ❖ [Chapter 3 \(*Implementing a Cluster Configuration*\)](#): This chapter provides instructions to install, setup and configure the applicable software for install nodes and additional nodes in a cluster configuration. It also includes instructions for configuring and setting up the web servers and database connections.
- ❖ [Chapter 4 \(*Content Server Configurations*\)](#): This chapter provides summaries and results from tested versions of the cluster configuration (an install node server and one or more additional node servers) with Content Server.

PRE-IMPLEMENTATION TASKS AND CONSIDERATIONS

OVERVIEW

This section covers the following topics:

- ❖ [Requirements](#) (page 2-1)
- ❖ [Recommendations and Prerequisites](#) (page 2-2)
- ❖ [Configuring Automation Activity](#) (page 2-3)



Important: You should review the *Content Server Clustering Concepts Guide* before setting up a cluster. That document contains important planning and installation considerations, and architecture scenarios.

Due to the complex interaction of cluster-related software and hardware components, you will need to retain Oracle's Consulting Services to assist you in planning and implementing your clustered Content Server configuration.

REQUIREMENTS

At least two servers are required for a cluster, and all servers must map to the same share point on a file server. The hostname specified in a cluster configuration should be the hostname used by browsers when they connect to the system, and therefore the hostname of the load balancing machine.

Clustering is possible with all operating systems and web servers that support Content Server. The following configurations have been tested and found to be stable clustering environments:

Server OS	Web Server	Database
UNIX: (SuSe Linux, Sun Solaris, HP-UX, IBM AIX, Red Hat Linux)	One of the following: Sun ONE, Apache, IBM HTTP Server	A third-party database such as IBM DB2, Oracle, Sybase, Informix (SQL Server)

Content ServerContent Server

RECOMMENDATIONS AND PREREQUISITES

Before implementing a clustered system, the following are strongly recommended:

- ❖ Advanced knowledge of and experience with Content Server.
- ❖ Knowledge and familiarity with your site's current infrastructure and operations. Installing a clustered Content Server may require changes in corporate policies or procedures.
- ❖ Installation of Content Server with a cluster environment in mind; more planning and effort will be required to implement a cluster configuration with a Content Server system that is already deployed. See the *Clustering Concepts Guide* for best practices and information needed before setting up a cluster system.
- ❖ The install program currently uses the Oracle 9i JDBC libraries. Before installing a cluster node, it is recommended that you install the Oracle 10 JDBC libraries by copying the ojdbc14.jar file to the <install_dir>/shared/classes directory.

Hardware Overview

The hardware configuration for each site is based on existing infrastructure. Content Server is not itself cluster-aware but it is an application that uses a site's existing cluster-aware infrastructure. Therefore no one specific hardware configuration can work for all sites.

In general, you can set up elements of the Content Server to be distributed to different servers or file systems. These machines can be part of the consumption cluster, or can be separate machines that have the Content Server installed on proportioned drives.

You can also set one machine in the cluster as a master node and administer the other nodes through that system.

In clustered Content Server configurations, the clock in each machine must be synchronized because several pieces of functionality require coordinating the time settings (for example, the work file queues are shared by all the nodes in the cluster). To ensure uninterrupted performance, the clock synchronization must be configured to have a maximum disparity of five seconds.

CONFIGURING AUTOMATION ACTIVITY

Configuration variables can be used to control automation activity for a load balancing or failover clustering scenario. These flags allow the refinery, indexing, and archiving functions to be disabled. For more information about these variables, see [Appendix A \(Configuration Settings\)](#). With the current release of Content Server, enhancements and additional configuration settings allow for more flexibility in load balancing configurations.

These configuration options can be used in situations such as load balancing with content servers, when some content servers are used only for consumption and others for contribution. In this scenario, turning off all automated activity for the consumption servers improves performance and manageability of those servers, while servers dedicated to conversion, indexing, and replication can remain unexposed to the web.

IMPLEMENTING A CLUSTER CONFIGURATION

OVERVIEW

This section covers the following topics:

- ❖ [Installing a Cluster](#) (page 3-3)
- ❖ [Reconfiguring an Updated Cluster](#) (page 3-5)
- ❖ [Remote Administration](#) (page 3-5)
- ❖ [Configuring the Cluster](#) (page 3-6)
- ❖ [Backing Up Clustered Content Servers](#) (page 3-6)
- ❖ [Cluster Installation Details](#) (page 3-7)



Important: You should review the *Content Server Clustering Concepts Guide* before setting up a cluster. That document contains important planning and installation considerations, and architecture scenarios.

Due to the complex relationship between existing infrastructure and cluster configuration, detailed installation instructions cannot be provided here. Contact Consulting Services for installation assistance.

The following steps are required to install a new cluster:

1. Install the server software on the first node server and start the required services. See [Setting Up the Install Node Server](#) (page 3-2).
2. Configure the server software on the nodes to build the cluster. See [Configuring Additional Nodes](#) (page 3-4).

SETTING UP THE INSTALL NODE SERVER

Use the following procedures to install the server software on the install node server in the cluster configuration.

1. Set up the file server with a shared mount point. Unless specified otherwise, the shared mount point is referred to as the *<shared_drive>* throughout the remainder of this manual.
2. Install the web server if one is not already installed.
3. Install Content Server to the shared mount point on the file server (see the *Content Server Installation Guide for UNIX* for complete content server software installation instructions).
 - ❖ When prompted to supply the **IP address filter**, you must include the IP addresses separated by the pipe symbol (|) for all the machines in the cluster.
 - ❖ When prompted to provide the **destination directory** of the core components, specify the directory name of the share point on the file server.
4. After installation, change the value of IntradocDir for all nodes to a UNC path or a mapped drive. This value should be changed in the following locations:
 - ❖ *<install_dir>/admin/bin/intradoc.cfg*
 - ❖ *<install_dir>/bin/intradoc.cfg*
 - ❖ *<shared_dir>/stellent/admin/data/servers/servers.hda*

To use the UNC path method to specify a shared resource do the following.

5. Open the intradoc.cfg file located in *<install_dir>/bin* and set the IntradocDir variable as follows:

```
IntradocDir=//<servername>/shared_directory
```
6. Do this for every server in the cluster, pointing to the main content server instance.

INSTALLING A CLUSTER

For Content Server release 7.5.1 and later, an Installer script can be run to set up your cluster configuration. This script configures the content server in the appropriate specified directory and sets up the configuration files. For details about the tasks performed by the Installer, see [Cluster Installation Details](#) (page 3-7).

Using the Installer

With the Installer command, you can set up and configure the nodes in your cluster. Run it for each node, using a different name for each node in the cluster.



Important: The Installer command is not interactive. It does not show a progress meter so you can gauge the progress of the installation. This command may require several minutes to complete.

The Installer is actually a series of commands that execute in sequence. It should be entered on one continuous line on the command line.

```
Installer
--set-ClusterNodeIntradocDir=<stubdir>
--set-ClusterNodeName=<nodename>
--set-ClusterNodeAddress=<ip.of.node>
--set-ClusterBinDirRule=<local|shared>
ConfigureClusterNode
ConfigureAdminClusterNode
```

You must specify the following configuration values for the variables on the Installer command line:

- ❖ **ClusterNodeIntradocDir** *stubdir*: Enter the name of the local installation directory where the content server will be installed. For example, /stellent-nodes/StellentNode1. Omit this parameter if you are using a high performance clustering file system, such as Veritas or AFS.
- ❖ **ClusterNodeName** *nodename*: Enter the name of the cluster node or the DNS alias of the node. For example, StellentNode1.
- ❖ **ClusterNodeAddress** *ip.of.node*: Enter the IP address of the cluster node or the node name.
- ❖ **ClusterBinDirRule**: Specify if the file system is *local* or *shared*. If you are using a high-performance clustering file system, this must be set to *shared*.

Copy the filter plug-in definition file (systemfilters.hda) from the shared directory to the stub directory. After doing so, change the IHS /usr/HTTPServer/conf/httpd.conf file to load the Stellent module (filter) from the local stub directory, as in the following example for AIX and Apache. Each line is on a single line in the .conf file:

```
% LoadModule IdcApacheAuth
    <stellent_stub>/shared/os/aix/lib/IdcApacheAuth.so
% IdcUserDB stellent <stellent_stub>/data/users/userdb.txt
```

Add the following lines to both the /bin/intradoc.cfg and the /admin/bin/intradoc.cfg files using a text editor. See [Appendix A \(Configuration Settings\)](#) for more information about other configuration variables that can be used:

```
DisableSharedCacheChecking=true
UserPublishCacheDir=<stellent-stubs>/data/users
SharedDir=<stellent-stubs>/shared
ArchiverDoLocks=true
```

You should then configure the web server to point to the Content Server and if needed, configure the load balancer.

Configuring Additional Nodes

As described in [Installing a Cluster](#) (page 3-3), you can use the Installer command to set up additional nodes. The Installer command will copy the appropriate content to the right directories and assign the correct values to the variables.

After setting up the nodes, open the intradoc.cfg file from the stellent-stubs/bin and stellent-stubs/admin/etc/ directories. Add the following lines in a text editor:

```
DisableSharedCacheChecking=true
ClusterNodeName=<node_name>
ClusterGroup=<cluster_name>
SocketServerAddress=<server_IP_address>
```

node_name is the identifier for the node and *cluster_name* is the identifier for the cluster. All nodes should have a different node name, but should have the same cluster group name.

For some infrastructure environments, you need to copy the filter plug-in definition file (systemfilters.hda) from the shared directory to the stub directory. Contact Oracle Consulting Services for details about other configuration settings which might be needed for your environment.

Start the Content Server, Admin Server, and IHS processes then verify that the log and PID files are created in the <stellent_stub>/admin/etc and <stellent_stub>/etc directories.

In the second cluster node, you can now see two listings for the content server. One is the local node and the second is the remote node. You can administer either content server from here. If you navigate back to the first cluster, you will see that its configuration has been updated so that it can control both nodes.

RECONFIGURING AN UPDATED CLUSTER

Prior to Content Server version 7.5.1, cluster installations were not admin server aware. Therefore you may have a misconfigured admin server. To reconfigure an updated cluster, perform these steps:

1. Stop all admin server cluster nodes.
2. Rename the `/netshare/stellent/admin/data/servers` directory to `/netshare/stellent/admin/data/old-servers`.
3. Restart all admin server cluster nodes.
4. Re-add all cluster nodes using the instructions in [Configuring Additional Nodes](#) (page 3-4).

REMOTE ADMINISTRATION

You can use the Admin Server to remotely administer Content Servers that are not part of a clustered group. Use these steps to set up remote administration:

1. Navigate to the Admin Server for one of the systems to be administered.
2. Click **Add Existing Server**. Select **Remote Server**, and click **Submit**.
3. Enter the IP address of the computer that is running the other Admin Server.
4. Enter the port that the other Admin Server is listening on (typically 4440). Click **Next**.

The Admin Server now connects to the remote Admin Server, exchanging configuration data. You can administer both content servers from either Admin Server.



Note: A change to one Admin Server's configuration (such as the addition of a new local node) is not automatically communicated to all remote Admin Servers. To re-synchronize information between Admin Servers, you must choose **Add Existing Server** and select **Remote Server** again. Re-specify the information to the Admin Server and click **Next**. The information will now be synchronized between the two servers again.



Note: Admin applets must talk directly to a specific node to function correctly. You must access the admin applets for a particular node from the node directly, and not through the Load Balancer.

CONFIGURING THE CLUSTER

After installation, the cluster must be configured.

- ❖ **Configuring the Web Servers:** Web server configuration for a cluster is the same as that for a regular Content Server installation. See the *Content Server Installation Guide for UNIX* for details.
- ❖ **Securing the File Servers:** File server security will vary according to the hardware and software installed at the site. You should make sure to configure your file system security so that all servers can access the files and set other security measures according to the needs of your site.

BACKING UP CLUSTERED CONTENT SERVERS

You should back up your clustered content servers just as you would a single content server machine. See the *Content Server System Migration Guide* for details about archiving and backups.

- ❖ Back up the primary content server installation directories and those containing data, databases, and data collections.
- ❖ Back up applicable stub_dir on each server. However, because the information in the stub_dir does not change, the backup only needs to be completed once.
- ❖ It is not practical to shut down all parts of the system before performing a backup, although this would ensure that the backup reflects a single point in time and prevents activity in the system that could corrupt the backup records. Try to perform your backups at times of least system activity, realizing that you may not be able to obtain an exact copy.
- ❖ The content server administrator should coordinate with the database administrator so that a database backup can be performed while the clustered system is shut down.

Each site should determine the appropriate backup procedures that are necessary for their data, infrastructure, and corporate culture.

CLUSTER INSTALLATION DETAILS

The following steps describe the process that takes place when you use the Installer script. **This is background information only**; the steps described here are performed by the Installer and do not need to be performed manually to configure a cluster. It is provided as an explanation of the process that occurs.

1. The Content Server is installed on the first node server in a directory on the shared file system (for example, `c:/netshare/stellent`).
2. The web server is installed.
3. The value of `SocketHostAddressSecurityFilter` in the `/netshare/stellent/config/config.cfg` file is modified to include the IP address of the server that was specified with the `ClusterNodeName` configuration value (in this example, named `StellentNode1`).
4. If this is a local cluster, a node directory is created on `/stellent-nodes/StellentNode1`. If this is a shared cluster, this step is omitted.

5. The following directories are copied to the local node directory:

`/netshare/stellent/bin`

`/netshare/stellent/admin/bin`

On a shared directory, the directories are copied to the following shared location:

`/netshare/stellent/bin-Stellent1`

`/netshare/stellent/admin/bin-Stellent1`

6. The following information from the Installer command line is added to the `intradoc.cfg` files in the above directories:

`ClusterNodeName=<the name specified with the Installer>`

`ClusterNodeAddress=<the address specified with the Installer>`

`ClusterBinDirRule=<local or shared, as specified with the Installer>`

The remaining steps pertain to local clusters only.

7. The following directories from the network shared directory are copied to the local cluster directory (named `stellent-nodes/StellentNode1` in this example):

`netshare/stellent/etc/`

`netshare/stellent/admin/etc`

8. The `/stellent-nodes/StellentNode1/etc/config` file is created and the following information is added to the first line of the file:

```
INTRADOCMS_HOME="value for ClusterNodeIntradocDir"
```

Using the previous example, this would be:

```
INTRADOCMS_HOME="/stellent-nodes/StellentNode1/"
```

9. The `/stellent-nodes/StellentNode1/admin/etc/config` file is created and the following information is added to the first line of the file:

```
INTRADOCMS_HOME="value for ClusterNodeIntradocDir/admin"
```

Using the previous example, this would be:

```
INTRADOCMS_HOME="/stellent-nodes/StellentNode1/admin"
```

10. In the config files created in steps 8 and 9 the following lines are added:

```
INTRADOCMS_ETC="${INTRADOCMS_HOME}/etc" INTRADOCMS_PID="${INTRADOCMS_ETC}/pid"  
INTRADOCMS_LOG="${INTRADOCMS_ETC}/log"  
INTRADOCMS_HISTORY="${INTRADOCMS_ETC}/history"  
INTRADOCMS_STARTFILE="${INTRADOCMS_ETC}/.starting"
```

CONTENT SERVER CONFIGURATIONS

OVERVIEW

This chapter provides summaries and results from tested versions of cluster configurations with Content Server. In this guide, this combination of servers is referred to as a *standard cluster configuration*. Each tested version combines the standard cluster configuration with another Universal Content Management product or third-party product.

This chapter covers these topics:

- ❖ [Clusters and Oracle Products](#) (page 4-1)
- ❖ [Clusters and HTTP Servers](#) (page 4-3)

CLUSTERS AND ORACLE PRODUCTS

The information in this section identifies any important pre-installation or configuration considerations that should be addressed before specific products are installed. For example, some products are contingent upon other Universal Content Management or third-party products to run or function properly. Also, some installation or configuration options generally available with various Universal Content Management products may not be supported in a cluster configuration.

CoreWebDAV and Folders Component

Prior to Content Server version 10gR3 the Core WebDAV (Web-Based Distributed Authoring and Versioning) server was only available with the Folders component.

To include CoreWebDAV support in a cluster configuration:

1. Ensure that the component has been installed with the Content Server 10g Release 3 (10.1.3.3.0) For the standard cluster configuration, there are no additional requirements needed.
2. Enable the component.

Content Publisher

Install the Content Publisher components (Site Builder and Site Server) on the first node server according to the instructions in the *Content Publisher Installation & Configuration Guide*. During the installation process on the first node server, you will need to provide the name of the content server that will provide source documents for your Web publications and also the name of the HTTP server hosting that content server.

When prompted for the content server and HTTP server names, you will need to point to the load balancing machine. This is because in the cluster configuration, the hostname used by browsers when they connect to the system is the hostname of the load balancing machine.

Inbound Refinery

Inbound Refinery does not perform any conversions. It runs in the background and manages all conversion add-ons (filters) at the input side of the content server system. The types of files that are converted and the result of the conversion depends on the specific conversion modules (add-ons or filters) installed on the Inbound Refinery computer.

The optional conversion add-ons (filters) are distributed individually and installed separately. Install Inbound Refinery on the first node server according to the instructions in the *Inbound Refinery, PDF Converter, and XML Converter Installation Guide for UNIX*. (For the standard cluster configuration, there are no additional requirements aside from the instructions in the referenced installation guide.)

Dynamic Converter

To install the Dynamic Converter component and extract the conversion filters, refer to the instructions provided in the *Dynamic Converter Installation Guide*. (For the standard cluster configuration, there are no additional requirements aside from the instructions in the referenced installation guide.)

CLUSTERS AND HTTP SERVERS

Standard cluster configurations support the following web servers:

- ❖ SunOne or IBM HTTP Server
- ❖ Apache

Depending on which web server you decide to use in your cluster configuration, there may be some important pre- or post-installation considerations that apply. Some web server features may actually conflict with Content Server functionality. Or, there may be authentication or other security issues related to specific web servers. For more detailed information, refer to the *Content Server Installation Guide for UNIX* and the *Content Server Security Integration Guide*.

Aside from any potential implementation or security considerations, there are no other requirements in addition to the instructions in the referenced installation or security guides.



CONFIGURATION SETTINGS

OVERVIEW

This appendix covers the following topics:

- ❖ [Introduction](#) (page A-1)
- ❖ [Local Directory Variables](#) (page A-2)
- ❖ [Automation Variables](#) (page A-3)
- ❖ [Locking and Tracing Variables](#) (page A-4)

INTRODUCTION

In recent releases of Content Server, several configuration variables have been introduced to control automation activity for a load balancing or failover clustering scenario. These flags allow the refinery, indexing, and archiving functions to be disabled. Other enhancements and additional configuration settings allow for more flexibility in load balancing configurations.

The configuration variables described in this appendix can be used to set up local data references, locking, automation, and tracing in a clustered environment. These flags can be set in one of the following files:

- ❖ A component configuration file
- ❖ The `<install_dir>/config/config.cfg`
- ❖ The `<install_dir>/bin/intradoc.cfg`

LOCAL DIRECTORY VARIABLES

The following variable settings are available to help maximize the content server performance in a cluster environment.

UserPublishCacheDir

This setting allows the web server plug-in to refer to data in a local directory (the "stub" directory) rather than the file system where the content server is installed. This configuration is only used to configure the system.hda file. All files published by the content server for consumption by the web server plug-in will be written to files in this directory.

Use the format:

```
UserPublishCacheDir=<stub_dir>/data/users
```

where <stub_dir> is the path of the stub directory (such as C:/stellent/).

SharedDir

This setting allows the content server to refer to code and resource files in a local directory (the "stub" directory) rather than the file system where the content server is installed.

Use the format:

```
SharedDir=<stub_dir>/shared
```

where <stub_dir> is the path of the stub directory (such as C:/stellclust/).

DisableSharedCacheChecking

This setting stops the file stamp checking on all loaded resources. This means that resource files defined in components will only be refreshed when the content server is restarted. This flag can be used to limit file system activity. It might be useful if the content server file system is on a network share. This flag does not stop timestamp checking on *hcst* and *hcsp* pages.

Use the format:

```
DisableSharedCacheChecking=true
```



Note: This setting is particularly important in production environments where file checking and time stamp overhead can seriously degrade performance. In a development environment this setting can be optional.

IdcServerSocketQueueDepth

This setting specifies the depth of the TCP/IP socket queue. It can be used to tune Content Server performance, particularly in a clustered environment. Decreasing this value might limit Content Server throughput. For example, under very heavy loads, performance can be adversely affected because then the queue is exhausted, the TCP/IP may perform random retries. The retries could limit the server throughput.

Use the format:

```
IdcServerSocketQueueDepth=200
```

AUTOMATION VARIABLES

In a load balancing/failover scenario, automation on more than one server in the cluster can remain enabled because of recent content server synchronization enhancements. There is no need to disable the automation feature unless system performance is compromised due to one of the following scenarios:

- ❖ If you are scaling beyond three or four servers in the cluster. In this case, all but a core set of servers can be set to have no automation. (Generally, semi-static systems that do not perform extensive search queries or database retrievals do not benefit by scaling beyond four servers.)
- ❖ If the shared drive between servers is on an NFS mount and has noticeable performance issues, automation can be disabled. In this case, the lag time in file name and content changes can cause the content server synchronization logic to fail.

The following variables can be used to control automation activity in a clustered environment.

NoAutomation

True = The content server will perform no automation activity (refinery conversion, indexing, or archiving).

IsAutoQueue

False = The content server will not process documents converted by the Inbound Refinery.

IsAutoSearch

False = The content server will not index documents automatically or automatically export documents to an archive.

IsAutoArchiver

False = The content server will not automatically import or transfer archives.

LOCKING AND TRACING VARIABLES

Locking and tracing variables are used to control the content server indexing and archiving processes.

- ❖ Locking mechanisms ensure that automated indexing, imports, and transfers can be performed by only one server at a time when multiple servers share an instance. A lock is created by placing a lock file in an appropriate directory. For a lock to be active, the timestamp on the lock file must change periodically; otherwise any contending content server will treat the lock as expired.
- ❖ Indexing uses a background monitoring thread that creates an indexing lock file.
- ❖ Archiving has a similar locking algorithm that can be activated at the file system level, but locking does not have to be enabled for archive files.
- ❖ Trace settings define how information about the automated processes and locking functions is recorded.

ArchiverDoLocks

True = Locking is enabled for archive files. (If only one server is doing automated imports or transfers in a load balancing configuration, this flag does not need to be set to true.)

ArchiverTraceLocks (deprecated)

True=Traces when locks are set and release on an archive file.

In place of this variable, use the System Audit Information page and select **archiverlocks** in the Edit Active Console Output Tracing section of the page.

ArchiverTraceLockMaintenance (deprecated)

True=Traces archive lock file timestamp maintenance for any outstanding active locks and timestamp changes on archive locks maintained by other servers.

In place of this variable, use the System Audit Information page and select **archiverlocks** and **Full Verbose Tracing** in the Edit Active Console Output Tracing section of the page.

ArchiverTouchMonitorInterval

This setting defines the interval that a content server waits before it treats an archive lock file with an unchanged timestamp as having expired. The content server will also update the timestamp for any locks it is maintaining twice during this interval. The default is 12000 milliseconds (12 seconds).

With smaller values, the locks can alternate between content servers more quickly. With larger values, it is more likely that high activity on the network drive could keep a touch request delayed for many seconds, possibly causing a content server to falsely report a lock as having expired. Falsely expired locks can cause duplicated archiving effort, but should not have any long-term impact on the state of the content server.

IndexerTraceStartStop (deprecated)

True = Traces the starting and stopping of index cycles.

In place of this variable, use the System Audit Information page and select `indexermonitor` in the Edit Active Console Output Tracing section of the page.

IndexerTraceCycleListChanges (deprecated)

True = Traces changes in the list of active indexing cycles that is maintained by the content server. At certain points, each indexing cycle needs an exclusive lock, at which time it must wait for the other indexing cycles to become inactive. This flag can help audit the success of such requests.

In place of this variable, use the System Audit Information page and select **indexermonitor** and **Full Verbose Tracing** in the Edit Active Console Output Tracing section of the page.

IndexerTraceLockFileChanges (deprecated)

True = Traces changes in the indexer lock and suspension files.

In place of this variable, use the System Audit Information page and select **indexermonitor** and **Full Verbose Tracing** in the Edit Active Console Output Tracing section of the page.

IndexerTouchMonitorInterval

This setting defines the interval that a content server waits before it treats an index lock file with an unchanged timestamp as having expired. The content server will also update the timestamp for any locks it is maintaining twice during this interval. The default is 12000 milliseconds (12 seconds).

LOCAL DIRECTORY VARIABLES FOR THE STUB_DIR

The following variable settings are available to help maximize the content server performance in a cluster environment.

UserPublishCacheDir

This setting allows the web server plug-in to refer to data in a local directory (the `stub_dir`) rather than the file system where the content server is installed. This configuration is only used to configure the `system.hda` file. All files published by the content server for consumption by the web server plug-in will be written to files in this directory.

Use the format:

```
UserPublishCacheDir=<stub_dir>/data/users
```

where `<stub_dir>` is the path of the stub directory (such as `/A1/users/projects/stellclust/`).

SharedDir

This setting allows the content server to refer to code and resource files in a local directory (the `stub_dir`) rather than the file system where the content server is installed.

Use the format:

```
SharedDir=<stub_dir>/shared
```

where `<stub_dir>` is the path of the stub directory (such as `/A1/users/projects/stellclust/`).

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  
*
```

Third Party Licenses

- * 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
* 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:

1. 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/>.

Third Party Licenses

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



A

- ArchiverDoLocks
 - Locking and tracing variable, A-4
- ArchiverTouchMonitorInterval
 - Locking and tracing variable, A-5
- ArchiverTraceLockMaintenance
 - Locking and tracing variable, A-5
- ArchiverTraceLocks
 - Locking and tracing variable, A-4
- Audience, 1-2
- Automation Variables, A-3
- Automation variables
 - Configuration settings, A-3
 - IsAutoArchiver, A-4
 - IsAutoQueue, A-3
 - IsAutoSearch, A-4
 - NoAutomation, A-3

B

- Backing up
 - Content servers, 3-6

C

- Cluster configurations
 - With Content Publisher, 4-2
 - With CoreWebDAV, 4-2
 - With Dynamic Converter, 4-3
 - With Inbound Refinery, 4-2
- cluster hardware
 - overview, 2-2
- cluster installation, 3-2
- Clustered content servers
 - Backing up, 3-6
- Configuration settings
 - ArchiverDoLocks, A-4
 - ArchiverTouchMonitorInterval, A-5

- ArchiverTraceLockMaintenance, A-5
- ArchiverTraceLocks, A-4
- Automation variables, A-3
- DisableSharedCacheChecking, A-2
- IdcServerSocket QueueDepth, A-3
- IndexerTouchMonitorInterval, A-6
- IndexerTraceCycleListChanges, A-5
- IndexerTraceLockFileChanges, A-5
- IndexerTraceStartStop, A-5
- IsAutoArchiver, A-4
- IsAutoQueue, A-3
- IsAutoSearch, A-4
- Local directory variables, 2-3, A-1, A-2
- Locking and tracing variables, A-4
- NoAutomation, A-3
- SharedDir, A-2, A-6
- UserPublishCacheDir, A-2, A-6

- Content Publisher, 4-2

- Content servers
 - Backing up, 3-6
- Conventions, 1-2
- CoreWebDAV, 4-2

D

- DisableSharedCacheChecking
 - Local directory variable, A-2
- Dynamic Converter, 4-3

H

- Hardware configuration
 - Time synchronization, 2-3
- HTTP servers, 4-3

I

- IdcServerSocket QueueDepth
 - Local directory variable, A-3

Index

- Inbound Refinery, 4-2
- IndexerTouchMonitorInterval
 - Locking and tracing variable, A-6
- IndexerTraceCycleListChanges
 - Locking and tracing variable, A-5
- IndexerTraceLockFileChanges
 - Locking and tracing variable, A-5
- IndexerTraceStartStop
 - Locking and tracing variable, A-5
- Install node
 - Server software installation, 3-2
- Installation, 3-2
- installing node software, 3-4
- IsAutoArchiver
 - Automation variable, A-4
- IsAutoQueue
 - Automation variable, A-3
- IsAutoSearch
 - Automation variable, A-4

L

- Local Directory Variables, A-2
- Local directory variables
 - Configuration settings, 2-3, A-1, A-2
 - DisableSharedCacheChecking, A-2
 - IdcServerSocket QueueDepth, A-3
 - SharedDir, A-2, A-6
 - UserPublishCacheDir, A-2, A-6
- Locking and tracing variables
 - ArchiverDoLocks, A-4
 - ArchiverTouchMonitorInterval, A-5
 - ArchiverTraceLockMaintenance, A-5
 - ArchiverTraceLocks, A-4
 - Configuration settings, A-4
 - IndexerTouchMonitorInterval, A-6
 - IndexerTraceCycleListChanges, A-5
 - IndexerTraceLockFileChanges, A-5
 - IndexerTraceStartStop, A-5
- Locking variables, A-4

N

- NoAutomation
 - Automation variable, A-3

O

- Operating systems
 - Supported, 2-2

R

- Recommendations
 - Pre-implementation considerations, 2-2
- Requirements, 2-1

S

- Server software
 - Installing on install node, 3-2
- SharedDir
 - Local directory variable, A-2, A-6
- Solaris
 - Supported file server OS, 2-2
- Symbols, 1-2
- Synchronized clocks
 - Hardware requirement, 2-3

T

- Time synchronization
 - Hardware requirement, 2-3
- Tracing variables, A-4

U

- UserPublishCacheDir, A-2, A-6
 - Local directory variable, A-2, A-6

W

- Web Servers
 - Configuring, 3-6
- Web servers
 - Supported, 2-2