

ContentDirector™ 2.0 Administration Guide

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Information e-mail: info@netapp.com
Information Web: <http://www.netapp.com>

Part number 210-00073
November 2001

Copyright and trademark information

Copyright information

Copyright © 1994–2001 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 1998–2001 The OpenSSL Project. All rights reserved.

No part of this book covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Portions of this product are derived from the Berkeley Net2 release and the 4.4-Lite-2 release, which are copyrighted and publicly distributed by The Regents of the University of California.

Copyright © 1980–1995 The Regents of the University of California. All rights reserved.

Portions of this product are derived from NetBSD, which is copyrighted by Carnegie Mellon University.

Copyright © 1994, 1995 Carnegie Mellon University. All rights reserved. Author Chris G. Demetriou.

Permission to use, copy, modify, and distribute this software and its documentation is hereby granted, provided that both the copyright notice and its permission notice appear in all copies of the software, derivative works or modified versions, and any portions thereof, and that both notices appear in supporting documentation.

CARNEGIE MELLON ALLOWS FREE USE OF THIS SOFTWARE IN ITS “AS IS” CONDITION. CARNEGIE MELLON DISCLAIMS ANY LIABILITY OF ANY KIND FOR ANY DAMAGES WHATSOEVER RESULTING FROM THE USE OF THIS SOFTWARE.

Software derived from copyrighted material of The Regents of the University of California and Carnegie Mellon University is subject to the following license and disclaimer:

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notices, this list of conditions, and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University 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 REGENTS 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 REGENTS 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.

Portions of the software were created by Netscape Communications Corp.

The contents of those portions are subject to the Netscape Public License Version 1.0 (the "License"); you may not use those portions except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/NPL/>.

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is Mozilla Communicator client code, released March 31, 1998.

The Initial Developer of the Original Code is Netscape Communications Corp. Portions created by Netscape are Copyright © 1998 Netscape Communications Corp. All rights reserved.

Software derived from copyrighted material of Network Appliance, Inc. is subject to the following license and disclaimer:

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. The use and purchase of this product do not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Network Appliance.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp and the Network Appliance design are registered trademarks of Network Appliance, Inc. in the United States, Canada, and the European Union. Network Appliance is a registered trademark of Network Appliance, Inc. in Monaco and a trademark of Network Appliance, Inc. in the United States and Canada. FAServer is a registered trademark of Network Appliance, Inc. in the United States and the European Union. NetCache is a registered trademark of Network Appliance, Inc. in the European Union and Japan, and a trademark of Network Appliance, Inc. in the United States. SnapCopy is a registered trademark of Network Appliance, Inc. in the European Union and a trademark of Network Appliance, Inc. in the United States. WAFL is a registered trademark of Network Appliance, Inc. in the European Union and a trademark of Network Appliance, Inc. in the United States and Canada. FilerView and SecureShare are registered trademarks of Network Appliance, Inc. in the United States. Data ONTAP is a trademark of Network Appliance, Inc. in the United States and Canada. Snapshot is a trademark of Network Appliance, Inc. in the United States and the European Union. ApplianceWatch, BareMetal, ContentDirector, ContentReporter, DataFabric, SecureAdmin, Serving Data by Design, Smart SAN, SnapManager, SnapMirror, SnapRestore, and Web Filer are trademarks of Network Appliance, Inc. in the United States.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the United States and/or other countries.

Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries.

RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

	Preface	.ix
Chapter 1	Overview of ContentDirector	1
	About ContentDirector	2
	Understanding how ContentDirector maps to your network	3
	Resources in the management console	5
	User types and tasks	6
	Where to go next	7
Chapter 2	How to Start and Stop ContentDirector	9
	Starting and stopping ContentDirector on Windows NT	11
	Starting and stopping ContentDirector on Windows 2000	13
	Starting and stopping ContentDirector on Solaris systems	15
	Viewing the management console on a remote machine	17
Chapter 3	Learning about the User Interface	19
	Information about the user interface	20
	What you see	21
	Context dependent features	22
	Summary of the main screen panels	23
	Views	25
	The Users view	26
	Jobs view	28
	Sources view	30
	Destinations view	31
	Policies view	33
	The Network view	35
	The menu bar and the toolbar	36
	The Status window	38

Customizing the display	39
Customizing the screen layout	40
Customizing status window content	41
Configuring job-related displays	44

Chapter 4

Configuring Users and User Rights	47
Overview of user rights	48
Components of user rights	49
Overview of user rights components	50
Domains	51
User groups	52
Users	53
User rights	54
Planning your user rights scheme	55
User rights information worksheet	56
Implementing your user rights scheme	58
Creating domains	59
Creating user groups	61
Creating domain administrators	64
Creating regular users	66
Granting instance rights	68

Chapter 5

Preparing for Content Distribution	71
Setting up content distribution resources	72
Setting content distribution preferences	73
Creating sources	77
Creating destination groups	80
Creating policies	82
Creating jobs	91
Reducing local disk space requirements	97
Modifying content distribution resources	100
Adding and deleting destinations	101
Modifying and deleting sources	104
Modifying and deleting destination groups	107
Modifying and deleting policies	109
Modifying and deleting jobs	111

Chapter 6	Distributing Content	.113
	Learning about content distribution	.114
	Technical description of the content distribution process	.115
	Stages of the distribution event	.116
	An example of running a job	.119
	Distributing a file.	.120
	Determining job status	.122
	Interpreting job status from the Status window	.123
	Using Job History	.125
	Deleting content from a destination	.127
	Changing master keys to ensure data security	.128
	Generating a new master key.	.129
	Changing master keys on the CDA on NetCache	.130
Appendix A	Error Code Descriptions	.137
Appendix B	Values in the Jobs View Status Field.	.147
Appendix C	Tree Display Icons.	.149
Appendix D	Default Mime Types.	.151
Appendix E	Mapped Drive Permissions	.165
Appendix F	Resolving Port Conflicts	.167
	Reconfiguring the RMI port	.168
	Reconfiguring the transport port	.170
	Glossary	.173
	Index	.179

Preface

Introduction

This guide describes how to configure, manage, and operate Network Appliance[™] ContentDirector 2.0.

Audience

This guide is for system administrators and others who will configure or use ContentDirector. It assumes you are familiar with your operating system, your network, and the basic concepts of how data moves over a network. If you plan to distribute streaming media content, you also need a basic understanding of streaming media protocols and service. This guide does not cover basic system or network administration topics, such as IP addressing, routing, and network topology; it emphasizes the characteristics of the ContentDirector software.

Terminology and conventions

The term “ContentDirector” refers to a complete configuration including all components of the software.

This guide uses the term “type” to mean pressing one or more keys on the keyboard. It uses the term “enter” to mean pressing one or more keys, and then pressing the Enter key.

Keyboard conventions

When describing key combinations, this guide uses the hyphen (-) to separate individual keys. For example, “Ctrl-D” means pressing the “Control” and “D” keys simultaneously. Also, this guide uses the term “Enter” to refer to the key that generates a carriage return, although the key is named “Return” on some keyboards.

Typographic conventions

The following table describes typographic conventions used in this guide.

Convention	Type of information
<i>Italic type</i>	Words or characters that require special attention. File names. Placeholders, for information you must supply. Book titles in cross-references.
Monospaced font	Command names. Information displayed on the system console or other computer monitors. The contents of files.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless you must type it in uppercase letters.

Special messages

This guide contains special messages that are described as follows:

Note

A note contains important information that helps you install or operate the system efficiently.

Caution

A caution contains instructions that you must follow to avoid loss of data.

About this chapter

This chapter provides a general product overview and introduction to basic features.

Topics in this chapter

This chapter contains the following topics:

- ◆ “[About ContentDirector](#)” on page 2
- ◆ “[Understanding how ContentDirector maps to your network](#)” on page 3
- ◆ “[Resources in the management console](#)” on page 5
- ◆ “[User types and tasks](#)” on page 6
- ◆ “[Where to go next](#)” on page 7

About ContentDirector

Function and key features

ContentDirector automates the distribution of content over a network from a central source to multiple destinations. The network can be a local area network, intranet, or internet. Reusable, user-defined rules govern content distribution tasks, including the location of source files and destinations, scheduling, and how the content is handled over the network.

Advantages of automated distribution

Rules-driven content distribution offers several advantages over manual content distribution:

- ◆ Pro-active distribution of content before it is needed - “push” instead of “pull”
- ◆ Bandwidth savings: content is sent at wire speed, at off-peak times
- ◆ Version control and replication over groups of destinations
- ◆ Consistent, accurate, timely content distribution without supervision
- ◆ Simple graphical interface allows users without systems-level expertise to perform routine tasks
- ◆ Repeatability
- ◆ Standardization across multiple departments or organizations

Components of distributed architecture

ContentDirector has a distributed architecture based on the following components:

- ◆ The *core application* runs at a central location and provides command and control over ContentDirector resources
- ◆ The *management console* provides a graphical user interface to the core application. You can run the management console directly from the system where the core application is resident, or remotely via a browser applet.
- ◆ ContentDirector *agents* (CDAs) are installed on nodes at the periphery of the network. Agents act as content destinations or as pass-through points to control content flow.

Understanding how ContentDirector maps to your network

Ways ContentDirector maps to your network

ContentDirector maps to your network in two different ways:

- ◆ Physically, through *transport groups* to move data over the network
- ◆ Logically, through *destination groups* to target specific content to individual destinations

Understanding the difference between transport and destination groups is critically important to using ContentDirector efficiently.

Transport groups

ContentDirector maps physically to your network through *transport groups*. Transport groups are configured during the installation process, and are defined by how devices are physically grouped. ContentDirector uses them to move content to the periphery of the network. Each transport group consists of a distribution point or *transport group leader*, and destinations or *transport group members*. The group leader passes content from the core application to the group members. Group leaders can also function as group members.

Destination groups

Superimposed on the transport groups are logical *destination groups*. Destination groups are collections of destinations that all receive the same content. They preposition content to where viewers are expected to need it.


A destination group can include destinations from any transport group. Destination groups do not need group leaders, because they deal only with the final destination. The task of deciding how the content gets there is left to the transport groups.

Users define destination groups to suit their needs through the ContentDirector management console.

How content flows through ContentDirector architecture

The flow of content in ContentDirector starts at the content source, passes through the core application and transport group leaders, and ends at the destination groups.

The following diagram illustrates the content flow through destination groups and transport groups.



Resources in the management console

User-defined resources

The management console enables you to control ContentDirector by defining the following resources:

- ◆ *Users and user rights*: access control over ContentDirector resources
- ◆ *Policies and jobs*: distribution rules that define what, where, when, and how content is moved
- ◆ *Sources*: source locations for content
- ◆ *Destination groups*: multiple destinations that receive the same content

Auto-discovered resources

When ContentDirector first starts, it contains the following resources:

- ◆ Transport group leaders and group members that are auto-discovered on your network
- ◆ Default policy for distributing content
- ◆ Default user domain (container for user rights)
- ◆ Default system administrator

You create all other resources through the management console.

User types and tasks

How user type corresponds to tasks

You log onto ContentDirector as one of the following types of users:

- ◆ System administrator: full access to all resources
- ◆ Domain administrator: full access to all resources in your domain
- ◆ Regular user: access to resources defined by your user group

In the ContentDirector user rights scheme, your tasks often depend on your user type. The task flow after installing ContentDirector can be divided into initial setup tasks and continuing . System and domain administrators perform most initial setup tasks, while regular users perform continuing administration.

Task flow

The following table describes the most efficient task flow for getting started with ContentDirector. An asterisk denotes tasks that must be performed by a system administrator. A double asterisk denotes tasks that must be performed by either a system or domain administrator.

Step	Phase	Task
1	Initial setup	Starting
2		Creating hosted sites*
3		Creating domains and domain administrators*
4		Creating user groups and users**
5		Creating policies
6		Creating jobs
7		Creating destination groups
8	Continuing administration	Running and monitoring jobs
9		Modifying resources as necessary to accommodate changes in strategy or network topology

Where to go next

First-time users

All first-time users should see the following chapters before logging onto ContentDirector:

- ◆ [“How to Start and Stop ContentDirector”](#) on page 9
- ◆ [“Learning about the User Interface”](#) on page 19

All others

Once you have logged onto ContentDirector and become familiar with the management console, use the table below to locate where you need to go next in this guide.

If you need to...	Then look at...
Create domains, user groups or users	“Configuring Users and User Rights” on page 47
Create sources	“Creating sources” on page 77
Create: ◆ Policies ◆ Jobs ◆ Destination groups	◆ “Creating policies” on page 82 ◆ “Creating jobs” on page 91 ◆ “Creating destination groups” on page 80
Run jobs	“An example of running a job” on page 119

About starting and stopping ContentDirector

Starting and stopping ContentDirector consists of starting and stopping the core application services, the ContentDirector Agent (CDA) services, and the management console, which can be initiated through an application or a Java applet.

The procedures for starting and stopping the ContentDirector services and the management console application vary according to platform.

The core application, management console, and applet run on the following platforms:

- ◆ Windows NT
- ◆ Windows 2000
- ◆ Solaris

The CDA runs only on NetCache appliances. You start each CDA after enabling and configuring it on NetCache. See the *ContentDirector Installation Guide* for detailed instructions.

Starting and stopping services

The core application services start automatically when you complete the installation process and any time you reboot the system on which they reside. If you do not want to have the services running in the background and do not want to reboot the system, you can start and stop them manually. See the instructions specific to each operating system in this chapter for details.

Starting and stopping the management console

The management console is loaded manually. You can run it as an application from the core application machine or as Java applet from a remote machine through a Web browser. This distinction determines the procedure for starting and stopping the management console, but the user interface is the same in both instances.

Chapter topics

The chapter is primarily organized by platform. The exception is starting and stopping the user interface with the Java applet. Because you run the applet from a browser, the procedures for doing so are determined by the specific browser being used. An example is given at the end of the chapter.

This chapter discusses the following topics.

- ◆ [“Starting and stopping ContentDirector on Windows NT”](#) on page 11
- ◆ [“Starting and stopping ContentDirector on Windows 2000”](#) on page 13
- ◆ [“Starting and stopping ContentDirector on Solaris systems”](#) on page 15
- ◆ [“Viewing the management console on a remote machine”](#) on page 17

Starting and stopping ContentDirector on Windows NT

Starting core application services

To start the core application services on Windows NT systems, complete the following steps.

Step	Action
1	From the desktop, select Start> Control Panel > Services.
2	Select ContentDirector.
3	Click the Start button. Result: The word “Started” appears in the Status column.

Stopping core application services

To stop the core application services on a Windows NT platforms, complete the following steps.

Step	Action
1	From the desktop, select Start > Control Panel > Services.
2	Select ContentDirector.
3	Click the Stop button. Result: The word “Started” disappears in the Status column.

Starting the management console application

To start the management console application on Windows NT systems, complete the following steps

Step	Action
1	From the Desktop, click the Start button.
2	Select Content Director group from the Start menu.

Step	Action
3	Select ContentDirector Program. Result: The splash screen appears briefly, followed by the login screen.
4	Enter the username and password in the login screen. Note_____ The login ID will determine your user rights when operating ContentDirector. When you first install ContentDirector, the default user name is “super” and password is the same. _____ Result: The main screen of the user interface is displayed.

Stopping the management console application

To stop the management console application on Windows NT systems, complete the following steps.

Step	Action
1	From the main screen of the user interface, click the File button.
2	Select Exit from the drop-down menu. Result: The Exit Confirmation screen is displayed.
3	Click Yes.

Starting and stopping ContentDirector on Windows 2000

Starting the core application services

To start the core application services on Windows 2000 systems, complete the following steps.

Step	Action
1	From the desktop, select Start > Control Panel > Administrative Tools > Services.
2	Right click ContentDirector.
3	Click Start. Result: The word “Started” appears in the Status column.

Stopping the core application services

To stop the core application services on Windows 2000 systems, complete the following steps.

Step	Action
1	From the desktop, select Start > Control Panel > Administrative Tools > Component Services.
2	Right click ContentDirector.
3	Click Stop. Result: The word “Started” disappears from the Status column.

Starting the management console application

To start the management console application on Windows 2000 systems, complete the following steps.

Step	Action
1	From the Desktop, click the Start button.
2	Select Content Director group from the Start menu.
3	Select ContentDirector Program. Result: The splash screen appears briefly, followed by the login screen.
4	Enter the username and password in the login screen. Note _____The login ID will determine your user rights when operating ContentDirector. When you first install ContentDirector, the default user name is “super” and password is the same. _____ Result: The main screen of the user interface is displayed.

Stopping the management console application

To stop the management console application on Windows 2000 systems, complete the following steps.

Step	Action
1	From the main screen of the user interface, click the File button.
2	Select Exit from the drop-down menu. Result: The Exit Confirmation screen is displayed.
3	Click Yes.

Starting and stopping ContentDirector on Solaris systems

Starting the core application services

To start the core application services on Solaris systems, complete the following steps.

Step	Action
1	To start core application services, enter the following command: <code>/etc/init.d/ccm start</code>

Stopping the core application services

To stop the core application services on Solaris systems, complete the following steps.

Step	Action
1	To stop core application services, enter the following command: <code>/etc/init.d/ccm stop</code>

Starting the management console application

To start the management console application on Solaris systems, complete the following steps.

Step	Action
1	Enter the following command: <code>/opt/NTAPccm/bin/gui start</code> Result: The splash screen appears briefly, followed by the login screen.

Step	Action
2	Enter the username and password in the login screen. Note_____ The login ID determines your user rights when operating ContentDirector. When you first install ContentDirector, the default user name is “super” and password is the same. _____ Result: The main screen of the user interface is displayed.

Stopping the management console application

To stop the management console application on Solaris systems, complete the following steps.

Step	Action
1	From the main screen of the user interface, click the File button.
2	Select Exit from the drop-down menu. Result: The Exit Confirmation screen is displayed.
3	Click Yes

Viewing the management console on a remote machine

About viewing the management console from a remote machine

You view the ContentDirector management console on a remote machine by accessing a Java applet on the core application server through a Web browser on the remote system. The browser that you use determines the procedure for starting and stopping the user interface from a remote machine.

Requirements

The applet requires version 1.3 of the Java Virtual Machine (JVM) plug-in. This plug-in is available from the ContentDirector installation CD. If you do not install it from CD, the applet downloads the JVM automatically from a SUN Systems Web site when you first load the applet from the browser..

Caution

On Windows NT and 2000, the administrative user who installed and configured ContentDirector must be logged into Windows on the core application machine for ContentDirector to correctly display and access mapped drives.

Using the applet: An example using MS Internet Explorer

To view the management console using the Java applet on a Windows NT system using the Internet Explorer browser, complete the following steps.

Step	Action
1	Open the Internet Explorer browser.
2	Type the URL in the Address box of the browser menu. The syntax for the URL is the following: <code>http://<IP address or hostname>:<port>/ContentDirector</code> where the default port is 8080.
3	Click the Go button. Result: The splash screen appears briefly, followed by the login screen.

Step	Action
4	Enter the username and password in the login screen. Note_____ The login ID determines your user rights when operating ContentDirector. When you first install ContentDirector, the default user name is “super” and password is the same. _____ Result: The main screen of the user interface is displayed.

Closing the applet view of the management console

To close the applet view of the management console on a Windows NT system using the Internet Explorer browser, complete the following steps.

Step	Action
1	From the main screen of the user interface, click the File button.
2	Select Exit from the drop-down menu. Result: The Exit Confirmation screen is displayed.
3	Click Yes.

About this chapter

The purpose of this chapter is to present a quick overview of the ContentDirector user interface in order to help you navigate through the menus and find the information that is available through the user interface.

Chapter topics

This chapter discusses the following topics:

- ◆ [“Information about the user interface”](#) on page 20
- ◆ [“Summary of the main screen panels”](#) on page 23
- ◆ [“Views”](#) on page 25
- ◆ [“The menu bar and the toolbar”](#) on page 36
- ◆ [“The Status window”](#) on page 38
- ◆ [“Customizing the display”](#) on page 39

Information about the user interface

About the user interface

The user interface provides three types of functionality.

- ◆ Menus and other points of entry into the user interface that enable you to operate ContentDirector
- ◆ Components that provide information about the configuration and status of many aspects of ContentDirector
- ◆ On-line and context sensitive help

What you can see and what you can do through the user interface, are, to some extent, determined by your users rights and what components have already been created.

What you see

What you see initially

When you first log into the user interface immediately after installing ContentDirector, some of the screen components are blank. These refer to components that you have not yet created or selected and for which there are no defaults. These include the following:

- ◆ Jobs
- ◆ Content groups
- ◆ Sources

Others contain defaults or information derived from the installation and the CDA discovery process. These include the following:

- ◆ Default policy
- ◆ Super user
- ◆ Default user group
- ◆ Default user domain
- ◆ Destination devices with CDA enabled and configured
- ◆ Network segments

Users rights and what you see

When you login to the ContentDirector, your user ID determines your access rights. The components that cannot be created or changed by you are greyed out.

When you first log in as a new user, you can see and access all of the preexisting jobs, content groups, and policies to which the domain administrator has given you access rights. Subsequently, any components that you create are accessible to you.

Context dependent features

Context dependent menus

When you right click many of the elements in the ContentDirector main screen, the menu for that component is displayed. The menu is appropriate to the item selected. In the Users view, for example, if a user is selected, the menu consists of the selections New, Modify, and Delete. If no user is selected, the menu consists only of the selection New, the other two are greyed out.

Context dependent help

Many screens have a Help button that invokes context-sensitive help for the screen.

Summary of the main screen panels

Parts of the ContentDirector main screen

When you log into the ContentDirector user interface, the first thing that you see is the main screen. The main screen consists of six parts.

View toolbar: contains six icons: Users, Jobs, Sources, Destinations, Policies, and Network. These icons point to views that provide status information for the various components.

Selected view: contains the view selected in the view toolbar. There are two types of views:

- ◆ **Tables (Users, Jobs, and Policies)**
Double clicking a record within a table opens the modify dialog for that item if you have the requisite user privileges.
- ◆ **Trees (Sources, Destination Groups, and Network)**
Trees illustrate the topology of the component in question in the left panel and status information about it in the right panel.

Main menu : enables you to create and manage all of the components of ContentDirector.

Main toolbar : enables you to create a new user, job, policy, or destination group, as well as initiate operations pertaining to a specific job.


Status window: provides information from the log files about actions occurring in ContentDirector and the user interface. You can configure the level of detail that is displayed.

Status bar : displays username, transport protocol, core application server name, and database server name.

There is some overlapping functionality between the main menu and the main toolbar. Both enable you to create new components.

An illustration of the main screen

The illustration below presents the elements of the ContentDirector main screen.



Views

About views

Views provide information about the current status or configuration of ContentDirector components. Within any given view, you can see only those view components to which your user's rights give you access.

There are two types of view:

Tree displays: Topology trees contain icons representing each object in the tree. The icons specify the type of object and its status.

The Sources, Destinations, and Network views are tree views. These consist of the following:

- ❖ A left panel that displays the topology of the objects referenced by the view.
When you right click a selected object, a menu is displayed through which you can modify or delete the object. Objects that are not configurable do not display a menu.
You can select a specific object within the tree by left clicking it. This determines the content of the right panel
- ❖ A right panel that displays the detailed information about the specific object selected in the left panel. This includes (among other things) the date and time created or last modified, which reflect when Content Director's auto-discovery process recorded the presence of the component.

When the focus is on a tree display and you click Refresh, the core application polls the devices of the component in question to refresh the topology.

Tables: The Users, Jobs, and Policies views consist of tables. When the focus is on a table and you click Refresh, the core application polls the database to update the table.


The Users view

About the Users view

When you login to ContentDirector, you identify yourself as a particular user. Your user ID determines your users rights and the resources to which you have access. Users are members of domains and users groups. You find information about the configuration of users, user groups, and domains in the Users view. The user rights attached to your login ID determine which of the tables of the Users view are visible to you.

An illustration of the Users view

The Users view is shown in the following illustration.



Tables in the Users view

The User view contains three tables. You open a table by clicking its tab at the top of the view. The Users view consists of the following tables:

The Users table : is visible to system administrators, domain administrators. and users. This table consists of the following fields:

- ❖ Username: login ID
- ❖ FullName: the actual name of the user (optional)

- ❖ Privileges: the privileges attached to the type of user (user, domain administrator, system administrator)
- ❖ E-mail: the users e-mail address (optional)
- ❖ Domain: the domain to which the user belongs.

The User Groups table : is visible to system administrators and domain administrators. This table consists of the following fields:

- ❖ UserGroup Name
- ❖ Domain: the name of the domain to which the user group belongs
- ❖ # Users: the number of users in the group.

The Domains table : is visible only to system administrators. This table consists of the following fields:

- ❖ Name: the domain name
- ❖ # Users: the number of users in the domain
- ❖ # UserGroups: the number of user groups in the domain

Note


When you first log into ContentDirector after installing, the only user is the default login account “Super.” Super has system administration privileges. A system administrator cannot be a member of a domain or user group.

Information about the Jobs view

When you create a job, you define all of the parameters that define and control the distribution event that is initiated when you run the job. The information is stored in the database, and jobs can be run more than once. The Jobs view provides configuration and status information about all of the jobs to which you have access.

An illustration of the Jobs view

The following illustration presents the Jobs view.



Configuration information in the Jobs view table

You define configuration information through the user interface when you create or modify a job. The fields in the Jobs view table are blank when you log into the user interface for the first time.

The following fields contain configuration information:

- ◆ Job Name
- ◆ Creator
- ◆ Schedule
- ◆ Quorum

Note

The Quorum field refers to the commit quorum. The commit quorum is the percent of destination servers or NetCaches that must receive a job assignment to complete the distribution event.

**Status information
in the Jobs view
table**

Status information is updated as a job runs. The following fields display status information.

- ◆ Status
- ◆ Last Started
- ◆ Last Completed

The status information displayed in the Jobs view is refreshed periodically. You can configure the refresh rate and the time period for which jobs are displayed through the Preferences menu. See “Configuring job-related displays” on page 44..

The Status field provides information about the state of the last or the currently running distribution event for the job. This includes such states as completed, aborted, creating the assignment, etc. For detailed information, see Appendix B, “[Values in the Jobs View Status Field](#),” on page 147.

While a job is running, a colored diamond appears next to the status designation: The color indicates how the job is progressing.

- ◆ Green indicates that all is well.
- ◆ Yellow indicates that there is a problem.
- ◆ Red indicates that there is a a serious problem and the job assignment will probably fail.

Views


Sources view

Information in the Sources view

ContentDirector creates jobs from content sources that are stored on a filer or a conventional hard drive. The Sources view is a tree display that contains information about the topology, types of transfer, the locations, addresses, and creation date of the sources from which the core application obtains content for job assignments. The protocol of the server is specified, for example, MMS or RTSP.

An Illustration of the Sources view

The following is an illustration of the Sources view.




Destinations view

Information in the Destinations view

The target devices for content distribution are NetCache appliances. The Destinations view contains information about target devices and destination groups.


An illustration of the Destinations tab

The Destinations tab displays the topological position and other information about each destination device as illustrated below.



An illustration of the Destinations Groups tab

The Destination Groups tab displays the topology of destination groups. When the tree control of a specific destination group is expanded it lists the destinations for that group. See the illustration below.



Views

Policies view

About the Policies view


A policy is a component of a job that contains a subset of the parameters that control the distribution event. Its definition is independent of the job definition, and it can be used in multiple jobs. The Policies view displays current values of the policies listed in the table.

The check boxes are for display purposes, not user input.

When ContentDirector is first installed, you see only the default policy.

Illustration of Policy view

The following illustration presents the Policy view



Policy view fields

The following list contains the policy view fields:

- ◆ Policy Name: the name you assigned to the policy when you created it
- ◆ Encryption: whether encryption is enabled
- ◆ Verification: whether verification is enabled

- ◆ Quorum: the percent of the target devices that must receive the content for a job assignment for the core application to commit the content and complete the distribution event
- ◆ Retries: the number of times that a failed distribution event will be rerun before the final status is designated as failed
- ◆ Repushes: the number of times ContentDirector tries to re-send a job to destinations that were not part of the quorum for a successful job
Repush applies only to jobs with quorums of less than 100%.
- ◆ Compression: the level of data compression


The Network view

Information about Network view

When you install ContentDirector on a destination device, you specify its position in a transport group consisting of a group leader that communicates with the core application and one or more group members that receives content. The Network view displays the topology of each transport group and the status of each component within it.

An illustration of Network view

The following is an illustration of the Network view.



The menu bar and the toolbar

Selections on the menu bar

The menu bar (the main menu bar) enables you to create, configure, and modify of all of the ContentDirector components except the Transport Group. It provides access to drop down menus.

The menu bar contains the following items:

- ◆ The File menu provides access to the Preferences dialog, enables you to log out and log in as another user, or exit the program.
- ◆ The View menu of the menu bar provides a means of changing the View displayed in the View panel. See the illustration below.
- ◆ The Users menu enables you to manage users, user Groups, and Domains.
- ◆ The Jobs menu provides access to job related functionality. this includes defining and running jobs.
- ◆ The Sources menu enables you to create, modify, and delete sources.
- ◆ The Destinations menu enables you to delete destination devices and to create, modify, and delete content groups.
- ◆ The Policies menu enables you to create, modify, and delete policies.
- ◆ The Tools menu enables you to manage e-mail contacts, set the browser path and file, and generate a new master key.
- ◆ The Help menu opens on-line help and displays the ContentDirector license key, build number, and serial number.

Selections on the toolbar

You use the items on the toolbar to create new job components and initiate operations related to a selected job. It also contains a button for creating a new user.

The toolbar contains the following icons:

- ◆ User: creates new users
This button enables you to the create of new users, but not domains or user groups.
- ◆ Job: creates new jobs
- ◆ Policy: creates new policies
- ◆ Group (Destination): creates new destination groups
- ◆ Refresh: refreshes the data in the fields or tree nodes of the selected view

- ◆ Abort: aborts the selected job if it is currently running
- ◆ History: displays the history of job assignments of the selected job
- ◆ Run: runs the selected job manually
- ◆ Help: opens on-line help

The Status window

information about the Status window

The Status window provides the following information:

- ◆ Information about the user interface such as changes to the refresh rate for the Jobs view table or modifications to a user's account
- ◆ Information about transport communication between a group leader and the core application such as reports sent from a group leader to the core application.
- ◆ Information about the content distribution process for pending or running jobs, such as error codes, group leader IP addresses, and assignment IDs

You can configure the level of detail reported in the Status window through the Preferences window. See “[Information in the Status window](#)” on page 41.

Customizing the display

Elements you can customize

You can customize the ContentDirector user interface by altering the screen layout and by configuring parameters that determine the content of the Status window and job-related displays.

Customizing the screen layout

Elements of screen layout that can be customized

The screen layout can be customized in the following ways.

- ◆ Screen size and individual panels can be resized by clicking and dragging edges
- ◆ The order of fields in tables can be rearranged by dragging and dropping.
- ◆ The banner can be enabled and disabled through the Tools menu.

Customizing status window content

Information in the Status window

The information in the Status window comes mostly from the log files, which are located in the \CCM\logs directory. The amount and types of information taken from the log files is configurable. Some information in the status window is written directly from the GUI and cannot be found in the log files.

Configuring the Status window

You can configure the status window from the Preferences menu. The parameters that can be modified are the following.

- ◆ The level of detail that is selected determines the amount and type of information that goes to the Status window from the log file. The default is “terminal.”
- ◆ The capacity of the screen writes determines the number of characters written to the Status window before it is cleared. the default is 10,000.
- ◆ The frequency at which the screen is refreshed determines how often (in seconds) information is written to the Status window. The default is 30 seconds.

To Configure these parameters complete the following steps.

Step	Action
1	From the main menu bar, select File > Preferences. Result: The Preference screen appears
2	Click the Status window tab.

Step	Action
3	Select the desired level of detail from Level drop down menu. The choices are as follows: ◆ Terminal displays only messages that warn of problems that could cause something to fail. ◆ Warning displays all of the terminal messages plus messages concerning non-critical problems. ◆ Info displays all of the terminal and warning messages as well as messages containing information such as the status of any job that is running. ◆ Debug displays all of the above messages as well as messages containing information needed by the program developers to troubleshoot a job.
4	Click OK.
5	Enter the buffer size in the buffer size field.
6	Click OK.
7	Enter the refresh rate in the refresh rate field.
8	Click OK.

Clearing the status window

ContentDirector writes to the Status window and to log files. You can clear the status window so that it only displays new data, starting with the next screen refresh.

To clear the status window manually, complete the following steps.

Step	Action
1	From the main menu bar, click View.
2	Click Clear Status Window.

Disable autoscrolling in the status window

To make reading the logging text easier, you can disable autoscrolling in the status window.

To disable autoscrolling, complete the following steps:

Step	Action
1	Right-click anywhere on the status window.
2	Click Autoscroll Off.

Configuring job-related displays

Parameters you can Configure

You can configure the following parameters in job-related displays:

- ◆ Number of assignments that you view in the job History window
- ◆ Length of time during which a job is displayed in the Jobs view table of the main screen
- ◆ Refresh rate of the fields in the jobs view table

Note

A higher refresh rate uses more resources than a lower one.

Setting the number of assignments viewed in Job History

To set the number of assignments to be viewed in the Job History window, complete the following steps.

Step	Action
1	In the Main menu bar, select File > Preferences.
2	Click the jobs tab.
3	Enter the desired value in the “job assignments to display in history dialog” box.
4	Click OK.

Setting the length of time jobs are displayed in the Jobs view

To set the length of time jobs are displayed in the Jobs view complete the following steps

Step	Action
1	In the Main menu bar, select File > Preferences.
2	Click the jobs tab.

Step	Action
3	Enter the desired value in the “Display jobs created in the last (days)” dialog.
4	Click OK.

Refresh rate of the fields in the Jobs view table

To configure the Refresh rate of the fields in the Jobs view table, complete the following steps.

Step	Action
1	In the Main menu bar, click File > Preferences.
2	Click the Jobs tab.
3	Enter the desired value in the “Refresh rate (seconds)” dialog.
4	Click OK.

When you first log in to the ContentDirector user interface, the main screen appears with the Jobs view selected. Because you have not yet created a job, the table is empty.

About this chapter

This chapter describes the user rights features in ContentDirector and how to configure them.

Topics in this chapter

This chapter contains the following topics:

- ◆ [“Overview of user rights”](#) on page 48
- ◆ [“Components of user rights”](#) on page 49
- ◆ [“Planning your user rights scheme”](#) on page 55
- ◆ [“Implementing your user rights scheme”](#) on page 58
- ◆ [“Granting instance rights”](#) on page 68

Overview of user rights

About user rights

ContentDirector 2.0 offers a comprehensive user rights scheme that allows multiple, secure domains within a single installation of ContentDirector. Each domain contains its own distinct, multilayered set of user rules. These rules create a virtual interface to your physical network: rather than displaying all content, resources, and options, it filters them through user rights and displays only those available to the current user.

Advantages of user rights

With this user rights scheme, ContentDirector can distribute content for several companies using a single shared installation:

- ◆ A service provider can create separate domains for each company, each domain restricted to the company's resources.
- ◆ The owner company uses the domain to create and manage its own user rights, resources, and content.
- ◆ Within the company domain, smaller groups of users and their resources can also define mutually invisible, secure environments.

ContentDirector user rights are both highly customizable and flexible:

- ◆ Different users can have widely varying levels of access to the various resources of the system. You can make resources available to some users and not to others.
- ◆ User groups define user rights, so that multiple users can conveniently share the same rights.
- ◆ You can assign user rights for a specific resource at the time of its creation.

Components of user rights

Overview of components

ContentDirector enables you to control user rights through four components: domains, user groups, users, and user rights. The interaction of these components creates your virtual interface.

For detailed information

The following section discuss the individual components of ContentDirector user rights and how they interact:


- ◆ [“Domains”](#) on page 51
- ◆ [“User groups”](#) on page 52
- ◆ [“Users”](#) on page 53
- ◆ [“User rights”](#) on page 54

Overview of user rights components

Interaction of user rights components

User rights components are like nesting containers:

- ◆ Domains hold user groups
- ◆ User groups hold users and user rights



How user rights create a virtual interface

Domains and user groups are user-created, logical components that do not need to map to the physical network. This flexibility allows them to correspond to any criteria: for example, company, department, or event. Your virtual interface can transparently map users to physical and logical resources according to your needs.

Domains

Purpose of domains	The purpose of domains is to set the upper limit of user rights for their constituent user groups. Domains do this in two ways: ◆ Define access to content sources, destinations, and folders ◆ Grant resource creation rights
Interaction between domains	Domains are not visible to each other. User rights cannot cross over domains: a user group in one domain cannot see or access resources belonging solely to another domain. This limits the rights of user groups, because each user group belongs to one and only one domain.
Where to set domain boundaries	A domain is advisable whenever there are two or more collections of user groups that need to be completely separate and invisible to each other. Even if these collections of user groups share the same physical resources, such as servers or networks, ContentDirector can maintain them in separate virtual spaces through separate domains.
Example use of domains	A typical use of domains is for a content distribution network (CDN) that offers services to several companies. Companies A, B, and C each have their own domain. Even though they share the same physical network, each company's information in ContentDirector is secure from the other companies.

User groups

Definition of a user group

A user group is a container for users and their rights. Rights belong to the group, rather than to individual users. All users in the group share the same rights.

Ways user groups set user rights

User groups set user rights in two ways:

- ◆ Restrict access to domain sources, servers, and folders
- ◆ Define ability to create content distribution resources

Each user group can access only the sources, servers, and folders included in its domain. The user group can further filter access to them.

The user group also defines whether its users can create each of the following types of resources: jobs, policies, destination groups, and contacts.

Rules for users in multiple user groups

Users can belong to multiple user groups, but only under a single domain. They cannot belong to groups outside that domain. Users belonging to multiple user groups have the combined rights of all the groups. If one group provides access to a resource, the user has access.

When to use user groups

You need to define a user group whenever there are one or more users within a domain who need a specific level of access to a specific set of domain resources.

Example use of user groups

A typical use of user groups is for a company distributing content internally for several different departments. Human Resources, Marketing, and Customer Service each have their own user group. The personnel responsible for the Human Resources intranet site cannot distribute content to the Marketing intranet site, and vice versa.

Users

Types of users

ContentDirector defines three types of users:

- ◆ System administrators
- ◆ Domain administrators
- ◆ Regular users

System administrators

System administrators have access to all resources in ContentDirector. Only system administrators can create other system administrators, domains, domain administrators, and sources.

Domain administrators

Domain administrators are like system administrators within their respective domains. They have access to all resources within the domain.

Domain administrators can coordinate the activities of users within their domain, to ensure optimal timing of bandwidth usage and physical resources.

Regular users

Regular users' rights are defined by the user groups to which they belong, with one exception: Regular users always have full access to resources they create, no matter how the group access rights and memberships are altered.

Regular users typically constitute the majority of ContentDirector users.

User rights

- Types of user rights** ContentDirector has two types of user rights:
- ◆ Rights for a group of users, or class rights
 - ◆ Rights to a specific resource, or instance rights

Class rights Class rights enable a user group to access a resource. A system or domain administrator sets these rights for each user group while creating the group. Class rights are the standard method for setting up rights in ContentDirector.

Instance rights Instance rights are attached to specific resources instead of user groups. They allow regular users to access resources created by other regular users.

Planning your user rights scheme

Information you need to plan your user rights scheme

To plan your user rights scheme, you need the following information:

- ◆ What groups of users belong to the same organizations, and constitute domains
- ◆ Within each domain, what users share the same resources, and constitute user groups
- ◆ Who fills the roles of system administrator, domain administrator, and regular user

The User Rights Planning Worksheets that follow can help you plan your user rights scheme.

User rights information worksheet

Domain worksheet Use the worksheet below to plan your domains.

Domain Name	
Domain owner (company, department, or other group)	
Domain administrator(s)	
Content sources/folders by URL	
Destinations by device name	
Resource creation rights	Yes/No
Jobs	
Polices	
Destination groups	
Contacts	

User group worksheet

Use the worksheet below to plan your user groups. Use the Content sources and Destinations entries to further filter access to domain resources, if desired.

User Group Name	
Domain	
User group owner (department or other group)	
Content sources/folders by URL	
Destinations by device name	
Resource creation rights	Yes/No
Jobs	
Polices	
Destination groups	
Contacts	

Implementing your user rights scheme

Information you need before you start

Before implementing your user rights scheme, you need to complete the plan suggested in [“Planning your user rights scheme”](#) on page 55. Be sure you have a printout of your plan available before you start.

Phases of implementation

The phases of implementing your user rights scheme correspond to its major components:

- ◆ Domains
- ◆ User groups
- ◆ Users

Because you set up user rights when you create user groups, they do not require a separate phase.

For detailed information

The following section discuss how to implement your ContentDirector user rights scheme:

- ◆ [“Creating domains”](#) on page 59
- ◆ [“Creating user groups”](#) on page 61
- ◆ [“Creating domain administrators”](#) on page 64
- ◆ [“Creating regular users”](#) on page 66

Creating domains

Set a guideline on domain workspaces

One important part of setting up a domain is selecting a workspace. The workspace is an area on hard disk that ContentDirector uses for temporary files and backup copies of distributed content.

You have the option of either using the default workspace, or setting up a separate workspace for each domain. To use ContentDirector efficiently, set a rule to consistently do one or the other.

Each approach has advantages and disadvantages:

- ◆ Using the default workspace is faster and easier, but does not organize or secure the backup files as well.
- ◆ Setting a separate workspace for each domain is more complex, but it:
 - ❖ Helps maintain security by keeping backup copies within the domain
 - ❖ Organizes administration and maintenance by storing each domain’s backup copies separately

The size of the workspace should be at least four times as large as the space required by the content you plan to distribute. For strategies for reducing disk space requirements, see “[Reducing local disk space requirements](#)” on page 97.

Creating domains

To create a domain in ContentDirector, complete the following steps.

Step	Action
1	Log in as the default system administrator using the following username/password combination: super/super.
2	On the top menu bar of the ContentDirector main screen, click Users.
3	Click Domains.
4	Click New. Result: The New Domain screen opens. It contains four tabs: Main, Sources, Folders, and Destinations. The Folders tab becomes active after you complete the Sources tab.

Step	Action
5	On the Main tab, ◆ Enter the domain name. ◆ Use the default ContentDirector workspace, or click Browse to select another workspace for this domain.
6	Click the Sources tab and select the desired sources: ◆ To permit access to all sources in ContentDirector, deselect the check box, and skip to the next step. To restrict access to sources, select the check box. ◆ To grant access to a specific source, click its name in the Available panel, then click Add. The source name moves to the Visible to Domain panel. To remove a source from the Visible to Domain panel, click its name, then click Remove.
7	Click the Folders tab. To restrict access to specific folders within the selected sources: 1. Deselect the check box under View All. 2. Click Browse and browse the local file system for the desired folders. 3. Click OK.
8	Click the Destinations tab. ◆ To permit access to all servers in ContentDirector, leave the checkbox blank, and skip the next two steps. ◆ To restrict access to servers, click the checkbox. ◆ To grant access to a server, click its name in the Available panel, then click Add. The server name moves to the Visible to Domain panel. ◆ To remove a server from the Visible to Domain panel, click its name, then click Remove.
9	Click OK to save the domain. The new domain appears on the the Users view in the main screen, on the Domains tab.

Creating user groups

Prerequisite for creating user groups

You must create at least one domain before you can create user groups.
For instructions on creating domains, see“[Creating domains](#)” on page 59.

Creating user groups

To create a user group, complete the following steps.

Step	Action
1	Log into ContentDirector as the default system administrator using the following username/password combination: super/super.
2	On the top menu bar of the ContentDirector main screen, click Users.
3	Click User Groups.
4	Click New. Result: The New User Group screen that appears consists of four tabs: Main, Sources, Folders, and Destinations. Only the Main tab is active when the screen first displays.
5	Click the dropdown menu next to the Domain field and choose the domain to which this user group belongs.
6	In the Group field, enter the name for the user group.
7	Under Creation Allowed For: select the checkboxes for resources that members of this group can create.

Step	Action
8	Click the Sources tab. ◆ If the Restrict Sources checkbox is checked and inactive, sources are restricted by the domain. Click a source name on the Available panel, then click Add to grant access to a source for this user group. ◆ If the Restrict Sources box is not checked, sources are not restricted by the domain. Your choices are: ❖ Leave it as is to grant access to all sources for this group ❖ Select it to restrict source access for this group If you choose to restrict access, you need to select which sources are available. To select sources, click a source name on the Available panel, then click Add.
9	Click the Folders tab. All available folders are listed in the Visible Folders panel. ◆ If access to folders is unrestricted by the domain, the Restrict checkbox is deselected. Select the Restrict checkbox to restrict access to folders. Click Browse and select the folders. ◆ If access to folders is restricted by the domain, the Restrict checkbox is selected and inactive. Click the button under Browse and select a subfolder to restrict access to that subfolder only.
10	Click the Destinations tab. ◆ If the Restrict Destinations checkbox is checked and inactive, destinations are restricted by the domain. Click a destination name on the Available panel, then click Add to grant access to a destination for this user group. ◆ If the Restrict Destinations box is not checked, destinations are not restricted by the domain. Your choices are: ❖ Leave it as is to grant access to all destinations for this group ❖ Select it to restrict destination access for this group If you choose to restrict access, you need to select which destinations are available. To select destinations, click a destination name on the Available panel, then click Add.

Step	Action
11	Click OK to save the user group. The new user group appears on the the Users view in the main screen, on the User Groups tab.

Implementing your user rights scheme

Creating domain administrators

Prerequisite for creating domain administrators

You must create a domain before you can create a domain administrator. For instructions on creating domains, see [“Creating domains”](#) on page 59.

Each domain requires an administrator

Each domain must have a domain administrator. After creating a domain, create a new user to be the domain administrator of that domain.

To create a domain administrator, complete the following steps.

Step	Action
1	Log in as the default system administrator using the following username/password combination: super/super.
2	On the menu bar near the top of the ContentDirector main screen, click Users.
3	Click New. The New User screen consists of three tabs: Account, Personal, and Group. The Account and Group tabs are required. The Personal tab contains optional information.
4	On the Account tab, fill in the Username, Password, and Confirm password fields. Note _____ Passwords must have a minimum of five characters. _____
5	At the Type field, click the drop-down menu and select Domain Administrator.
6	Select the check box next to User Must Change Password At Next Logon if you want the new administrators to immediately change their password.
7	Select Account Disabled to disable the account.

Step	Action
8	Click the Personal tab and fill in optional personal information, if desired.
9	Click the Groups tab, click the drop-down menu next to the Domain field, and select the domain.
10	Click OK to save. Result: The new domain administrator appears in the Users view of the main screen, on the Users tab.

Creating regular users

Prerequisites for creating regular users

You must create a domain and a user group before you can create regular users.

For instructions on creating domains, see [“Creating domains”](#) on page 59. For instructions on creating user groups, see [“Creating user groups”](#) on page 61.

To create a regular user, complete the following steps.

Step	Action
1	Log into ContentDirector as a system administrator or domain administrator.
2	On the menu bar at the top of the main screen, click Users.
3	Click New. Result: The New User screen that appears consists of three tabs: Account, Personal, and Group. The Group tab is inactive until a user type is selected. The Account and Group tabs are required. The Personal tab contains optional information.
4	On the Account tab, fill in the Username, Password, and Confirm password fields. Note _____ Passwords must have a minimum of five characters. _____
5	At the Type field, click the drop-down menu and select User.
6	Select the check box next to User Must Change Password At Next Logon if you want the new users to immediately change their password.
7	Select Account Disabled to disable the account.
8	Click the Personal tab and fill in optional personal information, if desired.

Step	Action
9	Click the Groups tab.
10	Click the drop-down menu next to the Domain field, and select the domain.
11	Select a user group in the Available panel, then click Add to make this user a member of the user group. If you are adding this user to multiple groups, repeat this step for each group.
12	Click OK to save. Result: The new user appears in the Users view of the main screen, on the Users tab.

Granting instance rights

About instance rights

Instance rights enable regular users to access resources created by other regular users in the same domain.

Instance rights are necessary because regular users cannot typically access resources created by other regular users. This is true even if all the users are members of the same user group. A domain administrator must alter the instance rights of the resource so that other regular users can access it. System administrators cannot grant instance rights, only domain administrators can do so.

Which resources have instance rights


Policies, jobs, destination groups, and contacts all have instance rights.

How instance rights differ from standard user rights

Standard user rights, or *class rights*, enable user groups to access a whole category of resources, for example, policies. Instance rights instead grant access to a specific resource, for example, Policy 1. The instance rights of Policy 1 are a property of the policy, rather than of the user groups that can access it.

Which rights take precedence

Instance rights take precedence over user group rights. In the following diagram, User Group A has access rights to Policy 1, even though the group rights do not grant access to policies.



Granting instance rights

To grant instance rights for a resource, complete the following steps.

Step	Action
1	Log into ContentDirector as a domain administrator.
2	Access the Modify screen for the desired resource : ◆ For policies, click the Policies button on the left menu bar, then double-click the policy to modify. ◆ For jobs, click the Jobs button on the left menu bar, then double-click the job to modify. ◆ For destination groups, click the Destinations button on the left menu bar, click the destination groups tab, right-click the destination group, then click Modify. ◆ For contacts, click Options on the top menu bar, click contacts, then double-click the contact to modify.
3	Click the Access tab.

Step	Action
4	Set access for user groups: ◆ To enable access for a specific user group, to the right of the user group name select the Read check box. ◆ To enable access for all user groups, click the Set All Access button. ◆ To disable access for a specific user group, to the right of the user group name deselect the Read check box. ◆ To disable access for all user groups, click the Clear All Access button.
5	Click OK to save your changes to the resource and return to the main screen.

About this chapter

This chapter describes how to prepare to distribute content ContentDirector by creating and modifying content distribution resources.

Topics in this chapter

This chapter contains the following topics:

- ◆ [“Setting up content distribution resources”](#) on page 72
- ◆ [“Reducing local disk space requirements”](#) on page 97
- ◆ [“Modifying content distribution resources”](#) on page 100

Setting up content distribution resources

Information to gather before you start

Before you can prepare for content distribution, you need to gather the following basic information:

- ◆ Names and locations of source directories and files
- ◆ Identity and location of destinations
- ◆ Which destinations receive the same content
- ◆ How often the content is sent
- ◆ Content type
- ◆ Approximate maximum size of jobs

Phases of preparing to distribute content

Preparing to distribute content has the following five phases:

- ◆ Set global and user preferences
- ◆ Create sources
- ◆ Create destination groups
- ◆ Create policies
- ◆ Create jobs

For the most efficient task flow, perform these steps in the sequence shown above.

For detailed information

The following sections discuss how to prepare for content distribution in ContentDirector:

- ◆ [“Setting content distribution preferences”](#) on page 73
- ◆ [“Creating sources”](#) on page 77
- ◆ [“Creating destination groups”](#) on page 80
- ◆ [“Creating policies”](#) on page 82
- ◆ [“Creating jobs”](#) on page 91

Setting content distribution preferences

Purposes of setting preferences Preferences create standard settings for many fields in ContentDirector screens. They serve two purposes:

- ◆ Save time when creating new resources
- ◆ Help ensure that fields are filled correctly and consistently

Types of preferences Global content distribution preferences are set by a system administrator. They apply to all users in ContentDirector. Only a system administrator can change or override global preferences.

User content distribution preferences are set by any user. They apply only to the user who set them. That user can manually override them in any instance where they apply.

Setting global preferences ContentDirector has three global content distribution preferences: network segment refresh rate, job workspace, and content mime types. These settings are visible only to system administrators.

To set these preferences, complete the following steps.

Step	Action
1	Log on as a system administrator.
2	Select File > Preferences.

Step	Action
3	Click the Jobs tab. The workspace setting at the bottom of this tab is a global preference. This workspace is used by all ContentDirector jobs, unless a system administrator changes the individual job workspace. You can use the current workspace preference, or select a different workspace. To select a different workspace, under Workspace, click Browse and choose the location in the local file system. Note Moving the workspace to a different drive reduces local disk space usage. See “Reducing local disk space requirements” on page 97 for details. You can select a Network Appliance filer or a mapped our mounted drive as your workspace. The core application machine must have administrative or root access to the drive that contains the workspace. For detailed instructions, see Appendix E, “Mapped Drive Permissions,” on page 165.
4	Click the Network tab. In the field next to Network segment refresh rate, enter in minutes how often you want transport group leaders to report to the core application with a list of group members and their statuses. Setting too frequent an interval for the network segment refresh rate can slow down ContentDirector, because the reports use application resources. Set the refresh rate to an interval that can catch changes in the status of group leaders without excessively slowing other functions, for example, once a minute. Note The default rate of 0 means that no automatic refresh is performed. Fill in the refresh rate unless you want users to perform all network segment reports manually.

Step	Action
5	Click the Cache tab. The scrolling list contains mime types recognized by ContentDirector. Mime identifiers are sent with these files to ensure they are correctly recognized by the destination and by end user applications. For a printed list of mime types included with ContentDirector at shipment, see Appendix D, “Default Mime Types,” on page 151. ◆ To add mime types, click New. ◆ To delete or modify, click the file extension on the scrolling list, then click Delete or Edit. ◆ To undo your changes to this list and restore the standard list of mime types, click Restore Default.
6	Click OK to save your preferences.

Setting user preferences

To set user preferences, complete the following steps.

Step	Action
1	Log in as the user who will use these preferences.
2	Click File > Preferences.

Step	Action
3	Click the Job tab. Under file types to exclude and include, you can make the job screen reflect your strategy for types of files to distribute: ◆ To distribute any type of file, leave both fields blank. ◆ To exclude only a few types of files, and distribute all others, enter the types to exclude in the File Types To Exclude field. ◆ To include only a few types of files, and exclude all others, enter the types to include in the File Types To Include field. ◆ To enter a file type, type its extension into the desired field. ◆ If you enter multiple types, separate them by a semicolon, with no space. For example: <code>.exe;.html;.jpeg;.zip</code> Note_____ Be sure to include file types you plan to distribute in your mime types. For more information, see “Setting global preferences” on page 73.
4	Click the Policy tab. Under File Types Not Compressed, add file types you want to exclude from compression in jobs where compression is activated. These are typically file formats that are already compressed. The default list includes common graphics, media, and archive formats. If you enter multiple types, separate them by a semicolon, with no space. For example: <code>.gif;.jpg</code>
5	Under SMTP Host, enter the name of the default mail server to use if you plan to automatically email contacts about the final status of jobs.
6	Under Sent By, enter the default name to write in the From field when sending job status emails.
7	Click OK to save your preferences.

Creating sources

About sources

Sources are directories on local, mapped, or mounted drives that contain the content you want to distribute to NetCache destinations. Each source is designated in ContentDirector by a Universal Resource Locator (URL), for example: <http://www.oursite.com>. The URL enables NetCaches to recognize and store the content.

Recommendation on setting up sources

Set up sources as directories whenever possible, rather than as specific files. This allows you to eject files from the destination NetCache appliances by deleting them from the source directory in ContentDirector. If you select individual sources as files, ContentDirector cannot trigger an eject by NetCache.

Who can create sources

Only system administrators can create sources.

Source protocols

ContentDirector recognizes the following protocols for sources:

Protocol	Name
http://	HTML
mms://	Microsoft Media Streaming
rtsp://	Real-Time Streaming Protocol

Using a filer as a source location

For increased performance and storage efficiency, consider using a Network Appliance™ filer as a source location instead of a conventional server:

- ◆ Data input and output are much faster on the filer than on a conventional server.
- ◆ Filer snapshots can replace the shadow copies of content that ContentDirector would otherwise store in the workspace.

Before you can use a filer as a source location for sources, you need to mount the filer and create a separate volume for each source. See your filer documentation for more information.

Note

Using a filer as a source location limits the number of jobs you can create for each source. Filers have a limit of 31 snapshots per volume. Each job uses three snapshots while it is running, and two while it is idle. This means for each source you can have a maximum of ten jobs running simultaneously, or 15 jobs in an idle state. If you exceed this limit, the job will fail.

Note

When using a filer as a source location, the filer must open permissions so the core application server can rsh to it. To enable these permissions, create the following entry on the filer in */etc/hosts.equiv*: **<hostname or IP address of core application server> <username>** where the username is the one specified in the advanced dialog for filers. Note the space between the first and second parts of the entry.

Filer permissions for ContentDirector

To set up ContentDirector sources on a filer, your operating system must grant ContentDirector permission to access the filer. Under Solaris the core application machine must have root access. Under Windows there are several possibilities for setting up permissions, depending on your domain configuration and how you run ContentDirector. For detailed instructions, see Appendix E, “[Mapped Drive Permissions](#),” on page 165.

Creating a source

To create a source, complete the following steps.

Step	Action
1	Log in as a system administrator.
2	On the menu bar at the top of the ContentDirector main screen, click Sources > New.
3	Next to Scheme, click the drop-down menu and select the protocol.

Step	Action
4	In the Host name field, enter a unique name for this source. The host name must have at minimum a domain, for example: http://oursite. It may contain a full domain name with a directory path, for example: http://www.oursite.com/products/new. Note _____ Path designations for sources cannot be strings greater than 1028 characters in length.
5	To select a source location on a filer, click Filer. If you do not plan to use a filer, skip this step. ◆ Enter the host name, volume name, and remote username. ◆ To use filer snapshots instead of workspace copies, click Optimize for Filer. ◆ Click OK to save the filer settings and return to the New Source screen.
6	Click Browse and select the source folder for the site content, then OK to select the location. Note _____ If you selected a filer as a source location, all content that you plan to distribute for that source must be on the filer for jobs to succeed.
7	Select the directory that includes all the content for this source. You can split up the content into subdirectories as necessary when you create jobs.
8	Click OK to save the new source.
9	The new source appears in the Network view of the main screen, on the Sources tab.

Creating destination groups

What destination groups are

A destination group is a collection of one or more destinations that receive content from the same source in ContentDirector. All content distribution in ContentDirector must use destination groups.

Who can create destination groups

Any administrator can create destination groups, and regular users can also create them if their user group rights permit it.

When to create destination groups

Although it is more efficient to create your destination groups before creating jobs, it is not required. If you create a job for a new combination of destinations that is not yet designated as a destination group, ContentDirector prompts you to create a new destination group.

Destination groups as part of your virtual interface

As with user rights, destination groups allow you to map your resources to a customized virtual interface. Because destination groups are purely logical and do not need to map to the physical network, they are completely flexible. A destination group can correspond to any group of destinations. For example, one destination group can consist of an office in a remote location, while another includes a whole department that has employees scattered across several locations.

Information you need to gather before you start

Before you create destination groups, you need to gather the following information for each group:

- ◆ Name of source that holds the content for each group
- ◆ Name of each destination that will receive content from this source

Creating destination groups

To create a destination group, complete the following steps.

Step	Action
1	On the menu bar on the ContentDirector main screen, click Destinations > Destination Groups > New.
2	Fill in the Group Name field.
3	Click the drop-down menu next to Source and select the source to provide content to this destination group.
4	On the Available Destinations panel, select a server, then click Add to select a server as a destination in this group. To remove a server from this group, on the Current Destinations panel click the destination name, then click Remove.
5	Click OK to save the new destination group. Result: The new destination group appears on the Destinations view of the main screen, in the destination groups tab.

Creating policies

About policies

Policies are content distribution rules that enable you to control how jobs handle content. Every job references a policy.

Function of policies

Policies control the following aspects of content distribution:

- ◆ Creation of lists of contacts to Email about job status
- ◆ Data security, integrity, and compression
- ◆ Criteria for distribution success
- ◆ Instructions to NetCache specifying when to make content available for serving

Synchronizing time for time-sensitive features

To use time-sensitive features such as Min Age in policies, you need to synchronize the time zone and time. This means the time zone and time must be exactly the same on the specified devices, even if they are separated geographically:

- ◆ If you are running CD at the core application server, you need to synchronize the core application and the NetCache destination.
- ◆ If you are running CD remotely through the applet, you need to synchronize the machine running the applet, the core application server, and the NetCache destination, and restart the machine running the applet.

Who should create policies

Any administrator can create policies, and regular users can also create them if their user group rights permit it.

Creating policies requires technical understanding of how networks handle content, as well as knowledge of your organization's goals and guidelines for distributing content. For this reason, only qualified personnel should create policies.

The default policy

ContentDirector ships with a default policy configured and ready to use. You can modify the default policy as desired. To modify the default policy, complete the following steps:

Step	Action
1	On the left side of the main screen, click the Policies button.
2	Double-click the listing for Default Policy.
3	Make any desired changes.
4	Click OK.

Before you start

The Policy tab of the Preferences screen contains settings that populate fields in each new policy. Before creating policies, be sure to set these preferences appropriately for your needs. For more information, see “[Setting content distribution preferences](#)” on page 73.

Tabs in the Policy screen

Each policy consists of Options, Contacts, Network, Quorum, Cache, and Compression tabs. The following instructions describe how to create a new policy by filling in each tab.

Opening the New Policy screen and filling in the Options tab

To fill in the Options tab, complete the following steps to fill in the tabs.

Step	Action
1	On the menu bar at the top of the ContentDirector main screen, click Policies > New.
2	In the Name field, enter a descriptive name for this policy.
3	The Import policy field enables you to select an existing policy to use a template for the new policy. If you do not want to use a template, skip this step. To import a policy, click the drop-down menu and select a policy.

Filling in the Contacts tab

The Contacts tab enables you to set up personnel who can be notified automatically of the final status of jobs.

Note

You can set up a contact to email Network Appliance autosupport when jobs fail. Network Appliance will examine the information sent by ContentDirector and contact you about any significant failure. See the procedure below for details.

To fill in the Contacts tab, complete the following steps.

Step	Action
1	Click the Contacts tab. If fields in this tab are already populated by your preferences, skip those steps.
2	In the SMTP Host field, enter the name of the default mail server to use when emailing contacts about the final status of jobs.
3	In the Sent by field, enter the default name to write in the From field when sending job status Emails.
4	To add new contacts, click New.
5	In the New Contact screen enter the name, Email address, and optional comments. To use Network Appliance autosupport for your ContentDirector installation, enter the following email address: autosupport@netapp.com
6	Under Options, select the check box specifying when you want ContentDirector to Email this contact automatically. For the autosupport contact, select Contact If Job Fails.
7	Click OK to save the new contact.
8	To add more contacts, repeat Steps 4 through 7. Note Contacts added by system administrators are visible to all users. Contacts added by any other users are visible only to the user who added them.

Step	Action
9	You can include or exclude contacts from the list of contacts to be Emailed by this policy: ◆ To include a contact in this policy, click the contact name under Available Contacts, then click Add. ◆ To exclude a contact from this policy, click the contact name under Current Contacts, then click Remove.

Filling in the Network tab

The Network tab contains settings that manage the integrity of the content as it is sent over the network.

To fill in the Network tab, complete the following steps.

Step	Action
1	Next to the Security field, leave the setting on None, or click the dropdown arrow and select a security setting: Encryption, Verification, or Both: ◆ None: Validates content, but does not encrypt or verify it. ◆ Encryption: Encrypts content. ◆ Verification: Verifies that content did not change during distribution. ◆ Both: Uses both encryption and verification. Note_____ Enabling any additional security can lengthen the time it takes to complete a job.
2	If you are distributing content via satellite, select the checkbox next to Enable Forward Error correction. Under Data Packet Loss Rate use the default setting. If your network does not include a satellite link, skip this step. Note_____ If your satellite link uses a terrestrial backchannel for control, setting the Data Packet Loss Rate is optional.

Filling in the Quorum tab

Quorums set criteria for job success based on the number of destinations that receive content.

To fill in the Quorum tab, complete the following steps.

Step	Action
1	Click the Quorum tab. At the Commit Quorum field, click the drop-down menu and set the percentage of targeted destinations you need to receive a job in order for ContentDirector to flag the job as successful. The default percentage is 100. For example, if you have 100 destinations targeted for a job, and need eighty of them to receive the content for the job to be a success, choose 80 percent. You can use the repush feature later to send the job to the remaining twenty destinations. A commit quorum of 0% requires at least 1 destination to process the assignment successfully for the job to succeed. For example, for a job with 0% commit quorum that pushes to 100 destinations, if 99 destinations fail to process the assignment, the job still succeeds. The job fails only if all 100 destinations fail to commit. Note_____ If you have multiple NetCache appliance destinations that receive the same content, a commit quorum of less than 100 percent can help increase the efficiency of content distribution.
2	At the Availability Quorum field, click the dropdown arrow and set the percentage of targeted destinations you need to be ready before ContentDirector starts sending a job. The default percentage is 100. The availability quorum saves time and bandwidth by preventing ContentDirector from starting a job if the minimum percentage of destinations is not available.

Step	Action
3	Under Retry Information: ◆ Select the check box next to Enable Retry Of Failed Jobs to enable the retry feature. Retry attempts to re-send a job to the whole destination group if the first attempt fails to meet the commit quorum. ◆ In the Number field, enter the number of times you want ContentDirector to retry. The default number of retries is zero. ◆ In the Interval Time field, enter the amount of time for ContentDirector to wait between job retries. Use the format hours: minutes: seconds. For example, for ten minutes, enter 00:10:00. The default time is five minutes. The interval time is used only if you entered a value in the Number field.
4	Under Repush Information: ◆ Select the check box next to Enable Repush To Failed Servers to enable the repush feature. A quorum of less than 100% permits a job to succeed even though not all destinations received it. Repush attempts to re-send the job to destinations that did not receive it the first time. ◆ In the Number field, enter the number of times you want ContentDirector to attempt to repush. The default number of retries is zero. ◆ In the Interval time field, enter the amount of time for ContentDirector to wait between repush attempts. Use the format hours: minutes: seconds. For example, for ten minutes, enter 00:10:00. The default time is five minutes. The interval time is used only if you entered a value in the Number field.

Filling in the Cache tab

The Cache tab settings determine how NetCache destinations handle content. These settings do not apply to Solaris or Windows client destinations.

To fill in the Cache tab, complete the following steps.

Step	Action
1	Click the Cache tab.
2	Select the Expires check box to enable an expiration age (time to live), and enter the age in days, hours, minutes, and seconds. When content reaches the Expires age, the NetCache must retrieve fresh content from the distribution server.
3	Select the Max Age check box to set a maximum age for content, and enter the maximum age in days, hours, minutes, and seconds. The default setting is 30 days. Max Age applies only to http:// content, not to mms:// or rtsp://. Note _____The Max Age setting has the same function as Expires. It is valid in HTML 1.1 only, where it takes precedence over Expires. If the destination or client web browser does not support HTML 1.1, it uses the Expires setting. The safest approach is to set both Expires and Max Age to the same value.
4	Select the Min Age check box to enable an earliest date and time content may be served, and then select the date and time as month/day/year and hours:seconds AM or PM.
5	Select the Must Revalidate Header check box to require destinations to revalidate headers when refreshing stale content.
6	At Lock Files In Cache, enter the duration of locking in days, plus any additional time in hours:minutes:seconds. The maximum value is: 999 days, 23 hours, 59 minutes, and 59 seconds. The default setting is 30 days.

Step	Action
7	Under Completion Criterion, select a value that determines what portion of the content sent to the cache must be published to the cache for the job to be successful: ◆ All: all content must be published – any publishing failure results in a failed status ◆ Best Effort: no minimum is set, and status is always successful ◆ % Bytes: specifies the percentage of bytes in the job that must be published ◆ % Files: specifies the percentage of files in the job that must be published The completion quorum interacts with the commit quorum. Each NetCache destination sends a status report of success or fail to the core application. The core application sums all the status reports and compares them to the commit quorum to determine whether the job succeeded.

Filling in the Compression tab

The Compression tab contains settings that enable compression of content as it is sent over the network. Compressing content lengthens the time to prepare the job, but shortens the time required to distribute it.

The default setting for compression is none. If you are not using compression, you can skip this step.

To fill in the Compression tab, complete the following steps.

Step	Action
1	Click the Compression tab.
2	Next to Compression type, click the drop down menu and select None, Low, Medium, or High compression.

Step	Action
3	File Types Not Compressed lists types of files to exclude from compression in jobs where compression is activated. If this section is not already populated by your preferences, add files to exclude from compression by clicking the file extension on the Available panel, then clicking Add. To remove files from the File types not compressed list, click the file extension on the Current panel, then click Remove. This means that the file type is compressed when ContentDirector distributes it. To add new file types to the Available panel, click New, enter the name and a description for the file type, and then click OK.
4	Click OK to save the policy. Result: The new policy appears in the Policies view of the main screen.

Creating jobs

About jobs

All content distribution through ContentDirector is defined and regulated by jobs. Jobs are sets of parameters that determine the content, destinations, and schedule for a distribution task. You can schedule a job to run periodically and perform the same distribution task at regular intervals. Each distribution task is governed by a single job.

Any administrator can create jobs, and regular users can also create them if their user group rights permit it.

Importance of good job design

Because jobs define what, where, and when you distribute, they are key to realizing your objectives with ContentDirector. Consider two primary factors when designing a job:

- ◆ Sizing
- ◆ Scheduling

Sizing considerations

Planning how to distribute content requires that you decide on criteria for sizing jobs. Be sure to take into account the following criteria:

- ◆ The wire speed at which you can distribute content
- ◆ How content storage is split over multiple destinations

For optimal efficiency, split your content into smaller jobs rather than combining it into larger jobs. This makes the best use of your system resources, and minimizes the impact of any distribution problems.

Scheduling considerations

Consider the following factors when scheduling jobs:

- ◆ Bandwidth conservation

A common objective when scheduling jobs is to conserve bandwidth during peak hours of network traffic. ContentDirector's scheduling function enables you to time content distribution for off-peak periods such as night, early morning, or weekends.

- ◆ **Frequency of updates**
Another important factor is how frequently the content needs to be updated. The interval of the schedule should reflect this frequency, and can range from once a month to every few minutes.
- ◆ **Ensuring that related content is available simultaneously**
If you split related content across multiple jobs, be sure to schedule the jobs so that the content will be available simultaneously.
Simultaneous jobs to the same destination are queued, so it is most efficient to schedule them sequentially.

Before you start

The Job tab of the Preferences screen contains settings that populate fields in each new job. Before creating jobs, be sure to set these preferences appropriately for your needs. For more information, see [“Setting content distribution preferences”](#) on page 73 .

Opening the New Job screen and filling in the Details tab

Each job consists of Details, Content, Filters, and Schedule tabs.

To create a job, complete the following steps to fill in the tabs.

Step	Action
1	At the top of the ContentDirector main screen, click Jobs.
2	Click New.
3	In the Job Name field, enter a descriptive name for this job.
4	In the Import job field, click the drop down menu and select an existing job to use as a template. If you do not want to use a template, skip this step.
5	In the Description field, enter an optional description of this job.
6	In the Policy field, click the drop down menu and select the policy for this job, or create a new policy.

Step	Action
7	If you have logged into ContentDirector as a system administrator, you can set the workspace for this job. Otherwise, skip this step and the next one. You can move the workspace to save space on the local drive. For more information, see “Reducing local disk space requirements” on page 97. Note_____ You must have write and delete rights to the drive where the workspace is located.
8	Next to Workspace, click the Browse button and select a new location for the workspace.
9	Click OK.

Filling in the Content tab

The Content tab designates the source and the destinations for this job.

To select the content source, complete the following steps.

Step	Action
1	Click the Content tab.
2	In the Source field, click the drop-down menu and select the source that provides the content for this job.
3	Click Add to select the source directory. In the Browse screen, select the source location for the content on the local file system.
4	Click the Destination tab.

Step	Action
5	You can select whole destination groups or individual devices as destinations: ◆ To select a whole destination group, at the Select Destinations in Group field, click the drop-down menu and select the group ◆ To select individual destinations instead of a whole group, click the destination name on the Available Servers panel, then click Add ◆ To remove a destination, click the its name on the Current Servers panel, then click Remove.

Filling in the Filters tab

The Filters tab enables you to choose a subset of the source content you designated on the Content tab. You can select only changed files, choose whether to include subdirectories or not, and select file types to include or exclude. All settings on the Filters tab are optional.

To fill in the Filters tab, complete the following steps.

Step	Action
1	Click the Filters tab.
2	Select the check box for Include Only Files That Have Changed to send only files that have been modified or added since the last time the job was run.
3	Select the check box for Include All Subdirectories to send all subdirectories within the location you specified in the Content tab. Deselect the check box to send only the top level directory.

Step	Action
4	Under From Source Location, set specific types of files to include or exclude, if your preferences have not already set them: ◆ To distribute any type of file, leave the File Names panel blank. ◆ To exclude only a few types of files, and distribute all others: click Exclude, click File Types, select the types of file to exclude, then click OK. ◆ To include only a few types of files and exclude all others, select Include, click File Types, select the types of file to include, then click OK. ◆ To clear your list of selected file types, click Clear.

Filling in the Schedule tab

The Schedule tab enables you to choose the frequency, interval, and corresponding specifics for when to run the job.

To fill in the Schedule tab, complete the following steps.

Step	Action
1	Click the Schedule tab.
2	At the Execute This Job Field, click the drop-down menu and select Manual or the interval for running this job. To set a time, type in the time in 12-hour (civilian) format, followed by AM or PM.

Step	Action
3	To set a date, click the drop-down menu next to the date field. Use the left and right arrows to set the month. Double-click a day to select the day and close the calendar. ◆ Manual: a user must initiate running the job from the core application. ◆ One time only: runs the job once. Select the start time and date. ◆ Every N minutes: runs the job at intervals of up to one hour. Set the start time, start date, and end date, choose to run every day or weekdays only, and set the interval in minutes. ◆ Hourly: runs the job at intervals of up to 24 hours. Set the start time, start date, and end date, choose to run every day or weekdays only, and set the interval in hours. ◆ Daily: runs the job once a day. Set the start time, start date, and end date, and choose to run every day or weekdays only. ◆ Weekly: runs the job once a week. Set the start time, start date, and end date, and choose the days of the week on which to run the job. ◆ Monthly: runs the job once a month. Set the start time, start date, and end date. Under Interval, set the interval for the job: ❖ To run the job every n days, click Day and select the number of days. ❖ To run the job once a week on a selected day, click Every, then click the drop-down menu next to the day and select the day. Under Of the months, select the months in which you want the job to run.
4	Click OK to save the job. Result: The new job appears in the Jobs view of the main screen.

Reducing local disk space requirements

About disk space requirements for content distribution

ContentDirector has local disk space requirements that are related to the content distribution process. These requirements are dynamic depending on your content distribution practices. They are separate from installation disk space requirements.

What you need to know to estimate requirements

To estimate your local disk space requirements, you need to know:

- ◆ The size of the largest job you need to run
- ◆ Whether you plan to run jobs concurrently, because the disk space required is additive across jobs – storage for jobs does not overlap
- ◆ Whether you are running CD optimized for filer – this reduces the amount of local disk space required by a factor of (1 x the size of the content being pushed)

Components of local disk space usage

There are three components of local disk space usage for content distribution:

- ◆ Shadow: a copy of the content pushed the last time a job was successfully run
- ◆ Assignment: content and header data for each job currently running.
The size of the assignment file is not the same as the size of the raw content. It depends on the following factors:
 - ❖ The size and number of the source files. In general, many small files require more overhead than a few large files.
 - ❖ For jobs that send only file deltas (changes), the number and scope of the changes made to the source file since the job was last run
 - ❖ For jobs with compression enabled, the level of compression
- ◆ Temporary directory: temporary copy of content for each job currently running

You can relocate or eliminate all three of these components.

Approximate ratios for estimating space requirements

Local disk space requirements fluctuate with the cycle of job processing. The peak is for a job that has already been run at least once, and is just starting to run again. At that time, the maximum local disk space used is:

- ◆ approximately four times the size of the content being pushed, if you implement *no* space-saving measures
- ◆ approximately the same as the size of the content being pushed, if you implement *all* space-saving measures

These ratios apply to all jobs, regardless of size.

Keep in mind that the actual space required depends on what is being added, changed, or deleted by a particular distribution event.

Strategies for reducing local disk space needs

There are three strategies for reducing local disk space needs:

- ◆ Move the shadow and temporary directories to a different drive by moving the workspace
- ◆ Eliminate the shadow directory entirely by using an optimized filer for the content source
- ◆ Move the assignment directory to a different drive

Move the workspace

The workspace contains the shadow directory and temporary directories. If you log onto ContentDirector as a system administrator, you can move the workspace either as a global preference or for each job. For details see “[Setting content distribution preferences](#)” on page 73 and “[Creating jobs](#)” on page 91.

Use an optimized filer for content source

You can eliminate the shadow directory by locating the content source on an optimized filer. ContentDirector integrates with DataONTAP and uses snapshots for content backup, instead of the shadow. For details see “[Creating sources](#)” on page 77 and Appendix E, “[Mapped Drive Permissions](#),” on page 165.

Relocating the assignment file

ContentDirector builds a temporary assignment file before running each job. The assignment file contains all the content and header information for the job. The default location for this file is under the install directory, in `\ContentDirector\Ccm\jakarta-tomcat\webapps\ContentDirector\Assignments`.

The assignment directory is different from the workspace, and its location is not affected by your choice of workspace.

To relocate the assignment directory, complete the following steps:

Caution_____

Caution: Do not make any other changes to these files other than the ones listed below. Be sure to use the correct case, as properties files are case-sensitive.

Step	Action
1	Stop the ContentDirector services.
2	Make backup copies of the files: ◆ ContentDirector\Ccm\config\cm.properties ◆ ContentDirector\jakarta-tomcat\conf\server.xml
3	Edit the file ContentDirector\Ccm\config\cm.properties.
4	Set the "cm_webserver_dir" parameter to the "root" of the new location, for example, for directory "z:/": cm_webserver_dir=z:/ Note _____ For each job, ContentDirector creates a subdirectory under this root location named after the internal jobID. For example, for job 5005, the workspace is z:/5005
5	Save and close cm.properties.
6	Edit the file ContentDirector\jakarta-tomcat\conf\server.xml.
7	Add a new alias for ContentDirector, for example: <Context path="/ContentDirector" docBase="z:/ContentDirector" debug="0" reloadable="true"> </Context>
8	Save and close server.xml.
9	Restart ContentDirector services.

Modifying content distribution resources

Reasons to modify resources

Changes in your site content, network topology, content distribution strategies, or other factors might necessitate changes to your content distribution settings.

For detailed information

The following sections discuss how to modify content distribution resources in ContentDirector:

- ◆ [“Adding and deleting destinations”](#) on page 101
- ◆ [“Modifying and deleting sources”](#) on page 104
- ◆ [“Modifying and deleting destination groups”](#) on page 107
- ◆ [“Modifying and deleting policies”](#) on page 109
- ◆ [“Modifying and deleting jobs”](#) on page 111

Adding and deleting destinations

Adding destination servers

To add destination servers, add the servers to your physical network and confirm their connectivity, then install or activate ContentDirector. Follow the instructions in the *ContentDirector Installation Guide* to activate and configure the embedded ContentDirector agent.

ContentDirector recognizes each new destination server as soon as it starts the ContentDirector Agent. If you create a group leader and group members, ContentDirector recognizes a new network segment.

Adding new destinations to destination groups

To add a new destination to an existing destination group, modify the destination group from the ContentDirector main screen and add the new destination as a member. For more information see “[Modifying and deleting destination groups](#)” on page 107.

When you add a new destination to an existing destination group, ContentDirector gives you the option of automatically replicating on the newly added destination all the jobs that are currently associated with the destination group.

To add a new destination to a destination group, complete the following steps.

Step	Action
1	At the Destinations view of the main screen, click the Destination Groups tab.
2	Right-click the destination group to which you want to add the new destination. .
3	Click Modify.
4	Under Available Destinations, click the name of the destination you want to add.
5	Click Add.
6	Click OK.

Step	Action
7	At the Replicate Options screen, you can select the upper checkbox to enable Replicate Job Content To New Destinations to replicate all jobs currently associated with this destination group to the new destination. If you choose not to replicate the content on the new destination, leave the checkbox disabled. Result: The lower checkbox becomes active.
8	Select the lower checkbox to enable Remove New Destinations From Group if Replicate Fails. This option forces ContentDirector to check the status of the replication attempt after it runs, and removes the newly added destination from the destination group if the replication failed.

Deleting network segments and destinations

ContentDirector enables system administrators to delete network segments or destinations.

To delete network segments or destinations, complete the following steps.

Caution

Deleting a network segment deletes *all* destinations on that segment from ContentDirector. Deleting any network segment or destination means that you can no longer distribute content to that location. Accidentally deleting the wrong transport group leader in ContentDirector means that you must restart the group leader machine to recover it. Be sure that you identify the correct network segment or destination before deleting it.

Step	Action
1	Log in as a system administrator.
2	On the ContentDirector left menu bar, click Networks.
3	Click the Network Segment tab.
4	Right-click the network segment or destination you want to delete.
5	Click Delete.

Recovering a deleted transport group member

If you accidentally delete the wrong transport group member, recover it by completing the following steps:

Step	Action
1	On the toolbar near the top of the ContentDirector main screen, click Refresh. Result: The refresh function re-enters the deleted transport group member into ContentDirector. Note_____ The above instructions work only with transport group members. If you have accidentally deleted the wrong transport group leader in ContentDirector, you must restart the group leader machine to recover it. _____

Modifying content distribution resources

Modifying and deleting sources

Who can modify sources Only a system administrator can modify sources.

Reasons to modify sources The only reason to modify a source is that the source files have moved to a new location. This is because you can change only the location of the source files in a source. You cannot change the name or scheme of the source without deleting the source and creating a new one.

Caution—————
ContentDirector displays a warning if you modify the volume for any source that is optimized for filer. ContentDirector deletes the snapshots for all jobs that are associated with this source and that have one or more successful runs, and handles them as new jobs.

Modifying sources To modify a source, complete the following steps.

Step	Action
1	Log in as a system administrator.
2	On the left tool bar on the top of ContentDirector main screen, click Sources.
3	Right-click the source you want to modify.
4	Click Modify.

Step	Action
5	To change the filer source location, click Filer. If you are not using a filer for the source location, skip this step. ◆ Enter the new host name, volume name, and remote username. ◆ Click Optimize for Filer to replace workspace copies with filer snapshots. Note_____ Before you can use a filer as a source location for sources, you must mount the filer and create a separate volume for each source. See your filer documentation for more information. _____
6	To change the source folder on either a filer or conventional file server, click Browse. Select the source folder, then click OK to select the location.
7	Click OK to save your changes to the source.

Who can delete sources

Only a system administrator can delete sources.

Deleting sources

Deleting a source removes the logic from ContentDirector that enables you to distribute the source content. It does not remove the content files from your disk drive.

Note_____

If you delete a source that is referenced by jobs and destination groups, you must modify the jobs and destination groups to reference a new source, or delete them from ContentDirector. For more information see “[Modifying and deleting destination groups](#)” on page 107 and “[Modifying and deleting jobs](#)” on page 111.

To delete a source, complete the following steps.

Step	Action
1	Log in as a system administrator.
2	On the left tool bar on the top of ContentDirector main screen, click Sources.
3	Right-click the source you want to delete.
4	Click Delete.

Modifying and deleting destination groups

Who can modify destination groups

Any system administrator, and the domain administrator for the domain where the destination group resides, can modify the destination group. Users can also modify destination groups if their group user rights permit them to do so.

Reasons to modify destination groups

You might want to change the properties of a destination group if:

- ◆ You have added a new destination server to your network, and want to include it in an existing destination group
- ◆ You have deleted a source and substituted a new source, and want to redirect an existing destination group to the new source

All properties of a destination group can be changed, except its name.

Modifying destination groups

To modify a destination group, complete the following steps.

Step	Action
1	At the ContentDirector main screen, on the left tool bar click Destinations.
2	Click the destination groups tab.
3	Right-click the destination group you want to modify.
4	Click Modify.
5	Make the desired changes.
6	Click OK to save your changes and return to the main screen.

Who can delete destination groups

Any system administrator, and the domain administrator for the domain where the destination group resides, can delete the destination group. Users can also delete destination groups if their group user rights permit them to do so.

Deleting destination groups

To delete a destination group, complete the following steps.

Step	Action
1	On the ContentDirector left tool bar click Destinations.
2	Click the destination groups tab.
3	Right-click the destination group you want to delete.
4	Click Delete.
5	Click Yes on the confirmation screen.

Modifying and deleting policies

Who can modify policies

Any system administrator, and the domain administrator for the domain where the policy resides can modify the policy. The user who created the policy can also modify it.

Modifying policies requires technical understanding of how networks handle content, as well as knowledge of your organization's goals and guidelines for distributing content. For this reason, only qualified personnel should modify policies.

Reasons to modify policies

You might want to change the properties of a policy if:

- ◆ You need to change the list of contacts who are automatically Emailed about job status
- ◆ Your organization changes strategies about how jobs should handle content

You can change all properties of a policy, except its name.

To modify a policy, complete the following steps.

Step	Action
1	At the ContentDirector main screen, on the left tool bar click Policies.
2	Select the policy you want to modify.
3	Double-click the policy to open it.
4	Make the desired changes.
5	Click OK to save your changes and return to the main screen.

Who can delete policies

Any system administrator, and the domain administrator for the domain where the policy resides, can delete the policy. The user who created the policy can also delete it.

Which policies you can delete

You can delete any policy that is not attached to jobs, except the default policy. If the policy is attached to jobs, ContentDirector displays an error if you attempt to delete it. Before deleting the policy, you must modify the jobs to reference another policy.

The default policy cannot be deleted.

Deleting policies

To delete a policy, complete the following steps.

Step	Action
1	At the ContentDirector main screen, on the left tool bar click, click Policies.
2	Right-click the policy you want to delete.
3	Click Delete.
4	Click Yes at the confirmation screen.

Modifying and deleting jobs

Who can modify jobs

Any system administrator, and the domain administrator for the domain where the job resides, can modify the job. The user who created the job can also modify it.

Reasons to modify jobs

You might want to change the properties of a job if:

- ◆ You need to add or change content
- ◆ The list of destinations has changed

If you add a destination to a job that has already run, ContentDirector gives you the option to immediately replicate the job to the newly added destination. If you choose to not replicate immediately, the job is sent to the new destination at the next scheduled run time.

- ◆ The schedule has changed

You can change all the properties of a job, except its name.

Modifying jobs

To modify a job, complete the following steps.

Step	Action
1	At the ContentDirector main screen, on the left tool bar click Jobs.
2	Double-click the job you wish to modify to open it.
3	Make the desired changes.
4	Click OK to save your changes and return to the main screen.

Who can delete jobs

Any system administrator, and the domain administrator for the domain where the job resides, can delete the job. The user who created the job can also delete it.

Deleting jobs

To delete a job, complete the following steps.

Step	Action
1	On the menu bar at the left side of the ContentDirector main screen, click Jobs.
2	In the Status column of the Jobs view, confirm that the job you want to delete is in one of the following states: ◆ Idle ◆ Completed ◆ Failed ◆ Aborted Note_____ If the job is running, you must abort it first before deleting it. _____
3	Select the job.
4	Right-click the job.
5	Click Delete.
6	To delete all content that this job has sent to servers, select the check box for Delete Previously Distributed Job Content from Servers. To leave the content on the servers, deselect the check box.
7	Click Yes at the confirmation screen.

About this chapter

This chapter discusses how the ContentDirector distribution process works, how to distribute content, and how to monitor and verify the distribution process using ContentDirector.

The material in this chapter starts with the distribution event. You define the components that control the distribution event ahead of time. These components make up the job. After the job is defined and the content is at the source locations, running a distribution event is simple and straightforward. Because jobs are predefined, easily modified, and reusable, ContentDirector makes the distribution of content flexible and controllable.

Topics in this chapter

This chapter discusses the following topics.

- ◆ [“Learning about content distribution”](#) on page 114
- ◆ [“Technical description of the content distribution process”](#) on page 115
- ◆ [“An example of running a job”](#) on page 119
- ◆ [“Determining job status”](#) on page 122
- ◆ [“Deleting content from a destination”](#) on page 127
- ◆ [“Changing master keys to ensure data security”](#) on page 128

Learning about content distribution

About this topic

This topic provides background information about the content distribution process. The information helps you understand what happens during the content distribution event and interpret the feedback from the user interface.

Summary of the distribution process

When ContentDirector distributes content to the destination devices, core application reads the original files in the source locations and processes the data according to the specifications defined in the job. The core application creates an assignment consisting of the following:

- ◆ The processed data
- ◆ Information about the data processing and about the distribution event itself

The core application pushes the assignment to the transport group leaders on the destination devices, and the group leaders send the assignment to those transport group Members that are included in the destination group for the job.

The ContentDirector agents (on the transport group devices) use the information sent with the processed data to verify and unpack the data from the assignment to the correct locations. The processed data is converted back to its original form and files on the destination devices.

Technical description of the content distribution process

How content distribution is initiated

The actions that make up the distribution event are initiated by the core application. A job is not the actual distribution event. It is a set of predefined parameters that control the content distribution event. A job is created through the user interface, and the information is stored in the database. After a job is created, the database has all the information needed to run the job, including the scheduling information, but the distribution event cannot start until the core application initiates the process.

To do this, a scheduler service polls the database. When the scheduler finds a ready-to-run indicator for a job, it notifies the core application, which initiates the content distribution process as defined in the job.

Note

When a job is run manually from the user interface (as opposed to running at a scheduled time), the user interface communicates the ready-to-run status to the database and notifies the scheduler.

The directory structure on the source

When a job is first run, the core application creates a subdirectory under the workspace named with the job ID. Under the job ID, two permanent directories and a temporary directory are created:

- ◆ The temporary directory contains a mapping of the source. It is the location from which the data for the assignment is retrieved for transport to the destination devices. The temporary directory is created each time an assignment is created. It is deleted after the distribution event is finished.
- ◆ The Shadow directory contains a mapping of the source. The backup copies of the data are stored there.
- ◆ The JournalFiles directory contains a file for each assignment run for a job. The file contains the command buffer for that assignment.

Stages of the distribution event

Distribution stages in the user interface

After the distribution event starts, the user interface displays the successive stages of the event in the Status field of the Jobs view table. The distribution event has the following stages:

- ◆ Directory creation
- ◆ Channel creation
- ◆ Assignment creation
- ◆ Delivery of the command buffer
- ◆ Delivery of data and data processing files
- ◆ Processing the assignment

The headings of the topics that follow include the stage designations that appear in the Status field of the Jobs view.

Note

These indicators might go by too quickly to be seen if the job is small.

Directory creation (Running)

When the distribution event begins, the first entry in the Status field is Running. During this phase of the process a directory labeled with the job ID and the temporary directory are created in the Workspace. Access to the Workspace is also verified.

Channel creation (ChanCreate)

ChanCreate appears next in the Status field, indicating that the distribution channels are being created. The core application gets the job ID, the destination group ID, and the destination group name and sends them to the group leaders. The group leaders set up the ports and addresses for the transport channels to the group members that are included in the destination group.

Assignment creation (AsgnCreate)

After the channels are established, the core application creates the assignment. An assignment consists three types of data.

- ◆ Data files to be pushed and corresponding directory structure
When the assignment is created, the core application retrieves the original data files from the source with the directory structure specifying their location. The core application processes this data according to the specifications of the job definition with regard to such things as compression and security. For this reason, the data in the assignment is not in the same form as that in the original data files.
- ◆ Assignment processing information
When the assignment is created, the core application attaches the information about the data processing specified by the job. ContentDirector agents use this information on the destination devices to interpret the assignment and unpack the data correctly.
- ◆ Command buffer
The command buffer contains the commands that initiate the processing of the assignment on the destination devices and information about the source location of the data. The command buffer essentially “tells” the ContentDirector agents what is about to be sent in the data delivery stage and the type of processing that occurs. The command buffer is sent before the rest of the assignment and processing information
After the assignment is created, the core application tests for the availability quorum. If enough devices on the back end are online and ready to receive data, the next phase of the process begins.

Delivery of the command buffer (CmdDeliver)

After the assignment is created on the front end and the availability quorum is met, the core application sends the command buffer to the destination devices. It goes first to the transport group Leaders. The group leader, in turn, sends it to those transport group members that are part of the destination group.

Before the rest of the assignment is sent, the group leaders report their success rates to the core application. Based on these reports, the core application tests for the commit quorum. If the quorum is met, the core application embarks on the next step in the process.

Delivery of data and data processing files (DataDeliver)

After the core application sends the command buffer, the assignment data files with their accompanying processing information are pushed to the group leaders, which send it to the group members.

Processing the assignment (Processing)

When a ContentDirector agent (on the group member) receives an assignment, it runs the commands in the command buffer. This initiates the assignment processing, which includes such things as verifying the integrity of the data, unencryption, decompression, parsing, and so on.

When the processing is completed, the group leader is notified by the group member, and the group leader sends a report to the core application. The core application tests for the commit quorum. If the commit quorum is met, the commit process occurs. The result is that the files in the temporary directory are now in the Shadow directory, and all temporary directories are removed. Finally, the distribution event is designated as Completed in the user interface.

An example of running a job

About running a job This example of running a job includes the steps for running a small job and the steps for aborting the job. The job is defined to push a single media file (MMS).

Prerequisites for running a job To run a job, you need to create the following components.

- ◆ The source defined with the correct schema, in this case, MMS
- ◆ The policy (if other than the default policy) configured with the appropriate NetCache parameters
- ◆ The destination group that specifies the group members to receive the content
- ◆ The job that specifies the components listed above, the source, and other job parameters

Things to check before you start It is advisable to check the following things before you run the job:

- ◆ Adequate storage capacity on the core application machine and the servers of the destination group

In addition, you can check the “Cache Total Space” and the “Cache Locked Space” in the Network Segment screen of the Network view by selecting the cache in question.

- ◆ Transport groups are functioning properly

This can be determined through Network view of the user interface.

Distributing a file

About distributing a file

You can set the distribution event for a job to be run on a schedule, manually, or once only in the job definition. Any job to which you have access can be run manually through the user interface. You can abort a distribution event after it is initiated.

Steps for distributing a file

To push a file to a Netcache Appliance, complete the following steps.

Step	Action
1	In the Jobs view table, select the job you want to run.
2	Click Run on the lower menu bar.
3	Click Yes at the prompt. After a few moments, the values in the Status field of the jobs table change as the event progresses through the various stages. New output appears in the Status window. When the job finishes running, Completed or Failed, appear in the Status field, and ASSIGNMENT FINISH appears in the Status window. For information about the job status that can be found in the user interface. See “Interpreting job status from the Status window” on page 123.

About aborting a distribution event

You can abort a distribution event while it is in progress. When you do so, the core application removes any files and directories that were created or copied on the source or destination devices. You cannot abort a distribution event after it has reached the processing stage.

Aborting a distribution event

To abort a distribution event, complete the following steps.

Step	Action
1	Click the Abort icon on the lower menu bar.
2	Click Yes at the prompt. Result: “aborting” appears in the Status field of the Jobs view table, after which the job status is listed as failed or aborted. In the History table for the job, the abort is listed as a separate assignment and is designated Completed in the history Status field. The distribution event assignment is listed as failed or aborted.

Determining job status

About job status information

As a job is run, information about the distribution event from transport group members, transport group leaders, and the core application is saved in log files. The core application accesses the information through the database. The user interface displays some of this information dynamically as the distribution event progresses. You can, also, find a static record of the final disposition of a job assignment in the History dialog box.

Where to look

You can find information about job status at three locations in the ContentDirector user interface.

- ◆ The Status field in the Jobs view table
For information about the Status field in the Jobs view table “[Stages of the distribution event](#)” on page 116
- ◆ The Status window
- ◆ The History dialog box

Interpreting job status from the Status window

About the Status window

The information in the Status window is used, primarily, by Network Appliance Technical Support. However two lines of information indicate the nature of the assignment and the status and outcome of the job. These lines are easily identified and come at the beginning and the end of the distribution event.

Note

You can copy and paste lines from the status window to the clipboard.

Text at the beginning of the job

You can determine the beginning of a job in the Status window by identifying the following lines.

```
AssignmentID: 50129 [TaskPublish] *****
```

```
AssignmentID: 50129 [TaskPublish] ***** ASSIGNMENT START (100)
```

```
AssignmentID: 50129 [TaskPublish] *****
```

Informational text about the job

The assignment ID (50129) identifies the assignment and corresponds to the ID number in the Snapshot directory for the job in the Workspace.

The number in the parentheses to the right of ASSIGNMENT START identifies the type of assignment. The value 100 designates the content distribution event. The core application creates numerous assignment types. Some are initiated by the core application itself. Others are user-initiated assignment types. These include the following.

- ◆ Content distribution event: 100
- ◆ Abort: 106

Text at the end of the job

You can determine the end of a job in the Status window by identifying the following lines.

```
AssignmentID: 50129 [TaskPublish *****  
AssignmentID: 50129 [TaskPublish] ***** ASSIGNMENT FINISH (291/3000)  
AssignmentID: 50129 [TaskPublish] *****
```

Information at the end of the job

The two numbers to the right of ASSIGNMENT FINISH are the job status code (291) and the error code for processing by the core application (3000).

Job status codes include the following.

- ◆ 290: Job failed
- ◆ 291: Job succeeded
- ◆ 269: Job aborted

Core application error codes include the following.

- ◆ Error state none: 3000
- ◆ Error state unknown: 3001
- ◆ Assignment creation failed: 3002
- ◆ Availability quorum failed: 3003
- ◆ Commit quorum failed: 3004
- ◆ Channel creation failed: 3005
- ◆ Workspace update failed: 3006
- ◆ Failed to initiate distribution event: 3007
- ◆ Assignment timed out: 3008

Using Job History

About Job History

The Job History table lists assignments that were created for a given job. These include content distribution assignments (publish) and abort assignments. You find details of the selected assignments in a tree display.

Fields in the Job History table

The Job History table includes the following fields:

- ◆ Start: the time and date that the distribution started
- ◆ Finish: time and date that it finished
- ◆ Status: current status. This field is the same as the Status field in the Jobs view table. It is updated as the distribution event progresses.
- ◆ Type: the assignment type
- ◆ Duration: the time that the event took to complete


Using Job History

To view the Job History table, complete the following steps.

Step	Action
1	Select the job of interest in the Jobs view table.
2	Click History in the lower menu bar. Result: The Job History table appears.
3	To view the history for another job, click the job drop-down menu.
4	Select the desired job.

An illustration of the Assignment Details display

The following illustration presents detailed information of the assignments selected in the Job History table.



Using the Assignment Details display

To view the Assignment Details display, complete the following steps

Step	Action
1	In the Jobs view table, select the desired assignments.
2	Click the Details button. Result: The Assignment details display appears.

Deleting content from a destination

About deleting content

You can delete content that you distributed through ContentDirector by deleting it from the source directory. The next time the job for that source runs, the files you deleted from the source are ejected from NetCache.

The requirements for deleting content are:

- ◆ You must delete the content from the source directory in ContentDirector.
- ◆ The content must have originally been distributed by CD, not added to the NetCache by other means.
- ◆ The content must not be locked in the cache.

Note

You cannot trigger an eject from NetCache for sources that specify only of a single file. The source must specify a directory to use this method. For this reason it is recommended to always create sources as directories rather than selecting specific files.

Procedure

To delete content from the source and from NetCache, complete the following steps:

Step	Action
1	Use the local file system on the core application machine to delete files from the source.
2	Run the job manually to trigger an immediate eject from NetCache, or wait until the next scheduled distribution event for the job. See “An example of running a job” on page 119 for detailed instructions.

Changing master keys to ensure data security

About master keys

ContentDirector uses a master key file at both ends of each data transfer to authenticate the identity of the sender and receiver. The installation process creates a default master key on the core application and destination machines.

For more secure content distribution, you can create a new, unique master key for your ContentDirector network. The process consists of the following steps:

- ◆ Generating the new master key at the core application
- ◆ Activating the new master key at the core application
- ◆ Transferring the new master key to destinations

How the process works

You create a master key through the Tools menu of the management console. When you do so, the core application generates a new master key with the name `master_key_new <digits>`, and places it in the config directory of the core application directory.

When you create a master key, it does not become the current key automatically. To function, the master key file must be renamed `master_key` on the core application and the transport group leaders and group members. After the new key is renamed `master_key` on the core application machine, it becomes the operative key and no content distribution can take place until the new key is transferred to the group leaders on the destination devices and is renamed `master_key`.

The method you use for transferring the new master key to destinations depends on the destination type and the degree to which security is a concern:

- ◆ For destinations where security is not a primary concern, you can use the NetCache Manager utility to transfer the master key.
- ◆ For destinations where security is a primary concern, you can copy the new master key from the core application server to a diskette and deliver it manually.

Generating a new master key

Generating a new master key

To create a new master key, complete the following steps.

Step	Action
1	On the main menu, select Tools > Generate Master Key. Result: A new master key is generated.
2	Click OK. Result: The following message is displayed. The master key is successfully created and stored in the file: <path>/core application/config/master_key_<digits>. Make sure to copy new Master Key to each Group Member.

Activating the new master key on the core application

To activate a new master key on the core application, complete the following steps.

Caution—
After activating the new master key on the core application, the old master key becomes invalid. You must also activate the new master key on group members before you can distribute content.

Step	Action
1	Rename the old master key file to back it up.
2	Change the name of the new master key file to master_key.

Changing master keys on the CDA on NetCache

ContentDirector contains two master key files

The ContentDirector core application generates two separate master key files that can be used to install the master key on the CDAs. These two files are

- ◆ master_key.sh

The master_key.sh file is a shell script that automatically installs the master key on the CDA. Install the master_key.sh file by using the NetCache automatic configuration feature on the CDA. See [“Installing the master key from the NetCache Manager utility”](#) on page 131 and [“Installing the master key from the CLI”](#) on page 133 for more information.

- ◆ master_key

The master_key file is a Java Serialized Object file that can be used to install the master key on a CDA if you want to avoid the security risks involved in transferring the master_key.sh file over the network. See [“If you want a more secure way of installing the master key”](#) on page 134 for more information.

Note

To upload or download files using the Secure Sockets Layer (https://), you must have purchased and installed the SecureAdmin™ software. See your Network Appliance sales representative for more information.

About using the automatic configuration feature

The primary function of the NetCache automatic configuration feature (also called *multiple cache management*) is to automate the periodic propagation of configuration settings to multiple target NetCache appliances. Although the automatic configuration feature can be used to install the ContentDirector master key, you must exercise caution because you can alter preexisting configuration rules.

Before performing the procedure described in [“Installing the master key from the NetCache Manager utility”](#) on page 131, consider the following recommendations:

- ◆ When you enter the URL of the master_key.sh file in the Automatic Configuration Files Box, do not overwrite any preexisting configuration rules.

- ◆ After you have downloaded the master_key.sh file, remove the MasterKey.sh URL from the Automatic Configuration Files box so that the action will not be repeated.

For more information about the automatic configuration feature, see the multiple cache management information in the *Routine Administration Guide*.

Installing the master key from the NetCache Manager utility

To use the NetCache automatic configuration feature to install the master key onto the CDA, complete the following steps.

Step	Action
1	Verify that the master_key.sh file has been copied from the config directory on the core application to an FTP or Web server that is accessible to the remote CDA. For example, the Tomcat Web server, which resides on the ContentDirector core application, can be used to serve the MasterKey.sh file. Record the URL of the master_key.sh file.
2	From the NetCache Manager utility, select Setup > ContentDirector > General.
3	Under ContentDirector Agent Enable, disable the ContentDirector service.
4	Click Commit Changes.
5	Access the NetCache CLI on the CDA for which you are installing the master_key.sh file.
6	Enter the following NetCache CLI command to clear any preexisting local master key: <pre>clear config.cd.master_key</pre>
7	From the NetCache Manager utility, select Setup tab > System > Automatic Configuration.

Step	Action
8	Enter the URL of the master_key.sh file (refer to Step 1) in the Automatic Configuration Files box. Caution_____ Do not overwrite any preexisting automatic configuration rules. _____ Use the following syntax: <pre>{http https ftp}://[user:passwd@]ccm_srvr:ccm_port/directory/master_key.sh</pre> In the preceding URL, <i>user</i> is the user name, <i>passwd</i> is the password, <i>ccm_srvr</i> is the FTP or Web server host name or IP address, and <i>ccm_port</i> is the Content Control Manager port (the default is port 8080). Note_____ You need only provide a user name and password if the server requires authentication. _____ For example: <pre>http://Tomcat:8080/ContentDirector/master_key.sh</pre>
9	In the Get Configuration Files Now option, click Manual Update to immediately install the master key. Result: The master key is installed on the CDA.
10	Enter the following NetCache CLI command to ensure that the master key was successfully installed: <pre>show config.cd.master_key</pre>
11	From the NetCache Manager utility, select Setup > ContentDirector > General.
12	Under ContentDirector Agent Enable, enable the ContentDirector service.
13	Click Commit Changes.
14	For security reasons, delete the master_key.sh file from the ContentDirector directory on the Tomcat Web server.

Installing the master key from the CLI

To install the master_key.sh file from the NetCache CLI, complete the following steps.

Step	Action
1	Copy the master_key.sh file to a temporary directory on an FTP or HTTP server. Ensure that the master_key.sh URL is accessible to the remote CDA. For example: <code>http://tomcat:8080/ContentDirector/master_key.sh</code>
2	From the NetCache Manager utility, select Setup > ContentDirector > General.
3	Under ContentDirector Agent Enable, disable the ContentDirector service.
4	Click Commit Changes.
5	Access the NetCache CLI on the CDA for which you are installing the master_key.sh file.
6	Use the NetCache CLI <code>add config.admin.config.list</code> command to import the master_key.sh file from its temporary FTP or HTTP location. Use the following syntax: <pre>add config.admin.config.list before local.cfg {http https ftp}://[userid:passwd@]temp_srvr:port/direct orypath/master_key.sh</pre> In the preceding syntax, <i>userid</i> is the user name, <i>passwd</i> is the password, <i>temp_srvr</i> is the FTP or HTTP server host name or IP address, and <i>port</i> is the FTP or HTTP server port. For example: <pre>add config.admin.config.list before local.cfg http://tomcat:8080/ContentDirector/master_key.sh</pre>

Step	Action
7	Enter the following NetCache CLI command to confirm that the master key was installed successfully: <pre>show config.cd.master_key</pre> If the master key installation is successful, a master key code number is displayed.
8	Use the NetCache CLI <code>delete</code> command to delete the <code>master_key</code> URL from the NetCache appliance configuration file download list. For example: <pre>delete config.admin.config.list http://tomcat:8080/ContentDirector/master_key.sh</pre>
9	From the NetCache Manager utility, select Setup > ContentDirector > General.
10	Under ContentDirector Agent Enable, enable the ContentDirector service.
11	Click Commit Changes.
12	Delete the <code>master_key.sh</code> file from its temporary transfer location on the FTP or HTTP server.

If you want a more secure way of installing the master key

If you are concerned about the security risks involved in transferring the `master_key.sh` file over the network, you can install the Java Serialized Object `master_key` file. The `master_key` file can be copied to disk and manually transported to the local network, where it can be copied directly to the CDA.

To install the `master_key` file, complete the following steps.

Step	Action
1	From the NetCache Manager utility, select Setup > ContentDirector > General.
2	Under ContentDirector Agent Enable, disable the ContentDirector service.
3	Click Commit Changes.

Step	Action
4	Enter the following NetCache CLI command to clear any preexisting local master key: <code>clear config.cd.master_key</code>
5	Copy the master_key file (from disk) to a location on the local network that is accessible to the CDA.
6	Using the NetCache CLI <code>cp</code> command, copy the master_key file to the following directory on the NetCache appliance: /tmp The syntax of the <code>cp</code> command is as follows: <code>cp source_file destination_file</code>
7	Enter the following NetCache CLI command to set the master key: <code>setmk</code>
8	Enter the following NetCache CLI command to ensure that the master key was successfully installed: <code>show config.cd.master_key</code>
9	From the NetCache Manager utility, select Setup > ContentDirector > General.
10	Under ContentDirector Agent Enable, enable the ContentDirector service.
11	Click Commit Changes.

Note

After the `setmk` command is executed, the master_key file is automatically deleted from the /tmp directory location on this appliance.

About this appendix

This appendix contains the following three tables:

- ◆ Core application error messages sorted by code number, and cross-referenced by detailed history text and log text
- ◆ General error codes and their corresponding messages, sorted numerically by code number
- ◆ General error messages and their corresponding codes, sorted alphabetically by message (the same data as the second table, but sorted by message)

Core application error messages and codes, sorted by code number

The following table contains core application messages and codes from the job detailed history and from the logs . Data is sorted numerically by code.

Code	Job Detailed History	Error Log	Possible Cause(s)
3000		IConstantsFES.ERR_CCM_NONE	No Error
3001	Unknown error.	FES.ERR_CCM_UNKNOWN	Unspecified error condition.
3002	Assignment Creation Failed.	FES.ERR_CCM_ASSGN_CREATE_FAILED	Error determining assignment changes. Files not readable. Disk accessibility.
3003	Availability Quorum Failed.	FES.ERR_CCM_PREQUORUM_FAILED	Not enough nodes reported online to meet availability quorum. Nodes offline/unreachable.
3004	Commit Quorum Failed.	FES.ERR_CCM_QUORUM_FAILED	Not enough nodes processed to commit assignment. Nodes offline/unreachable. Nodes failed for local reason.
3005	Transport Channel Creation Failed.	FES.ERR_CCM_CHANNELS_FAILED	Problem attempting to setup multicast channels. GLs offline/unreachable.

Code	Job Detailed History	Error Log	Possible Cause(s)
3006	Workspace Update Failed.	FES.ERR_CCM_WORKSPACE_UPDATE_FAILED	Failure updating shadow after publish. Possible disk space, permission issues.
3007	Initialization Failed.	FES.ERR_CCM_INIT_FAILED	Could not create assignment directory. Problems accessing database.
3008	Assignment Timed Out.	FES.ERR_CCM_ASSGN_TIMEOUT	GLs not reporting progress.
3009	User Aborted.	FES.ERR_CCM_USER_ABORT	User aborted assignment.
3010	Disk Space Availability Failure	FES.ERR_CCM_LOW_SPACE	Running low on disk space.
3011	No Destinations for Assignment.	FES.ERR_CCM_NODESTS	Core application found no nodes to push to. Possible database problem, destination group issue.
3012	No Journal file found for Repush.	FES.ERR_CCM_REPUSH_NOSNAP	Could not find a JournalFile to perform repush. Failed to access db, file was deleted by user.
3013	No Target found for Repush.	FES.ERR_CCM_REPUSH_NOTARG	Could not find a node to repush to. Failed to access db, db has no info.
3020	Removal of Temporary Channels Failed.	FES.ERR_CCM_TMPCHANCLEANUP	Failed to remove a temporary mcast channel. Not a fatal error but a warning. Could not contact GL.
3024	A Required file was not found.	FES.ERR_CCM_FILENOTFOUND	Disk accessibility, file removed by another process.
3025	The Shadow Directory Does Not Exist.	FES.ERR_CCM_NOSHADOW	Failed to create, disk accessibility.
3026	The Directory does not Exist	FES.ERR_CCM_DIR_NOEXIST	Failed to create, disk accessibility.

Code	Job Detailed History	Error Log	Possible Cause(s)
3027	The Specified Object is not a Directory.	FES.ERR_CCM_DIR_NOTDIR	A file by that name already exists in the file system and it is not a directory.
3028	No Read Permission to this Directory.	FES.ERR_CCM_DIR_NOTREAD	Disk accessibility - can't read.
3029	No Write Permission to this Directory.	FES.ERR_CCM_DIR_NOTWRITE	Disk accessibility - can't write.
3030	Could not create Workspace.	FES.ERR_CCM_WORKSPACE_NOCREATE	Disk accessibility - can't write/read.
3031	Insufficient Access to Workspace.	FES.ERR_CCM_WORKSPACE_INSUFFICIENTPERMS	Disk accessibility - can't write/read.
3032	No Snapshot Detected.	FES.ERR_CCM_NOFILERSNAP	Snapshots not enabled, out of snapshots.
3033	Could not find the assignment dir.	FES.ERR_CCM_NOASSIGNDIR	Disk accessibility, can't read/write. Configuration problem.
3034	The CDFES service was shutdown.	FES.ERR_CCM_SERVICESHUTDOWN	The CDFES service was shutdown. Network or servlet issues.
3035	Could not communicate with GrpLdr.	FES.ERR_CCM_TRANSPORTException	Could not communicate with GL. GL down/unreachable. Servlet problem.

General error messages and codes, sorted by code number

The following table contains messages from the job detailed history and codes from the logs and Status window of the main screen. Data is sorted numerically by code.

Code	Message	Possible Cause(s)
13	No disk space.	Running low on disk space.
14	Time out expired.	No response from target. Possible network/connectivity issues.
15	Invalid parameters.	Possible configuration problem.
17	Not enough memory.	Running low on memory.
18	DNS error.	DNS misconfigured, not available.
6001	Failed to open URL.	Attempt to open a socket for content assignment (CA) fetch failed.
6002	URL connection failed.	An I/O Exception raised while attempting to issue HTTP GET for content assignment (CA) fetch.
6003	File size is zero.	Caller passed a zero (or negative) file size parameter to assignment fetch logic (sendCA()). Value must be positive and greater than zero.
6004	Content is written short.	The thread that GCDP started to fetch a given content assignment (CA) received less than the expected number of bytes.
6005	Unknown exception thrown for receiver.	An unanticipated exception was raised during execution of assignment fetch logic (sendCA()).
6006	Http get failed.	Non-success status was returned in HTTP header in response to HTTP GET. It should be "200 OK" but was not. Content assignment cannot be fetched.
6007	No fetch because of no output stream.	MCTP returned a null output stream reference. A valid non-null such reference is required for assignment fetch logic (send)CA()) to transfer the incoming HTTP stream up to MCTP for multicast distribution.
6008	Remote service unreachable for receiver.	Cannot make contact with remote service for MCTP to issue either an asynchronous completion of a request or an unsolicited status notification to core application.
8002	Channel not found.	Core application attempts to push job to channel that GL does not have. Possible problem with database relating group to GL channels or possibly problem with core application creating/communicating channels to use to GL.

Code	Message	Possible Cause(s)
8003	Content assignment not found.	Core application asks to abort or asks status on assignment that GL does not have any knowledge of. GL might not be able to contact core application to report status and has aborted or completed already without informing core application.
8005	No report.	
8007	Back End Server is down.	Cache is not responding to ping requests and was eliminated from push because of this make sure CDA is started on cache.
8019	MCTP client timed out.	GM doesn't return status information during data transmission in about 50 sec. The GM is probably not functioning.
8020	MCTP retransmission count exceeded.	GM cannot get all the data within a data burst during data transmission after several retransmissions. Probable reason is that channel loss rate is extremely high.
8021	MCTP client window stale.	The data delivered to a GM are not processed in timely manner.
8022	MCTP client stream closed.	GM discontinued data reception. Any assignment processing error can cause this.
12003	Content verification failed.	When verification is turned on, this means it failed. Corrupted assignment.
12004	Security alert, possibly MasterKey mismatch.	Masterkey mismatch.
12007	I/O Error while reading/writing to cache.	Socket connection to cache was closed for some reason. Something in the push header indicates that the object is not cacheable, such as .cgi, ? in filename, streaming push with non-streaming object, wrong content-type for streaming object.
12013	Content Interpreter thread died unexpectedly	Any unexpected event. Check the CDLog.Systemerr.log for more detail.
12015	Content assignment not supported.	Abort an assignment that can't be aborted, such as commit.
12017	Command buffer not found.	Assignment header is corrupted, or no disk space on cache to extract command buffer.

Code	Message	Possible Cause(s)
12021	Unknown assignment command.	Command buffer corrupted, or core application is a totally new version with new commands that the GM does not understand.
12022	CI/DP registration failed	Obsolete cd.zip file present on NetCache. Delete cd.zip if exists in /etc/java in cache.
12023	Cache quorum not met.	Completion criterion is set for quorum and it is not met. Check xml log for exception details.
12024	Group Member type unknown.	Config incorrect, check registry.
12025	Command buffer invalid.	Command buffer corrupted, or core application is a total new version with new commands and arguments that GM does not understand.
12026	Failed connecting to cache.	Cannot make connection to cache on loop back.
12027	System error.	Occurs only on upgrade GM from core application, failed exec command on cache.
12028	I/O Error on Streaming, incorrect number of bytes sent	GL or core application aborted the assignment.
12030	Content assignment header not found for payload.	PUB failed for that assignment but the payload still sent. Check disk space on cache.
12031	Content event unknown.	Core application is a totally new version with new events that GM does not understand
12400	Failed to cache object, Bad request.	Return code 400 from cache, check push header in xml file log.
12500	Failed to cache object.	Return code 500 from cache, check log.
12510	Cache content fill error.	Return code 510 from cache, many causes, such as disk space, non-cacheable object, corrupted streaming file.
15001	Property file not found.	One or both of two possible virtual site property files is not where system expects it to be.

Code	Message	Possible Cause(s)
15007	Unknown exception thrown.	An unanticipated exception was raised during execution of cache registry gathering and reporting logic.
15008	Number format exception thrown.	Unanticipated non-numeric data encountered during execution of cache registry gathering and reporting logic.
15009	Property file data invalid.	Ill-formatted or unexpected data found in one or both of two possible virtual site property files.
15010	Gigo exception thrown.	A call to an AgentStorage interface method was made before AgentStorage data is fully populated and ready to return valid data.

General error messages and codes, sorted by message

The following table contains messages from the job detailed history and codes from the logs and Status window of the main screen. Data is sorted alphabetically by message.

Message	Code	Possible Cause(s)
No disk space.	13	Running low on disk space.
Time out expired.	14	No response from target. Possible network/connectivity issues.
Invalid parameters.	15	Possible configuration problem.
Not enough memory.	17	Running low on memory.
DNS error.	18	DNS misconfigured, not available
Failed to open URL.	6001	Attempt to open a socket for content assignment (CA) fetch failed.
URL connection failed.	6002	An I/O Exception raised while attempting to issue HTTP GET for content assignment (CA) fetch.
File size is zero.	6003	Caller passed a zero (or negative) file size parameter to assignment fetch logic (sendCA()). Value must be positive and greater than zero.

Message	Code	Possible Cause(s)
Content is written short.	6004	The thread that GCDP started to fetch a given content assignment (CA) received less than the expected number of bytes.
Unknown exception thrown for receiver.	6005	An unanticipated exception was raised during execution of assignment fetch logic (sendCA()).
Http get failed.	6006	Non-success status was returned in HTTP header in response to HTTP GET. It should be "200 OK", but was not. Content assignment cannot be fetched.
No fetch because of no output stream.	6007	MCTP returned a null output stream reference. A valid, non-null such reference is required for assignment fetch logic (send)CA()) to transfer the incoming HTTP stream up to MCTP for multicast distribution.
Remote service unreachable for receiver.	6008	Cannot make contact with remote service for MCTP to issue either an asynchronous completion of a request or an unsolicited status notification to core application.
Channel not found.	8002	Core application attempts to push job to channel that GL does not have. Possible problem with database relating group to GL channels or possibly problem with core application creating/communicating channels to use to GL.
Content assignment not found.	8003	Core application asks to abort or asks status on assignment that GL does not have any knowledge of. GL might not be able to contact core application to report status and has aborted or completed already without informing core application.
No report.	8005	
Back End Server is down.	8007	Cache is not responding to ping requests and was eliminated from push because of this make sure CDA is started on cache

Message	Code	Possible Cause(s)
MCTP client timed out.	8019	GM doesn't return status information during data transmission in about 50 sec. The GM is probably not functioning.
MCTP retransmission count exceeded.	8020	GM cannot get all the data within a data burst during data transmission after several retransmissions. Probable reason is that channel loss rate is extremely high.
MCTP client window stale.	8021	The data delivered to a GM are not processed in timely manner.
MCTP client stream closed.	8022	GM discontinued data reception. Any assignment processing error can cause this.
Content verification failed.	12003	When verification is turned on, this means it failed. Corrupted assignment.
Security alert, possibly MasterKey mismatch.	12004	Masterkey mismatch.
I/O Error while reading/writing to cache.	12007	Socket connection to cache was closed for some reason. Something in the push header indicates that the object is not cacheable, such as .cgi, ? in filename, streaming push with none streaming object, wrong content-type for streaming object.
Content Interpreter thread died unexpectedly	12013	Any unexpected event. Check the CDLog.Systemerr.log for more detail.
Content assignment not supported.	12015	Abort an assignment that can't be aborted, such as commit.
Command buffer not found.	12017	Assignment header is corrupted, or no disk space on cache to extract command buffer.
Unknown assignment command.	12021	Command buffer corrupted, or core application is a totally new version with new commands that the GM does not understand.
CI/DP registration failed	12022	Obsolete cd.zip file present on NetCache. Delete cd.zip if exists in /etc/java in cache.
Cache quorum not met.	12023	Completion criterion is set for quorum and it is not met. Check xml log for exception details.

Message	Code	Possible Cause(s)
Group Member type unknown.	12024	Config incorrect, check registry.
Command buffer invalid.	12025	Command buffer corrupted, or core application is a total new version with new commands and arguments that GM does not understand.
Failed connecting to cache.	12026	Cannot make connection to cache on loop back.
System error.	12027	Occurs only on upgrade GM from core application, failed exec command on cache.
I/O Error on Streaming, incorrect number of bytes sent	12028	GL or core application aborted the assignment.
Content assignment header not found for payload.	12030	PUB failed for that assignment but the payload still sent. Check disk space on cache.
Content event unknown.	12031	Core application is a totally new version with new events that GM does not understand
Failed to cache object, Bad request	12400	Return code 400 from cache, check push header, it should be in the xml file log.
Failed to cache object.	12500	Return code 500 from cache, check log.
Cache content fill error.	12510	Return code 510 from cache, many causes, such as disk space, non- cacheable object, corrupted streaming file.
Property file not found.	15001	One or both of two possible virtual site property files is not where system expects it to be.
Unknown exception thrown.	15007	An unanticipated exception was raised during execution of cache registry gathering and reporting logic.
Number format exception thrown.	15008	Unanticipated non-numeric data encountered during execution of cache registry gathering and reporting logic.
Property file data invalid.	15009	Ill-formatted or unexpected data found in one or both of two possible virtual site property files.

Status field values The following table describes the Status field values.

Status	Description
Idle	The job has not been run yet.
Failed	The job failed.
Completed	The job succeeded.
Aborted	The user aborted the job while it was running.
Pending	The job is scheduled to run and will run shortly.
Expired	The job's end schedule date expired.
Running	The job is initiated and the subdirectories in the Workspace are being created.
ChanCreate	1st stage: creating the distribution channels.
AsgnCreate	2nd stage: creating the assignment.
CmdDeliver	3rd stage: delivering the command buffer.
DataDeliver	4th stage: delivering the job assignment data (content).
Processing	5th stage: the destination devices are processing the job assignment. The Commit process occurs during this stage.
Aborting	The job encountered errors and is automatically aborting.

Descriptions of tree display icons

The following table describes the tree display icons.

Icon	Description
	General
	All Network Segments
	All destination servers, all content groups, sources, and network segments (transport groups)
	Unknown device
	Source
	Destination group
	No license
	Transport Group Members
	Online, responding, version matches
	NetCache appliance with host site
	Version mismatch

Icon	Description
	Offline
	Not responding
	Transport Group Leaders
	Online, responding, version and master key match
	Version or master key mismatch
	Offline

About default Mime types

ContentDirector includes an extensive list of default mime types. If you are distributing content of a type not included in the defaults, you can add new mime types in global Preferences. See “[Setting content distribution preferences](#)” on page 73 for detailed instructions.

List of default mime types

Following is a list of the default mime types included with ContentDirector 2.0 as shipped:

123	application/vnd.lotus-1-2-3
ai	application/postscript
aif	audio/x-aiff
aifc	audio/x-aiff
aiff	audio/x-aiff
asd	application/astound
asf	video/x-ms-asf
asn	application/astound
asp	application/x-asap
asx	video/x-ms-asf
au	audio/basic
avb	application/octet-stream
AVB	application/octet-stream
avi	video/msvideo
bcpio	application/x-bcpio
bin	application/octet-stream
bld	application/bld
bld2	application/bld2
bmp	image/bmp

bpk	application/octet-stream
bz2	application/x-bzip2
c	text/plain
cal	image/x-cals
cc	text/plain
ccad	application/clariscad
ccn	application/x-cnc
cco	application/x-cocoa
cdf	application/x-netcdf
cgi	magnus-internal/cgi
class	application/octet-stream
clp	application/x-msclip
cmx	application/x-cmx
cod	image/cis-cod
cpio	application/x-cpio
crd	application/x-mscardfile
cs	application/x-csh
csm	chemical/x-csml
csml	chemical/x-csml
css	text/css
dcm	x-lml/x-evm
dcr	application/x-director
dcx	image/x-dcx
dir	application/x-director
doc	application/msword
dot	application/x-dot
drw	application/drafting

dvi	application/x-dvi
dwf	drawing/x-dwf
dwg	image/vnd
dxg	image/x-dxg
dxr	application/x-director
ebk	application/x-expandedbook
emb	chemical/x-embl-dl-nucleotide
embl	chemical/x-embl-dl-nucleotide
eps	application/postscript
es	audio/echospeech
esl	audio/echospeech
etc	application/x-earthtime
etx	text/x-setext
evm	x-lml/x-evm
evy	application/x-envoy
exe	application/octet-stream
f90	text/plain
fh4	image/x-freehand
fh5	image/x-freehand
fhc	image/x-freehand
fif	image/fif
fm	application/x-maker
fpx	image/x-fpx
FPX	image/x-fpx
fvi	video/isisvideo
gau	chemical/x-gaussian-input
gdb	x-lml/x-gdb

gif	image/gif
gps	application/x-gps
gtar	application/x-gtar
gz	x-gzip
h	text/plain
hdf	application/x-hdf
hdml	text/x-hdml
hh	text/plain
hlp	application/winhlp
hqx	application/mac-binhex40
htm	text/html
html	text/html
ice	x-conference/x-cooltalk
ico	image/x-icon
ief	image/ief
ifs	image/ifs
iges	application/iges
igs	application/iges
ins	application/x-NET-Install
ipx	application/x-ipix
it	audio/x-mod
itz	audio/x-mod
ivr	i-world/i-vrml
jpe	image/jpeg
jpeg	image/jpeg
jpg	image/jpeg
js	application/x-javascript

jwt	application/jwt
lak	x-lml/x-lak
latex	application/x-latex
lcc	application/fastman
lcl	application/x-digitalloca
lcr	application/x-digitalloca
lgh	application/lgh
lml	x-lml/x-lml
lmlpack	x-lml/x-lmlpack
lsf	video/x-ms-asf
lsx	video/x-ms-asf
lzh	application/x-lzh
LZH	application/x-lzh
m	text/plain
m13	application/x-msmediaview
m14	application/x-msmediaview
m3u	audio/x-mpegurl
man	application/x-troff-man
map	magnus-internal/imapemap
mbd	application/mbedlet
mct	application/x-mascot
mdb	application/x-msaccess
mdz	audio/x-mod
me	application/x-troff-me
mi	application/x-mif
mid	audio/midi
midi	audio/midi

mif	application/x-mif
mil	image/x-cals
mio	audio/x-mio
mng	video/x-mng
mny	application/x-msmoney
moc	application/x-mocha
mocha	application/x-mocha
mod	audio/x-mod
mof	application/x-yumekara
mol	chemical/x-mdl-molfile
mop	chemical/x-mopac-input
mov	video/quicktime
movie	video/x-sgi-movie
mp2	audio/x-mpeg
mp3	audio/x-mpeg
mpe	video/mpeg
mpeg	video/mpeg
mpg	video/mpeg
mpp	application/vnd.ms-project
mps	application/x-mapserver
mpv2	video/x-mpeg2
mrl	text/x-mrml
ms	application/x-troff-ms
mts	application/metastream
nc	application/x-netcdf
ndb	x-lml/x-ndb
nif	application/x-nif

nmz	application/x-scream
npx	application/x-netfpx
oda	application/oda
pac	audio/x-pac
paе	audio/x-epac
pan	application/x-pan
pbm	image/x-portable-bitmap
pcx	image/x-pcx
pda	image/x-pda
pdb	chemical/x-pdb
pdf	application/pdf
pfr	application/font-tdpfr
pgm	image/x-portable-graymap
pict	image/x-pict
pl	application/x-perl
pm	application/x-perl
png	image/png
pnm	image/x-portable-anymap
pot	application/vnd.ms-powerpoint
ppm	image/x-portable-pixmap
pps	application/vnd.ms-powerpoint
ppt	application/vnd.ms-powerpoint
pqf	application/x-cprplayer
pqi	application/cprplayer
proxy	application/x-ns-proxy-autoconfig
prt	application/pro_eng
ps	application/postscript

ptlk	application/listenup
pub	application/x-mspublisher
qcp	audio/vnd.qcelp
qt	video/quicktime
qti	image/x-quicktime
qtif	image/x-quicktime
r3t	text/vnd.rn-realtex3d
ra	audio/x-pn-realaudio
ram	audio/x-pn-realaudio
ras	image/x-cmu-raster
rf	image/vnd.rn-realflash
rgb	image/x-rgb
rlf	application/x-richlink
rm	audio/x-pn-realaudio
rnx	application/vnd.rn-realplayer
roff	application/x-troff
rp	image/vnd.rn-realpix
rpm	audio/x-pn-RealAudio-plugin
rt	text/vnd.rn-realtex
rte	x-lml/x-gps
rtf	application/rtf
rtg	application/metastream
rtx	text/richtext
rv	video/vnd.rn-realvideo
rwc	application/x-rogerwilco
s3m	audio/x-mod
s3z	audio/x-mod

scd	application/x-msschedule
sdf	application/e-score
sea	application/x-stuffit
set	application/set
sh	application/x-sh
shar	application/x-shar
shtml	magnus-internal/parsed-html
shw	application/presentations
sit	application/x-stuffit
skd	application/x-Koan
skm	application/x-Koan
skp	application/x-Koan
skt	application/x-Koan
slc	application/x-salsa
smi	application/smil
smp	application/studiom
snd	audio/basic
sol	application/solids
spc	text/x-speech
spl	application/futuresplash
spr	application/x-sprite
sprite	application/x-sprite
spt	application/x-spt
src	application/x-wais-source
step	application/STEP
stl	application/SLA
stp	application/STEP

sv4cpio	application/x-sv4cpio
sv4crc	application/x-sv4crc
svf	image/vnd
svg	image/svg+xml
svh	image/svh
svr	x-world/x-svr
swf	application/x-shockwave-flash
swfl	application/x-shockwave-flash
t	application/x-troff
tad	application/octet-stream
TAD	application/octet-stream
talk	text/x-speech
tar	application/x-tar
taz	application/x-tar
tbp	application/x-timbuktu
tbt	application/timbuktu
tcl	application/x-tcl
tex	application/x-tex
texi	application/x-texinfo
texinfo	application/x-texinfo
text	text/plain
tgz	application/x-tar
tif	image/tiff
tiff	image/tiff
tki	application/x-tkined
tkined	application/x-tkined
toc	application/toc

tr	application/x-troff
trk	x-lml/x-gps
trm	application/x-msterial
tsi	audio/tsplayer
tsv	text/tab-separated-values
ttf	application/octet-stream
ttz	application/t-time
txt	text/plain
unv	application/i-deas
ustar	multipart/x-ustar
uu	x-uuencode
uue	x-uuencode
vda	application/vda
vdo	video/vdo
viv	video/vivo
vivo	video/vivo
vmi	application/x-dremacast-vms-info
vms	application/x-dremacast-vms
vox	audio/voxware
vqe	audio/x-twinvq-plugin
vqf	audio/x-twinvq
vql	audio/x-twinvq
vre	x-world/x-vream
vrt	x-world/x-vrt
vrw	x-world/x-vream
vts	workbook/formulaone
wav	audio/x-wav

wax	audio/x-ms-wax
wbk	application/MSBinder
wbmp	image/vnd.wap.wbmp
wi	image/wavelet
wm	video/x-ms-asf
wma	audio/x-ms-wma
wmf	application/x-msmetafile
wml	text/vnd.wap.wml
wmlc	application/vnd.wap.wmlc
wmls	text/vnd.wap.wmlscript
wmv	video/x-ms-wmv
wmx	video/x-ms-wmx
wpt	x-lml/x-gps
wri	application/x-mswrite
wrl	x-world/x-vrml
WRL	x-world/x-vrml
wrz	x-world/x-vrml
WRZ	x-world/x-vrml
wv	video/wavelet
wvx	video/x-ms-wvx
wxl	application/x-wxl
xbm	image/x-xbitmap
xdm	application/x-xdma
xdma	application/x-xdma
xdw	application/vnd.fujixerox.docuworks
xla	application/vnd.ms-excel
xlc	application/x-excel

xll	application/x-excel
xlm	application/x-excel
xls	application/x-excel
xlt	application/vnd.ms-excel
xlw	application/x-excel
xm	audio/x-mod
xml	text/xml
xmz	audio/x-mod
xpm	image/x-xpixmap
xul	text/xul
xwd	image/x-xwindowdump
xyz	chemical/x-xyz
yz1	application/x-yz1
z	x-compres
zac	application/x-zaurus-zac
zip	application/zip

When you need mapped drive permissions

You need mapped drive permissions when you are using either a mapped drive or a Network Appliance filer to store a source or a workspace.

Operating system considerations

The procedure you need to follow varies between Solaris and Windows:

- ◆ If you are using Solaris, you only need to give the machine running the core application root access to the mounted drive.
- ◆ If you are using Windows, there are several possible scenarios depending on how you are running the ContentDirector service. See “[Windows scenarios](#)” for details.

Windows scenarios

The simplest scenario for using a filer for ContentDirector sources or workspaces under Windows fulfills *both* of the following conditions:

- ◆ The filer and the machine on which the ContentDirector core application resides are both in the same domain.
- ◆ The ContentDirector service runs as a user in that domain, with permissions on the filer.

There are three other scenarios that are more complex:

- ◆ ContentDirector runs as a local system account.
Because the local system account runs with a NULL session ID, you must enable special privileges on the server side to allow access to a NULL session ID. For instructions on configuring the NULL user in DataONTAP, see “Using a NetApp Filer with Lotus Domino Server R5 for Windows NT/2000,” section 7.1.1, on the Network Appliance NOW site, now.netapp.com.
- ◆ The filer and the machine running the ContentDirector core application are in separate domains.
The client where the ContentDirector core application is resident must access the filer as the user *FilerDomain\username*.
- ◆ The filer is not in any domain, but in a workgroup instead.
You must make a registry change to enable plain text passwords on the client where the ContentDirector core application resides. For instructions on enabling plain text plain text passwords for filers, see solution 3.0.1234115.2574126 on the Network Appliance NOW site,

now.netapp.com. For information on running applications as services under Windows NT 4.0, see *Srvany.doc* in the Windows NT Resource Kit.

About port conflicts Conflicts can arise in the RMI and transport ports when other applications are installed or activated on the core application machine. To reconfigure these ports, you can run a script under Windows or edit properties files under Solaris.

For detailed information The following sections describe how to resolve port conflicts in ContentDirector:

- ◆ [“Reconfiguring the RMI port”](#) on page 168
- ◆ [“Reconfiguring the transport port”](#) on page 170

Reconfiguring the RMI port

About the RMI port

Remote method invocation (RMI) is a protocol used for communication between remote Java objects.

The ports ContentDirector uses for RMI are assigned to 1099 by default. Other applications might also bind to port 1099, causing port conflicts. MS Outlook is one example of this. If port 1099 is bound by another application at the time of the ContentDirector installation process, the conflict is resolved silently by the installation utility. If a port conflict arises after installation, follow the instructions below to resolve it.

Reconfiguring the RMI port

Two scripts are available for reconfiguring the RMI port.

- ◆ chgports allows you to specify the port to use for RMI.
- ◆ chgreg automatically changes the port to the next available port and displays the new port number

The location of these scripts depends on the operating system of the core application machine:

- ◆ Under Windows, the ContentDirector\Ccm directory
- ◆ Under Solaris, the /opt/NTAPccm/bin directory

To reconfigure the RMI port, complete the following steps.

Step	Action
1	Stop the ContentDirector service.
2	Navigate to ContentDirector\Ccm (for Windows) or /opt/NTAPccm/bin (for Solaris).

Step	Action
3	At the prompt, run either chgports or chgreg: To run chgports, type: <pre>chgports -rn <port number></pre> where <port number> is the new port, then press Enter. To run chgreg, type: <pre>chgreg</pre> then press Enter.
4	Start the ContentDirector service.

Reconfiguring the transport port

About the transport port

The ports used for transport between the core application and the transport group leaders are assigned to 8080 by default. Other applications might also bind to port 8080, causing port conflicts. RealNetworks' RealPlayer is one example of this.

To reconfigure the transport port, you need to change it at the core application and at each CDA.

Reconfiguring the transport port on the core application

Two scripts are available for reconfiguring the transport port.

- ◆ `chgports` allows you to specify the port to use for transport.
- ◆ `chgtom` automatically changes the port to the next available port and displays the new port number

The location of these scripts depends on the operating system of the core application machine:

- ◆ Under Windows, the `ContentDirector\Ccm` directory
- ◆ Under Solaris, the `/opt/NTAPccm/bin` directory

To reconfigure the transport port, complete the following steps.

Step	Action
1	Stop the ContentDirector service.
2	Navigate to <code>ContentDirector\Ccm</code> (for Windows) or <code>/opt/NTAPccm/bin</code> (for Solaris).
3	At the prompt, run either <code>chgports</code> or <code>chgtom</code>: To run <code>chgports</code>, type: <pre>chgports -tn <port number></pre> where <code><port number></code> is the new port, then press Enter. To run <code>chgtom</code>, type: <pre>chgtom</pre> then press Enter.

Step	Action
4	Start the ContentDirector service.
5	Change the ContentDirector port on each CDA to match the new port. See “ Reconfiguring the transport port on the CDA ”.

Reconfiguring the transport port on the CDA

When you change the transport port at the core application, you must also change it at each CDA to match the new port. To change the port at the CDA, complete the following steps.

Step	Action
1	From the NetCache Manager utility, select Setup > ContentDirector > Advanced.
2	Under Core Application Port, enter the new port number you configured at the core application.
3	Click Commit Changes.

Glossary

appliance	A device that performs a single, well-defined function and is simple to install and operate, for example, a NetCache appliance or NetApp filer.
authentication	The process of verifying the identity of a user who is logging in to a computer system.
cache	<i>Verb.</i> To store data temporarily for expedited access. <i>Noun.</i> The location in which data is stored temporarily. There are a variety of types of caches, including disk cache, memory cache, and Web cache.
client	A workstation or PC in a client-server architecture; that is, a computer system or process that requests services from and accepts the responses of another computer system or process.
console	A physical or virtual terminal that is used to monitor and control a NetApp appliance.
ContentDirector	A Network Appliance scalable software solution that automates the distribution of content from the center to the edge of networks.
ContentDirector agent (CDA)	A component of ContentDirector that acts as distribution or destination point at the edge of the network, for content pushed from the center of the network.
content provider	Any company in the business of aggregating or producing content for delivery via the Internet.

core application	The controlling component of ContentDirector at the center of the network that schedules, selects content, and creates policies for the distribution of content to the edge.
Data ONTAP	A software product designed exclusively for NetApp appliances that optimizes file service by combining patented file-system technology and a microkernel design dedicated to network data access. Developed by Network Appliance.
destination	A node at the edge of the network that receives content through ContentDirector and serves it to clients.
destination group	A group of network edge nodes that are targeted to receive content from the same source in ContentDirector. Destination groups are purely logical and do not need to map to physical groups of nodes.
domain	A subset of ContentDirector content sources and destinations, associated with specific groups of users and visible only to them and administrators.
Domain Name System (DNS)	An Internet service for finding computers on an IP network.
domain administrator	A ContentDirector user with administrative rights to a specific set of program resources.
edge device	Any device, for example a NetCache appliance, that is deployed at the periphery of a LAN and WAN environment to interface traffic from the edge of one network to the other.
file server	A computer that provides file-handing and storage functions for multiple users.
filer	An appliance that attaches to a computer network and is used for data storage. NetApp filers operate using Data ONTAP software.

firewall	A server or device deployed between a LAN and the external WAN to control all inbound and outbound traffic, thus the enhancing security of the network behind the firewall.
fully qualified domain name (FQDN)	The complete name of a specific computer on the Internet, consisting of the computer's host name and its domain name.
group leader	A network node that acts as a ContentDirector distribution point, or passthrough, for content near the edge of a network. Group leaders transmit content to group members.
group member	A network node that acts as a ContentDirector destination point for content at the edge of a network. Group members receive content from group leaders.
IP address	The unique numerical address of a computer that is attached to a TCP/IP network, for example, 198.95.226.66.
job	A re-usable set of parameters in ContentDirector that specify the content and schedule for content distribution events.
MIME	Multipurpose Internet Mail Extensions. A specification that defines the mechanisms for specifying and describing the format of Internet message bodies. An HTTP response containing the MIME Content-Type header allows the HTTP client to invoke the application that is appropriate for the data received.
MMS	Microsoft Media Streaming. The streaming media control protocol used in Windows Media Technologies 3 and 4 streams.
mount	To make the file system on a particular device accessible by attaching it to the file hierarchy.

NetCache	The proprietary software that runs on NetCache appliances. Developed by Network Appliance.
NetCache appliance	An appliance designed specifically to serve Internet and intranet content quickly and reliably. Developed by Network Appliance.
NetCache Manager	The browser interface through which a NetCache appliance is managed.
network interface card (NIC)	A printed circuit board that connects a computer or other node to a network. Also known as a “network adapter.”
NOW	NetApp On the Web. The secure customer service Web site for Network Appliance.
NTFS permissions	The permissions used in NTFS-style security to control access to files and directories.
policy	A re-usable set of rules in ContentDirector governing how content is handled during distribution events, criteria for success, and email notifications of job status.
quorum	Threshold criteria for the success of a distribution event in ContentDirector.
regular user	A non-administrative ContentDirector user. Regular users have only the rights granted by membership in a group of users, defined by an administrator.
Remote Shell	A program that enables a user on one system to execute a program on another system. Remote Shells are usually established with the <code>rsh</code> command and are not interactive.

resource	An object in ContentDirector that is necessary for using the program. Users can create, modify, use and delete resources. Resources include: jobs, policies, destination groups, destinations, sources, checkpoints, contacts, users, and user groups.
rsh	The command used to establish a Remote Shell connection.
RTSP	Real Time Streaming Protocol. An open standard for delivery of real-time media over the Internet.
Snapshot	A Data ONTAP feature that creates an online, read-only copy of the entire file system—called a “snapshot”—that protects against accidental deletions or modifications of files without duplicating file contents.
source	Directory on a local disk drive at the center of the network, containing content to be distributed through ContentDirector to NetCache appliances at the edge .
stream	A media file that is transmitted in a continuous flow over a network. Streams are of two types, live and on-demand.
streaming media	A term used to describe media files that begin playing while they are being transmitted over the network to the media player on the client computer.
streaming server	A server configured to serve streaming media to clients.
system administrator	A ContentDirector user with administrative rights to all program resources.
user group	A list of non-administrative ContentDirector users, and the rights that they share.

Time To Live (TTL)	The time limit on how long a copy of an object remains in the cache before NetCache verifies the object with the originating Web server.
volume	An independent file system with its own RAID groups.

Index

A

- access tab
 - purpose 68
- assignment file
 - relocating 98
- autosupport
 - configuring 84

C

- compression
 - enabling, in policy 89
- contacts
 - creating 84
- content distribution
 - global preferences 73
- ContentDirector
 - components 2
 - flow of content through 3
 - key features 2
 - master key
 - file differences 130
 - installing the MasterKey.sh file 131
 - installing with automatic configuration feature 130
 - secure installation 134
 - overview of 2
 - SecureAdmin for 130
 - starting and stopping 9
 - remote viewing of user interface 17
 - requirements for remote viewing 17
 - services on Solaris 15
 - services on Windows 2000 13
 - services on Windows NT 11
 - user interface on Solaris 15
 - user interface on Windows 2000 13
 - user interface on Windows NT 11
 - starting services
 - on Windows NT 11

D

- default policy 83
- Destination groups
 - definition 3
- destination groups
 - creating 80
 - definition 80
- destinations
 - recovering deleted 103
 - replicating content on new 101
- disk space
 - estimating 97
 - reducing use 98
- domain administrators
 - creating 64
 - prerequisite 64
 - definition 53
- domains
 - creating 59
 - interaction between 51
 - purpose 51
 - workspace 59

E

- email contacts
 - creating 84
- encyption
 - enabling, in policy 85

F

- file deltas
 - sending only 94
- file header
 - controlling, in policy 87
- file types
 - restricting 94
- filer
 - as source location 77
 - rsh permission requirement 78

- filtering content
 - in job 94
- filters
 - for content 94
- forward error correction
 - enabling, in policy 85

H

- host name
 - requirements 79

I

- instance rights
 - compared to class rights 54, 68
 - granting 69
 - precedence over class rights 68
 - why required 68

J

- job history
 - number of assignments in 44
- job status
 - color coding 29
 - notifying personnel of 84
- jobs
 - aborting 120
 - cannot delete running job 112
 - creating 92
 - running 120
 - scheduling 95
 - scheduling considerations 91
 - setting criteria for success 89
 - sizing considerations 91
- jobs view
 - setting length of history 44
 - setting refresh rate 45

M

- master key
 - purpose of 128
- MasterKey file
 - installing 131

- mime types
 - adding new 75

P

- passwords
 - requirements 66
- policies
 - creating 83
 - default policy 83
- preferences
 - setting global 73
 - types of 73
- protocols
 - supported by CD 77

Q

- quorums
 - setting, in policy 89

R

- regular users
 - creating 66
 - prerequisite 66
 - definition 53
- resources
 - auto-discovered 5
 - user-defined 5
- rsh
 - on filer source location 78

S

- satellite links
 - error correction for 85
- sources
 - creating 78
- status window
 - clearing 42
 - disable autoscrolling 42
 - error codes 124
 - information source for 41
 - job status codes 124
 - setting buffer size 41

- setting detail level 41
- setting refresh rate 41
- source of information 41
- system administrators
 - definition 53

T

- time
 - synchronizing 82
- time zones
 - synchronizing 82
- Transport groups
 - definition 3

U

- user groups
 - creating 61
 - prerequisite 61
 - relationship with user rights 52
 - users in multiple 52
- user interface

- at startup 21, 27
- job status color coding 29
- main screen panels 23
- user preferences
 - setting 75
- user rights
 - types of 54
- Users
 - user types and tasks 6
- users
 - in multiple user groups 52

V

- verification
 - enabling, in policy 85

W

- workspace
 - default 74
 - for domain 59
 - setting preference 74

