

http://www.cisco.com/en/US/products/sw/conntsw/ps491/products_installation_and_configuration_guides_list.htm Go MAR JUL FEB
11 captures 21 Mar 2005 - 12 Jan 2014 2007 2008 2013 About this capture



Cisco Application and Content Networking System (ACNS) Software

Configuration Guides

Some links below may open a new browser window () to display the document you selected.

View documents by topics:

Cisco ACNS Software Version 5.5

- [Cisco ACNS Software Configuration Guide for Centrally Managed Deployments, Release 5.5](#)
- [Cisco ACNS Software Configuration Guide for Locally Managed Deployments, Release 5.5](#)
- [Configuring Cisco Access Routers and the NME-WAE Network Module for ACNS Deployments](#)

Cisco ACNS Software Version 5.4

- [Cisco ACNS Software Configuration Guide for Centrally Managed Deployments, Release 5.4](#)
- [Cisco ACNS Software Configuration Guide for Locally Managed Deployments, Release 5.4](#)

Cisco ACNS Software Version 5.3

- [Cisco ACNS Software Configuration Guide for Centrally Managed Deployments, Release 5.3](#)
- [Cisco ACNS Software Configuration Guide for Locally Managed Deployments, Release 5.3](#)

Cisco ACNS Software Version 5.2

- [Cisco ACNS Software Configuration Guide for Centrally Managed Deployments, Release 5.2](#)
- [Cisco ACNS Software Configuration Guide for Locally Managed Deployments, Release 5.2](#)

Cisco ACNS Software Version 5.1

- [Cisco ACNS Software Caching and Streaming Configuration Guide, Release 5.1](#)
- [Cisco ACNS Software Deployment and Configuration Guide, Release 5.1](#)

Cisco ACNS Software Version 5.0

- [Cisco ACNS Software Caching Configuration Guide, Release 5.0](#)
- [Cisco ACNS Software Deployment and Configuration Guide, Release 5.0](#)

Cisco ACNS Software Version 5.0.3

- [Creating Manifest Files for Cisco ACNS Software, Release 5.0.3](#)

Cisco ACNS Software Version 4.2

- [Cisco ACNS Software Caching Configuration Guide, Release 4.2](#)

Cisco ACNS Software Version 4.1.1

- [Cisco ACNS Software Caching Configuration Guide, Release 4.1](#)

[Contacts](#) | [Feedback](#) | [Help](#) | [Site Map](#)

© 1992-2008 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

http://cisco.com/en/US/docs/app_ntwk_services/waas/acns/v55/configuration/central/ Go JAN FEB NOV
4 captures 06
6 Feb 2009 - 17 Jan 2013 2008 2009 2010 About this capture



❑ Cisco ACNS Software Configuration Guide for Centrally Managed Deployments, Release 5.5

- ❑ Index
- ❑ Title Pages
- ❑ Preface
- ❑ Chapter 1: Understanding the ACNS 5.5 Network
- ❑ Chapter 2: Getting Started
- ❑ Chapter 3: Using the Content Distribution Manager GUI
- ❑ Chapter 4: Setting Up Content Request Routing in the ACNS Network
- ❑ Chapter 5: Configuring the ACNS Network for Content Distribution
- ❑ Chapter 6: Configuring the ACNS Network for Content Acquisition
- ❑ Chapter 7: Creating and Managing Programs
- ❑ Chapter 8: Configuring Caching Services
- ❑ Chapter 9: Configuring Streaming Media Services
- ❑ Chapter 10: Managing Caching Performance and Differentiated Services
- ❑ Chapter 11: Viewing Content Replication Status
- ❑ Chapter 12: Configuring Login Authentication, Configuration Authorization, and Accounting
- ❑ Chapter 13: Working with Device Configurations
- ❑ Chapter 14: Configuring Network Interfaces
- ❑ Chapter 15: Configuring Request Authentication and Authorization
- ❑ Chapter 16: Configuring Request Processing Services
- ❑ Chapter 17: Creating and Managing IP Access Control Lists
- ❑ Chapter 18: Configuring Group Authorization Using an Access Control List
- ❑ Chapter 19: Using the Transaction Logs
- ❑ Chapter 20: Configuring Platform and System Settings
- ❑ Chapter 21: Monitoring and Troubleshooting the ACNS Network
- ❑ Chapter 22: Monitoring with SNMP
- ❑ Appendix A: Creating Manifest Files
- ❑

Cisco ACNS Software Configuration Guide for Centrally Managed Deployments, Release 5.5

Click the links on the left to view the individual chapters in HTML format.

 [Download the complete book](#) (PDF - 19 MB)

http://cisco.com/en/US/docs/app_ntwk_services/waas/acns/v55/configuration/central/
Go
 JAN
 FEB
 NOV
 06
 2008
 2009
 2010
 4 captures
 6 Feb 2009 - 17 Jan 2013
 About this capture

© 1992-2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)



Cisco ACNS Software Configuration Guide for Centrally Managed Deployments

Release 5.5

Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Text Part Number: OL-9136-01



THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0601R)

Cisco ACNS Software Configuration Guide for Centrally Managed Deployments, Release 5.5
© 2006 Cisco Systems, Inc. All rights reserved.



Preface	25
Document Objectives	25
Audience	25
Organization	26
Conventions	27
Related Documentation	28
Obtaining Documentation	28
Cisco.com	28
Product Documentation DVD	29
Ordering Documentation	29
Documentation Feedback	29
Cisco Product Security Overview	29
Reporting Security Problems in Cisco Products	30
Obtaining Technical Assistance	31
Cisco Technical Support & Documentation Website	31
Submitting a Service Request	31
Definitions of Service Request Severity	32
Obtaining Additional Publications and Information	32

PART 1

Introducing ACNS Software for Centrally Managed Deployments

CHAPTER 1

Understanding the ACNS 5.5 Network 1

ACNS Network Overview	2
Content Distribution Managers	2
Content Engines	2
Content Routers	2
IP/TV Program Manager	3
Functional Description	3
Content Delivery Network Component	3
Content Caching Device	3
Deployment Scenarios	4
Content Overview	5
On-Demand Content	5
Pre-Positioned Content	6

Pre-Positioned Versus Preloaded Content	6
Live Content	6
Creating and Managing Content in an ACNS Network	6
Content Request Interception and Router Redirection Overview	7
WCCP Transparent Interception	7
Direct Proxy Routing	7
Dynamic Proxy Autoconfiguration	8
Content Router Routing	8
Content Services Overview	8
Content Pre-Positioning	9
Content Caching and Access Control	9
User Authentication and Content Filtering	9
About Device Groups	10
Content Acquisition and Distribution Architecture	11
Understanding Basic Concepts and Terms	11
Acquiring Content	13
Distributing Content Through Channels	13
Content Replication Overview	14
Unicasting	14
Multicasting	14
ACNS Network Topology Considerations	15

CHAPTER 2

Getting Started	1
Overview	1
Initial Device Configuration Overview	2
Initial Network Configuration Overview	4
Summary of Initial Configuration Tasks	6
Enable a Content Distribution Manager	6
Enable a Content Router (optional)	6
Enable Content Engines	7
Configure Content Engine Disk Space	7
Activate All Devices	7
Enable Software Services	8
Configure One or More Request Routing Methods	8
Using the Quick Setup Method	8
Quick Device Configuration Using Autoregistration	8
Quick Device Configuration Using the Setup Utility	10
About Selecting Static IP Addresses or Using Interface-Level DHCP	12
Quick Network Configuration Using the Quick Start Tool	12

Initial Device Configuration Using the Standard CLI	14
Configuring the Device Mode	14
Configuring the Device Network Settings	15
Configuring the Content Distribution Manager	16
Configuring Disk Space	19
Registering a Device with the Content Distribution Manager	20
Initial Device Configuration Summary	21
Initial Network Configuration Using the Standard Content Distribution Manager GUI	22
Activating Devices in the Content Distribution Manager GUI	22
Activating All Inactive Content Engines	24
Where to Go Next	25

CHAPTER 3**Using the Content Distribution Manager GUI 1**

Launching the Content Distribution Manager GUI	1
Content Distribution Manager GUI Primary Feature Tabs	2
Understanding the Interface	3
Window Elements	3
Navigation	5
The Device Home Window	6
Taskbar Icons	8
Using Online Help	11
Logging On Following a Session Timeout	11
Exiting the Content Distribution Manager GUI	12

PART 2**Configuring the ACNS Network for Content Services****CHAPTER 4****Setting Up Content Request Routing in the ACNS Network 1**

Configuring WCCP Transparent Routing	2
About WCCP Services	5
About Dynamic WCCP Redirection Services	6
About WCCP Service for HTTPS	6
About WCCP Custom Web Cache Service	7
Configuring WCCP General Settings for the Content Engine	7
About Client IP Spoofing	9
About WCCP Slow Start	9
About WCCP Clean Shutdown	9
Configuring WCCP Service Settings for the Content Engine	10
About Load Balancing	13

Creating WCCP Router Lists	14
Modifying WCCP Router Lists for the Content Engine	14
Viewing WCCP Router Lists for the Content Engine	15
Assigning Unique Router Lists to Content Engines in a Device Group	16
Configuring WCCP Port Lists for the Content Engine	17
Configuring WCCP Service Masks for the Content Engine	18
Configuring WCCP Transparent Routing Bypass Options	22
Configuring Authentication Traffic Bypass	22
Configuring WCCP Bypass Settings	23
Creating WCCP Bypass List Entries	25
Enabling WCCP on the Router	26
Basic WCCP Router Configurations for Web Caching	26
Web Cache Service Configuration with Clients and Content Engine on the Same Subnet	26
Configuring the Content Engine for Web Cache Service—Clients and Content Engine on the Same Subnet	27
Configuring the Router for Web Cache Service—Clients and Content Engine on the Same Subnet	28
Configuration Examples—Web Cache Service with Clients and Content Engine on the Same Subnet	28
Web Cache Service Configuration with Clients and Content Engine on Different Subnets	29
Configuring the Content Engine for Web Cache Service—Clients and Content Engine on Different Subnets	30
Configuring the Router for Web Cache Service—Clients and Content Engine on Different Subnets	30
Configuration Examples—Web Cache Service with Clients and Content Engine on Different Subnets	31
Configuring Direct Proxy Routing	32
Using a Proxy Autoconfiguration File Server	32
Creating PAC File Templates	33
Registering a PAC File Template and Assigning Coverage Zone Information	35
Configuring a Content Engine as a PAC File Server	38
Configuring Client Proxy Autoconfiguration Settings	39
Configuring Content Request Routing Using a Content Router or Routing Content Engine	41
Coverage Zones and Coverage Zone Files	42
Registering Coverage Zone Files	43
Choosing the Coverage Zone	43
Using the Default Coverage Zone	44
Registering and Applying a User-Defined Coverage Zone	44
Coverage Zone File Examples	47
Scenario 1: No Network Address Translation Firewall	47
Scenario 2: Content Engine Behind a NAT Firewall	48

Scenario 3: Multiple Content Engines Behind a NAT Firewall	49
Scenario 4: NAT Firewall with Multiple IP Addresses	51
Configuring a Content Engine as a Routing Content Engine	51
Configuring Dynamic Content Routing Using a Content Router	52
Enabling Dynamic Content Routing	52
Changing the Content Engine Metric Value for Dynamic Content Routing in the Coverage Zone File	53
Configuring Load-Based Content Routing Using a Content Router	55
Enabling Load-Based Routing on the Content Router	55
Configuring Content Engine Threshold Limits for Load-Based Routing	56
Content Router Request Overflow Routing	58
Defining Alternate Domains for Request Overflow Routing	59
Viewing the Overflow Domain List	59
Troubleshooting Content Router Configurations	59
Where to Go Next	60

CHAPTER 5**Configuring the ACNS Network for Content Distribution 1**

Preparing for Channel Distribution of Content	3
Creating and Modifying Locations	3
Configuring Channel Elements	6
Creating and Modifying Content Providers	6
Deleting Content Providers	7
Creating and Modifying Websites	8
Deleting Websites	10
Configuring Channels for Channel Distribution	10
Creating and Modifying Channels	10
Creating a Channel	11
Modifying a Channel	15
Designating the Root Content Engine	16
Adding and Removing Content Engines from Channels	16
Adding and Removing Device Groups from Channels	17
Priming a Live Channel Using the Channels Interface	18
Deleting Channels	19
Viewing Channels	19
Distributing Content Through Unicast Replication	20
Configuring Location Leader Preference and Forwarder Selection Probability	21
Configuring the Forwarder Lookup Level	22
Configuring the Content Metadata Poll Interval Multiplier	22
Configuring Unicast Content Distribution Through a NAT Firewall	23

Commands Related to Distribution Settings	24
Distributing Content Through Multicast Replication	25
About IP Multicasting	25
About Multicast Clouds	26
About Configuring Carousel Passes	27
About Multicast Logging Enhancements	28
About the NACK Interval Multiplier	28
About the File Resend or Stop Feature	29
About Configuring Multicast Forward Error Correction	29
About Configuring the Multicast Sender Delay Interval	29
Configuring PGM and File Transmission Parameters Using Multicast Expert Mode	30
About Configuring Channel Options for Content Replication	32
Enabling Content Engines for Multicasting	33
Configuring Multicast Cloud Properties	34
Modifying a Multicast Cloud	39
Deleting Multicast Clouds	40
Viewing Multicast Clouds	40
Assigning Content Engines to a Multicast Cloud	40
Assigning and Removing Multicast Clouds from Channels	43
Verifying the Multicast Configuration	45
Commands Related to Multicasting	45
Troubleshooting the Distribution Hierarchy	47
Where to Go Next	48

CHAPTER 6

Configuring the ACNS Network for Content Acquisition 1

Preparing to Acquire Pre-positioned Content	2
About the Acquirer	2
Set Up the Origin Server	3
Using HTTP and HTTPS	4
Using FTP	5
Create a Channel	5
Assign a Root Content Engine	6
Assigning a Backup Root Content Engine	6
Choosing a Content Acquisition Method for a Channel	6
Acquiring Pre-Positioned Content Using the Content Distribution Manager GUI	9
Adding Content Items to a Channel Using the Quick Crawl Utility	11
Adding a Crawl Task to a Channel	13
Configuring Content Acquisition Rules for a Crawl Task	15
Configuring Advanced Content Serving Settings	17

Modifying Multiple Content Items	22
Deleting Multiple Content Items	23
Acquiring Pre-Positioned Content Using an Externally Hosted Manifest File	23
Create a Manifest File	24
Specify Channel Acquisition and Distribution Properties	24
Distribution Priority	24
Item Priority	25
Channel Quota	25
Manifest File Update Interval	26
Configure Manifest File and Proxy Information for the Channel	26
Generate the Publishing URL	30
Assign a Playserver	33
Acquiring Content through a Proxy Server	35
Configuring Host and Proxy Server Settings Using the Content Distribution Manager GUI	35
Configuring an HTTP Proxy Server Using the CLI	37
Configuring a Proxy Server Using the Manifest File	38
About the <i>noProxy</i> Attribute	39
Configuring a Proxy Server and Port for the Manifest File	40
Acquiring Content Using Indexed Web Servers	40
Authentication Support	42
Acquiring Content that Requires NTLM Authentication	42
Proxy Authentication Support	43
Configuring Authentication for a Proxy to Fetch the Manifest File	43
Configuring Authentication for a Proxy Using the Manifest File	44
Configuring Authentication for an HTTP Proxy Using the Content Distribution Manager GUI	44
Configuring Authentication for a Proxy Using the CLI	45
Configuring Authentication for a WCCP Proxy Using the Content Distribution Manager GUI	46
Modifying WCCP Proxy Authentication Settings	47
Configuring Authentication for a WCCP Proxy Using the CLI	47
Verifying the Results	48
Troubleshooting Content Acquisition	48
Retry and Refresh Mechanisms	52
Updating Channel Content	53
Bandwidth Control	54
Configuring Acquisition and Distribution Default Bandwidth Settings	54
Displaying a Graphical Representation of the Acquisition and Distribution Bandwidth Settings	56
Configuring Acquisition and Distribution Bandwidth Settings for Scheduled Times	57
Scheduling Bandwidth for Streaming Acquisition	59
Error Codes	59

ACNS Unified Name Space Errors	59
Acquirer Error Codes	61
Where to Go Next	62

CHAPTER 7

Creating and Managing Programs 1

About Programs	1
Program Prerequisites	2
Playing Hinted and Nonhinted IP/TV Programs over an ACNS Network	3
Configuring Multicast IP Addresses for IP/TV Programs	3
About Live Events	3
Configuration Requirements for Managed Live Events	4
Live Stream Interruptions	4
About Root Content Engine Failover and Fallback Support for Cisco Streaming Engine Live Programs	5
About Root Content Engine Failover and Fallback Support for a Windows Media Live Program	6
About ACNS 5.5 Support for Windows Media Server-Side Playlists	6
About Scheduled Rebroadcasts	7
About TV-Out Programs	8
About Export Programs	8
Viewing Programs	8
Viewing and Modifying API Programs	10
Previewing a Program	11
Defining a Program	11
Creating a Live Program	13
Choosing a Live Channel	13
Adding a New Live Channel	14
Assigning Content Engines and Device Groups	15
Priming a Live Channel Using the Programs Interface	16
Configuring Live Stream Settings	18
Configuring Live Stream Settings for a Windows Media Live Program	18
Configuring Live Stream Settings for a Cisco Streaming Engine Live Program	20
About Unicast Publishing URLs	22
About Multicast Publishing URLs	23
About NSC Files for Multicasting Windows Media Server-Side Playlists	23
Scheduling Live Programs	23
Creating a Scheduled Rebroadcast Program	25
Choosing the Media Files	25
Assigning Content Engines	28
Assigning Device Groups	28

Configuring Streaming Settings	29
Scheduling Rebroadcast Programs	30
Creating a TV-Out Program	31
Enabling the TV-Out Feature	31
Supported MPEG File Specifications	32
Selecting Media Files and Overlay Images for Playlists	33
Adding Media Files to a Playlist	33
Adding an Overlay Image to a Playlist	34
Editing a Overlay Image	35
Assigning an Overlay Image to a Media File	36
Scheduling TV-Out Programs	36
Assigning Content Engines to Playlists	37
Assigning Device Groups to Playlists	38
Viewing the TV-Out Playback Guide	39
Creating an Export Program	40
Selecting Export Channels for Set Top Box Playback	41
Assigning Content Engines to an Export Channel	43
Assigning Device Groups to an Export Channel	44
Viewing Export Playlist Settings	45
Copying a Program	45
Managing Multicast Addresses for Programs	46
Configuring Program Multicast Addresses	46

PART 3**Configuring ACNS Software Applications****CHAPTER 8****Configuring Caching Services 1**

Content Engine Proxy Caching Overview	1
Network Protocols and Caching	2
Configuring HTTP Settings	2
Configuring HTTP Connection Settings	2
Configuring HTTP Cache Settings	6
Configuring HTTP Cache Freshness Settings	10
Configuring Authenticated HTTP Cache Settings	13
Specifying a Reauthentication Interval	19
Configuring Advanced HTTP Cache Settings	19
Adding or Modifying Additional HTTP Request Methods for the Content Engine	23
Configuring HTTP Cache Vary User Agent	25
Configuring HTTP Destination Port Restrictions	26
Configuring HTTPS Settings	27

Configuring HTTPS Certificates	27
Filtering HTTPS Certificates for the Content Engine	29
Configuring HTTPS Certificate Groups	31
Filtering HTTPS Certificate Groups for the Content Engine	32
Configuring HTTPS Keys	34
Filtering HTTPS Keys for the Content Engine	35
Configuring HTTPS Proxy Settings	37
Configuring HTTPS Servers	42
Filtering HTTPS Servers for the Content Engine	44
Transparent HTTPS Caching Using SSL	47
Configuring HTTP and HTTPS Outgoing Proxy Exclusion Settings	48
Configuring FTP Settings	50
FTP-over-HTTP Caching	51
Configuring FTP-Over-HTTP Connection Settings	52
Configuring FTP-Over-HTTP Cache Freshness	55
Native FTP Caching	57
Configuring the Content Engine for Transparent Native FTP Caching	58
Configuring Native FTP Connections	58
Configuring the Client Side of Nontransparent Native FTP Caching	60
Configuring Native FTP Cache Freshness	61
Enabling Inetd FTP Service	61
Configuring Native FTP Custom Messages	62
About Writing FTP Custom Welcome Messages	62
Configuring FTP Message Download Settings	62
Viewing the Content Engine FTP Message Configurations	63
Resetting a Native FTP Custom Message	63
Configuring Native FTP Custom Message Upload Settings	64
Configuring TFTP Settings	65
Configuring TFTP General Settings	65
Configuring TFTP Directory Settings	69
Configuring DNS Caching Name Service for WCCP Transparent Caching	71
Configuring DNS Caching Server Settings for the Content Engine	71
Cacheable Resource Records in the Content Engine DNS Cache	74
Configuring DNS Server Bindings for the Content Engine	75
Configuring DNS Address Record Mappings for the Content Engine	76
Configuring DNS Canonical Name Record Mappings for the Content Engine	77
Configuring the DNS Server for HTTP Proxy Caching	78
Configuring ICP Settings	80
ICP Client Settings	80

ICP Remote Client Settings	82
ICP Server Settings	83
ICP Remote Server Settings	84

CHAPTER 9**Configuring Streaming Media Services 1**

Content Engine Streaming Overview	1
Live Streaming for the Non-Multicast Enabled Network	2
Unicast Live Stream Splitting	2
Hybrid Unicast-to-Multicast Live Streaming	3
Video On Demand Streaming	4
Demand-Cached Video	5
Pre-Positioned Video	6
Licensing and Enabling Streaming Servers	6
Using the Real-Time Streaming Protocol Gateway	7
Configuring the RTSP Gateway	8
Enabling RealProxy	10
Enabling RealSubscriber	12
Enabling the Cisco Streaming Engine	14
Using Windows Media Technologies	15
Enabling Windows Media Services	15
Configuring Windows Media Proxy Settings for the Content Engine	18
About Fast Start	24
About Fast Cache	25
Configuring the Windows Media Player to Use the Content Engine as a Proxy	26
Configuring Windows Media Live Splitting for Unicast and Multicast Transmissions	26
Live Stream Splitting Scenarios	27
Configuring Unicast-In Multicast-Out	28
Configuring Multicast-In Multicast-Out	31
Configuring Multicast-In Unicast-Out	34
Configuring Unicast-In Unicast-Out	36
Configuring Windows Media Multicast Station Schedules	38
Configuring Weighted Load Balancing for Live Stream Splitting	38
Related CLI Commands	40
Configuring Bandwidth Settings for Content Services	40
Configuring Content Services Default and Maximum Bandwidth Settings	41
Configuring Scheduled Bandwidth Settings	42
Displaying a Graphical Representation of the Content Service Bandwidth Settings	45
Configuring the Windows Media Incoming Bandwidth Bypass List	47

CHAPTER 10

Managing Caching Performance and Differentiated Services 1

- Configuring TCP Settings for Increased Cache Performance 1
- Configuring TCP Memory Limits 6
- About Explicit Congestion Notification 7
- About Congestion Windows 7
- About the Retransmit Time Multiplier 7
- About TCP Slow Start 8
- About TCP-Over-Satellite Extensions 8
- Configuring IP Differentiated Services 9
- Enabling the MTU Discovery Utility 10

CHAPTER 11

Viewing Content Replication Status 1

- Viewing the System-Wide Replication Status by Channel 1
- Viewing the Replication Status for a Specific Channel 5
- Viewing the Detailed Replication Status of Content Items for a Specific Channel 9
- Viewing the Detailed Replication Status of Content Items for a Channel by Device 10
- Viewing the Detailed Replication Status of Content Items Across All Content Engines in a Channel 12
- Viewing the System-Wide Replication Status by Content Engine 13
- Viewing the Replication Status for a Specific Content Engine 14
- Viewing the Detailed Replication Status of Content Items for a Selected Channel by Content Engine 15
- Using CLI Commands for Replication Status 17

PART 4

Managing ACNS Network Devices

CHAPTER 12

Configuring Login Authentication, Configuration Authorization, and Accounting 1

- Understanding Authentication, Authorization, and Accounting 1
- Understanding Login Authentication and Authorization 2
- Understanding User Accounts and Privilege Profiles 2
 - Accessing a Device for the First Time 2
- Understanding AAA Database Management 3
 - Default Login Authentication and Authorization Configuration 3
 - Understanding Login Authentication Failover 4
 - Understanding Local Login Authentication and Authorization 5
 - Understanding RADIUS Authentication and Authorization 6
 - Understanding TACACS+ Authentication and Authorization 6
- Configuring Devices for Login Authentication and Authorization 7
 - Configuring Remote Authentication Server Settings for the Device 7
 - Configuring RADIUS Server Settings 7

Configuring TACACS+ Server Settings	9
Configuring Login Authentication and Configuration Authorization Schemes for the Device	12
Changing the Admin Account Password for a Device	13
Configuring User Accounts for Centrally Managed Devices	14
Creating User Accounts for Centrally Managed Devices	14
Modifying and Deleting User Accounts for Centrally Managed Devices	14
Managing User Accounts and Privilege Profiles for the Content Distribution Manager	15
Creating and Managing User Accounts for the Content Distribution Manager	16
Creating New User Accounts Using the Content Distribution Manager GUI	17
Modifying and Deleting User Accounts	19
Changing Your CLI User Password	20
Viewing User Accounts	20
Creating and Managing Roles	20
Modifying and Deleting Roles	21
Viewing Role Settings	22
Assigning Roles to User Accounts	22
Configuring and Managing Domains	23
Adding an Entity to a Domain	24
Modifying and Deleting Domains	25
Viewing Domains	26
Assigning Domains to User Accounts	26
Configuring AAA Accounting	27
Viewing Audit Trail Logs	30

CHAPTER 13

Working with Device Configurations 1

Working with Content Distribution Managers	1
Viewing the Content Distribution Manager Configuration	2
Configuring a Warm Standby Content Distribution Manager	2
Switching a Content Distribution Manager from Warm Standby to Primary	3
Deleting a Warm Standby Content Distribution Manager	5
Working with Device Locations	5
Creating and Modifying Locations	5
Deleting Locations	8
Viewing the Location Tree	8
Working with Content Engines	8
Modifying Content Engine Properties	9
Rebooting a Content Engine	12
Deleting a Content Engine	12
Viewing Content Engine Channel Assignments	14

Commands Related to Content Acquisition and Distribution	14
Configuring the Content Engine Common Internet File System Server	17
Configuring CIFS Server Settings	17
Configuring CIFS Server Website Access Control Settings	19
Customizing the CIFS Share Name	21
Working with Content Routers	22
Modifying Content Router Properties	22
Deleting a Content Router	24
Configuring Services on Content Routers	25
Working with Device Groups	26
Creating a Device Group	26
Assigning Content Engines to the Device Group	29
Enabling Device Group Overlap	30
Deleting a Device Group	30
Viewing Device Groups	30
Working with Baseline Groups	31
Configuring a Baseline Group for a Service	31
Assigning Content Engines to the Baseline Group	33
Configuring the Service Settings for the Baseline Group	34
Forcing the Baseline Group Service Settings on All Content Engines in the Baseline Group	34
Switching the Baseline Group for a Service	35
Viewing Configured Service Settings for a Device Group	35
Local Device Configuration Versus Central Management	35
Repurposing an ACNS Network Device	36

CHAPTER 14

Configuring Network Interfaces 1

Configuring Multiple Network Interfaces and Standby Interfaces	1
Configuring a Standby Interface Using the Content Distribution Manager GUI	3
Configuring the Interface Priority Setting	4
Configuring Multiple IP Addresses on a Single Interface	6
Modifying Fast Ethernet or Gigabit Ethernet Interface Settings	7
Configuring the Fibre Channel Interface	9
Modifying Fibre Channel Interface Settings for the Content Engine	10
Configuring EtherChannel	11
Configuring Port Channel Settings Using the Content Distribution Manager GUI	11
Configuring a Load-Balancing Method for Interfaces within a Channel Group Using the Content Distribution Manager GUI	13
Configuring Interfaces for DHCP	13

Configuring IP Addresses for Network Address Translation	14
Configuring Static IP Routes	15

PART 5**Managing Website Access****CHAPTER 15****Configuring Request Authentication and Authorization 1**

Understanding Request Authentication and Authorization for Content Engines	1
About Proxied Authentication and Pass-Through Authentication Modes	2
Content Delivery Protocols Supported	3
Authentication Mechanisms Supported	3
About Authenticating Windows Media Requests	6
About Authenticating and Authorizing HTTP Requests	6
Understanding the Content Engine Authentication Cache	7
Using ACLs for Group-Based Authorization	12
Using Rules for Group-Based Authorization	12
Using Microsoft Active Directory for Group Searches	13
Using Digest Authentication	14
About Authenticating Native FTP Requests	15
Considerations When Configuring Native FTP Proxied Authentication	16
Configuring Request Authentication on Centrally Managed Content Engines	17
Configuring Authentication Server Settings	17
Configuring LDAP Server Settings	17
Configuring NTLM Server Settings	23
Configuring the NTLM Server LDAP Memory Cache Settings	28
Support for No Default NTLM Domain	30
About NTLM Load Balancing for HTTP Request Authentication	30
About NTLM Failover for HTTP Request Authentication	31
Configuring NTLM Allowed Domains for HTTP Request Authentication	31
Configuring RADIUS Server Settings	32
Configuring TACACS+ Server Settings	34
Setting the Authentication Scheme for Request Authentication	36

CHAPTER 16**Configuring Request Processing Services 1**

Configuring ICAP	1
Overview of ICAP	1
About ICAP Services	2
About ICAP Services and Vectoring Points	2
About ICAP Performance	3
ICAP Vendors Supported	3

Maximum File Size Supported	4
Enhancements to ICAP Support in ACNS 5.5 Software	4
Configuring ICAP Settings	4
Configuring ICAP Services	7
Configuring an ICAP Server for ICAP Service	10
About ICAP Processing and Java Applets	12
Displaying Information About an ICAP Configuration	13
Displaying Statistics for ICAP Services	13
Configuring Service Rules	13
Enabling Rule Settings	14
Configuring Service Rules	14
Configuring the Rules Template Using CLI Commands	23
Configuring a Pattern List	23
Adding a Pattern to an Existing Pattern List	24
Associating an Action with an Existing Pattern List	24
Verifying an Action Performed on a Pattern List	25
Configuring URL Filtering	25
Configuring URL Filter Settings Using the Content Distribution Manager GUI	26
Configuring URL Filters Using the CLI	30
Configuring Custom Blocking Messages Using the CLI	31
Creating a Text File URL List for URL Filtering	32
Configuring URL Filtering with the N2H2 Server	33
N2H2 Features Supported	34
Enabling N2H2 Filtering Using the Content Distribution Manager GUI	35
Enabling N2H2 Filtering Using the CLI	36
N2H2 Configuration and Restrictions	39
Configuring URL Filtering with Websense Enterprise Software	39
Using a Websense Enterprise Server	39
Configuring Websense URL Filtering with External Websense Servers	40
Using the Integrated Websense Server	41
Configuring Ports for the Integrated Websense Server	42
About Websense Server Failover	42
Enabling Websense Filtering and Configuring a Failover Scheme	43
Configuring Websense Server Failover and URL Filtering Using the CLI	45
About Websense Server Components	48
Installing Websense Server Components for the Content Engine	50
Specifying RADIUS Agent Settings for the Content Engine Using the CLI	52
Specifying eDirectory Agent Settings for the Content Engine Using the CLI	52
Websense Status and Statistics Commands	53

Websense Configuration and Restrictions	55
Configuring URL Filtering with SmartFilter Software	55
About the SmartFilter Control List	56
SmartFilter Software and the Action No-Auth Command Rule Interaction	56
Configuring Content Engines to Bypass URL Filtering for Specific HTTP and HTTPS Requests	57
Execution Order of Rule Actions	58
Configuring the Content Engine GUI for Secure or Nonsecure Access	59
Setting Up the Content Engine to Interoperate with Third-Party Policy Servers	60
About URL Signature Keys	62
Setting QoS Policy Server Settings for Content Engines	62
Creating New URL Signature Keys for Content Engines	63
Creating Rule Patterns and Rule Actions for Directing Requests to a Policy Server	64

CHAPTER 17**Creating and Managing IP Access Control Lists 1**

Introducing IP ACLs	1
Creating or Modifying an IP ACL	2
Creating a New IP ACL	3
Adding Conditions to an IP ACL	3
Modifying, Deleting, or Reordering a Condition	11
Associating an IP ACL with an Application	12
Applying an IP ACL to an Interface	13
Deleting an IP ACL	14

CHAPTER 18**Configuring Group Authorization Using an Access Control List 1**

Configuring an Access Control List for Group Authorization	1
Enabling the Access Control List	2
Using the CLI to Configure and Enable an ACL for Group Authorization	2
Modifying an Access Control List	3
Changing the Currently Applied Settings	4
Repositioning Settings in the Access Control List	5
Viewing Access Control List Settings	5

CHAPTER 19**Using the Transaction Logs 1**

Understanding Transaction Log Formats	1
Squid-Style Transaction Logging	1
Extended Squid Log Format	2
Apache-Style Transaction Logging	2
Custom Format Transaction Logging	2

W3C Customizable Logging Format	6
Transaction Logging and NTLM Authentication	7
Usage Guidelines for Log Files	8
Understanding Working Logs	8
Archive Working Log	8
Sanitizing Transaction Logs	8
Exporting Log Files	9
Exporting Transaction Logs to External FTP Servers	9
Restarting Export After Receiving a Permanent Error from the External FTP Server	9
Exporting Transaction Logs to External SFTP Servers	10
Secure File Transfer Protocol Access for Nonadministrative Users	10
Enabling Transaction Logging with the Content Distribution Manager GUI	10
Using WMT Transaction Logging	15
Log Formats Accepted by Windows Media Services	16
Configuring WMT Transaction Logging	17
Using Real-Time Transaction Logging	19
Configuring Real-Time Transaction Logging	20
About Specifying the Transaction Log Entry Type when Logging to a Remote Syslog Host	21

PART 6

System Administration and Configuration

CHAPTER 20

Configuring Platform and System Settings 1

Configuring Remote Login Access Control Settings	1
Enabling the Telnet Service on the Content Engine	1
Configuring Secure Shell Settings	2
About the Secure Shell Host Key	4
About Secure Shell Daemon Password Guesses	4
Configuring Message of the Day Settings for the Content Engine	5
Configuring Exec Timeout Settings	6
Configuring Line Console Carrier Detection	6
Configuring Network Services	7
Configuring CDP Settings	7
Enabling RCP Services on the Content Engine	8
Configuring the Content Router for Interoperability with a Layer 4 Switch	9
Configuring Date and Time Settings	11
Configuring NTP Settings	11
Configuring Device Clock and Time Zone Settings	12
Configuring System-Wide Distribution QoS Settings Based on Channel Priority	16

Setting DSCP Values for QoS Packets	18
Modifying System Default Properties	18
Configuring Faster Detection of Offline Content Engines	21
About Faster Detection of Offline Content Engines	22
Printing ACNS Network Data	23
Creating Proxy Error Messages	23
Configuring Proxy Error Message Download Settings for the Content Engine	24
Configuring Proxy Error Message Upload Settings for the Content Engine	26
Enabling Offline Operation of Network Devices	27

CHAPTER 21**Monitoring and Troubleshooting the ACNS Network 1**

Monitoring System Status	1
Device Alarms	2
Troubleshooting Devices Using the System Status Bar	3
Using the Content Distribution Manager GUI show Command Tool	4
Content Alarms	4
Troubleshooting Content Replication Issues Using the System Status Bar	5
Monitoring System Events Using the System Message Log	5
Configuring System Event Logging Using the Content Distribution Manager GUI	5
About Multiple Hosts for System Logging	7
Configuring System Event Logging Using CLI Commands	8
Viewing the System Message Log	9
Alarm Overload Detection	10
Configuring Alarm Overload Detection Settings for the Content Engine	11
Monitoring Device Status	12
Monitoring Device Performance	12
Viewing Content Engine Statistics	14
Viewing Device Group Statistics	15
Viewing Routing Statistics	16
Viewing Streaming Statistics	16
Incoming Statistics	17
Outgoing Statistics	17
Multicast Statistics	18
Bandwidth Savings Statistics	18
Revised Command Syntax	18
Clearing the Statistics	18
Configuring Report Options	19
Bytes Served Report	19

Bandwidth Efficiency Gain Report	19
Streaming Sessions Report	19
CPU Utilization Report	20
Monitoring Specified HTTP URLs	20
Enabling Kernel Debugger	23

CHAPTER 22

Monitoring with SNMP 1

Supported SNMP Versions	1
Configuring SNMP Traps	2
Configuring SNMP Traps Using the CLI	5
Configuring SNMP Community Settings	5
Configuring SNMP Group Settings	6
Configuring SNMP User Settings	8
Configuring SNMPv2 View Settings	9
Configuring SNMP Host Settings	10
Configuring SNMP Asset Tag Settings	11
Supported MIBs	12
Key SNMP CLI Commands	13

PART 7

Appendices

APPENDIX A

Creating Manifest Files 1

Manifest File User Guidelines	1
Overview	2
Quick Start	2
Writing XML Tags	2
Important Manifest Tags	3
Writing a Single-Item HTTP Manifest File	4
Writing a Single-Item FTP Manifest File	4
Writing an FTP Crawler Manifest File	5
Writing an HTTPS Crawler Manifest File	5
Validating Manifest Files	5
Migrating from ACNS 4.x Software to ACNS 5.x Software	5
Working with Manifest Files	6
Specifying a Single Content Item	6
Specifying a Crawl Job	7
Writing Common Regular Expressions	10
Scheduling Content Acquisition	10

Specifying Shared Attributes	10
Specifying a Crawler Filter	11
Specifying Content Priority	13
Generating a Playserver List	13
Specifying Attributes for Content Serving	16
Specifying Metadata for Content Serving	18
Specifying Time Values in the Manifest File	18
Refreshing and Removing Content	19
Specifying Live Content	20
More Sample Manifest Files	21
Sample 1	22
Sample 2	24
Sample 3	25
Sample 4	27
Sample 5	27
Downloading the Sample Manifest Files	28
Manifest Validator Utility	28
Running the Manifest Validator Utility	28
Understanding Manifest File Validator Output	31
Correcting Manifest File Syntax	32
Manifest File Reference	33
Manifest File Structure and Syntax	33
CdnManifest	37
playServerTable	38
playServer	39
options	40
server	42
host	42
proxyServer	44
item	44
crawler	54
item-group	58
matchRule	61
match	61
contains	63
http-meta-data	63
Configuring the Freshness of Pre-Positioned Content	64
XML Schema	65
Manifest XML Schema	65
PlayServerTable XML Schema	69

Default PlayServerTable Schema	70
Manifest File Time Zone Tables	71
Manifest File Automated Scripts	78
Installing Perl on Your Workstation	79
Obtaining the Perl Scripts	79
Listing Website Content Using the Spider Script	80
Selecting Live and Pre-Positioned Content Using the Manifest Script	82
Creating a Rules File for the Spider and Manifest Scripts	85
Using the ACNS Manifest Generator Tool	86
Supported Platforms	86
Installing the ACNS Manifest Generator	87
Understanding the ACNS Manifest Generator User Interface	87
Creating a Basic Manifest File Using the ACNS Manifest Generator	88
Elements and Tags Not Supported in the ACNS Manifest Generator	97

APPENDIX B

IP Multicast Addressing 1

Unusable IP Multicast Addresses	1
Insecure Services	4
Copying Files Between Servers and Clients	4
Limited Scope Addresses	4
Source-Specific Multicast Addresses	5
GLOP Addresses	5
Layer 2 Multicast Addresses	5
Ethernet MAC Address Mapping	6

APPENDIX C

Mapping ACNS Software CLI Commands to the Content Distribution Manager GUI 1

CLI Commands Supported in the Content Distribution Manager GUI	1
CLI Commands Not Supported in the Content Distribution Manager GUI	27

INDEX



Preface

This preface describes who should read the *Cisco ACNS Software Configuration Guide for Centrally Managed Deployments, Release 5.5*, how it is organized, and its document conventions.

This preface contains the following sections:

- [Document Objectives, page 25](#)
- [Audience, page 25](#)
- [Organization, page 26](#)
- [Conventions, page 27](#)
- [Related Documentation, page 28](#)
- [Obtaining Documentation, page 28](#)
- [Documentation Feedback, page 29](#)
- [Cisco Product Security Overview, page 29](#)
- [Obtaining Technical Assistance, page 31](#)
- [Obtaining Additional Publications and Information, page 32](#)

Document Objectives

This configuration guide describes how to use Cisco Application and Content Networking System (ACNS) software for setting up and managing content delivery and content caching services in a centrally managed environment.

Audience

This guide is intended for network administrators and content managers. The person responsible for managing the Content Distribution Manager, Content Routers, and Content Engines should be experienced with the following:

- IP network configuration
- Domain Name Server (DNS) configuration

The content manager should be familiar with the responsibilities of a webmaster and should have a thorough knowledge of Extensible Markup Language (XML).

Organization

This document is organized in the following manner:

Chapter	Title	Description
Chapter 1	Understanding the ACNS 5.5 Network	Describes the basic concepts of the Cisco ACNS 5.x network, the function and position of the different devices, topology considerations, and how content is retrieved, stored, and distributed to the end user.
Chapter 2	Getting Started	Describes how to initially configure your ACNS devices and activate your devices on the ACNS network.
Chapter 3	Using the Content Distribution Manager GUI	Describes how to navigate within the Content Distribution Manager GUI.
Chapter 4	Setting Up Content Request Routing in the ACNS Network	Describes the three request routing methods supported in ACNS 5.5 software: WCCP transparent request interception, direct proxy routing, and Content Router routing.
Chapter 5	Configuring the ACNS Network for Content Distribution	Provides information about configuring ACNS network elements for channel distribution and multicasting.
Chapter 6	Configuring the ACNS Network for Content Acquisition	Outlines the tasks necessary for acquiring pre-positioned content in the ACNS network and contains information on acquisition and distribution bandwidth control and acquisition authentication support.
Chapter 7	Creating and Managing Programs	Describes how to create and manage live, WMT rebroadcast, TV-out, and export programs (for export to set top boxes).
Chapter 8	Configuring Caching Services	Describes how to configure caching services on Content Engines.
Chapter 9	Configuring Streaming Media Services	Describes how to configure streaming media services on Content Engines.
Chapter 10	Managing Caching Performance and Differentiated Services	Describes how to configure TCP parameters and how to configure Type of Service (ToS) and differentiated services code point (DSCP) settings.
Chapter 11	Viewing Content Replication Status	Explains how to monitor content replication status from the Content Distribution Manager GUI.
Chapter 12	Configuring Login Authentication, Configuration Authorization, and Accounting	Explains how to configure login authentication settings and external authentication servers for the Content Engine.
Chapter 13	Working with Device Configurations	Explains how to modify ACNS device properties.
Chapter 14	Configuring Network Interfaces	Describes how to configure additional interfaces using options for redundancy, load balancing, performance optimization, and so forth.

Chapter	Title	Description
Chapter 15	Configuring Request Authentication and Authorization	Explains how to configure request authentication and authorization for centrally-managed ACNS networks.
Chapter 16	Configuring Request Processing Services	Describes how to configure content filtering services for your device.
Chapter 17	Creating and Managing IP Access Control Lists	Describes the procedure for applying IP ACLs to devices and management services using the Content Distribution Manager GUI.
Chapter 18	Configuring Group Authorization Using an Access Control List	Describes how to configure group authorization using an access control list.
Chapter 19	Using the Transaction Logs	Describes how to use the transaction logs.
Chapter 20	Configuring Platform and System Settings	Explains how to configure platform and system settings.
Chapter 21	Monitoring and Troubleshooting the ACNS Network	Provides information on monitoring and troubleshooting devices and content replication. Includes using the system message logs, transaction logs, and SNMP.
Chapter 22	Monitoring with SNMP	Explains how to configure SNMP traps, recipients, community strings and group associations, user security model groups, and user access permissions.
Appendix A	Creating Manifest Files	Describes how to create and import manifest files.
Appendix B	IP Multicast Addressing	Discusses IP multicasting and assigning IP multicast addresses.
Appendix C	Advanced Routing Configurations	Explains how to pre-position and serve content acquired from multiple origin servers using either a WCCP-enabled router or a Content Router.
Appendix C	Mapping ACNS Software CLI Commands to the Content Distribution Manager GUI	Maps ACNS software CLI commands to the corresponding Content Distribution Manager GUI window.

Conventions

This document uses the following conventions:

Convention	Description
boldface font	Commands and keywords are in boldface .
<i>italic</i> font	Arguments for which you supply values are in <i>italics</i> .
[]	Elements in square brackets are optional.
{x y z}	Alternative keywords are grouped in braces and separated by vertical bars.
[x y z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.

Convention	Description
string	An unquoted set of characters. Do not use quotation marks around the string, or the string will include the quotation marks.
screen font	Terminal sessions and information the system displays are in <code>screen font</code> .
boldface screen font	Information you must enter is in boldface screen font .
<i>italic screen font</i>	Arguments for which you supply values are in <i>italic screen font</i> .
^	The symbol ^ represents the key labeled Control—for example, the key combination ^D in a screen display means hold down the Control key while you press the D key.
< >	Nonprinting characters, such as passwords, are in angle brackets.

Notes use the following conventions:



Note

Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the publication.

Related Documentation

The following documentation provides additional information about Cisco ACNS software:

- *Documentation Guide for Cisco ACNS Software, Release 5.5*
- *Release Notes for Cisco ACNS Software, Release 5.5*
- *Cisco ACNS Software Command Reference, Release 5.5*
- *Cisco ACNS Software Configuration Guide for Locally Managed Deployments, Release 5.5*
- *Cisco ACNS Software Upgrade and Maintenance Guide, Release 5.x*

Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. Cisco also provides several ways to obtain technical assistance and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Cisco.com

You can access the most current Cisco documentation at this URL:

<http://www.cisco.com/techsupport>

You can access the Cisco website at this URL:

<http://www.cisco.com>

You can access international Cisco websites at this URL:

http://www.cisco.com/public/countries_languages.shtml

Product Documentation DVD

The Product Documentation DVD is a comprehensive library of technical product documentation on a portable medium. The DVD enables you to access multiple versions of installation, configuration, and command guides for Cisco hardware and software products. With the DVD, you have access to the same HTML documentation that is found on the Cisco website without being connected to the Internet. Certain products also have .PDF versions of the documentation available.

The Product Documentation DVD is available as a single unit or as a subscription. Registered Cisco.com users (Cisco direct customers) can order a Product Documentation DVD (product number DOC-DOCDVD= or DOC-DOCDVD=SUB) from Cisco Marketplace at this URL:

<http://www.cisco.com/go/marketplace/>

Ordering Documentation

Registered Cisco.com users may order Cisco documentation at the Product Documentation Store in the Cisco Marketplace at this URL:

<http://www.cisco.com/go/marketplace/>

Nonregistered Cisco.com users can order technical documentation from 8:00 a.m. to 5:00 p.m. (0800 to 1700) PDT by calling 1 866 463-3487 in the United States and Canada, or elsewhere by calling 011 408 519-5055. You can also order documentation by e-mail at tech-doc-store-mkpl@external.cisco.com or by fax at 1 408 519-5001 in the United States and Canada, or elsewhere at 011 408 519-5001.

Documentation Feedback

You can rate and provide feedback about Cisco technical documents by completing the online feedback form that appears with the technical documents on Cisco.com.

You can submit comments about Cisco documentation by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Customer Document Ordering
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Cisco Product Security Overview

Cisco provides a free online Security Vulnerability Policy portal at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

From this site, you will find information about how to:

- Report security vulnerabilities in Cisco products.
- Obtain assistance with security incidents that involve Cisco products.
- Register to receive security information from Cisco.

A current list of security advisories, security notices, and security responses for Cisco products is available at this URL:

<http://www.cisco.com/go/psirt>

To see security advisories, security notices, and security responses as they are updated in real time, you can subscribe to the Product Security Incident Response Team Really Simple Syndication (PSIRT RSS) feed. Information about how to subscribe to the PSIRT RSS feed is found at this URL:

http://www.cisco.com/en/US/products/products_psirt_rss_feed.html

Reporting Security Problems in Cisco Products

Cisco is committed to delivering secure products. We test our products internally before we release them, and we strive to correct all vulnerabilities quickly. If you think that you have identified a vulnerability in a Cisco product, contact PSIRT:

- For Emergencies only—security-alert@cisco.com

An emergency is either a condition in which a system is under active attack or a condition for which a severe and urgent security vulnerability should be reported. All other conditions are considered nonemergencies.

- For Nonemergencies—psirt@cisco.com

In an emergency, you can also reach PSIRT by telephone:

- 1 877 228-7302
- 1 408 525-6532



Tip

We encourage you to use Pretty Good Privacy (PGP) or a compatible product (for example, GnuPG) to encrypt any sensitive information that you send to Cisco. PSIRT can work with information that has been encrypted with PGP versions 2.x through 9.x.

Never use a revoked or an expired encryption key. The correct public key to use in your correspondence with PSIRT is the one linked in the Contact Summary section of the Security Vulnerability Policy page at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

The link on this page has the current PGP key ID in use.

If you do not have or use PGP, contact PSIRT at the aforementioned e-mail addresses or phone numbers before sending any sensitive material to find other means of encrypting the data.

Obtaining Technical Assistance

Cisco Technical Support provides 24-hour-a-day award-winning technical assistance. The Cisco Technical Support & Documentation website on Cisco.com features extensive online support resources. In addition, if you have a valid Cisco service contract, Cisco Technical Assistance Center (TAC) engineers provide telephone support. If you do not have a valid Cisco service contract, contact your reseller.

Cisco Technical Support & Documentation Website

The Cisco Technical Support & Documentation website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The website is available 24 hours a day, at this URL:

<http://www.cisco.com/techsupport>

Access to all tools on the Cisco Technical Support & Documentation website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a user ID or password, you can register at this URL:

<http://tools.cisco.com/RPF/register/register.do>



Note

Use the Cisco Product Identification (CPI) tool to locate your product serial number before submitting a web or phone request for service. You can access the CPI tool from the Cisco Technical Support & Documentation website by clicking the **Tools & Resources** link under Documentation & Tools. Choose **Cisco Product Identification Tool** from the Alphabetical Index drop-down list, or click the **Cisco Product Identification Tool** link under Alerts & RMAs. The CPI tool offers three search options: by product ID or model name; by tree view; or for certain products, by copying and pasting **show** command output. Search results show an illustration of your product with the serial number label location highlighted. Locate the serial number label on your product and record the information before placing a service call.

Submitting a Service Request

Using the online TAC Service Request Tool is the fastest way to open S3 and S4 service requests. (S3 and S4 service requests are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Service Request Tool provides recommended solutions. If your issue is not resolved using the recommended resources, your service request is assigned to a Cisco engineer. The TAC Service Request Tool is located at this URL:

<http://www.cisco.com/techsupport/servicerequest>

For S1 or S2 service requests, or if you do not have Internet access, contact the Cisco TAC by telephone. (S1 or S2 service requests are those in which your production network is down or severely degraded.) Cisco engineers are assigned immediately to S1 and S2 service requests to help keep your business operations running smoothly.

To open a service request by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411 (Australia: 1 800 805 227)

EMEA: +32 2 704 55 55

USA: 1 800 553-2447

For a complete list of Cisco TAC contacts, go to this URL:

<http://www.cisco.com/techsupport/contacts>

Definitions of Service Request Severity

To ensure that all service requests are reported in a standard format, Cisco has established severity definitions.

Severity 1 (S1)—An existing network is down, or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Severity 2 (S2)—Operation of an existing network is severely degraded, or significant aspects of your business operations are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Severity 3 (S3)—Operational performance of the network is impaired, while most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Severity 4 (S4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- The *Cisco Product Quick Reference Guide* is a handy, compact reference tool that includes brief product overviews, key features, sample part numbers, and abbreviated technical specifications for many Cisco products that are sold through channel partners. It is updated twice a year and includes the latest Cisco offerings. To order and find out more about the Cisco Product Quick Reference Guide, go to this URL:

<http://www.cisco.com/go/guide>

- Cisco Marketplace provides a variety of Cisco books, reference guides, documentation, and logo merchandise. Visit Cisco Marketplace, the company store, at this URL:

<http://www.cisco.com/go/marketplace/>

- *Cisco Press* publishes a wide range of general networking, training and certification titles. Both new and experienced users will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press at this URL:

<http://www.ciscopress.com>

- *Packet* magazine is the Cisco Systems technical user magazine for maximizing Internet and networking investments. Each quarter, Packet delivers coverage of the latest industry trends, technology breakthroughs, and Cisco products and solutions, as well as network deployment and troubleshooting tips, configuration examples, customer case studies, certification and training information, and links to scores of in-depth online resources. You can access Packet magazine at this URL:

<http://www.cisco.com/packet>

- *iQ Magazine* is the quarterly publication from Cisco Systems designed to help growing companies learn how they can use technology to increase revenue, streamline their business, and expand services. The publication identifies the challenges facing these companies and the technologies to help solve them, using real-world case studies and business strategies to help readers make sound technology investment decisions. You can access iQ Magazine at this URL:

<http://www.cisco.com/go/iqmagazine>

or view the digital edition at this URL:

<http://ciscoinq.texterity.com/ciscoinq/sample/>

- *Internet Protocol Journal* is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:

<http://www.cisco.com/ipj>

- Networking products offered by Cisco Systems, as well as customer support services, can be obtained at this URL:

<http://www.cisco.com/en/US/products/index.html>

- Networking Professionals Connection is an interactive website for networking professionals to share questions, suggestions, and information about networking products and technologies with Cisco experts and other networking professionals. Join a discussion at this URL:

<http://www.cisco.com/discuss/networking>

- World-class networking training is available from Cisco. You can view current offerings at this URL:

<http://www.cisco.com/en/US/learning/index.html>



PART 1

Introducing ACNS Software for Centrally Managed Deployments





Understanding the ACNS 5.5 Network

The Cisco Application and Content Networking System (ACNS) software solution expedites the delivery of strategic applications from the corporate data center to branch offices and local points of sale. Cisco ACNS software combines the technologies of caching, pre-positioning, and live and on-demand streaming to accelerate delivery of web applications, object files, live events, and video, enabling organizations of all sizes to reduce costs, drive up productivity, and increase revenues.

Cisco ACNS software runs on the Cisco Content Engine hardware platform. Content Engines in an ACNS network save bandwidth and protect networks by acting as forward proxies that authenticate, filter, and cache all traffic between an organization's internal network and the Internet. Content Engines also perform reverse proxy caching to reduce the load on expensive back-end web servers.



Note

In the ACNS 5.3.3 software release, support for the Wide Area Application Engine (WAE) was added. The WAE platform operates with either WAAS, ACNS, or WAFS software. When WAAS software is installed, the WAE appliance can function as either a Central Manager or as an Application Acceleration Engine. When ACNS software is installed, the WAE appliance functions as a Content Engine or one of the other ACNS device modes (Content Router or Content Distribution Manager). When WAFS software is installed, the WAE appliance functions as a File Engine.

This chapter describes the basic concepts of the Cisco ACNS 5.x network. It describes the function and relationships of the different device modes that make up the ACNS network and explains how content is acquired, distributed, stored, and delivered to the end user.

This chapter contains the following sections:

- [ACNS Network Overview, page 1-2](#)
- [Content Overview, page 1-5](#)
- [Content Request Interception and Router Redirection Overview, page 1-7](#)
- [Content Services Overview, page 1-8](#)
- [Content Acquisition and Distribution Architecture, page 1-11](#)
- [Content Replication Overview, page 1-14](#)
- [ACNS Network Topology Considerations, page 1-15](#)

ACNS Network Overview

Cisco ACNS software Release 5.5 is the basis for the network solution described in this publication. Cisco ACNS software operates in four different device modes. These device modes are Content Engine, Content Distribution Manager, Content Router, and Cisco IP/TV Program Manager. An ACNS network consists of at least one Content Distribution Manager, one or more Content Engines, and one or more optional Content Routers.

Cisco ACNS software allows content services to be configured, reconfigured, and monitored centrally from the Content Distribution Manager graphical user interface (GUI), Cisco's web-based application management tool.

Content Distribution Managers

The primary role of a Content Distribution Manager is to perform centralized content and device management. In the ACNS 5.5 network, the Content Distribution Manager manages both content acquisition and distribution and also manages policy settings and configurations on individual Content Engines. Through the Content Distribution Manager GUI, the network administrator can specify what content is to be distributed and to whom. The Content Distribution Manager also allows the administrator to monitor network nodes and apply changes, such as software upgrades, to groupings of nodes from a central location.

Content Engines

The primary role of a Content Engine in the ACNS network is to serve client requests for content. Content Engines also play a major role in content request routing and in channel distribution of content. The ACNS network deploys Content Engines either inside an enterprise firewall on an internal network, or at the edge of the enterprise network.

Content Engines can be managed centrally through the Content Distribution Manager or locally as separate standalone content caches. (Refer to the *Cisco ACNS Software Configuration Guide for Locally Managed Deployments* publication.)

Content Routers

The primary role of a Content Router in the ACNS network is to redirect client requests for content to the closest Content Engine containing that content. In the ACNS 5.5 network, Content Routers use the Domain Name System (DNS) to ensure that the Content Router receives all requests for content. Once the Content Router receives these requests, it can redirect the end user to the content. (See the [“Configuring Content Request Routing Using a Content Router or Routing Content Engine”](#) section on page 4-41 for more information.)

An alternative to Content Router routing is to use an enterprise router that supports Web Cache Communication Protocol (WCCP) and is configured to intercept and route requests for content. WCCP detects client requests and routes the request to a Content Engine within the same network. WCCP does not require the presence of a Content Router. (See the [“Configuring WCCP Transparent Routing”](#) section on page 4-2.)

IP/TV Program Manager

IP/TV Program Manager Release 5.x is a Linux-based application running on the Cisco Content Engine hardware. IP/TV Program Manager is accessed from a browser and is used by the system administrator or broadcast administrator to set up and manage IP/TV scheduled or on-demand programs, channels, recordings, and file transfers among IP/TV Servers.

You can access IP/TV Program Manager from Netscape 4.5x through 4.7x, or Microsoft Internet Explorer 5.x or later. The browser must have Java and JavaScript enabled.

Functional Description

The Content Engine operates either as a component of a Cisco Application and Content Networking System (ACNS) or as a standalone content-caching device and is generally positioned on the WAN edge between your enterprise and the Internet.

Content Delivery Network Component

Cisco Content Engines with Cisco ACNS software installed are the content delivery components of a larger Cisco ACNS network solution, which includes content routing and content distribution and management. Content Engines with Cisco ACNS 5.x software are deployed in conjunction with a Content Distribution Manager to create a centrally managed ACNS network.

The ACNS solution offers accelerated content delivery, hosting, content replication, video streaming, and other content-based services. The ACNS solution addresses the need to distribute and receive high-bandwidth, media-rich content across the Internet or an intranet without performance losses or content delivery delays.

Content Caching Device

Cisco Content Engines operating with Cisco ACNS software accelerate content delivery and optimize bandwidth usage by transparently caching frequently accessed content and fulfilling content requests locally rather than traversing the Internet or intranet to a distant server farm each time a request is made. The Content Engine works in tandem with a router to handle web traffic, including user requests to view pages and graphics (objects) on World Wide Web servers—whether internal or external to your network.

To deploy Cisco Content Engines with Cisco ACNS software within your existing network, your network must support Cisco IOS software and the Web Cache Communication Protocol (WCCP). WCCP transparently redirects HTTP requests to a Content Engine, and the Content Engine responds to those requests.

For example, when a user requests an object from a web server, the router first sends the request to a Content Engine. If the Content Engine has a copy of the requested object in storage, the Content Engine sends the user the object. Otherwise, the Content Engine simultaneously obtains the requested objects from the web server, stores a copy of the objects (caches them), and forwards the objects on to the user.

You can also configure your Content Engine as a proxy server that acts as a network gateway device, which is optimized to retrieve content on behalf of web clients. Direct proxy routing is known as nontransparent caching because the web clients and media players in the network are configured to explicitly point to the Content Engine that is acting as the proxy server.

**Note**

Unlike transparent caching, nontransparent caching is possible even if your network does not support Cisco IOS software and the WCCP.

In nontransparent caching cases, the Content Engine sends the content to the web client from its local storage if the requested content is already available at that location (cache hit). If the requested content is not already stored in the Content Engine's local cache (cache miss), the Content Engine retrieves the requested content from the origin server, stores a local copy of the content if the content is cacheable, and sends the requested content to the web client. When the Content Engine receives subsequent requests for the same content, it sends the content from its local storage.

By caching web objects, the Content Engine can speed the completion of user requests when more than one user wants to access the same objects. Caching in this manner also reduces the amount of traffic between your network and the Internet, potentially improving your overall network performance and optimizing your bandwidth usage.


Deployment Scenarios

All types of devices do not need to be present in the ACNS network for Cisco ACNS software to function. For example, your ACNS 5.5 software might be deployed in any one of the following scenarios:

- One Content Engine only
- Many standalone Content Engines
- One Content Distribution Manager and many managed Content Engines
- One Content Distribution Manager and many managed Content Engines and Content Routers

Cisco ACNS 5.5 software is especially beneficial for deployment in an enterprise or Internet service provider (ISP) network where there is a requirement to both pre-position and cache content. Typically, such a network uses topologies that contain both high-speed and low-speed connections. Cisco ACNS 5.5 software allows companies to use off-peak hours to move moderate or high-bandwidth content to the edge of their networks over low-speed links.

[Figure 1-1](#) shows a typical enterprise intranet. Branch offices at the enterprise edge connect to the enterprise data center using low-speed (56-kbps, 128-kbps or 512-kbps) links. In addition, branch offices in the same region may connect to a regional headquarter data center, which then connects to corporate headquarters using limited-bandwidth links.

Figure 1-1 Typical Enterprise ACNS Network Topology

In this example, each branch office has a Content Engine. The corporate data center has a Content Distribution Manager that is managed by the central Information Technology (IT) staff. Content Routers might also be deployed in the data center as well.

Content Overview

Content is the fundamental element of the ACNS network; it represents all the data that the ACNS network handles. Content can be classified on the basis of how it is acquired, distributed, or served. The ACNS 5.5 network serves three classes of content: on-demand content, pre-positioned content, and live content.

This section explains some of the terminology that is used to discuss content in this publication.

On-Demand Content

On-demand content is acquired and cached by the Content Engine, and then it is delivered by the Content Engine in response to a client request. When the first client request is made for a piece of content, the Content Engine pulls the data from the origin web server and serves it to the client. The content is also stored (or cached) on the Content Engine. Subsequent requests for the same content are served to the client directly from the Content Engine cache.

Pre-Positioned Content

Pre-positioned content is a means of distributing content to populate Content Engines in a centrally managed ACNS network environment. Bandwidth-intensive content objects, such as Java applets, Flash animations, Shockwave programs, and other file formats can be managed and scheduled for distribution to Content Engines during off-peak hours.

The ACNS network administrator uses the Content Distribution Manager GUI to configure a number of channels to specify the destination Content Engines for content delivery and assigns a Content Engine (called a root Content Engine) that is generally (but not necessarily) located in close proximity to an origin web server to fetch the content according to attributes that are set out in an XML file called a manifest file. The manifest file (see [Appendix A, "Creating Manifest Files"](#)) is used to specify the location from which the pre-positioned content objects should be fetched and the frequency with which the content should be checked for updates. The root Content Engine distributes the content to other Content Engines that serve as content caches and distribution gateways along the enterprise edge.

Pre-Positioned Versus Preloaded Content

Whereas pre-positioned content is associated with a centrally managed ACNS network environment, preloaded content (cache preloading) is configured on a Content Engine-by-Content Engine basis and is not generally associated with a centrally managed ACNS network.

Content Engines can be configured to preload specific content items using HTTP or FTP. Websites are scanned several link levels down for content. Preloaded content can be configured with specified bandwidth limits for better control of network usage.

Content preloading is done by configuring the Content Engine to create a cache request for all the content located at the origin web server that stores the primary content. At a time specified in the **pre-load schedule** command, the Content Engine verifies that its content is still current and updates any content that has changed.

Live Content

Live content is acquired as a live streaming broadcast from either a satellite or terrestrial broadcasting source. The ACNS network administrator configures the policies associated with obtaining the live multimedia stream, such as the program listing URL (playlist), the maximum bit rate, and so forth, as well as the distribution policies, such as priority, schedule, and maximum bandwidth.

Creating and Managing Content in an ACNS Network

Responsibilities for creating and managing content in an ACNS network are divided between webmasters and network administrators. This publication addresses both areas of responsibility.

Creating websites and adding content to the ACNS network are the responsibilities of the webmaster, whereas managing the content and the network is the responsibility of the ACNS network administrator. Webmasters are generally not involved with network topologies, configurations, WAN bandwidths, or Content Engine locations.

Enterprise companies can have multiple webmasters who are responsible for content by business unit, geography, content type, and so forth. Webmasters are authorized to add content to the ACNS network by creating one or more manifest files for channels associated with their website. Webmasters define and schedule acquisition jobs through the manifest file. Content is fetched according to the time defined in the manifest file. (See [Appendix A, “Creating Manifest Files.”](#))

To associate one or more manifest files with the same logical content, the website and all of its associated channels must be registered and predefined in the Content Distribution Manager GUI. A network administrator is the person responsible for registering and defining content providers, websites, and channels through the Content Distribution Manager GUI. The administrator points the channel to a manifest file, assigns Content Engines to the channel, and designates a root Content Engine. (See [Chapter 5, “Configuring the ACNS Network for Content Distribution.”](#))

The ACNS network administrator also assigns channels and allocates disk space. The administrator allocates both the total amount and the per-channel amount of disk space for the website and assigns each channel to a set of device groups. The administrator also schedules bandwidth allocation for each root Content Engine. (See [Chapter 13, “Working with Device Configurations,”](#) and the *Cisco ACNS Software Update and Maintenance Guide*.)

Content Request Interception and Router Redirection Overview

Cisco ACNS 5.5 software supports three methods of request handling to deliver content to the end user:

- WCCP transparent interception
- Direct proxy routing (with or without dynamic proxy autoconfiguration)
- DNS-style routing using a Content Router



Note

An ACNS 5.5 network can combine both WCCP and routing by using a Content Router.

WCCP Transparent Interception

The Content Engine supports WCCP Version 2 redirected requests. When configured for WCCP routing, the Content Engine receives requests from its assigned router and compares the content request against the content currently stored in its cache. If the Content Engine has the requested content (cache hit), the request is served from this cached storage. If the Content Engine does not have the content (cache miss), the content is requested from the origin web server and served to the requesting client through the Content Engine. No proxy configuration is required at the end user's browser. (See the [“Configuring WCCP Transparent Routing”](#) section on page 4-2.)

Direct Proxy Routing

In this mode, end user web browsers need to be explicitly configured to use the IP address or host name of the Cisco Content Engine, and there is no need for additional hardware, such as Layer 4 switches, Content Routers, or WCCP-enabled routers to intercept user requests.

If the Content Engine has the requested content, the request is serviced from this cached storage. If the content is not in the Content Engine's storage, the content is requested from the origin web server and served to the requesting client by way of the Content Engine. (See the [“Configuring Direct Proxy Routing”](#) section on page 4-32 for more information.)

Dynamic Proxy Autoconfiguration

The dynamic proxy autoconfiguration (PAC) routing feature uses a Content Engine that is enabled as a PAC file server. This routing method can be deployed in an ACNS 5.5 network that contains a Content Distribution Manager and at least one Content Engine. If end users are behind a NAT device or firewall and the administrator wants to direct the end users to different proxies based on their IP addresses, then there must be a PAC file server behind the NAT device as well. You do not need a Content Router or a routing-enabled Content Engine.

When using the PAC file server feature, the administrator writes one PAC file with ACNS software-defined macros that serves as a template. The PAC file server substitutes the closest proxies for the macros, where the closest proxy is determined by coverage zone information.

The PAC feature differs from proxy mode configuration in the following ways:

- Dynamic PAC supports multiple PAC files.
- The user can specify the names of the PAC files instead of being limited to preset names.
- The PAC file server dynamically adds a set of the nearest proxies to a PAC file template based on coverage zone information and the source IP address of the requester.
- The Content Distribution Manager acquires the PAC file templates and distributes them to Content Engines that are PAC file servers, rather than having the Content Engines themselves acquire the PAC file templates.

Content Router Routing

Content Router routing uses HTTP or RTSP redirection through a Content Router. The Content Router determines which Content Engine is best suited to deliver the desired content to the client by comparing the source IP address of the client end system against a table of address ranges assigned to Content Engines, known as the *coverage zone*. The coverage zone provides information on the proximity of client end systems to Content Engines based on each client's IP address. The Content Router can then choose the closest, best-suited Content Engine to serve the website request to the client. (See the [“Configuring Content Request Routing Using a Content Router or Routing Content Engine”](#) section on page 4-41 for more information.)

If several Cisco Content Engines are located behind a Network Address Translation (NAT) device or firewall, the Content Router can issue the redirect response to a routing-enabled Content Engine behind the NAT device. The routing Content Engine has the ability to issue a redirect request to another Content Engine that is also located behind the NAT device or firewall. This Content Engine looks at its coverage zone file and issues a subsequent redirect response to the appropriate Content Engine that is configured to serve the client's request. (To enable routing on a Content Engine, see the [“Configuring a Content Engine as a Routing Content Engine”](#) section on page 4-51.)

Content Services Overview

The ACNS 5.5 network combines two types of content services using a single software base: content pre-positioning and content caching with filtering and access control. This section explains the objectives of these services and provides examples that are based on some typical needs and requirements of an enterprise network.

Content Pre-Positioning

When pre-positioning content, the administrator can specify a group of content objects to be “pushed” to the remote branch offices in off-peak times, such as at night. This capability is advantageous because after the content is pre-positioned at a nearby Content Engine in the branch office, for example, it can be accessed using a higher-bandwidth connection than the typical low-bandwidth link between branch offices and remote corporate headquarters. Pre-positioning can be particularly advantageous for streaming objects.

For example, the training department at <http://www.bigcorp.com> might prepare an online training course that includes multiple videos and slides. The course might be posted at an internal website. If there is no ACNS network, an employee from a remote office taking the course online must wait a long time for the video and slides to load. If, however, the administrator uses an ACNS network to pre-position the videos and slides on a Content Engine at the local branch office, then the employee can take the course as if the website were on the local LAN.

Content Caching and Access Control

Content caching, filtering, and access control services address the need to accomplish the following objectives:

- Reduce bandwidth usage by caching frequently accessed Internet content.
- Authenticate, monitor, log, and control web access from end users to implement corporate policies regarding work environment and system security.

Both objectives are made possible because of the Content Engine’s special position as an “in-line” device between end users and the Internet.

For example, BigCorp Enterprise wants to make sure that the following policy, security, and budgetary objectives are met:

- Only full-time employees can access the web.
- No employee can access objectionable content through the corporate network.
- Viruses such as “Code Red” are prevented from being propagated inside the corporate network.
- Monthly payments to the company’s ISP are reduced.

BigCorp Enterprise can achieve all of these objectives by placing Content Engines at the network edge, where the corporate network interfaces with the Internet, and by placing Content Engines at remote branch offices. The Content Engines can then intercept content requests made by end users and enforce the above policies.

A Content Router might also be placed at the BigCorp Enterprise data center to aid in redirecting requests to Content Engines in the ACNS network.

User Authentication and Content Filtering

Content Engines can be configured to perform a number of authentication and content filtering services. Once the Content Engine receives a request, it does a number of things, including the following tasks:

- Authenticates the client, asks the client to provide a username and password so that it can consult an authentication, authorization, and accounting (AAA) server, and checks whether the client is allowed to access the web.

- Passes the request through a “filter,” such as Websense or SmartFilter, to make sure that the requested object is not objectionable content.

ACNS 5.5 software relies on third-party software to implement content filtering. Supported filtering software includes Websense, N2H2, and SmartFilter.

- Passes the request through “rules,” which might rewrite certain headers, redirect the request, or otherwise manipulate the request.
- Checks to see whether the requested content is already in the Content Engine cache. If so, then the Content Engine serves the object directly, rather than from the origin server, thus saving bandwidth to the Internet.
- If the requested content is not already in the cache, then the Content Engine fetches the content from the Internet on the client’s behalf, and caches the content for future use if appropriate.

This functionality is typically called “proxy” functionality. Cisco ACNS 5.5 software supports Content Engine proxy functionality by supporting all web access protocols (including HTTP) and all streaming protocols.

**Note**

For more information, see the [“Authentication Support” section on page 6-42](#) and the [“Configuring Request Authentication on Centrally Managed Content Engines” section on page 15-17](#).

About Device Groups

To facilitate configuring content services for large numbers of Content Engines, ACNS 5.5 software supports device groups.

A *device group* is a set of similar devices (such as Content Engines) that share common qualities and capabilities. Some common qualities might include disk capacity, bandwidth capacity, or routing properties.

By grouping devices into device groups, you can configure settings and channel assignments for the entire group of devices at one time. Settings that can be configured through device groups include Content Engine properties, bandwidth settings, and content services.

**Note**

Content services that require a license key must be enabled on the Content Engine before you can configure service settings for a device group. You cannot enable software licenses on a device group. Licenses must be enabled on an individual Content Engine basis. (See the [“Licensing and Enabling Streaming Servers” section on page 9-6](#).)

Content Engines can be assigned to multiple device groups when the device group overlap feature is enabled in the Content Distribution Manager GUI. (See the [“Enabling Device Group Overlap” section on page 13-30](#).)

The Content Distribution Manager GUI allows you to change settings on individual devices to override the device group settings, and gives you the option of reverting back to the device group settings, if you wish.

To create and modify device groups, see the [“Working with Device Groups” section on page 13-26](#).

Content Acquisition and Distribution Architecture


This section discusses the ACNS network content acquisition and distribution architecture. It describes how content arrives at a Content Engine and discusses the role of content providers, websites, channels, and manifests in fetching and distributing content to the various Content Engines in the ACNS network.

Understanding Basic Concepts and Terms

Content is made available for distribution by *content providers*. A content provider is a person, group of persons, or organization that is responsible for providing content for an enterprise organization. Content is usually stored on file servers or web servers that are not part of the ACNS network of managed devices. In this publication, the server that stores the content is called the *origin server*. Content from the content provider is often placed on the origin server by a webmaster, who uses the content to build a *website* or multiple websites.

Content from a single website is mapped to a set of devices or device groups by means of a *channel*. A channel is a configuration used to define how content is to be acquired, distributed, and stored in advance of a client request. Content objects associated with a specific channel have a common domain (host) name; in other words, the content in a specified channel resides in a single location on an origin server. Each channel is restricted to a single website and maps domain names to devices one to one for Content Router DNS interception. (See [Figure 1-2](#).)

Figure 1-2 One-to-One Channel Mapping to Devices



When a Content Engine is assigned to a channel, it is also by definition subscribed to the channel's website. A Content Engine that is subscribed to a website can handle content requests for this website through a Content Router and cache the content from this website. A Content Engine must be assigned to a channel to accept and distribute *pre-positioned content* for that channel. Channel assignment is also required for the Content Router to be able to serve content from the website of that channel.

[Table 1-1](#) defines some important ACNS software concepts and terms.

Table 1-1 ACNS Software Concepts and Terms

Term	Definition
Locations	Groupings of well-connected (often physically connected) or interchangeable Content Engines with adjacent IP addresses and usually found in the same proximate geographic location. Locations organize and group Content Engines into virtual networks for distribution of content through channels. Locations are configured in the Content Distribution Manager GUI; each location can have zero to many Content Engines assigned to it. (See the “Creating and Modifying Locations” section on page 5-3.)
Level	Hierarchical arrangement of locations defined by the child-to-parent relationships between locations. - Each location can have zero to one location as its parent. Each location can have zero to many locations as its children. - Locations with no parents are placed at level 1. A location whose parent is at level 1 is placed at level 2, and so on. An ACNS 5.5 network can have up to 4 levels and support up to 200 Content Engines per location and up to 200 child locations per parent location. To configure location levels, see the “Creating and Modifying Locations” section on page 5-3.
Location tree	Set of locations organized in the form of a tree structure. The structure is defined by the location levels. An ACNS network can consist of one or many location trees.
Channels	Groups of content to which Content Engines can be assigned. Content Engines in the same location can be assigned to different channels. - When configuring a channel in the Content Distribution Manager GUI, the administrator specifies a list of Content Engines (which belong to the channel and which replicate the content of the channel) and designates a root Content Engine to acquire the content from the origin server and publish it. - For any given channel, there is only one publisher of content (the root Content Engine) and multiple receivers of that content (the Content Engines that are assigned to that channel). The location that contains the root Content Engine for a given channel is called the <i>root location</i>. A channel can have only one configured root Content Engine. Other Content Engines in the root location can act as backup publishers if the configured root Content Engine fails. To configure channels in the Content Distribution Manager GUI, see the “Creating and Modifying Channels” section on page 5-10.
Root Content Engine	Content Engine designated to download the content from the origin server and forward it to the Content Engines that are assigned to the content channel. You designate a Content Engine to be the root Content Engine for a channel through the Channels tab of the Content Distribution Manager GUI. (See the “Designating the Root Content Engine” section on page 5-16.)
Forwarder	Content Engine that has been chosen as a distribution host by another Content Engine from a another location.
Location Leader	Content Engine that acts as a forwarder to other Content Engines within its own location.
Content provider	Person, group, or organization that is responsible for providing content for an enterprise or organization.
Origin server	Server that stores the source content.

Table 1-1 ACNS Software Concepts and Terms (continued)

Term	Definition
Website	Collection of content objects from a single origin server. Websites can be classified as either routable or nonroutable. Routable websites are controlled and operated by the enterprise corporation where the content is owned. Routable website domain names are fully qualified domain names (FQDNs) that are recognizable to Content Routers. Routable websites support all ACNS software routing and edge intercept mechanisms: WCCP interception, direct proxy routing, and Content Router routing. (See the “Content Request Interception and Router Redirection Overview” section on page 1-7.) Nonroutable websites are not controlled or operated by the enterprise corporation, and the content is not owned by the enterprise. For example, www.cnn.com or www.yahoo.com are nonroutable websites. Nonroutable websites support WCCP interception and proxy configuration only.
Manifest file	XML file consisting of explicit item descriptions and crawler items that is used by the root Content Engine to acquire content from an origin server.
Multicast cloud	Group of multicast-enabled Content Engines configured to communicate multicast session information with one another.

Acquiring Content

Before content is distributed, it is acquired from the origin server by a *root* Content Engine. The root Content Engine is the one Content Engine that is authorized to go directly to the origin server for content. The root Content Engine then publishes the content to other Content Engines in the channel. You assign the root Content Engine by designating it as such in the Content Distribution Manager GUI. In general, the root Content Engine should be in the same general location as the origin server.

The root Content Engine is assigned a *manifest file* that specifies the content for a single channel. The manifest file is an XML file consisting of explicit item descriptions and crawler items (items to be acquired by crawling a website or FTP server) with prefix constraints or match rules. The manifest file is defined by the webmaster or the ACNS network administrator acting as an agent for the webmaster and controls both content acquisition and content distribution. The manifest file is assigned to one or more root Content Engines (a primary and a list of backups) through the Content Distribution Manager GUI.

Distributing Content Through Channels

Channels form logical routes for content to travel from an origin server through the root Content Engine to all the other Content Engines in the channel. Logical routes for content distribution are based on the device *location* hierarchy, or *location tree*. (See [Table 1-1](#) for an explanation of these terms.)

The content distribution route follows the general tree structure of the manually constructed ACNS network location tree, where content is distributed from the root of the tree (root Content Engine) toward the branches; however, only the Content Engines assigned to a particular channel can be a part of that channel distribution tree. The channel distribution tree is constructed specifically for each channel.

Content is forwarded from Content Engine to Content Engine through the channel distribution tree until all Content Engines in the channel have received their subscribed content.

Content Replication Overview

Content is forwarded (or replicated) either by unicast pull (transmission initiated by a client request for the content) or, if it is enabled, by multicast push (transmission initiated in accordance with a preconfigured program or schedule). Unicast content forwarding involves communication between a single sender and single receiver, whereas multicast replication involves communication between a single sender and a selected group of receivers.

Unicasting

Unicasting is a method of data transmission in which each Content Engine or client receives its own copy of a file or a data stream from a single sender. For example, unicasting is used when only one Content Engine downloads a set of files or a data stream. In another scenario, several clients request the file or the data stream, and each receives a new copy at the same time; for example, a proxy Content Engine receives the unicast stream and splits the stream, which is then fanned out to the multiple downstream clients that requested it. An advantage of unicasting is that each client can receive the entire data set at any time. The disadvantage of unicasting is that existing line bandwidths cannot be used optimally. (See the [“Configuring Windows Media Live Splitting for Unicast and Multicast Transmissions”](#) section on page 9-26 for information on configuring unicast transmissions.)

Multicasting

Multicasting allows efficient distribution of content to multiple Content Engines and is useful when many end users are interested in the same content. ACNS software supports Pragmatic General Multicast (PGM)-based multicast replication, using either satellite or multicast-enabled terrestrial infrastructures. (PGM is a reliable multicast protocol that enables PGM receivers to report loss of data and request retransmission by the PGM sender.)

In ACNS 5.x software, the administrator configures the ACNS network for multicasting by configuring a *multicast cloud* in the Content Distribution Manager GUI. The multicast cloud consists of one sender Content Engine, an optional backup sender (available in ACNS 5.1 software and later releases) for multicast-to-multicast failover, and at least one receiver Content Engine in a hub-spoke topology. All the Content Engines in one cloud share a unique advertising address, allowing them to communicate multicast session information. Content Engines that are assigned to the multicast cloud must be enabled for multicasting.

The administrator enables the multicasting feature in the Content Distribution Manager GUI by entering a license key (purchased from Cisco Systems) for each Content Engine that is to be multicast-enabled. (See the [“Enabling Content Engines for Multicasting”](#) section on page 5-33.)

The multicast-enabled Content Engines must then be assigned to one or more multicast-enabled channels. (See the [“Adding and Removing Content Engines from Channels”](#) section on page 5-16 and the [“Assigning and Removing Multicast Clouds from Channels”](#) section on page 5-43.) To implement multicast distribution, a multicast-enabled channel should have at least one multicast enabled cloud associated with it. In addition, the multicast-enabled Content Engines in the cloud must subscribe to the channel. You can enable a channel for multicasting when you first create the channel, or you can modify the channel properties for multicasting. (See the [“Creating and Modifying Channels”](#) section on page 5-10.)

When configuring the multicast cloud, the administrator specifies a range of addresses by entering a start IP address and an end IP address. Once a multicast cloud is configured, the multicast address range is used to provide each channel associated with it a unique data channel multicast address. When a channel

is assigned to a multicast cloud (see the [“Assigning and Removing Multicast Clouds from Channels” section on page 5-43](#)), an unused IP address is automatically selected from this range to ensure that the address is used by only one channel and by only one multicast cloud. Since different multicast clouds may be associated with the same channel, the multicast address used for each channel needs to be different in each multicast cloud.

For more information on IP multicasting, including assigning IP multicast addresses, see [Appendix B, “IP Multicast Addressing.”](#)

ACNS Network Topology Considerations

When you are planning your network topology, you need to understand what your objectives are. With your objectives clearly in mind, you can then determine the number and types of ACNS devices that you need to install and map out the best placement of each device in your network so that you can reach those objectives.

After your physical devices are installed in a centrally managed ACNS deployment, the network administrator configures the logical ACNS network topology by defining the device relationship parameters in the Content Distribution Manager GUI. In the GUI, the administrator names the device locations, assigns the parent-child relationship to the locations, creates channels, and assigns Content Engines to channels. Procedures for using the Content Distribution Manager GUI to set up device locations and relationships are covered in [Chapter 5, “Configuring the ACNS Network for Content Distribution.”](#)

Use the following table to list your objectives and prioritize them.

Objective	Priority
Accelerate access to website content	
Reduce bandwidth usage	
Control web access from end users	
Authenticate requests for content	
Monitor and log website activity	
Protect against virus attacks	
Optimize your installed infrastructure	

When planning your ACNS network, be prepared to answer questions such as these:

- What does my network look like?
- Where is my origin web server located?
- Where do I want the content to go?
- How many Content Engines do I need at each location to support fan-out of child locations and end users?
- How many Content Engines do I need for failover and redundancy?
- Which Content Engine models do I need to install to support the type of work and activities required?
- What routing mechanism do I want to use to direct user requests to Content Engines: direct proxy configuration, transparently using a WCCP-enabled router, HTTP or RTSP redirection using a Content Router or routing Content Engine?
- What type of content do I need to deliver: unicast or multicast live streaming content, video-on-demand, static web content?

Use the following table to record information about your environment so that you can determine the number and types of ACNS devices to install in your network. This information will also help you determine the best place to install each ACNS device within your existing network infrastructure.

Factors	Your Information
Existing infrastructure:	
Number of routers	
Number of switches	
Number of servers	
Location of origin servers	
Number of remote sites	
Location of each site	
Number of end users (client nodes) in each location	
Type of work or activity	
Expected bandwidth usage during peak hours	



Getting Started

This chapter discusses initial device configuration and basic ACNS network configuration for devices that make up an ACNS network. The following sections explain how to configure ACNS devices to operate and communicate in a centrally managed environment:

- [Overview, page 2-1](#)
- [Using the Quick Setup Method, page 2-8](#)
- [Initial Device Configuration Using the Standard CLI, page 2-14](#)
- [Initial Network Configuration Using the Standard Content Distribution Manager GUI, page 2-22](#)
- [Where to Go Next, page 2-25](#)

Overview

To get started, you must perform two kinds of initial configurations:

1. Device
2. Network

To configure initial device settings and network elements, you can choose between two configuration methods:

- Quick method
- Standard method

The first step in either method is to initially configure the devices. Initial device configuration requires the configuring of device network values, such as IP address, netmask, DNS host, and gateway IP address.

After your devices are up and running, you must register each device with a central management device, known as the Content Distribution Manager, to bring the device online in the Content Distribution Manager GUI. Devices can then be configured and managed from one central location through the Content Distribution Manager GUI.

Initial Device Configuration Overview

You must initially configure all of your ACNS devices before they can participate in the ACNS network. You can choose between three methods of initial device configuration:

- Autoregistration—Quick method

The autoregistration method performs the initial device configuration and network registration with the Content Distribution Manager automatically when the device is powered on for the first time. (See the [“Quick Device Configuration Using Autoregistration”](#) section on page 2-8.)


- Setup utility—Quick method

With the setup utility method, the administrator performs initial device configuration and network registration using the CLI-based interactive setup utility through a console connection. (See the [“Quick Device Configuration Using the Setup Utility”](#) section on page 2-10.)

- Command-line-interface—Standard method

Initial device configuration and network registration can also be performed using the standard command-line-interface (CLI) of the device. (See the [“Initial Device Configuration Using the Standard CLI”](#) section on page 2-14.)

[Figure 2-1](#) shows the task flow for initially configuring devices.

Figure 2-1 *Initially Configuring Devices in a Centrally Managed ACNS Network*

Initial Network Configuration Overview


After you have initially configured your ACNS devices, you must activate the devices and then configure a basic ACNS network. Device activation is done through the Content Distribution Manager GUI and allows only approved devices to participate in the ACNS network and communicate with other activated nodes.

You can choose between two methods of initial network activation and configuration:

- Quick Start tool—Quick method
- Content Distribution Manager GUI—Standard method

We recommend that you use the GUI-based Quick Start tool to activate your devices and to set up your basic ACNS network. Later you can modify this initial configuration through the standard Content Distribution Manager GUI or CLI. (See the [“Quick Network Configuration Using the Quick Start Tool” section on page 2-12.](#))

[Figure 2-2](#) shows the task flow for setting up a basic ACNS network using the Quick Start tool in the Content Distribution Manager GUI.

Figure 2-2 **Setting Up a Basic ACNS Network**

Summary of Initial Configuration Tasks

This section summarizes the initial configuration tasks for setting up and configuring your ACNS network for both the quick setup method and the standard configuration method, and it directs you to sections or other chapters that describe these tasks in more detail.

Enable a Content Distribution Manager

1. Power up a WAE appliance, and bypass autoregistration by pressing any key within the first 30 seconds during bootup.
2. Configure the device mode using the interactive setup utility or by using the **device mode** global configuration command in the CLI.
(To use the setup utility, see [Quick Device Configuration Using the Setup Utility, page 2-10.](#))
(For CLI configuration, see [Configuring the Device Mode, page 2-14.](#))
3. Configure a primary Content Distribution Manager, and enable management services using the interactive setup utility or the CLI.
(To use the setup utility, see the [Quick Device Configuration Using the Setup Utility, page 2-10.](#))
(For CLI configuration, see the [Configuring the Content Distribution Manager, page 2-16.](#))
4. Configure a standby Content Distribution Manager (optional).

Enable a Content Router (optional)

1. Power up the Content Router or Content Engine.
2. Configure the device mode using the interactive setup utility or by using the **device mode** global configuration command in the CLI.
(To use the setup utility, see [Quick Device Configuration Using the Setup Utility, page 2-10.](#))
(For CLI configuration, see [Configuring the Device Mode, page 2-14.](#))
3. Do you have a DHCP server configured?
 - Yes. Autoregistration automatically configures network settings and registers Content Routers with the Content Distribution Manager upon bootup.
(See [Quick Device Configuration Using Autoregistration, page 2-8.](#))
 - No. Disable autoregistration and manually configure the network settings, and register the device with the Content Distribution Manager using the interactive setup utility or the CLI.
(To use the setup utility, see [Quick Device Configuration Using the Setup Utility, page 2-10.](#))
(For CLI configuration, see [Configuring the Device Network Settings, page 2-15](#) and [Registering a Device with the Content Distribution Manager, page 2-20.](#))
4. Configure your DNS server for the Content Router.
(See [Configuring Content Request Routing Using a Content Router or Routing Content Engine, page 4-41.](#))

Enable Content Engines

1. Power up the Content Engine.
2. Do you have a DHCP server configured?
 - Yes. Autoregistration automatically configures network settings and registers Content Engines with the Content Distribution Manager upon bootup.
(See [Quick Device Configuration Using Autoregistration, page 2-8.](#))
 - No. Disable autoregistration and manually configure network settings, and register the device with the Content Distribution Manager.
Using either the interactive setup utility or the CLI, configure the following network settings:
Ethernet interface
IP domain name
Host name
IP name server
Default gateway
Primary interface
(To use the setup utility, see [Quick Device Configuration Using the Setup Utility, page 2-10.](#))
(For CLI configuration, see [Configuring the Device Network Settings, page 2-15](#), and [Registering a Device with the Content Distribution Manager, page 2-20.](#))

Configure Content Engine Disk Space

1. Determine which software services you wish to enable.
The software services that you have enabled will determine the types of file systems that you need to configure and the amount of disk space appropriate for each type of service.
2. Create disk space for conventional caching (cfs), pre-positioning (cdnfs), streaming services (mediafs), and system use (sysfs).
 - Use the Quick Start tool to choose a basic disk configuration.
(See [Quick Network Configuration Using the Quick Start Tool, page 2-12.](#))
 - Configure disk space using the CLI.
(See [Configuring Disk Space, page 2-19.](#))
 - Configure disk space using the Content Distribution Manager GUI File System Configuration window.
(See the *Cisco ACNS Software Update and Maintenance Guide*.)
3. Reboot the device for the disk configuration to take effect.

Activate All Devices

From the Content Distribution Manager GUI, activate all devices that are registered with the Content Distribution Manager using either the Quick Start tool or the standard Device Activation window.

(To use the Quick Start tool, see [Quick Network Configuration Using the Quick Start Tool, page 2-12.](#))

(To use the Device Activation window, see [Activating Devices in the Content Distribution Manager GUI, page 2-22.](#))

Enable Software Services

Choose basic software services in the Quick Start tool, or enable software services using the Content Distribution Manager GUI.

(For advanced caching, streaming, and pre-positioning configuration information, see the appropriate chapters found in the [“Managing ACNS Network Devices”](#) section of this guide.)

Configure One or More Request Routing Methods

1. Initially configure the request routing method or methods for your network.
 - Use the Quick Start tool (recommended method for getting your ACNS network up and running quickly).
 - Use the standard Content Distribution Manager GUI to configure your Content Engine for request routing.
2. Configure other devices, and set up your network infrastructure to route requests according to the routing methods used in your network.
3. For more information, see [Chapter 4, “Setting Up Content Request Routing in the ACNS Network.”](#)

Using the Quick Setup Method

This section explains how to quickly configure your ACNS devices and set up a basic ACNS network. Choose one of the following quick device configuration methods:

- [Quick Device Configuration Using Autoregistration, page 2-8](#)
- [Quick Device Configuration Using the Setup Utility, page 2-10](#)

After you complete the initial device configuration, use the following quick network configuration method to quickly set up a basic ACNS network:

- [Quick Network Configuration Using the Quick Start Tool, page 2-12](#)

Quick Device Configuration Using Autoregistration

Autoregistration automatically configures network settings and registers Content Engines with the Content Distribution Manager.

On bootup, devices running ACNS 5.x software (with the exception of the Content Distribution Manager itself) automatically discover the Content Distribution Manager and register with it. The administrator does not have to do any manual configuration on the device. Once the device is registered, the administrator approves the device and configures it remotely using the Content Distribution Manager GUI. (See the [“Activating Devices in the Content Distribution Manager GUI”](#) section on [page 2-22.](#))

Autoregistration uses a form of Dynamic Host Configuration Protocol (DHCP). For autoregistration to work, you must have a DHCP server that is configured with the host name of the Content Distribution Manager and that is capable of handling vendor class option 43.

**Note**

The form of DHCP used for autoregistration is *not* the same as the interface-level DHCP that is configurable through the **ip address dhcp** interface configuration command. (See the [“About Selecting Static IP Addresses or Using Interface-Level DHCP”](#) section on page 2-12 for an explanation of interface-level DHCP.)

The vendor class option (option 43) information needs to be sent to the ACNS network device in the format for encapsulated vendor-specific options as provided in RFC 2132. The relevant section of RFC 2132, Section 8.4, is reproduced here as follows:

The encapsulated vendor-specific options field should be encoded as a sequence of code/length/value fields of syntax identical to that of the DHCP options field with the following exceptions:

1. There should not be a “magic cookie” field in the encapsulated vendor-specific extensions field.
2. Codes other than 0 or 255 may be redefined by the vendor within the encapsulated vendor-specific extensions field but should conform to the tag-length-value syntax defined in section 2.
3. Code 255 (END), if present, signifies the end of the encapsulated vendor extensions, not the end of the vendor extensions field. If no code 255 is present, then the end of the enclosing vendor-specific information field is taken as the end of the encapsulated vendor-specific extensions field.

In accordance with the RFC standard, the DHCP server needs to send the Content Distribution Manager host name information in code/length/value format (code and length are single octets). The code for the Content Distribution Manager host name is 0x01. DHCP server management and configuration are not within the scope of the autoregistration feature.

The ACNS network device sends “CISCOCDN” as the vendor class identifier in option 60 to facilitate device groupings by customers.

Autoregistration DHCP also requires that the following options be present in the DHCP server’s offer to be considered valid:

- Subnet-mask (option 1)
- Routers (option 3)
- Domain-name (option 15)
- Domain-name-servers (option 6)
- Host-name (option 12)

In contrast, interface-level DHCP requires only subnet-mask (option 1) and routers (option 3) for an offer to be considered valid; domain-name (option 15), domain-name-servers (option 6), and host-name (option 12) are optional. All of the above options, with the exception of domain-name-servers (option 6), replace the existing configuration on the system. The domain-name-servers option is added to the existing list of name servers with the restriction of a maximum of eight name servers.

Autoregistration is enabled by default on the first interface of the device. The first interface depends on the Content Engine model as follows:

- For the CE-507, CE-507AV, CE-560, CE-560AV, CE-590, and CR-4430: FastEthernet 0/0
- For the CE-510, CE-511, CE-565, CE-566, CE-7305, CE-7325, and CE-7320: GigabitEthernet 1/0

If you do not have a DHCP server, the device is unable to complete autoregistration and eventually times out. You can disable autoregistration at any time after the device has booted and proceed with manual setup and registration.

To disable autoregistration, or to configure autoregistration on a different interface, use the **no auto-register enable** command in global configuration mode.

**Note**

Autoregistration is automatically disabled if a static IP address is configured or if interface-level DHCP is configured on the same interface as autoregistration. (See the [“About Selecting Static IP Addresses or Using Interface-Level DHCP”](#) section on page 2-12.)

The following example disables autoregistration on Fast Ethernet port 0/0:

```
CE(config)# no auto-register enable FastEthernet 0/0
```

Autoregistration status can be obtained by using the following **show** command:

```
CE# show status auto-register
```

You have completed the initial device configuration requirements by using the autoregistration feature. Now you must activate your Content Engines and Content Routers in the Content Distribution Manager GUI and configure any additional network interfaces. You can activate your Content Engines by using the Quick Start tool (see the [“Quick Network Configuration Using the Quick Start Tool”](#) section on page 2-12), or you can follow the steps in the [“Initial Network Configuration Using the Standard Content Distribution Manager GUI”](#) section on page 2-22.

Quick Device Configuration Using the Setup Utility

The setup utility is automatically invoked at system bootup when the startup configuration area in flash memory is empty (such as when the device is factory-fresh). This tool can also be launched using the **setup** command in EXEC mode.

The setup utility allows you to configure network settings and register with the Content Distribution Manager using an interactive dialog instead of the CLI. If you wish, you can disable autoregistration by responding to the setup utility prompt and configuring the first interface on your device (Ethernet 0/0 or GigabitEthernet 1/0) with either a static IP address or with interface-level DHCP. (See the [“About Selecting Static IP Addresses or Using Interface-Level DHCP”](#) section on page 2-12.) If you specify the configuration on the interface that uses autoregistration, autoregistration is automatically disabled. You can also use the setup utility to configure other interfaces if you have multiple interfaces on your device.

**Note**

The admin password prompt only appears at bootup when the startup configuration area is empty. The utility prompts you for the device mode only on Content Engine models that allow device mode changes.

To use the setup utility, follow these steps:

Step 1 Power up the device and open a console connection.

After the operating system boots, the following prompt appears:

```
ACNS boot:detected no saved system configuration
Do you want to enter basic configuration now?
hit RETURN to enter basic configuration:0028
```

At the appearance of this prompt, a 30-second countdown begins, during which you can respond to this prompt and initiate the setup utility. If you do not respond, the system continues booting and automatically registers with the Content Distribution Manager.

Step 2 Press **Enter**, and then enter responses as you are prompted.

In the following sample setup utility, administrator entries are shown in bold. The administrator has chosen not to enable DHCP.

```

admin password:
re-enter password:

What is the mode of the device (CE/CR/CDM/PM) [CE]:CE

Is this CE going to be managed by a
CDM (Content Distribution Manager) (y/n) [y]: y

Please choose an interface to configure from the following list:
1: GigabitEthernet 1/0
2: GigabitEthernet 2/0

Enter choice: 1

Do you want to enable DHCP on this interface (y/n) [n]: n

Please enter the IP address of this interface [209.165.200.231]: <Enter>

Please enter the netmask of this interface [255.255.255.224]: <Enter>

Please enter the default gateway [209.165.200.225]: <Enter>

Please enter the domain name server ip [172.30.2.133]: <Enter>

Please enter the domain name [cisco.com]: <Enter>

Please enter the hostname [CE565]: <Enter>

Please enter CDM (Content Distribution Manager) IP or Host name [172.16.226.212]: <Enter>

Do you want to apply the configurations (y/n) [y]: y

```

Step 3 Choose whether or not to save the configuration. You should see the following message:

```

This may take few moments. Please wait..

For your reference, configuration generated
during this setup session is available at /local1/setup_gen_config.txt.
You can view this file with the command, "type /local1/setup_get_config.txt".

Press any key to continue..

```

After you choose to save the configuration, the device performs hardware and software initialization tasks and then serves a login challenge that, when answered correctly, launches the CLI. At this point, the initial configuration is complete.

Step 4 To make sure that a primary interface has been configured, use the **show running-config EXEC** command. A primary interface should have been chosen automatically by the software during the initial startup; however, you can configure a primary interface manually by using the **primary-interface** command in global configuration mode.

You have completed the initial device configuration requirements by using the setup utility. Now you must activate your Content Engines and Content Routers in the Content Distribution Manager GUI and configure any additional network interfaces. You can activate your Content Engines by using the Quick Start tool. (See the [“Quick Network Configuration Using the Quick Start Tool”](#) section on page 2-12, or you may follow the steps in the [“Initial Network Configuration Using the Standard Content Distribution Manager GUI”](#) section on page 2-22.)

About Selecting Static IP Addresses or Using Interface-Level DHCP

During the initial configuration, you have the option of configuring a static IP address for the device or choosing DHCP. This section describes these two options.

DHCP is a communications protocol that lets network administrators manage their networks centrally and automate the assignment of IP addresses in an organization's network. When an organization sets up its computer users with a connection to the Internet, an IP address must be assigned to each device. Without DHCP, the IP address must be entered manually for each computer, and if computers move to another location in another part of the network, a new IP address must be entered. DHCP automatically sends a new IP address when a computer is connected to a different site in the network.

In the following sample setup utility, the administrator has chosen to enable DHCP. Administrator entries are shown in bold. (See also the [“Configuring Interfaces for DHCP” section on page 14-13.](#))

```

BOOT-100:sw-filesystems mounted, applying pending upgrades...
SW up-to-date

ACNS boot:detected no saved system configuration
  Do you want to enter basic configuration now?
    hit RETURN to enter basic configuration:0019
admin password:
re-enter password:
Is this CE going to be managed by a
CDM (Content Distribution Manager) (y/n) [y]: y

Please choose an interface to configure from the following list:
1: FastEthernet 0/0
2: FastEthernet 0/1

Enter choice: 1

Do you want to enable DHCP on this interface (y/n) [n]: y
Please enter CDM (Content Distribution Manager) IP or Host name [172.16.226.212]: <Enter>

Do you want to apply the configurations (y/n) [y]: y

```

Quick Network Configuration Using the Quick Start Tool

The Content Distribution Manager GUI Quick Start tool helps you set up a small and simple ACNS network with minimal time and effort. From the Quick Start tool you can select a set of Content Engines and provision them for caching, streaming, and pre-positioning of content. You can configure these services using a small subset of options. Advanced options are available from the main user interface. Additionally, the Quick Start tool allows you to choose one or more methods for content routing. As with content services, the tool only supports a small set of options for content routing.

[Table 2-1](#) describes the configuration tasks that can be accomplished using the Content Distribution Manager Quick Start tool.

Table 2-1 **Using the Content Distribution Manager Quick Start Tool**

Task	Description
Choose Content Engines from the device listing.	The tool creates a device group and a location. It activates any inactive devices.
Choose caching, streaming, and pre-positioned content services desired.	The tool enables the protocols for caching, streaming, and pre-positioning services and configures disk space allocation based on the types of services chosen.
Choose a content request routing method.	The choices are WCCP Version 2 routing, direct proxy routing, or content routing (the last requires a Content Router).

Before using the Quick Start tool, you must first complete the initial device configuration and setup as described for centrally managed deployments.

**Note**

Before you can access the Content Distribution Manager GUI, you must use the CLI to register yourself as a user. Use the **username** global configuration command to configure your name and password for GUI access.

To access the Quick Start tool, follow these steps:

- Step 1** In your web browser, enter the URL or IP address for the Content Distribution Manager. For example, enter the URL:
- `https://Name_of_Content_Distribution_Manager:8443`**
- Alternatively, enter the IP address:
- `https://IP_address_of_Content_Distribution_Manager:8443`**
- The Security Alert window appears.
- Step 2** Click **Yes** to accept the security certificate. The Login window appears.
- Step 3** Enter your username in the Username field. Enter your password in the Password field, and click **Login**.

**Note**

The default username is *admin* and the default password is *default*. If the defaults have been changed by another Content Distribution Manager administrator, you need to obtain the new username and password.

The System Home window appears.

- Step 4** Click the **Quick Start** button in the upper right corner of the window. The Quick Start tool opens in a separate window.
- Step 5** Follow the instructions in the Quick Start dialog boxes to complete a quick configuration.

Initial Device Configuration Using the Standard CLI

This section discusses initial device configuration using the standard CLI. You can choose this method to initially configure your devices *instead* of using one of the quick device configuration methods. You can also refer to these procedures when you need to troubleshoot the device network settings.

The procedures that follow assume that you have physically installed the hardware and powered up your device. After you physically install the hardware and power up your device, you can access the ACNS software through a console connection or a Telnet session. If you choose to use this method to initially configure your devices, we recommend that you complete the tasks in the order presented in this section.

**Note**

If you have upgraded your software from a previous ACNS software release to ACNS 5.x software, your network configuration is preserved. You do not need to reconfigure your network settings.

Configuring the Device Mode

All ACNS devices ship from the factory as Content Engines. Before configuring network settings for Content Distribution Managers and Content Routers using the CLI, you first need to change the device from a Content Engine to the proper device mode.

**Note**

Additional steps are required to change the device mode of a device that is already operating as part of your ACNS network. See the [“Repurposing an ACNS Network Device”](#) section on page 13-36.

Configuring the device mode is not a supported option on all hardware models. However, some hardware models can be configured to operate as any one of the three Content Networking device types. Devices that can be reconfigured using the **device mode** global configuration command are shipped from the factory by default as Content Engines.

The following hardware models support device mode configuration:

- CE-7305
- CE-7326
- CE-565
- CE-566
- WAE-611
- WAE-612

To change the device mode of your eligible Content Engine, you must also configure the disk space allocations, as required by the different device modes, and reboot the device for the new configuration to take effect. (See the [“Configuring Disk Space”](#) section on page 2-19.)

When you change the device mode of a Content Engine to a Content Router or a Content Distribution Manager, you need to configure the system file system (sysfs). We recommend that for a Content Router or a Content Distribution Manager, you configure sysfs at 100 percent. For example:

```
CR# disk config sysfs 100% cfs 0% mediafs 0% cdnfs 0%
```

If, however, you are changing the device mode of a Content Router or a Content Distribution Manager back to a Content Engine, you must configure disk space allocations for the caching (cfs), pre-positioning (cdnfs), streaming (mediafs), and system use (sysfs) file systems that are used on the Content Engine. (See the *Cisco ACNS Software Update and Maintenance Guide*.)

For example:

```
DeviceName# disk config sysfs 10% cfs 20% mediafs 10% cdnfs 60%
```

To configure the device mode, follow these steps:

-
- Step 1** Boot the device and access the device CLI through the console or a Telnet session.
- Step 2** Enter the **show device-mode current** EXEC command to view the current device mode.
- ```
DeviceName# show device-mode current
Current device mode: content-engine
```
- Step 3** Configure the new device mode by using the **device mode** command in global configuration mode.
- ```
DeviceName# configure
DeviceName(config)# device mode content-router
```
- Step 4** Save the configuration by issuing the **write memory** EXEC command or the **copy running-config startup-config** EXEC command.
- ```
DeviceName# write memory

DeviceName# copy running-config startup-config
```
- Step 5** Use the **restore factory-default preserve basic-config** EXEC command to reload the software and apply the device mode configuration. Your network settings are preserved.
- ```
DeviceName# restore factory-default preserve basic-config
```
- Step 6** Verify the new configuration. Check that the current and configured device modes are the same.
- ```
DeviceName# show device-mode configured
Configured Device mode: content-router
DeviceName# show device-mode current
Current device mode: content-router
```
- Step 7** Configure the disk space allocation based on the new device mode and needs of your network by using the **disk config** EXEC command.
- Step 8** Reload the device for the disk configuration to take effect.
- ```
DeviceName# reload
```

After the device reloads and you have verified the device mode configuration, you are ready to configure the network settings. (See the next section, “[Configuring the Device Network Settings](#).”) We recommend that you begin by configuring the Content Distribution Manager.

Configuring the Device Network Settings

To configure network settings using the CLI, follow these steps:

-
- Step 1** Power up the device and log on through the console. At the login prompt, enter the username **admin** and the password **default**.
- You must log in to the CLI with an ACNS system account that has superuser privileges.

- You must use a console connection to complete this initial configuration. After this initial configuration is complete, you can use Telnet sessions to access the CLI for subsequent configuration tasks, such as disk configuration.

Step 2 From the device CLI, enter global configuration mode.

```
CE-507# config
CE-507(config)#
```

Step 3 In global configuration mode, configure the device network settings by using the following commands:

- Configure your Ethernet interface for DHCP.

```
CE-507(config)# interface {FastEthernet | GigabitEthernet} slot/port ip address dhcp
```

The remainder of your network settings are configured automatically. You have finished configuring your network settings.

- Alternatively, configure the static IP address of your Ethernet interface.

```
CE-507(config)# interface {FastEthernet | GigabitEthernet} slot/port ip address
ip-address netmask
```

Step 4 Continue configuring the remaining network settings.

- a. Configure the IP domain name.

```
CE-507(config)# ip domain-name name1 name2 name3
```

- b. Configure the host name.

```
CE-507(config)# hostname name
```

- c. Configure the IP name server.

```
CE-507(config)# ip name-server ip-address
```

- d. Configure the IP default gateway.

```
CE-507(config)# ip default-gateway ip-address
```

- e. Configure the primary interface.

```
CE-507(config)# primary-interface {FastEthernet | GigabitEthernet} slot/port [dhcp]
```

You have completed configuration of network settings. To continue with the Content Distribution Manager setup, proceed to the next section, [“Configuring the Content Distribution Manager.”](#)

Configuring the Content Distribution Manager

The Content Distribution Manager has some unique setup requirements. You cannot autoregister this device. However, you can choose to initially configure it using the interactive setup utility or the CLI. This section describes how to initially configure a Content Distribution Manager using the CLI. You must configure the Content Distribution Manager before you can register any Content Engines or Content Routers with the Content Distribution Manager.

In ACNS 5.x software, Content Distribution Managers can operate in two different roles: primary and standby. The primary role is the default. You can have only one primary Content Distribution Manager that is active in your network; however, you can have any number of Content Distribution Managers operating in a standby role to provide redundancy and failover capacity. We recommend that you configure the primary Content Distribution Manager first.

**Note**

Primary and standby Content Distribution Managers must be running the same version of ACNS software. If they are not, the standby Content Distribution Manager detects this and shuts down the Centralized Management System (CMS) until the problem is corrected. We recommend that you upgrade your standby Content Distribution Manager first and then upgrade your primary Content Distribution Manager.

To configure the primary Content Distribution Manager using the CLI, follow these steps:

- Step 1** Create and initialize the management database, and enable ACNS services by using the **cms enable** global configuration command.

```
CDM-4630(config)# cms enable
```

- Step 2** Verify that services have been enabled by using the **show cms processes** EXEC command.

```
CDM-4630# show cms processes
Service cms_httpd is running
Service cms_cdm is running
CDM-4630#
```

- Step 3** Check the current running configuration by using the **show running-config** EXEC command.

```
CDM-4630# show running-config
```

A message similar to the following should appear:

```
-----
...
cms enable
!
...
-----
```

- Step 4** To save the configuration, use the **copy running-config startup-config** EXEC command.

```
CDM-4630# copy running-config startup-config
```

To configure a standby Content Distribution Manager role, follow these steps:

- Step 1** Configure the standby Content Distribution Manager role by using the **cdm role** global configuration command.

```
CDM-4630(config)# cdm role standby
```

- Step 2** Configure the standby Content Distribution Manager with the IP address of the primary Content Distribution Manager by using the **cdm ip** *{ip-address | hostname}* global configuration command.

```
CDM-4630(config)# cdm ip 10.1.1.90
```

This command associates the device with the primary Content Distribution Manager so that it can be approved as a part of the network.

- Step 3** Create and initialize the management database, and enable ACNS services by using the **cms enable** global configuration command.

```
CDM-4630(config)# cms enable
```

- Step 4** Verify that services have been enabled by using the **show cms processes EXEC** command.

```
CDM-4630# show cms processes
Service cms_httpd running
Service cms_ui running
Service cms_cdm running
```

- Step 5** Check the current running configuration by using the **show running-config EXEC** command.

```
CDM-4630# show running-config
```

A message similar to the following should appear:

```
-----
...
cms enable
!
...
-----
```

- Step 6** To save the configuration, use the **copy running-config startup-config EXEC** command.

```
CDM-4630# copy running-config startup-config
```

- Step 7** Activate the newly registered standby Content Distribution Manager in the primary Content Distribution Manager GUI in the same manner as you would activate a newly registered Content Engine or Content Router. (See the [“Activating Devices in the Content Distribution Manager GUI”](#) section on page 2-22.)

After you activate your standby Content Distribution Manager, the primary Content Distribution Manager notifies all registered Content Engines and Content Routers that a standby Content Distribution Manager exists in the system and sends each of these devices the information they need to contact the standby Content Distribution Manager when the primary Content Distribution Manager fails or becomes inactive.



Note You cannot log into the standby Content Distribution Manager GUI, and the standby Content Distribution Manager does not send configuration updates to any of the registered Content Engines or Content Routers. While it is operating as a standby Content Distribution Manager, its only function is to maintain an up-to-date copy of the primary Content Distribution Manager database. The standby accomplishes this by periodically polling the primary Content Distribution Manager for the latest configuration changes.

You have configured a primary Content Distribution Manager and an optional standby Content Distribution Manager. You are now ready to configure the network settings for all of your Content Engines and Content Routers. (See the [“Configuring the Device Network Settings”](#) section on page 2-15.)

To change a standby Content Distribution Manager to a primary Content Distribution Manager, follow these steps:

- Step 1** If your primary Content Distribution Manager is still operating, access the primary Content Distribution Manager through the console or through a Telnet Secure Shell session, and issue the following global configuration command:

```
CDM1(config)# cdm role standby
```



Note If your primary Content Distribution Manager is not operating because the hardware has failed or for any other reason, you can omit [Step 1](#).

Step 2 Access the standby Content Distribution Manager through the console or through a Telnet Secure Shell session, and issue the following global configuration command:

```
CDM1(config)# cdm role primary
```

You do not need to take any action on your Content Engines or Content Routers. These devices automatically detect the new primary Content Distribution Manager and begin receiving configuration updates from this Content Distribution Manager.



Note If you have recently made configuration changes to the primary Content Distribution Manager, we recommend that you wait at least one polling interval (System.datafeed.pollRate) before you change roles with the standby Content Distribution Manager, so that the standby has a record of the most recent configuration changes.



Note You cannot deploy a standby Content Distribution Manager behind a NAT firewall when the Content Engines and Content Routers that it manages are on the other side of the firewall.

Configuring Disk Space

If you are configuring a Content Engine for the first time, you need to create disk space for system use (sysfs), caching (cfs), streaming (mediafs), and pre-positioning (cdnfs) on the Content Engine by using the **disk config** command in EXEC mode. The sysfs should not be less than 3 GB.



Note For the Content Router or Content Distribution Manager device modes, we recommend that you allocate 100 percent of the disk space to the sysfs.

To configure disk space using the CLI, follow these steps:

Step 1 Exit configuration mode, if you have not already done so.

```
CE-507(config)# exit
CE-507#
```

Step 2 Enter the **disk config** EXEC command. For example:

```
CE-507# disk config sysfs 10% cfs 20% mediafs 10% cdnfs 60%
```

Step 3 Reload the Content Engine for the disk configuration to take effect.

```
CE-507# reload
```


**Tip**

For the new disk space configuration to take effect, you must first reboot the device; however, if you are going to change the device mode, you can wait to reboot until after you have configured the new device mode.

To update the disk configuration through the Content Distribution Manager GUI, follow the procedure in the “Updating Storage Capacity on Your Content Engines” section in the *Cisco ACNS Software Upgrade and Maintenance Guide*.

If you are attaching and configuring a Fibre Channel storage array, see the “Using a Fibre Channel Storage Array” section in the *Cisco ACNS Software Upgrade and Maintenance Guide*. Do not attempt to assign the Fibre Channel storage to the Content Engine and configure the file systems with a single reload of the Content Engine. If you do, the Fibre Channel storage assignment is recognized, but the disk configuration is not applied. An error message appears at bootup, similar to the following:

```
-----
ruby_disk:physical disk setup appears to have changed
ruby_disk:not applying 'disk config' changes. Please re-enter via CLI.
-----
```

If you encounter this error message, reenter your disk configuration and use the **reload** command on the Content Engine for the disk configuration to be applied.

Registering a Device with the Content Distribution Manager

Before a newly installed Content Engine or Content Router can be recognized by the Content Distribution Manager, it must be registered with the Content Distribution Manager. This section describes how to register a device using the standard CLI.

**Note**

To deregister a device, see the “Deleting a Content Router” section on page 13-24 or the “Deleting a Content Engine” section on page 13-12.

If you do not have a DHCP server in your network, or if for some reason your DHCP server is configured incorrectly, you can set up and register a device manually. To manually register your ACNS network devices, you must first disable autoregistration. This can be done by enabling the setup utility, or through the CLI after the device has booted by using the **no auto-register enable** command in global configuration mode.

To register a new device using the CLI, follow these steps:

- Step 1** Set the IP address or host name of the Content Distribution Manager with which the device is to be associated by using the **cdm ip** global configuration command.

```
CE-507# config
CE-507(config)# cdm ip {ip-address | hostname}
```

This command associates the device with the Content Distribution Manager so that the device can be approved as a part of the network.

After the device is configured with the Content Distribution Manager IP address, it presents a self-signed security certificate as well as other essential information, such as its IP address or host name, disk space allocation, and so forth, to the Content Distribution Manager.

- Step 2** To register the device, create and initialize the management database, and enable ACNS services, use the **cms enable** global configuration command.

```
CE-507(config)# cms enable
```

The CMS database table stores the device node configuration, allowing the node to be centrally managed from the Content Distribution Manager.

- Step 3** Use the **show cms info** EXEC command to verify that your device has been registered.

```
CE-507# show cms info
```

A message similar to the following should appear:

```
Registration information :
CDM address              = 10.1.1.90
Device Mode              = ce
Model                    = CE507
Node Id                  = 84
ce507#
```

- Step 4** To verify that the CMS is running, use the **show cms processes** EXEC command.

```
CE-507# show cms processes
```

A message similar to the following should appear:

```
Service cms_ce running
ce507#
```

- Step 5** Verify the current running configuration by using the **show running-config** EXEC command.

```
CE-507# show running-config
```

A message similar to the following should appear:

```
-----
...
cms enable
!
...
-----
```

- Step 6** To save the configuration, use the **copy running-config startup-config** EXEC command.

```
CE-507# copy running-config startup-config
```

The next step is to authorize the device to become part of the ACNS network by approving the device in the Content Distribution Manager GUI, as described in the section, “[Activating Devices in the Content Distribution Manager GUI](#).”

Initial Device Configuration Summary

In this section you accomplished the following:

- Changed the Content Engine device mode to enable one or more Content Distribution Managers
- Changed the Content Engine device mode to enable one or more Content Routers (optional)
- Configured network settings for your Content Distribution Managers, Content Engines, and Content Routers

- Configured one primary Content Distribution Manager and one or more standby Content Distribution Managers
- Configured disk space for various file systems
- Registered Content Engines and Content Routers with the Content Distribution Manager

You have completed the initial device configuration requirements. Now you must activate your Content Engines and Content Routers in the Content Distribution Manager GUI and configure any additional network interfaces. The next section (“[Initial Network Configuration Using the Standard Content Distribution Manager GUI](#)”) guides you through these steps.

Initial Network Configuration Using the Standard Content Distribution Manager GUI

If you have chosen *not* to use the Quick Start tool to initially configure a basic ACNS network, then the first task that you must perform after initially configuring your devices is to activate your devices in the Content Distribution Manager GUI.

This section describes the steps for activating devices in a centrally managed environment. The procedures for configuring other ACNS network elements using the standard Content Distribution Manager GUI are found in subsequent chapters of this guide. (See the “[Summary of Initial Configuration Tasks](#)” section on page 2-6 for chapter and section references.)

Activating Devices in the Content Distribution Manager GUI

Content Engines and Content Routers that are being newly added to the ACNS network need to be approved by the network administrator. The network administrator approves each device by making it active in the Content Distribution Manager GUI. This security feature prevents unauthorized devices from joining the network. This section discusses how to activate devices using the standard Content Distribution Manager GUI.

**Note**

Before you can access the Content Distribution Manager GUI, you must use the CLI to register yourself as a user. Use the **username** global configuration command to configure your name and password for GUI access.

To activate a Content Engine or Content Router using the Content Distribution Manager GUI, follow these steps:

- Step 1** Open the Content Distribution Manager GUI from your web browser by entering the URL or IP address of the Content Distribution Manager. For example, from your browser, enter one of these URLs:
- `https://name_of_Content_Distribution_Manager:8443`**
- or
- `https://IP_address_of_Content_Distribution_Manager:8443`**
- A window appears, requesting your username and password.
- Step 2** Enter the administrator username **admin** and the password **default**, and click **OK**. The Cisco Application and Content Networking System Home window appears.

- Step 3** Choose **Devices > Devices**. A list of all your registered Content Engines and Content Routers appears.
- Step 4** Click the **Edit** icon next to the device name. The Device Home window appears.
- Step 5** To display the entire table of contents, click the **Show All** button above the Contents pane.
- Step 6** In the Contents pane choose Device Activation. The Device Activation window appears. (See [Figure 2-3](#).)

Figure 2-3 Device Activation Window

The screenshot shows the Cisco Application and Content Networking System (ACNS) GUI in a Microsoft Internet Explorer browser. The main navigation bar includes 'Devices', 'Services', and 'System'. The 'Devices' tab is active, and the 'Device Activation, CONTENTENGINE' window is displayed. The window is divided into several sections:

- Contents**: A sidebar on the left with a tree view showing 'Device Home', 'Device Activation', 'Assignments', 'Prepositioning', 'Request Routing', 'Request Processing', 'Applications', 'General Settings', and 'Device Monitoring'. 'Device Activation' is selected.
- General Configuration**: Contains fields for 'Name' (CONTENTENGINE), 'Status' (Inactive), 'Activate' (checkbox), 'Replaceable' (checkbox), and 'cdnfs' (0 MB). There is a checkbox for 'Set Default Coverage Zone File' which is checked.
- Locality**: Contains a 'Select a Location' dropdown menu (set to 'ContentEngine-location'), a 'Create a new Location' checkbox, and a 'Parent of the new Location' dropdown menu (set to 'None').
- Request Routing**: Contains a checkbox for 'Enable Request Routing' and a 'Coverage Zone File' dropdown menu (set to 'None').
- NAT Configuration**: Contains a checkbox for 'Use CE's primary IP Address' (checked), a 'NAT Address' text field, and a 'Port' field (set to 443).
- Comments**: A text area with 'n/a'.

At the bottom, there is a 'Note: * - Required Field' and a 'Change submitted.' message. The 'Submit' and 'Cancel' buttons are at the bottom right.

- Step 7** Check the **Create a New Location** check box to create a default location for this Content Engine. Locations need to be configured before you can activate Content Engines and Content Routers and bring them online in the ACNS network.



Note This field is visible only if the Content Engine is in Inactive status. When checked, a default location for an inactive Content Engine is created. This option automatically creates a new location named <CE-name>-location and assigns the Content Engine to that location.

Alternatively, choose a location from the Location drop-down list. The list of locations that have been created using the Locations window are displayed in this drop-down list. If the location that has been chosen for the Content Engine contains a parent location, then the same location tree hierarchy is applied to the Content Engine.

This option allows you to choose an already created location. (If you have not already created a location, see the [“Working with Device Locations”](#) section on page 13-5.)


- Step 8** Choose a parent location for the default location from the Parent of the new Location drop-down list. This option allows you to choose an already created location as the parent for the newly created default location.
- Step 9** To activate the Content Engine or Content Router, check the **Activate** check box.
- Step 10** To save the settings, click **Submit**. When the device becomes active, the status changes from “Inactive” to “Online.”
- Other Content Engine properties can be configured now or at a later time. (To enable routing or to modify the default bandwidth settings, see the [“Modifying Content Engine Properties”](#) section on page 13-9.)

Activating All Inactive Content Engines

To activate all unactivated Content Engines as a group, follow these steps:


- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices listing window appears.
- Step 2** Click the **Activate all inactive CEs** icon. (See [Figure 2-4](#).) The Activate All Inactive Content Engines window appears.

Figure 2-4 Devices Listing Window



- Step 3** In the Activate All Inactive Content Engines window (see [Figure 2-5](#)), choose an existing location for all unactivated Content Engines by clicking the **Select an existing location for all inactive CEs** radio button. Then choose a location from the drop-down list.

Alternatively, choose to create a new location for each inactive Content Engine by clicking the **Create a new location for each inactive CE** radio button. Then specify a parent location for all newly created locations by choosing a location from the Select a parent location for all newly created locations drop-down list.

Figure 2-5 **Activate All Inactive Content Engines Window**

Step 4 To save the settings, click **Submit**.

Where to Go Next

You have completed the initial ACNS device configuration and have set up a basic ACNS network. Your Content Engines and Content Routers (optional) are identified on the network and are communicating with the Content Distribution Manager. Basic software services for conventional caching, caching and streaming, and pre-positioning content are enabled, and your disk file systems are configured to allow the various types of content to be stored on the Content Engine. Now you must finish setting up the request routing infrastructure so that client requests for content can be routed to the origin server where the content resides. Proceed to the next chapter, [Chapter 4, “Setting Up Content Request Routing in the ACNS Network.”](#)



Using the Content Distribution Manager GUI

The Content Distribution Manager is the central location from which much of the work of creating and managing ACNS networks and hosted content can be controlled. Functions related to managing your ACNS network are done using the Content Distribution Manager graphical user interface (GUI). All modifications made through the Content Distribution Manager GUI are propagated to the Content Routers and Content Engines that make up your ACNS network.

This chapter explains how to use the Content Distribution Manager GUI. It contains the following sections:

- [Launching the Content Distribution Manager GUI, page 3-1](#)
- [Content Distribution Manager GUI Primary Feature Tabs, page 3-2](#)
- [Understanding the Interface, page 3-3](#)
- [Using Online Help, page 3-11](#)
- [Logging On Following a Session Timeout, page 3-11](#)
- [Exiting the Content Distribution Manager GUI, page 3-12](#)

Launching the Content Distribution Manager GUI

To launch the Content Distribution Manager GUI from your system, follow these steps:


-
- Step 1** In your web browser, enter the URL or IP address for the Content Distribution Manager. For example, enter this URL:
- `https://Name_of_Content_Distribution_Manager:8443`**
- Alternatively, enter this IP address:
- `https://IP_address_of_Content_Distribution_Manager:8443`**
- Step 2** The Security Alert window appears. Click **Yes** to accept the security certificate. The Login window appears.
- Step 3** Enter your username in the Username field. Enter your password in the Password field, and click **Login**.
- User name: **admin**
Password: **<password>**

**Note**

The default username is **admin** and the default password is **default**. If the defaults have been changed by another Content Distribution Manager administrator, you need to obtain the new username and password.

The System Home window appears. (See [Figure 3-1](#).)

Figure 3-1 ACNS Content Distribution Manager GUI System Home Window



The ACNS 5.5 system home window provides current configuration information about devices and service baseline settings.

Content Distribution Manager GUI Primary Feature Tabs

[Table 3-1](#) describes the three primary feature tabs and their associated functions of the Content Distribution Manager GUI.

Table 3-1 **Content Distribution Manager GUI Feature Tabs**

Feature Tab	Description
Devices	Lets you configure and manage Content Distribution Managers, Content Engines, Content Routers, and device groups. Lets you create and manage ACNS network elements such as locations and location trees. Lets you view system logs and streaming and routing statistics for ACNS network devices.
Services	Lets you configure three service areas: Web, Video, and Platform. • Web service lets you create and modify ACNS network elements such as channels, content providers, websites, and multicast clouds. • Video service lets you configure programs and program addresses. • Platform service lets you modify baseline settings for a baseline group and customize the table of contents so that only certain device group configuration windows are available to the user. Lets you define a baseline configuration for each service.
System	Lets you configure routing and Content Distribution Manager configurations, and add a software update file. Lets you configure and manage user accounts, roles, domains, and passwords. Lets you view audit trail logs and system messages.


Understanding the Interface

This section describes the Content Distribution Manager GUI main window elements, the navigation scheme, and the taskbar elements.

Window Elements

The Content Distribution Manager GUI contains several window elements with which you need to familiarize yourself. (See [Figure 3-2](#).) (See [Table 3-2](#) for a description of each element.)

Figure 3-2 Content Distribution Manager GUI Elements



1	Contents pane	5	Window
2	Navigation bar	6	System status bar
3	Tabs	7	Submit and Cancel buttons
4	Taskbar	8	Tools

Table 3-2 Content Distribution Manager GUI Elements

Element	Description
Contents pane	Displays available suboptions, if required.
Navigation bar	Displays the menu options available for the selected tab. The Services tab and the System tab have two levels of menu options in the navigation bar. (See Figure 3-3.)
Tabs	Provides access to all the Content Distribution Manager product functionality. Tabs include Devices, Content, and System. (See Table 3-1 for a description of each tab.)
Taskbar	Displays available tasks for the selected option. (See Table 3-3 for a description of all the Content Distribution Manager GUI icons. Some of these icons appear in the taskbar.)
Window	Displays the area in which you perform application tasks.

Table 3-2 *Content Distribution Manager GUI Elements (continued)*


Element	Description
System status bar	Displays device and content alarms and presents the overall device and content health of the system.
Submit button	Saves the entries made in the window.
Cancel button	Returns to the previous window without saving entries. Use this button instead of the browser Back button.
Tools	Contains the following links: Home—Returns you to the System Home window. Quick Start—Launches the Quick Start tool. Help—Opens a new window that displays context-sensitive help information. Logout—Logs you out of the Content Distribution Manager.

**Note**

Use only the Content Distribution Manager GUI elements, which are described in [Table 3-2](#), to navigate within the GUI. Web browser buttons, such as the Back and Refresh buttons, are not supported and do not allow you to navigate to the correct GUI window.

Navigation

To navigate within the Content Distribution Manager GUI, begin by clicking one of three primary feature tabs: Devices, Services, or System. From each tab option, choose a first-level menu option in the upper navigation bar and a second-level menu option in the lower navigation bar, if available. Navigational directions in procedures are written following this sample scheme: **Services > Web > Channels**. (See [Figure 3-3](#).)

Figure 3-3 *Two-Tiered Options Menu in the Services Navigation Bar*

In most cases, after you reach the third navigational level (or the second level when only two levels exist), a basic listing window appears. In the listing window you can create a new element to add to the list by clicking the **Create New** icon in the task bar, or you can edit existing elements by clicking the **Edit** icon next to the element that you want to modify. A configuration window appears.

The configuration window displays a Contents pane on the left for further navigation. An **Expand** button expands the Contents pane to display all menu items that are nested under the top-level items. (See [Figure 3-4](#).)

Figure 3-4 Contents Pane for a Channel Configuration Window

Cisco Application and Content Networking System - Microsoft Internet Explorer

File Edit View Favorites Tools Help

CISCO SYSTEMS Application and Content Networking System

System Status Devices: 2 Devices, Critical Content: 18 Channels, Minor Home | Help | Logout

Devices Services System

Web Video Platform Channels Content Provider Websites Multicast Clouds Baseline Settings

Expand All

Contents

- Definition
- Channel Content
- Assign Multicast Cloud
- Assign Content Engines
- Assign Device Groups
- List all assigned Content Engines
- Replication Status
- Tools

Modifying Channel, chan2

Channel information

Name:

Request Routed FQDN:

Origin Server:

Content Provider: [Edit Content Provider](#)

WebSite: [Edit Website](#)

Live Channel:

Channel Quota: MB

Acquisition and Distribution Properties

Distribution Priority: Use null cipher for Distribution: ☐

* ☐ Multicast only ☐ Multicast/Unicast ☒ Unicast only

Root CE failover/fallback grace period: mins Never ☐ If there is only one CE in root location or if 'Never' is selected, failover/fallback will not occur.

Use system-wide settings for QoS for unicast data: ☒ System-wide QoS settings page is under the Systems Tab.

QoS value for unicast data: or Select DSCP from the drop-down or enter decimal value in the text box.

Comments

[Submit](#) [Cancel](#)

The Device Home Window

In the Device Home window, the Contents pane displays menu items in two views: basic or extended.


In basic view, the Contents pane displays menu items that are essential for basic configuration of the device only. In basic view, the toggle button reads **Show All**.

In extended view, the Contents pane displays all top-level menu items that are necessary for advanced configurations. To display extended view, click the **Show All** button. In extended view, this toggle button reads **Show Basic**. (See [Figure 3-5](#).)



Note

This toggle button is available for the Device table of contents only.

Figure 3-5 Device Table of Contents in Extended Mode

From the Device Home window, you can also perform the following tasks without needing to navigate through the table of contents:

- Activate the device, and set the device location (only available if the device is inactive)
- Modify channel assignments
- Modify device group assignments
- Upgrade the device software
- Telnet to the device CLI
- Assign the device to baseline groups
- Modify baseline group settings
- View detailed bandwidth efficiency gain report
- View detailed bytes served report

Taskbar Icons

Table 3-3 lists and describes the Content Distribution Manager icons that you can use to perform specific tasks in the various configuration windows.

Table 3-3 *Content Distribution Manager GUI Icons*

Taskbar Icon	Function
	Activates all inactive Content Engines
	Alerts you to a new system message
	Displays table of contents
	Displays all playlists
	Displays all devices
	Deactivates the device
	Runs Content Engine Cache GUI
	Updates application statistics
	Forces refresh of replication information
	Forces full database update
	Lets you view read-only items

Table 3-3 *Content Distribution Manager GUI Icons (continued)*

Taskbar Icon	Function
	Creates new items
	Edits items
	Deletes items
	Adds a content item for acquisition
	Deletes selected items
	Manages between host and proxy servers for content acquisition
	Saves to disk
	Exports a table to a comma-separated values (CSV) file
	Creates a filter table
	Displays a graph
	Applies the default settings to the device
	Overrides the group settings on the device

Table 3-3 *Content Distribution Manager GUI Icons (continued)*

Taskbar Icon	Function
	Lets you view all items
	Refreshes the table
	Reboots the device
	Shuts down the device
	Prints the current window
	Copies a program
	Deletes all conditions from an IP ACL
	Assigns all items to the entity
	Removes all items from the entity
	Identifies services for roles
	Switches the baseline groups
	Forces the group baseline settings

Table 3-3 *Content Distribution Manager GUI Icons (continued)*

Taskbar Icon	Function
	Indicates that the current transaction was successfully completed
	Indicates that user input is invalid and that the transaction did not finish

Using Online Help

ACNS 5.x software offers online help for the Content Distribution Manager GUI. There are several ways to access online help, depending on the information you want and the method you find most convenient.

If you want to read one or more help topics, click the **Help** link. A separate browser window appears that contains online help. From the online help Contents tab, click the name of a topic to display information about that topic. For information about a specific window, click the **Help** link while viewing that window.

Use the browser Back and Forward buttons to navigate between online help windows.

**Note**

Although the web browser Back and Forward buttons are not supported when navigating in the Content Distribution Manager GUI, these buttons *are* supported when navigating within online help.

Logging On Following a Session Timeout

The Content Distribution Manager web interface is equipped with a timeout feature that terminates idle sessions after a set length of time. (See the [“Modifying System Default Properties”](#) section on [page 20-18](#) for instructions on changing the length of time that must pass before your session times out.)

When a session has timed out, you immediately lose access to the Content Distribution Manager web interface, which is replaced with this message:

Your session has expired.

To log back on to the Content Distribution Manager, follow these steps:

Step 1

Close all web browser windows that were used to access the Content Distribution Manager by clicking the **Close** button to close the active window, or click the **File** menu and choose **Close** or **Exit**.

**Note**

Often when you log on to the Content Distribution Manager, the web interface is loaded into a separate instance of the web browser from the one used to point to the Content Distribution Manager URL. Make sure that both instances of the browser are terminated before attempting to log back on to the Content Distribution Manager.

- Step 2** Open a new instance of your preferred web browser and point it to the URL of your Content Distribution Manager. You are prompted to log on to the Content Distribution Manager web interface.
-

Exiting the Content Distribution Manager GUI

To exit the Content Distribution Manager GUI, follow these steps:

-
- Step 1** Click **Submit** to save any changes you made, or click **Cancel** to cancel the changes.
- Step 2** In the upper right corner of the Content Distribution Manager GUI, click **Logout**. You are prompted to confirm your decision.
- Step 3** Click **OK**. The Content Distribution Manager GUI returns to the Login window.
- Step 4** To exit the browser, choose **Close** from the browser File menu, or click the browser **Close** button.
-



PART 2

Configuring the ACNS Network for Content Services





Setting Up Content Request Routing in the ACNS Network

Cisco ACNS 5.5 software supports three methods of request handling to deliver content to the end user. These request handling methods are described in the following sections:


- [Configuring WCCP Transparent Routing, page 4-2](#)
- [Configuring Direct Proxy Routing, page 4-32](#)
- [Configuring Content Request Routing Using a Content Router or Routing Content Engine, page 4-41](#)
- [Troubleshooting Content Router Configurations, page 4-59](#)
- [Where to Go Next, page 4-60](#)



Note

For complete syntax and usage information for the CLI commands used in this chapter, refer to the *Cisco ACNS Software Command Reference, Release 5.5* publication.

[Figure 4-1](#) shows the basic tasks involved in configuring each of the three request handling methods.

Figure 4-1 *Choosing a Routing Method***Legend:**

 This step is performed externally to the ACNS device or central Content Distribution Manager GUI.

Configuring WCCP Transparent Routing

In the WCCP transparent routing method, requests for content made to an origin server are intercepted by a WCCP-enabled router. The WCCP-enabled router transparently redirects the request to a Content Engine containing the content. This type of transparent redirection allows for traffic interception on any port for traffic that traverses this router. WCCP contains many fail-safe mechanisms to ensure that the request interception remains entirely transparent to the end user.

To use a WCCP-enabled router, an IP address must be configured on the interface connected to the Internet, and the interface must be connected to the Content Engine. A Content Engine and a WCCP-enabled router cannot be separated by a firewall. The firewall handles only packet traffic toward the origin web server and does not handle packet traffic sent to the client by the Content Engine on behalf of the server.


You can configure only a single router to run the web cache service (WCCP Version 1 service). To use WCCP Version 1 with the Content Engine, you must point the Content Engine to a designated home router.

However, WCCP Version 2 enables a series of Content Engines, called a Content Engine cluster, to connect to multiple routers. This feature provides redundancy and a more distributed architecture for instances when a Content Engine needs to connect to a large number of interfaces. WCCP Version 2 also balances traffic load across a cache cluster and ensures fault-tolerant and fail-safe operation. As Content Engines are added to or deleted from a cache cluster, the WCCP-aware router or switch dynamically adjusts its redirection map to reflect the currently available caches.

When operating with WCCP Version 2, the Content Engine performs a clean shutdown after reboot, enabling WCCP Version 1, or disabling WCCP Version 2. A clean shutdown prevents broken TCP connections. (See the [“About WCCP Clean Shutdown”](#) section on page 4-9.)

Figure 4-2 shows how requests for content are routed using the WCCP transparent interception method.


Figure 4-2 Request Interception Using WCCP



1. A user requests a web page from a browser.
2. The WCCP-enabled router intercepts the request, and based on the TCP destination port number, the router determines whether it should transparently redirect the request to a Content Engine. Access lists are used to control which requests can be redirected.
3. If the Content Engine does not have the requested content, it sets up a separate TCP connection to the origin server to retrieve the content (3a in Figure 4-2). The content returns to, and is stored on, the Content Engine (3b in Figure 4-2).
4. The Content Engine sends the content to the client. Upon subsequent requests for the same content, the Content Engine transparently fulfills the request from its local storage.

There are some significant advantages to using a WCCP-enabled router for request routing interception in transparent mode:

- No end user configuration—Users do not have to point their browsers to the Content Engine.
- Fail-safe—Content Engines are automatically fault-tolerant and fail-safe. Any Content Engine failure does not cause denial of service to the end user.
- Scalable—Cache services can be scaled by deploying multiple Content Engines.
- Automatic bypass—Sites which depend on end user authentication or which fail to conform to HTTP standards will automatically bypass a transparent Content Engine.

Figure 4-3 Packet Flow Diagram

In [Figure 4-3](#), if the Content Engine does not have the data, a cache miss occurs, and the Content Engine sends the request to the origin server. As the Content Engine receives the data from the origin server, it saves the data and forwards the data to the client.

About WCCP Services

WCCP uses the concept of a service group to define caching related services for a WCCP-enabled router and WCCP-enabled Content Engines in a cluster. Service groups are identified by a service group number. Each WCCP service group number specifies a router list, single port list (containing up to eight ports), application type, hash parameters, password, and weight (for load balancing).

[Table 4-1](#) lists the service group numbers and describes the services supported by a WCCP-enabled router.

Table 4-1 WCCP Service Groups

Service Group Number	Description of Services
0	Web cache
60	FTP
70	HTTPS
80	HTTP, RTSP
81	MMST
82	MMSU
90–97	Dynamic or user-configurable
98	Custom
99	Reverse proxy



Note

All service group numbers listed in [Table 4-1](#) except for web cache services (service group 0) require WCCP Version 2.

A WCCP-enabled Content Engine supports various Internet protocols depending on the style of the request URL. If the WCCP-enabled Content Engine receives a server-style URL, only HTTP and RTSP protocols are supported (if the RTSP user agent criteria is met). If the WCCP-enabled Content Engine receives a proxy-style URL, the HTTP, HTTPS, FTP, and RTSP protocols are supported, as long as the respective proxy services are configured on the Content Engine. Proxy-style request URLs include the protocol and host name, whereas server-style request URLs do not. The RTSP protocol is supported for both server-style and proxy-style requests.

For proxy-style requests, the Content Engine supports up to eight incoming ports each for FTP, HTTPS, HTTP, MMS, and RTSP proxy modes. The incoming proxy ports can be the same ports that are used by transparent mode services. The incoming proxy ports can be changed without stopping any WCCP services running on the Content Engine or on other Content Engines in the Content Engine farm.


The Content Engine WCCP implementation currently allows global settings that apply to all WCCP services, such as healing parameters, slow start, and others. The multiple service model does not change that, and the settings in question remain global for the whole WCCP system. (See the [“Configuring WCCP General Settings for the Content Engine”](#) section on page 4-7.)

About Dynamic WCCP Redirection Services

You can configure the Content Engines to handle up to eight user-configurable or dynamic WCCP redirection services (service group numbers 90 to 97) on a Content Engine, provided that the services are also configured on the router.

With 8 dynamic services using a maximum number of 8 ports each, the maximum number of ports that can be specified for transparent redirection is 64. In Figure 4-4, the two Content Engines on the left handle only HTTP traffic through port 80 and are defined as members of service group 90. The two Content Engines on the right handle only Microsoft Media Server (MMS) requests and are defined as members of service group 91.

Figure 4-4 WCCP Service Group Feature



About WCCP Service for HTTPS

WCCP service for HTTPS traffic (service number 70) instructs the corresponding router to intercept port 443 TCP traffic (port 443 is the default port) and forward it to the Content Engine. This service has a special filtering feature that a normal WCCP service does not have; it uses an accept list to selectively redirect traffic to an application instead of redirecting all traffic to related applications. The accept list is a list of destination IP addresses. When the destination server's IP address in a client request matches any of the destination IP addresses in the list, WCCP redirects HTTPS traffic to the proper listening application. All other traffic is redirected to the corresponding router (normal bypass behavior). If bypass lists are configured to permit traffic from specified sources to bypass the Content Engine, traffic is not intercepted but passes directly to the origin server. If an address is present in both the accept list and bypass list, the Content Engine can return traffic to the WCCP-enabled router or switch and inform the router or switch to forward the packets as if the Content Engine were not present.

For HTTPS filtering feature to work, ensure that the **Accept all HTTPS traffic** check box (**Request Routing > WCCP > General Settings**) in the Content Distribution Manager GUI is checked so that all HTTPS traffic will be intercepted by the Content Engine. The Content Engine can negotiate an SSL connection with the client and serve HTTPS requests if the corresponding destination HTTPS server is configured on the Content Engine. Otherwise, HTTPS traffic is tunneled through the Content Engine

(Content Engine passes traffic from client to server without modifying anything, except for filtering). When the SmartFilter software or Websense filtering server is enabled on the Content Engine, the Content Engine will send URLs corresponding to tunneled HTTPS traffic to the filtering server. Filtering servers can decide to block or pass the traffic. If traffic needs to be blocked, the Content Engine closest to the connection with the client and the request is not served.

About WCCP Custom Web Cache Service

Custom web cache service (service number 98) causes the Content Engine to establish WCCP Version 2 redirection services automatically with a Cisco router on a user-specified port number. The Content Engine then performs transparent web caching for all HTTP requests over that port while port 80 transparent web caching continues without interruption. For custom web caching, service 98 must be enabled on the routers. WCCP Version 1 does not support custom web caching.

Transparent caching on ports other than port 80 can be performed by the Content Engine when WCCP is not enabled or when client browsers have previously been configured to use a legacy proxy server.

Configuring WCCP General Settings for the Content Engine

To use WCCP, the Content Engine must be properly configured. Use the following configuration guidelines to help you configure the Content Engine for edge interception using a WCCP-enabled router:

- Versions of software on the Content Engines must be compatible with those installed on the router.
- The Content Engines must not have their packets encrypted or compressed and should be part of the “inside” Network Address Translation firewall if one is present.
- Placing a Content Engine beyond a web cache redirect-enabled interface and along the route to the server will not cause the IP route cache to be populated with an entry.
- To use a WCCP-enabled router, an IP address must be configured on the router interface that is connected to the Internet, and this interface must be visible to the Content Engine on the network.



Note

ACNS 5.x software currently emphasizes the use of WCCP Version 2. However, WCCP Version 1 is still supported by the ACNS 5.x release.

To configure general WCCP settings for the Content Engine, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears, listing all the device types that are configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure WCCP general settings. The Device Home for Content Engine window appears.
- Step 3** To display the entire table of contents, click the **Show All** button above the Contents pane.
- Step 4** In the Contents pane, choose **Request Routing > WCCP > General Settings**. The WCCP Configuration Settings for Content Engine window appears. [Table 4-2](#) describes the fields in this window and provides the corresponding CLI global configuration commands.
- Step 5** From the **WCCP Version** drop-down list, choose the version of WCCP that the Content Engine should use.

Both of the WCCP versions allow transparent caching of web content. It is not necessary to disable WCCP Version 1 before enabling WCCP Version 2, or to disable WCCP Version 2 before enabling WCCP Version 1. Be sure the routers used in the WCCP environment are running a software version that supports the WCCP version configured on the Content Engine.

- Step 6** In the Home Router IP Address field, specify an IP address to configure a WCCP Version 1 router. To use WCCP Version 1 with the Content Engine, you must also point the Content Engine to a designated home router. This may also be the address of the IP default gateway.



Note Make sure that WCCP Version 1 is enabled on the router.

- Step 7** To keep the TCP flow intact, and to avoid overwhelming Content Engines when they come up or are reassigned new traffic, check the **Enable Flow Redirection** check box.
- This configuration works with WCCP Version 2 only. This feature also has a slow start mechanism which enables a Content Engine to manage an amount of data that is appropriate for its capacity.
- Step 8** To enable the Content Engine to send the client IP address for authentication on the origin web servers during the transparent redirection process, check the **Spoofing Client IP** check box. (For more information about IP spoofing, see the [“About Client IP Spoofing”](#) section on page 4-9.)
- Step 9** To enable the slow start capability on the Content Engine, check the **Slow Start** check box. (For more information about the WCCP slow start feature, see the [“About WCCP Slow Start”](#) section on page 4-9.)
- Step 10** In the Shutdown Delay field, specify the maximum amount of time (in seconds) the Content Engine waits to perform a clean shutdown of WCCP. (For more information about WCCP shutdown delay, see the [“About WCCP Clean Shutdown”](#) section on page 4-9.)
- Step 11** To enable the Content Engine to accept all HTTPS traffic by default, check the **Accept all HTTPS traffic** check box. Otherwise, the Content Engine will accept the traffic only if the HTTPS server is configured. (For more information about the WCCP service for caching HTTPS traffic, see the [“About WCCP Service for HTTPS”](#) section on page 4-6.)
- Step 12** To save the settings, click **Submit**.

Table 4-2 WCCP General Settings

GUI Parameter	Function	CLI Command
WCCP Version	Enables WCCP Version 1 or Version 2.	wccp version {1 2}
Home Router IP Address	IP address to configure a WCCP Version 1 router.	wccp home-router <i>ipaddress</i>
Enable Flow Redirection	Keeps the TCP flow intact.	wccp flow-redirect enable
Spoofing Client IP	Enables the Content Engine to send out the client’s IP address for authentication purposes on the origin web servers during transparent redirection process.	wccp spoof-client-ip enable
Slow Start	Enables the slow start capability.	wccp slow-start enable
Shutdown Delay	Maximum amount of time (in seconds) the Content Engine waits to perform a clean shutdown of WCCP.	wccp shutdown max-wait <i>seconds</i>
Accept all HTTPS traffic	Enables the Content Engine to accept all HTTPS traffic by default.	wccp https-cache accept-all

About Client IP Spoofing

In transparent caching, when the Content Engine contacts the origin web server, it uses the Content Engine's own IP address instead of the client's IP address, on behalf of which the Content Engine is making the request. With IP spoofing, the transparent redirection process enables the Content Engine to send out the client's IP address for authentication purposes on origin web servers.

When the Content Engine is performing normal WCCP caching without IP spoofing, the Content Engine connects the origin servers using its own IP address. Routers identify the requests as coming from the Content Engine using its own source IP address in the request and forwards the requests out onto the WAN. The router forwards requests coming from the clients on the LAN to the Content Engine.

When configured for IP spoofing, the Content Engine connects to the origin server using the client IP address instead of its own IP address. At this point, the router cannot identify requests coming from the Content Engine because the source IP address of the request is not that of Content Engine. To prevent these Content Engine requests from being returned to the Content Engine unresolved, the **ip wccp redirect exclude in** command must be applied to the WCCP-enabled router interface to which the Content Engine is connecting.

The **ip wccp redirect exclude in** command prevents the WCCP-enabled router from intercepting any requests coming in on this interface. To allow interception of client requests, the **ip wccp redirect out** command must be configured on those interfaces connected to the client and the WAN.

About WCCP Slow Start

Within a cluster of Content Engines, TCP connections are redirected to other Content Engines as units are added or removed. A Content Engine can be overloaded if it is reassigned new traffic too quickly or introduced abruptly into a fat pipe.

WCCP slow start performs the following tasks to prevent a Content Engine from being overwhelmed when it comes online or is reassigned new traffic:

- TCP flow protection when WCCP Version 2 is enabled and a Content Engine is introduced into the cluster
- TCP flow protection when WCCP Version 2 is disabled and a Content Engine is leaving the cluster
- Load assignment to the Content Engine in slow increments rather than a full load at bootup

Slow start is applicable only in the following cases:

- Initial bootup when there is no Content Engine yet present in the server farm
- When a new Content Engine is added to a cluster that is not handling the full load; for example, when there are some buckets that are being shed by the cluster

In all other cases slow start is not necessary and all the Content Engines can be assigned their share of traffic right away.

About WCCP Clean Shutdown

To prevent broken TCP connections, the Content Engine performs a clean shutdown of WCCP after a reload or WCCP version change. The Content Engine does not reboot until either all connections have been serviced or the configured maximum wait interval has elapsed.

During a clean shutdown, the Content Engine continues to service the flows it is handling but starts to bypass new flows. When the number of flows goes down to zero, the Content Engine takes itself out of the cluster by having its buckets reassigned to other Content Engines by the lead Content Engine. TCP connections can still be broken if the Content Engine crashes or is rebooted without WCCP being cleanly shut down. The clean shutdown can be aborted while in progress.

Configuring WCCP Service Settings for the Content Engine

To configure WCCP services, you need to configure four sets of settings:

- WCCP Service
- Dynamic Service Settings
- Load Balancing Hash
- Other Settings

To configure WCCP service settings for the Content Engine, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears, listing all device types that are configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure WCCP service settings. The Device Home for Content Engine window appears.
- Step 3** In the Contents pane, choose **Request Routing > WCCP > Services**. The WCCP Service Settings for Content Engine window appears.



Note The Content Engine settings or its associated device group settings will be applied to this window depending on the choice that you have made in the WCCP General Settings window.

- Step 4** To create a new service, configure router lists, and associate ports with WCCP Version 2 dynamic service, click the **Create New WCCP Service Setting** icon in the taskbar. The Creating New WCCP Service window appears.

You can configure a maximum of seventeen services. If seventeen services have already been defined, the icon is not visible in the taskbar.



Note All settings for a particular WCCP service can be configured only after it has been associated with a router list.

- Step 5** Under the WCCP Service section, follow these steps:
 - a. Choose the type of WCCP service from the Service Type drop-down list. (See [Table 4-3](#) for a description of the service types.)

Table 4-3 WCCP Service Types

Service Type	Description of Services	CLI Command
DNS	Allows a WCCP-enabled router to intercept and redirect DNS packets to a boomerang server, which will then accept the intercepted packets and process them, including reinserting the packets into the request stream.	wccp dns
FTP Native	Enables transparent interception of FTP traffic with WCCP Version 2.	wccp ftp-native
HTTPS	Enables WCCP flow redirection to a Content Engine configured as an HTTPS server.	wccp https
RTSP	Configures WCCP Version 2 Real-Time Streaming Protocol (RTSP) transparent interception.	wccp rtsp
Windows Media	Enables WCCP Version 2 Windows Media caching service.	wccp wmt
Windows Media RTSPU	Enables Windows Media RTSPU (port 5005) transparent interception.	wccp wmt-rtspu
Reverse Proxy	Enables WCCP Version 2 reverse proxy service. When WCCP reverse proxy is enabled, the router does load balancing in a cluster based on the source IP address (for example, the client's browser IP address).	wccp reverse-proxy
WebCache	Enables web cache service with WCCP Version 2. With web cache service, the router balances the traffic load within a Content Engine cluster based on the destination IP address (for example, web server IP address).	wccp web-cache
Custom WebCache	Enables the Content Engine to accept redirected HTTP traffic on a port other than 80.	wccp custom-web-cache
Dynamic Service Number (90–97)	Enables up to eight dynamic WCCP redirection services on the Content Engine.	wccp service-number

- b. Associate a router list with the WCCP service by choosing one from the Router List drop-down list. Only configured router lists are displayed in the drop-down list. Router lists can be configured, modified, and viewed by following the links in the Creating New WCCP Service window. For more information about router lists, see these sections:

- [Creating WCCP Router Lists, page 4-14](#)
- [Modifying WCCP Router Lists for the Content Engine, page 4-14](#)
- [Viewing WCCP Router Lists for the Content Engine, page 4-15](#)

- Step 6** Under the Dynamic Service Settings section, follow these steps. All settings under this section can be configured for dynamic web cache services only.
- a. To associate ports with specific WCCP Version 2 dynamic services, choose the port list number from the Port List drop-down list. All configured port lists will be displayed here. You can associate ports with user-configurable web cache services 90–97 only. For all other services, the Port List drop-down list is disabled.
 - b. To configure a port list for WCCP Version 2 dynamic services or if no port list already exists, click the **New Port List** button. The WCCP Port List Settings window appears. (For more information about creating port lists, see the [“Configuring WCCP Port Lists for the Content Engine”](#) section on page 4-17.)
 - c. To add or delete a port number to a configured port list, click the **Edit Port List** button. (For more information about modifying port lists, see the [“Modifying WCCP Port Lists for the Content Engine”](#) section on page 4-18.)
 - d. From the Application drop-down list, choose the application running on the Content Engine to which intercepted traffic must be redirected. (See [Table 4-4](#) for a description of the options in the drop-down list.)

Table 4-4 Application Options

Option	Description
cache	Redirects traffic to the caching application running on the Content Engine.
https-cache	Redirects traffic to the HTTPS caching application running on the Content Engine.
streaming	Redirects traffic to the streaming media application running on the Content Engine.

- e. To match the source port for redirection of traffic, check the **Match Source Port** check box.

- Step 7** Under the Load Balancing Hash section, follow these steps. The hash assignment method of load balancing cannot be defined for FTP, reverse proxy, and web cache WCCP Version 2 services. (See the [“About Load Balancing”](#) section on page 4-13.)

- a. To define the load-balancing hash of the destination IP address, check the **Destination IP** check box. This method is the default method of hash assignment for the following WCCP Version 2 services:
 - RTSP
 - Windows Media
 - Windows Media RTSPU
 - Custom web cache
 - Dynamic services

- b. To define load-balancing hash of the source IP address, check the **Source IP** check box. This method is the default method of hash assignment for HTTPS cache service.
- c. To define load-balancing hash of the destination port, check the **Destination Port** check box.
- d. To define load-balancing hash of the source port, check the **Source Port** check box. This method is the default method of hash assignment for DNS service.

Step 8 Under the Other Settings section, follow these steps:

- a. To force WCCP to use the configured assignment method only, check the **Use Selected Assignment Method** check box. This setting can be configured only for DNS, FTP native, HTTPS cache, reverse proxy, custom web cache, web cache, and Windows Media services. When applied, either of the two load balancing methods, hash assignment or mask assignment, can be used.
- b. To permit the Content Engine to receive transparently redirected traffic from a WCCP Version 2-enabled switch or router, if the Content Engine has a Layer 2 connection with the device, and the device is configured for Layer 2 redirection, check the **Layer2 Redirection** check box.
- c. In the Password field, specify the password to be used for secure traffic between the Content Engines within a cluster and the router for a specified service. Be sure to enable all other Content Engines and routers within the cluster with the same password.

Passwords must not exceed 8 characters in length. Reenter the password in the Confirm Password field.
- d. In the Weight field, specify the weight parameter that represents a percentage of the total load redirected to the Content Engine (for example, a Content Engine with a weight of 30 receives 30 percent of the total load). If the total of all weight parameters in a Content Engine cluster exceeds 100, the percentage load for each Content Engine is recalculated as the percentage that its weight parameter represents of the combined total.
- e. In the Port field, specify the port number other than port 80 that will accept redirected HTTP traffic (configurable only for custom web cache service). (See the [“About WCCP Custom Web Cache Service”](#) section on page 4-7.)

Step 9 To use the mask method for Content Engine assignment, check the **Use Mask Assignment** check box.

Step 10 To create a new WCCP service mask, click the **Create Mask** button. (For more information about creating service masks, see the [“Configuring WCCP Service Masks for the Content Engine”](#) section on page 4-18.)

Step 11 To modify or delete a configured service mask, click the **Edit Mask** button. (For more information about modifying service masks, see the [“Modifying WCCP Service Masks for the Content Engine”](#) section on page 4-20.)

Step 12 To view all configured service masks, click the **View Masks Configured for All Services** button. (For more information about viewing service masks, see the [“Viewing WCCP Service Masks for the Content Engine”](#) section on page 4-21.)

About Load Balancing

WCCP Version 2 has the capability to adjust the load being offered to individual Content Engines to provide more effective use of the resources available and at the same time help to ensure high quality of service to clients. Load balancing allows the set of hash buckets assigned to a Content Engine to be adjusted so that the load can be shifted from an overwhelmed Content Engine to other Content Engines that have available capacity.

With the hash assignment method, the hash parameters specify how traffic should be load balanced among the different service groups. Specifically, hashing maps items from one item set to another, such as mapping destination IP addresses or destination ports to different Content Engines in a service group from the WCCP-enabled router for load-balancing purposes. Only one load-balancing method can be used per Content Engine farm.

Creating WCCP Router Lists

You can create up to 8 different router lists with up to 32 IP addresses per list for use with WCCP Version 2 services. For example, you can specify one router list for WCCP Version 2 web cache service and another router list for reverse proxy at the same time without needing to reconfigure groups of routers or Content Engines. A router list must contain at least one IP address.

To create WCCP router lists, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, navigate to the Creating New WCCP Service window or the Modifying WCCP Service window. (See the [“Configuring WCCP Service Settings for the Content Engine”](#) section on page 4-10.)
 - Step 2** To create a new router list for WCCP Version 2, click the **New Router List** button. The Creating New WCCP Router window appears.
 - Step 3** In the Router IP Address field, specify the IP address of the router to be added to the list. You must enter at least one IP address. Otherwise, an error message is displayed on submit.

All IP addresses added must be unique within the router list. Otherwise, an error message is displayed on submit.
 - Step 4** To add an IP address to the list, click **Add**. This list represents the IP address of every router that will redirect traffic to this Content Engine for a particular service. If different routers will be used for different services, you must create more than one router list.

The window refreshes and the addresses are listed in numerical order. The order might not match the order in which IP addresses were entered. IP addresses will be displayed in red until settings are saved.
 - Step 5** To save the router list or to save any edits you have made to the router IP addresses, click **Submit**. You will be returned to the WCCP Service window.
-

Modifying WCCP Router Lists for the Content Engine

To add or delete an IP address from a router list or to delete a router list, follow these steps:

-
- Step 1** From the Creating New WCCP Service window or Modifying WCCP Service window, choose the router list number from the List Number drop-down list.
 - Step 2** Click the **Edit Router List** button. The Modifying WCCP Router List window appears.
 - Step 3** To add a router to the router list, enter the router’s IP address in the Router IP Address field, and click the **Add** button.
 - Step 4** To modify the router IP address, click the **Edit IP Address** icon next to the IP address that you want to modify. The IP address will be displayed in the Router IP Address field. Once you have changed the IP address, click the **Edit** button to add the modified IP address to the list. The window refreshes itself, and the changed IP address will be displayed in purple until settings are saved.

**Note**

The Edit button next to the Router IP Address field is displayed as Add to allow you to add any new IP address to the list. The Edit button toggles with the Add button.

Step 5 To save the settings, click **Submit**. The window refreshes itself, and the settings are saved.

Step 6 To delete an IP address from the router list, follow these steps:

- a. Click the **Delete IP Address** icon next to the listed IP address. The window refreshes itself, and the deleted IP address will be displayed in red with a strikethrough until settings are saved.
- b. To save the settings, click **Submit**. The window refreshes itself, and the deleted IP addresses are removed from display.

Step 7 To delete a router list, follow these steps:

- a. From the Modifying WCCP Router window, click the **Delete Router List** icon in the taskbar. The system displays a dialog box asking you to confirm that you want to permanently delete the router list configuration.

**Note**

When you delete a router list, the WCCP Version 2 services that have been configured to use this router list are also deleted. Make sure that the WCCP service is associated with a different router list, if required, before deleting the previously configured router list.

- b. To confirm your decision, click **OK**. The selected router list and the services with which it is associated are deleted.

Viewing WCCP Router Lists for the Content Engine

To view all router lists configured for a WCCP Version 2 dynamic service, follow these steps:

Step 1 To view all configured router lists for WCCP Version 2, click the **View All Router List** button from the Creating New WCCP Service window or the Modifying WCCP Service window. The WCCP Router List Configurations for Content Engine window appears.

The List Number column displays all configured router list numbers. The Router IP Addresses column displays a list of IP addresses that were added to the router list. From this window you can perform the following actions:

- Return to the previous window
- Edit an existing router list
- Create a new router list

Step 2 To return to the Creating New WCCP Service window, click the **Back** arrow in the Content Distribution Manager GUI taskbar.

Step 3 To configure a new router list and add IP addresses from this window, click the **Create New WCCP Router List Configuration** icon in the taskbar. The Creating New WCCP Router List window appears.

Step 4 To edit an existing router list, click the **Edit WCCP Router List Configuration** icon next to the router list that you want to modify. The Modifying WCCP Router window appears. You can add, update, or delete IP addresses.

Assigning Unique Router Lists to Content Engines in a Device Group


When a device group contains devices that are in different branch locations, the devices cannot share the same router lists. In ACNS 5.5 software, you can configure router lists for all devices in a device group from a single window in the Content Distribution Manager GUI Device Groups configuration area. Having a single window from which you can assign unique router lists to Content Engines within a device group eliminates the tedious process of visiting separate windows to configure router lists for each device that is in a different location.

To work with router lists for devices in a device group, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**.
- Step 2** Click the **Edit** icon next the device group with which you want to work.
- Step 3** In the Contents pane, choose **Request Processing > WCCP > Services**. The Modifying WCCP Service window appears.
- Step 4** To edit router lists for specific devices, navigate to the Modifying WCCP Router List window:
 - Click the **View All Router Lists** button, and then click the **Edit** icon for a router list.
 - Or choose a router list number from the drop-down menu, and click the **Edit Router List** button.

The Modifying WCCP Router List window appears. (See [Figure 4-5](#).)

Figure 4-5 Modifying WCCP Router List Window



- Step 5** Under the Device Specific Router Lists heading, click the **Edit** icon next to the name of the Content Engine for which you want to assign a unique router list. The Creating new Router List for Device window appears.



Note You must be in the Modifying WCCP Router List window for a device group for the Device Specific Router Lists section to be displayed in this window.

- Step 6** To add or remove routers from the router list on the device, enter an IP address in the Add Router field, and click either the **Add Router** button or the **Remove Router** button.
- Step 7** To save the settings, click **Submit**.

Configuring WCCP Port Lists for the Content Engine

With eight custom services that use a maximum number of eight ports each, the maximum number of ports that can be specified for transparent redirection is 64.

The legacy custom web cache and reverse proxy services (service numbers 98 and 99) can be configured with only one port each. If only one legacy service is configured, the total maximum number of transparent redirection ports is 57. If both legacy services are configured, the maximum port total is 50.

All ports receiving HTTP that are configured as members of the same WCCP service share the following characteristics:

- They have the same hash or mask parameters as configured using the Creating New WCCP Service window.
- The service on individual ports cannot be stopped or started individually (WCCP Version 2 restriction).

With Content Engines in a farm, the following restrictions apply:

- All Content Engines that use the same WCCP service are required to configure the same list of ports and the same hash or mask parameters.
- A Content Engine that tries to join the farm with the same WCCP service using a different list of ports or different hash or mask parameters is rejected by the router.
- To change the port list for a particular WCCP service, WCCP service must be stopped on all involved Content Engines, and then all must be restarted with the new parameters.

To configure WCCP port lists for the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure WCCP port lists. The Device Home for Content Engine window appears.
- Step 3** To display the entire table of contents, click the **Show All** button above the Contents pane.
- Step 4** In the Contents pane, choose **Request Routing > WCCP > Services**. The WCCP Service Settings for Content Engine window appears.
- Step 5** Click the **Create New WCCP Service Setting** icon in the taskbar to configure router lists for a new service. The Creating New WCCP Service window appears.

Alternatively, click the **Edit WCCP Service Setting** icon next to an existing WCCP service for which you want to configure router lists. The Modifying WCCP Service window appears.

- Step 6** Under the WCCP Service section, choose a WCCP dynamic redirection service from the Service Type drop-down list.
- Step 7** Under the Dynamic Service Settings section, click the **New Port List** button to configure a port list number to associate ports for a WCCP Version 2 dynamic service. The WCCP Port List Settings window appears.
- Step 8** Under the Ports column and against a port list number, specify the port number to be added to the list. You can configure a maximum of 8 port numbers per list.
- Each port number added must be unique among all configured port lists. Otherwise, an error message is displayed on submit.
- For a port list, port numbers need not necessarily be entered in successive fields under the Ports column. Blank fields in between will be removed on submit and all added port numbers will be reordered and displayed in ascending order when you visit the window the next time.
- Step 9** To save the settings, click **Submit**. The Creating New WCCP Service window or Modifying WCCP Service window appears depending on the window from which you traversed.

Modifying WCCP Port Lists for the Content Engine

To modify a port list, follow these steps:

- Step 1** From the Creating New WCCP Service window or Modifying WCCP Service window, click the **Edit Port List** button to configure a port list for WCCP Version 2 dynamic services. The WCCP Port List Settings window appears.
- Step 2** To modify port numbers in the port list, do the following:
- Choose a port list number and edit any of the port numbers that you wish to change in the Ports column for that row.
 - To save the settings, click **Submit**. The Creating New WCCP Service window or Modifying WCCP Service window appears depending on the window from which you traversed.
- Step 3** To delete a port list, do the following:
- Check the **Clear List** check box next to the port list that you want to delete. All added port numbers will be removed from the port list.
 - To save the settings, click **Submit**. The Creating New WCCP Service window or Modifying WCCP Service window appears depending on the window from which you traversed.

Configuring WCCP Service Masks for the Content Engine

You can configure up to 16 WCCP service masks. For FTP service, only IP masks can be configured, so port mask fields are disabled for entry.

Bit masks are specified as hexadecimal numbers. All the specified bit masks together cannot have more than 7 bits set. For example, a correct way of using three masks is 0xF (4 bits), 0x1 (1 bit), and 0x3 (2 bits) for a total of 7 bits. In this case, you cannot configure any additional mask other than 0x0. Otherwise, an error message is displayed.

An example of using four masks could be 0xA (2 bits), 0x7 (3 bits), 0x8 (1 bit), 0x1 (1 bit) for a total of 7 bits.

To configure WCCP service masks for a WCCP service configured on the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure WCCP service masks. The Device Home for Content Engine window appears.
- Step 3** To display the entire table of contents, click the **Show All** button above the Contents pane.
- Step 4** In the Contents pane, choose **Request Routing > WCCP > Services**. The WCCP Service Settings for Content Engine window appears.
- Step 5** Click the **Create New WCCP Service Setting** icon in the taskbar to configure router lists for a new service. The Creating New WCCP Service window appears.

Alternatively, click the **Edit WCCP Service Setting** icon next to an existing WCCP service for which you want to configure router lists. The Modifying WCCP Service window appears.
- Step 6** Under the WCCP Service section, choose a WCCP service from the Service Type drop-down list.
- Step 7** Under the Other Settings section, click the **Create Mask** button next to the Use Mask Assignment check box to configure a service mask and associate it with a WCCP Version 2 service. The Creating New WCCP Service Mask window appears.
- Step 8** Choose the type of WCCP service from the Service Type drop-down list. [Table 4-5](#) describes these service types.

Table 4-5 WCCP Service Types

Service Type	Description
DNS	Allows a WCCP-enabled router to intercept and redirect DNS packets to a boomerang server, which will then accept the intercepted packets and process them, including reinserting the packets into the request stream.
FTP Native	Enables transparent interception of FTP traffic with WCCP Version 2.
HTTPS	Enables WCCP flow redirection to a Content Engine configured as an HTTPS server.
RTSP	Configures WCCP Version 2 Real-Time Streaming Protocol (RTSP) transparent interception.
Windows Media	Enables WCCP Version 2 Windows Media caching service.
Windows Media RTSPU	Enables Windows Media RTSPU (port 5005) transparent interception.
Reverse proxy	Enables WCCP Version 2 reverse proxy service. When WCCP reverse proxy is enabled, the router does load balancing in a cluster based on the source IP address (for example, the client's browser IP address).
Web cache	Enables web cache service with WCCP Version 2. With web cache service, the router balances the traffic load within a Content Engine cluster based on the destination IP address (for example, web server IP address).

Table 4-5 WCCP Service Types (continued)

Service Type	Description
Custom web cache	Enables the Content Engine to accept redirected HTTP traffic on a port other than 80.
Dynamic Service Number (90–97)	Enables up to eight dynamic WCCP redirection services on the Content Engine.

- Step 9** In the Source IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, 0xFE000000) used to match the packet source IP address. The range is 0x00000000–0xFE000000. The default is 0x00000000.
- Step 10** In the Source Port Mask field, specify the port mask defined by a hexadecimal number (for example, 0xFE00) used to match the packet source port number. The port mask range is 0x00–0xFE. The default is 0x0.
- Step 11** In the Destination IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, 0xFE000000) used to match the packet destination IP address. The range is 0x00000000–0xFE000000. The default is 0x00001741.
- Step 12** In the Destination Port Mask field, specify the port mask defined by a hexadecimal number (for example, 0xFE00) used to match the packet destination port number. The port mask range is 0x00–0xFE. The default is 0x0.
- Step 13** To save the settings, click **Submit**. The Creating New WCCP Service window or the Modifying WCCP Service window appears, depending on the window from which you navigated.

Modifying WCCP Service Masks for the Content Engine

To modify a service mask for a WCCP service configured on the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure WCCP service masks. The Device Home for Content Engine window appears.
- Step 3** To display the entire table of contents, click the **Show All** button above the Contents pane.
- Step 4** In the Contents pane, choose **Request Routing > WCCP > Services**. The WCCP Service Settings for Content Engine window appears.
- Step 5** Click the **Create New WCCP Service Setting** icon in the taskbar to configure router lists for a new service. The Creating New WCCP Service window appears.
- Alternatively, click the **Edit WCCP Service Setting** icon next to an existing WCCP service for which you want to configure router lists. The Modifying WCCP Service window appears.
- Step 6** Under the WCCP Service section, choose a WCCP service from the Service Type drop-down list.
- Step 7** Under the Other Settings section, click the **Edit Mask** button next to the Use Mask Assignment check box to modify a service mask and associate it with a WCCP Version 2 service. The Modifying WCCP Service Mask window appears.
- The Edit Mask button toggles with the Create Mask button if a service mask has already been configured for a WCCP service.

-
- Step 8** Choose the type of WCCP service from the Service Type drop-down list. See [Table 4-5](#) for a description of the service types.
- Step 9** In the Source IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, 0xFE000000) used to match the packet source IP address. The range is 0x00000000–0xFE000000. The default is 0x00000000.
- Step 10** In the Source Port Mask field, specify the port mask defined by a hexadecimal number (for example, 0xFE00) used to match the packet source port number. The port mask range is 0x00–0xFE. The default is 0x0.
- Step 11** In the Destination IP Mask field, specify the IP address mask defined by a hexadecimal number (for example, 0xFE000000) used to match the packet destination IP address. The range is 0x00000000–0xFE000000. The default is 0x00001741.
- Step 12** In the Destination Port Mask field, specify the port mask defined by a hexadecimal number (for example, 0xFC00) used to match the packet destination port number. The port mask range is 0x0–0xFE. The default is 0x0.
- Step 13** To save the settings, click **Submit**. The Creating New WCCP Service window or the Modifying WCCP Service window appears depending on the window from which you navigated.
-

To delete a service mask, follow these steps:

-
- Step 1** From the WCCP Service Mask Settings for Content Engine window, click the **Edit WCCP Service Mask Setting** icon next to an existing WCCP service mask for which you want to change the mask settings. The Modifying WCCP Service Mask window appears.
- Step 2** Click the **Delete WCCP Service Mask** icon in the taskbar. The system displays a dialog box asking you to confirm whether you want to permanently delete the service mask configuration.
- Step 3** Click **OK** to confirm. The WCCP Service Mask Settings for Content Engine window appears.
-

Viewing WCCP Service Masks for the Content Engine

To view all masks configured for a WCCP service, follow these steps:

-
- Step 1** From the Creating New WCCP Service window or Modifying WCCP Service window, click the **View Masks Configured for All Services** button to view all configured service masks for WCCP services. The WCCP Service Mask Settings for Content Engine window appears. A list of all services for which masks have been configured to match the packet source IP address, source port number, destination IP address, or destination port number are displayed.
- Step 2** To return to the Creating New WCCP Service or Modifying WCCP Service window, click the **Back** (left arrow) icon in the taskbar.
- Step 3** To configure a new service mask, click the **Create New WCCP Service Mask Setting** icon in the taskbar. The Creating New WCCP Service Mask window appears.

Alternatively, click the **Edit WCCP Service Mask Setting** icon next to an existing WCCP service mask for which you want to change the mask settings. The Modifying WCCP Service Mask window appears.


Configuring WCCP Transparent Routing Bypass Options

One of the fundamental principles of transparent network request redirection is that the Content Engine must remain transparent to the end user at all times. A transparent caching solution in an ACNS network environment must not introduce any possible failure conditions or side effects in the network.

The Cisco ACNS transparent caching solution uses a WCCP-enabled router and various advanced techniques to ensure that the Content Engine remains transparent, even if web browsers are nonoperational or web servers are not HTTP-compliant.

If a Content Engine becomes overwhelmed with traffic, it can use the overload bypass feature to reroute the overload traffic. When the Content Engine is overloaded and the **bypass load** command is enabled, the Content Engine refuses additional requests and forwards them to the origin servers. If the load remains too high, more traffic is bypassed to the servers, and so on until the Content Engine can handle the load. The time interval between one bucket being bypassed and the next is set by the **out-interval** option. The default is 4 seconds. (See [Figure 4-6](#).)

Figure 4-6 Overload Bypass



When the first bucket bypass occurs, a set interval must elapse before the Content Engine begins to again service the bypassed buckets. The duration of this interval is set by the **time-interval** option. The default is 10 minutes.

When the Content Engine begins to service the bypassed traffic again, it begins with a single bypassed bucket. If the load is serviceable, it picks up another bypassed bucket, and so on. The time interval between picking up one bucket and the next is set by the **in-interval** option. The default is 60 seconds.

Configuring Authentication Traffic Bypass

Because of IP authentication, some websites do not allow the Content Engine to connect directly on behalf of the client. In order to preserve cache transparency and avoid disruption of service, the Content Engine can use authentication traffic bypass to automatically generate a dynamic access list for the selected client/server pairs. Authentication bypass triggers are also propagated upstream and downstream in the case of hierarchical caching.



Note

The bypass feature is only available when WCCP Version 2 is enabled in your local network.

Configuring WCCP Bypass Settings

To configure WCCP bypass settings, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**. The Device Group window appears. (Alternatively, choose **Devices > Devices**.)
- Step 2** Click the **Edit** icon next to the name of the device group (or Content Engine) that you want to configure. The Contents pane appears on the left.
- Step 3** To display the entire table of contents, click the **Show All** button above the Contents pane.
- Step 4** From the Contents pane, choose **Request Routing > WCCP > Bypass**. The WCCP Bypass Settings window appears. (See Figure 4-7.) Table 4-7 describes the fields in this window and provides the corresponding CLI global configuration commands.

Figure 4-7 WCCP Bypass Settings Window

The screenshot shows the Cisco Application and Content Networking System GUI. The main window is titled "WCCP Bypass Settings for Content Engine, CONTENTENGINE". The left sidebar shows a tree view with "Request Routing" expanded, and "WCCP" selected. The main content area displays the following settings:

WCCP Bypass Settings	
Current applied settings from Content Engine, CONTENTENGINE	
Bypass Enable:	☒
Authentication Bypass:	☒
Bypass Gateway:	
Bypass entry expiration time: *	20 (minutes) (1-1440)
Error Handling: *	transparent
Load Bypass Enable:	☒
Interval between bypassing buckets: *	4 (seconds) (4-600)
Time that a bucket is bypassed: *	6 (minutes) (1-1440)
Time interval between buckets coming back: *	2 (seconds) (2-600)

Note: * - Required Field

At the bottom right, there are "Submit" and "Cancel" buttons.

- Step 5** To enable the Content Engine to bypass incoming requests from clients, check the **Bypass Enable** check box.
- Step 6** To enable authentication bypass, check the **Authentication Bypass** check box.



Note

Some websites may not allow the Content Engine to connect directly on behalf of the client. To avoid a disruption of service when traffic is bypassed, the Content Engine can use authentication bypass to generate a dynamic access list for these client/server pairs. Authentication bypass triggers are also propagated upstream and downstream in an ACNS network environment.

- Step 7** In the Bypass Gateway field, specify the IP address of the bypass gateway to enable the Content Engine to return a bypassed packet through Layer 2 redirection to the configured redirecting switch. If no bypass gateway is configured, the bypass packet will be forwarded to the source MAC address that sent the packet.
- Step 8** In the Bypass entry expiration time field, enter the number of minutes that an idle client/server pair remains on the bypass access list. The default value is 20 minutes.
- Step 9** From the Error Handling drop-down list, choose a method for handling errors: **transparent**, **reset-connection**, or **send-cache-error**. (See [Table 4-6](#).)

Table 4-6 Error Handling Options

Option	Description
Transparent	The Content Engine will not send errors to the client but will bypass the client connections to the server.
Send-cache-error	The Content Engine sends an error page to the client.
Reset-connection	The Content Engine resets the TCP connection without specifying an error.

- Step 10** Check the **Load Bypass Enable** check box to enable traffic bypass.
- Step 11** In the Interval between bypassing buckets field, enter the amount of time (in seconds) between the bypassing of one bucket and the next. The default is 4 seconds.
- When the Content Engine is overwhelmed with traffic and the bypass option is enabled, it bypasses traffic one bucket at a time until it is no longer overloaded.



Note A bucket is defined as a certain subsection of the allotted hash assigned to each Content Engine in a Content Engine farm. If only one Content Engine exists in this environment, it has 256 buckets assigned to it.

- Step 12** In the Time that a bucket is bypassed field, enter the number of minutes that the Content Engine remains in bypass mode and does not attempt to pick up the bypassed load. The default is 10 minutes.
- Step 13** In the Time interval between buckets coming back field, enter the time between the pickup of each bucket (in seconds). The default is 2 seconds.
- Once the time interval allotted to bypass mode has elapsed, the Content Engine begins to pick up bypassed traffic one bucket at a time.
- Step 14** To save the settings, click **Submit**.
- A “Click Submit to Save” message appears in red next to the current settings line when there are pending changes to be saved after you have applied default or device group settings. To revert to the previously configured window settings, click **Reset**. The Reset button appears only when you have applied default or group settings to change the current device settings but the settings have not yet been submitted.

Table 4-7 WCCP Bypass Settings

GUI Parameter	Function	CLI Command
Bypass Enable	Enables the Content Engine to bypass incoming requests from clients.	—
Authentication Bypass	Enables the use of authentication bypass to generate a dynamic access list for client/server pairs.	bypass auth-traffic enable
Bypass Gateway	IP address of the bypass gateway.	bypass gateway <i>ipaddress</i>
Bypass entry expiration time	Number of minutes that an idle client/server pair remains on the bypass access list. The default value is 20 minutes.	bypass timer <i>minutes</i>
Error Handling	Specifies a method for handling errors.	error-handling { reset-connection send-cache-error transparent }
Load Bypass Enable	Enables traffic bypass.	bypass load enable
Interval between bypassing buckets	Time period (in seconds) between the bypassing of one bucket and the next. The default value is 4 seconds.	bypass load in-interval <i>seconds</i>
Time that a bucket is bypassed	Number of minutes that the Content Engine does not attempt to pick up the bypassed load.	bypass load time-interval <i>minutes</i>
Time interval between buckets coming back	Time (in seconds) between the pickup of each bucket once the Content Engine begins to pick up bypassed traffic again.	bypass load out-interval <i>seconds</i>

Creating WCCP Bypass List Entries

The Content Engine can use authentication bypass to generate a dynamic access list for client/server pairs. To generate these lists with the Content Distribution Manager GUI, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**. The Device Group window appears. (Alternatively, choose **Device > Devices**.)
- Step 2** Click the **Edit** icon next to the name of the device group (or Content Engine) that you want to configure. The Contents pane appears on the left.
- Step 3** To display the entire table of contents, click the **Show All** button above the Contents pane.
- Step 4** From the Contents pane, choose **Request Routing > WCCP > Bypass List**.
- Step 5** Click the **Create New WCCP Bypass List** icon. The Creating new WCCP Bypass List window appears.
- Step 6** Enter the IP address for the client in the Client Address field.

Step 7 Enter the IP address for the server in the Server Address field.

Step 8 Check **Submit** to save the settings.

To create a WCCP bypass list entry from the CLI, use the following global configuration command:

bypass static { *clientip* [*serverip* | **any-server**] | **any-client** *serverip* }



Note

You must not exceed 50 bypass list entries for any one Content Engine.

Enabling WCCP on the Router

To enable an interface to redirect web traffic to the Content Engine using WCCP Version 2, perform the following tasks beginning in global configuration mode on the router:

	Command	Purpose
Step 1	Router(config)# ip wccp enable	Enables the router to use WCCP.
Step 2	Router(config)# ip wccp redirect-list [<i>number</i> <i>name</i>]	(Optional) Specifies the redirect access list. Only packets that match this access list are redirected. If you do not configure this command, all packets are redirected.
Step 3	Router(config)# interface <i>type number</i>	Enters interface configuration mode.
Step 4	Router(config-if)# ip web-cache redirect	Configures the interface connected to the Internet to redirect web traffic to the Content Engine.
Step 5	Router(config-if)# ip route-cache same-interface	(Optional) If the client and a Content Engine are located on the same network, configures the router to use the fast switching path on the interface.
Step 6	Router(config-if)# end	Exits configuration mode.
Step 7	Router(config)# copy running-config startup-config	Saves the configuration.

Basic WCCP Router Configurations for Web Caching

When WCCP support is enabled on the router and on the Content Engines, the devices can communicate and deliver the services for which they are configured. To suspend these services, you can disable WCCP support on the router rather than powering off or otherwise disabling individual Content Engines. (For instance, use the **no ip wccp** router command to disable WCCP support on the router.)


Many WCCP Version 2 services also require a configuration of the appropriate **wccp** global configuration command. Refer to the *Cisco ACNS Software Command Reference, Release 5.x* publication and ACNS software release notes for more details.

Web Cache Service Configuration with Clients and Content Engine on the Same Subnet

In this scenario, the Content Engine and the requesting clients are on the same subnet, as shown in [Figure 4-8](#).

A router running WCCP Version 2 transparently redirects client HTTP traffic bound for router interface s0/0 to the Content Engine. The web cache service redirects HTTP traffic on port 80 only.

Figure 4-8 Web Cache Service with Clients and Content Engine on Same Subnet



Configuring the Content Engine for Web Cache Service—Clients and Content Engine on the Same Subnet

To configure the Content Engine for web cache service, perform the following steps while logged in to the ACNS software in global configuration mode:

	Command	Purpose
Step 1	ContentEngine(config)# wccp version 2	Ensures that the Content Engine is running WCCP Version 2.
Step 2	ContentEngine(config)# wccp router-list 1 10.10.10.1	Configures a router list.
Step 3	ContentEngine(config)# wccp web-cache router-list-num 1	Informs the routers in the specified router list that the Content Engine is accepting web cache service.
Step 4	ContentEngine(config)# exit	Exits global configuration mode.
Step 5	ContentEngine# copy running-config startup-config	Writes running configurations to nonvolatile memory.

Configuring the Router for Web Cache Service—Clients and Content Engine on the Same Subnet

To configure the router for web cache service, perform the following steps while logged in to the router in global configuration mode:

	Command	Purpose
Step 1	Router(config)# ip wccp web-cache	Instructs the router to run the web cache service.
Step 2	Router(config)# interface Ethernet0	Specifies which router interface to configure. In this scenario, Ethernet0 is the router interface to which the Content Engine and clients are connected.
Step 3	Router(config-if)# ip route-cache same-interface	Enables fast switching of redirected packets back through the interface on which they were received. Without this command, the router does not use the high-speed switching cache, and the packets are process-switched, a much slower method.
Step 4	Router(config)# interface Serial0	Specifies which router interface to configure. In this scenario, Serial0 is the router interface to the Internet.
Step 5	Router(config-if)# ip wccp web-cache redirect out	Instructs the router to redirect web cache traffic bound for the specified interface to Content Engines that accept web cache service. In this scenario there is only one router. Web cache traffic is defined as TCP port 80 traffic.
Step 6	Router(config-if)# exit	Exits interface configuration mode.

Configuration Examples—Web Cache Service with Clients and Content Engine on the Same Subnet

This example shows a Content Engine and router configured for web cache service with the topology illustrated in [Figure 4-8](#):

Content Engine

```

hostname Content_engine_4.2
!
clock timezone pst -8 0
!
ip domain-name cu.net
!
interface FastEthernet 0
ip address 10.10.20.10 255.255.255.0
no autosense
bandwidth 100
full-duplex
exit
interface FastEthernet 1
shutdown
exit
!
ip default-gateway 10.10.20.1
!
ip name-server 10.10.10.100
!
!
!
wccp router-list 1 10.10.20.1
wccp web-cache router-list-num 1

```

```
wccp version 2
!  
!  
!  
. . .
```

WCCP-Enabled Router

Building configuration...

Current configuration:


```
!  
version 12.0  
!  
hostname WCCP-Router  
!  
!  
ip wccp web-cache  
!  
interface Ethernet0  
 ip address 10.10.10.1 255.255.255.0  
 ip route-cache same-interface  
!  
interface Serial0  
 ip address 192.168.1.2 255.255.255.252  
 no ip directed-broadcast  
 no ip mroute-cache  
 ip wccp web-cache redirect out  
!  
end
```

Web Cache Service Configuration with Clients and Content Engine on Different Subnets

In this scenario, the Content Engine and the requesting clients are on different subnets, as shown in [Figure 4-9](#).

A router running WCCP Version 2 transparently redirects client HTTP traffic bound for router interface s0/0 to the Content Engine. The web cache service redirects HTTP traffic on port 80 only.

Figure 4-9 Web Cache Service with Content Engine and Clients on Different Subnets



Configuring the Content Engine for Web Cache Service—Clients and Content Engine on Different Subnets

To configure the Content Engine for web cache service, perform the following steps while logged in to the ACNS software in global configuration mode:

	Command	Purpose
Step 1	ContentEngine(config)# wccp version 2	Ensures that the Content Engine is running WCCP Version 2.
Step 2	ContentEngine(config)# wccp router-list 1 10.10.20.1	Configures a router list.
Step 3	ContentEngine(config)# wccp web-cache router-list-num 1	Informs the routers in the specified router list that the Content Engine is accepting web cache service.
Step 4	ContentEngine(config)# exit	Exits global configuration mode.
Step 5	ContentEngine# copy running-config startup-config	Writes running configurations to nonvolatile memory.

Configuring the Router for Web Cache Service—Clients and Content Engine on Different Subnets

To configure the router for web cache service, perform the following steps while logged in to the router in global configuration mode:

	Command	Purpose
Step 1	Router(config)# ip wccp web-cache	Instructs the router to run the web cache service.
Step 2	Router(config)# interface Serial0	Specifies which router interface to configure. In this scenario, Serial0 is the router interface to the Internet.

Step 3	Router(config-if)# ip wccp web-cache redirect out	Instructs the router to redirect web cache traffic bound for the specified interface to Content Engines that accept web cache service. In this scenario there is only one router. Web cache traffic is defined as HTTP packets on port 80.
Step 4	Router(config-if)# exit	Exits interface configuration mode.

Configuration Examples—Web Cache Service with Clients and Content Engine on Different Subnets

This example shows a Content Engine and router configured for web cache service with the topology illustrated in [Figure 4-9](#):

Content Engine Configuration

```
. . .
hostname Content_engine_4.2
!
clock timezone pst -8 0
!
ip domain-name cisco.com
!
exec-timeout 20
!
interface FastEthernet 0
ip address 10.10.20.10 255.255.255.0
no autosense
bandwidth 100
full-duplex
exit
interface FastEthernet 1
shutdown
exit
!
ip default-gateway 10.10.20.1
!
ip name-server 10.10.10.100
!
!
wccp router-list 1 10.10.20.1
wccp web-cache router-list-num 1
wccp version 2
!
!
!
...
```

WCCP-Enabled Router Configuration

Building configuration...

Current configuration:

```
!
hostname WCCP-Router
!
!
ip subnet-zero
!
ip wccp web-cache
!
interface Ethernet0
```

```

ip address 10.10.10.1 255.255.255.0
!
interface Ethernet1
ip address 10.10.20.1 255.255.255.0
!
interface Serial0
ip address 192.168.1.2 255.255.255.252
no ip directed-broadcast
no ip mroute-cache
ip wccp web-cache redirect out
!
end

```

Configuring Direct Proxy Routing

In a proxy configuration, the end user's browser must be configured with the IP address of the Content Engine; therefore, a proxy-style request arrives at the Content Engine after having been specifically routed to the Content Engine by the client.

Using a Proxy Autoconfiguration File Server

Cisco ACNS 5.5 software supports dynamic proxy autoconfiguration (PAC). This feature is a form of direct proxy routing that allows a browser to dynamically obtain proxy IP address and port configuration information from a special file called a PAC file template. The PAC feature is supported by the Microsoft Internet Explorer and Netscape Communicator browsers. Browsers must be manually configured for automatic proxy configuration.

The PAC feature can be configured through the Content Distribution Manager (see the [“Registering a PAC File Template and Assigning Coverage Zone Information”](#) section on page 4-35) and requires at least one Content Engine that is enabled as a PAC file server (see the [“Configuring a Content Engine as a PAC File Server”](#) section on page 4-38). The PAC file server dynamically chooses the proxy Content Engine that is closest to the client requesting the content, based on instructions contained in the PAC file template. (See the [“Creating PAC File Templates”](#) section on page 4-33.)

Dynamic PAC routing supports the following features:

- It supports multiple PAC files.
- The user can specify the names of the PAC files instead of being limited to preset names.
- The PAC file server dynamically adds a set of “nearest proxies” to a PAC file template based on coverage zone information and the source IP address of the requester.
- The Content Distribution Manager acquires the PAC file templates and distributes them to Content Engines that are PAC file servers rather than the Content Engines themselves acquiring the PAC file templates.



Note

This new feature is not an extension of the **proxy-auto-config** EXEC command feature supported in ACNS 5.0 software, nor does it replace that feature. Both features are supported in ACNS 5.5 software.

Creating PAC File Templates

The system administrator creates a PAC file template on a workstation and places it where the Content Distribution Manager can access the URL. The administrator then registers the PAC file template URL in the Content Distribution Manager GUI and assigns a coverage zone file to it.

The template is a complete PAC file except that the administrator inserts predefined macros where the PAC file server is to generate and insert dynamic information. The nearest proxies are determined by the metrics contained in the coverage zone file.

The Content Distribution Manager acquires the template and distributes it to all PAC file servers. For each Content Engine that is a PAC file server, the administrator must also configure the ports used for PAC file requests. Typically, port 80 is used.

The port can be configured in the CLI using the **http proxy incoming port** global configuration command. (There is no default.) Alternatively, you can use the Content Distribution Manager to configure the incoming proxy ports in the HTTP Connection Settings window (**Devices > Devices > Applications > Web > HTTP > HTTP Connections**).

The administrator then gives the PAC file URLs, including the desired servers, to end users. Users must configure their browsers to use the URL. When the browser first comes up, it requests the URL. The PAC file server dynamically replaces any ACNS software macros with the current information from the local coverage zone, if present, and from the default coverage zone based on the source IP address of the requester. The completed PAC file is returned to the requesting client.

When end users are behind a NAT device and the network administrator wants to direct end users to different proxies based on their IP address, then there must also be a PAC file server behind the NAT device for those users to use. The reason for this is because a PAC file server that is not behind the NAT would receive PAC requests for users behind the NAT with the source IP address of the NAT device rather than the source IP address of the end user.

Three macros are defined as follows:

- CE_NAME_n
- CE_IPADDR_n
- NEAREST_PROXIES_n

where *n* can be a value of 1 to 5.

For CE_NAME_n, the PAC file server substitutes the name of the *n*th closest proxy. For CE_IPADDR_n, the PAC file server substitutes the IP address of the *n*th closest proxy. If there is no *n*th closest proxy, the PAC file server substitutes the empty string "" in place of the macro. The PAC template creator must keep this in mind when creating the PAC file templates.

For the NEAREST_PROXIES_n macro, the PAC file server substitutes the string "PROXY <ip address>;" for up to *n* closest proxies where <ip address> is the IP address of the proxy. If the template writer includes a port number with the macro, then this port number is appended to each proxy IP address.

For example, NEAREST_PROXIES_2:8080 might expand into a string such as "PROXY 192.30.120.12:8080; PROXY 168.10.10.1:8080;". If *n* is larger than the number of available proxies, then the PAC file server substitutes fewer than *n* proxies. For example, NEAREST_PROXIES_2 might expand into "PROXY 192.30.120.12;" or even just "" when there are no nearest proxies. Again, the PAC template creator must keep this in mind when creating the PAC file templates.

PAC File Example

Assume that CE1 serves subnet 10.1.1.0/24 and CE2 serves subnet 10.1.2.0/24. All clients on subnet 10.1.1.0/24 are directed to CE1 as the closest proxy, and all clients on subnet 10.1.2.0/24 are directed to CE2 as the closest proxy.

The PAC file template would look like this:

```
Function FindProxyForURL(url, host)
{
    // handle plain host names
    if (isPlainHostName(host) {
        return "DIRECT";

    // handle exception list
        if (myIpAddress() == "10.1.1.2") {
            return "DIRECT";
        }

    // handle special domain list
        if (shExpMatch(host, "specialdomain.com")) {
            return "PROXY specialDomainProxy.foo.com";
        }

    // handle direct domain list
        if (shExpMatch(host, "directdomain.com")) {
            return "DIRECT";
        }

    // closest proxies + defaults
        p1 = CE_NAME_1;
        p2 = CE_NAME_2;
        if (p1 != "") {
            p1 = "PROXY " + p1 + ".proxy.foo.com:8080; ";
        }
        if (p2 != "") {
            p2 = "PROXY " + p2 + ".proxy.foo.com:8080; ";
        }
        return p1 + p2 + "DIRECT";
    }
}
```

The coverage zone file looks like this:

```
<?xml version="1.0"?>
<CDNNetwork>
<!--No coverage zone file from CDM is available-->
<!--Default coverage zones, if any, follow this comment-->
<coverageZone>
    <network>10.1.1.0/24</network>
    <CE>CE1</CE>
    <metric>20</metric>
</coverageZone>
<coverageZone>
    <network>10.1.2.0/24</network>
    <CE>CE2</CE>
    <metric>20</metric>
</coverageZone>
</CDNNetwork>
```

For a client source IP address of 10.1.1.20, the generated proxy.pac file would look like this:

```
function FindProxyForURL(url, host)
{
    // handle plain host names
```

```

if (isPlainHostName(host) {
    return "DIRECT";

// handle exception list
if (myIpAddress() == "10.1.1.2") {
    return "DIRECT";
}

// handle special domain list
if (shExpMatch(host, "specialdomain.com")) {
    return "PROXY specialDomainProxy.foo.com";
}

// handle direct domain list
if (shExpMatch(host, "directdomain.com")) {
    return "DIRECT";
}

// closest proxies + defaults
p1 = "CE1";
p2 = "CE2";
if (p1 != "") {
    p1 = "PROXY " + p1 + ".proxy.foo.com:8080; ";
}
if (p2 != "") {
    p2 = "PROXY " + p2 + ".proxy.foo.com:8080; ";
}
return p1 + p2 + "DIRECT";
}


```

Registering a PAC File Template and Assigning Coverage Zone Information

To register a PAC file template and assign coverage zone information, follow these steps:

-
- Step 1** In the Content Distribution Manger GUI, choose **System > Configuration > Routing > PAC File Registration**. The PAC File Registration window appears.
- Step 2** Click the **Create New PAC File** icon. The Registering New PAC File window appears.
- Step 3** Choose a file import method:
- **Upload**—The upload method allows you to upload a PAC file from any location that is accessible from your PC, by using the browse feature.
 - **Import**—The import method allows you to import the PAC file from an external HTTP, HTTPS, FTP, or CIFS server. (See [Figure 4-10](#).)

Figure 4-10 Registering New PAC File Window—Import Method



(For descriptions of the fields associated with the upload method, see [Table 4-8](#). For descriptions of the fields associated with the import method, see [Table 4-9](#).)

Table 4-8 Upload Method for PAC Files

Property	Description
File Import Method	Sets the file import method. Choose Upload to display the required fields to upload the coverage zone file.
PAC File Upload	Local directory path to the PAC file. Use the Browse button to help you locate the file.
Destination Filename	Name of the PAC file being distributed to the PAC file server. This field is filled in automatically with the filename from the local directory path. Note The PAC file with this destination filename is copied into the /state/pacFile directory on the PAC file server.
End User Request URL	URL that an end user would use to configure the browser to obtain the PAC file from a PAC file server. Note If a port number other than 80 is to be used to obtain the file, it must be specified in the URL. For example, if port 8080 is to be used, the URL would look like this: http://10.5.1.151:8080/sample.pac

Table 4-8 Upload Method for PAC Files (continued)

Property	Description
Coverage Zone File	Coverage zone file that is the basis for calculating the nearest proxy.
Use the default Coverage Zone	When checked, calculates the nearest proxy based on the default coverage zone. Note If only a coverage zone file is chosen from the drop down list, the nearest proxies are calculated using the information in the specified coverage zone file. If only Use the default Coverage Zone check box is checked, the nearest proxies are calculated using only the default coverage zone. If both a coverage zone file is specified and Use the default Coverage Zone check box is checked, the nearest proxies are calculated based on the union of the coverage zone file and default coverage zone.

Table 4-9 Import Method for PAC Files

Property	Description
File Import Method	Sets the file import method. Choose Import to display the required fields to import the PAC file.
PAC File URL	Full path of the origin URL (including the PAC filename) where the PAC file can be accessed.
Destination File Name	Name of the PAC file being distributed to the PAC file server. Note The PAC file with this destination filename is copied into the /state/pacFile directory on the PAC file server.
End User Request URL	URL that an end user would use to configure the browser to obtain the PAC file from a PAC file server. Note If a port number other than 80 is to be used to obtain the file, it must be specified in the URL. For example, if port 8080 is to be used, the URL would look like this: http://10.5.1.151:8080/sample.pac
Coverage Zone File	Coverage zone file that is the basis for calculating the nearest proxy.
Use the default Coverage Zone	When checked, calculates the nearest proxy based on the default coverage zone. Note If only a coverage zone file is entered, the nearest proxies are calculated using the information in the specified coverage zone file. If only the default coverage zone is indicated, the nearest proxies are calculated using only the default coverage zone. If both a coverage zone file is specified and the default coverage zone check box is checked, the nearest proxies are calculated based on the union of the coverage zone file and default coverage zone.

Table 4-9 Import Method for PAC Files (continued)

Property	Description
Update Interval (minutes)	Frequency with which the Content Distribution Manager looks for changes to the PAC file located at the PAC file URL. The default value is 10 minutes.
UserName	Name of the user to be authenticated to fetch the PAC file.
Password	User password for accessing the PAC file.
NTLM user Domain	NTLM user domain name for NTLM authentication.
Disable Basic Authentication	When checked, NTLM headers cannot be stripped off to allow fallback to the basic authentication method. If you leave this check box unchecked, NTLM authentication headers can be stripped to allow fallback to the basic authentication method, and the username and password information can be passed to the origin server in clear text with a basic authentication header.

Step 4 To save the settings, click **Submit**.

Configuring a Content Engine as a PAC File Server

PAC file servers look for changes in their PAC file server status, coverage zone information, general PAC information, and PAC template information. If necessary, a PAC file server updates its information so that it can correctly serve dynamic PAC files. It also installs the PAC file templates and coverage zone files in a known location. The administrator can configure any number of Content Engines as PAC file servers.

To configure a Content Engine as a PAC file server, follow these steps:

-
- Step 1** In the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure as a PAC file server. The Devices Home window appears.
- Step 3** To display the entire table of contents, click the **Show All** button above the Contents pane.
- Step 4** In the Contents pane, choose **Device Activation**. The Device Activation window appears. (See [Figure 4-11](#).)

Figure 4-11 Device Activation for Content Engine Window

The screenshot shows the Cisco Application and Content Networking System web interface in Microsoft Internet Explorer. The main window is titled "Device Activation, CONTENTENGINE". It features a left-hand navigation pane with a "Contents" menu. The main content area is divided into several sections: "General Configuration", "Locality", "Request Routing", "NAT Configuration", and "Comments".

General Configuration:

- Name: CONTENTENGINE
- Status: Online
- Activate: ☒
- cdnts: 0 MB
- Set Default Coverage Zone File: ☒

Locality:

- Location: ContentEngine-location

Request Routing:

- Enable Request Routing: ☐
- Coverage Zone File: None
- Enable PAC File Server: ☒

NAT Configuration:

- Use CE's primary IP Address: ☒
- NAT Address: [Empty field]
- Port: 443

Comments:

- n/a

At the bottom right, there are "Submit" and "Cancel" buttons. A note at the bottom left states: "Note: * - Required Field".

Step 5 In the Request Routing section, check the **Enable PAC File Server** check box.

Step 6 Click **Submit** to save the changes.

Configuring Client Proxy Autoconfiguration Settings

ACNS 5.x software supports proxy automatic configuration files (.pac files). A browser obtains proxy IP address and port configuration information from the proxy automatic configuration file when the browser's autoconfiguration URL field is configured with the Content Engine IP address, incoming port number, file directory, and .pac filename.



Note

You must configure disks /local1 or /local2 as a sysfs volume before downloading the autoconfiguration file to either of these two disk locations.

The Microsoft Internet Explorer and Netscape browsers support the proxy autoconfiguration feature. The browser must be manually configured for automatic proxy configuration.

This example demonstrates the URL syntax to enter in the proxy automatic configuration URL field of the browser:

http://ContentEngine-IPaddress:portnumber/proxy.pac

**Note**

Use a port number specified by the proxy incoming settings for configuring proxy incoming ports. For instance, if port 8080 is specified, then use 8080 as your port number in the example shown.

To configure proxy autoconfiguration, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**. If you have created device groups, the Device Group window appears. (Alternatively, choose **Devices > Devices**.)
- Step 2** To display the entire table of contents, click the **Show All** button above the Contents pane.
- Step 3** Click the **Edit** icon next to the name of the device group (or Content Engine) that you want to configure. The Contents pane appears on the left.
- Step 4** From the Contents pane, choose **Request Routing > Client Proxy Autoconfiguration**. The Client Proxy Auto Configuration Settings window appears. (See Figure 4-12.)

Figure 4-12 Client Proxy Auto Configuration Settings Window

The screenshot shows the Cisco Application and Content Networking System GUI. The main window is titled "Client Proxy Auto Configuration Settings for Device Group, DG2". On the left, there is a "Contents" pane with a tree view showing various configuration options. The "Client Proxy Auto Configuration" option is selected. The main area contains several configuration fields:

- Enable:** A checkbox that is currently unchecked.
- Remote FTP Server Host Name:** A text input field.
- Remote File (full pathname):** A text input field.
- Username:** A text input field.
- Password:** A text input field.
- Confirm Password:** A text input field.

At the bottom of the form, there is a note: "Note: * - Required Field". At the bottom right of the window, there are "Submit" and "Cancel" buttons. A red message at the top right of the form area states: "There are currently no Client Proxy Auto Config settings for this Group".

- Step 5** To enable proxy autoconfiguration, check the **Enable** check box.
- Step 6** In the Remote FTP Server Host Name Field, enter a host name or an IP address if you are downloading the .pac file from a remote FTP server.
- Step 7** In the Remote File field, enter the name of the autoconfiguration file to be accessed on the remote FTP server.
- Step 8** In the Username field, enter the username or ID needed to access to the FTP server.
- Step 9** In the Password field, enter the password needed to access the file on the remote FTP server.

Step 10 In the Confirm Password field, enter the password again.

Step 11 To save the settings, click **Submit**.

Configuring Content Request Routing Using a Content Router or Routing Content Engine

This routing method requires a Cisco Content Router or routing Content Engine. The Content Router or routing Content Engine uses HTTP or RTSP redirection to route requests to the Content Engine which the Content Router or routing Content Engine determines is best suited to deliver the desired content to the client. This determination is based on the requested FQDN, the client's IP address, and the responsiveness of a Content Engine. The Content Router compares the source IP address of the client end system against a table of address ranges assigned to Content Engines, known as the coverage zone. The Content Router or routing Content Engine then sends the client a redirect response pointing to the selected Content Engine. Subsequently, the Content Engine can be configured as a proxy caching server to serve content, or it can be configured to distribute pre-positioned content.

To configure the Content Engine for proxy caching, see [Chapter 8, "Configuring Caching Services."](#) To configure your network for pre-positioned content distribution, see [Chapter 5, "Configuring the ACNS Network for Content Distribution."](#)



Note

References to the Content Router in this section include the routing Content Engine. See the ["Configuring a Content Engine as a Routing Content Engine" section on page 4-51.](#)

Content Router request routing works in the following manner:

1. The client sends a Domain Name System (DNS) query for the IP address of the domain name to its local DNS proxy server.
2. The local DNS proxy sends the DNS query to a name server to resolve the domain name to an IP address.
3. The resolution process eventually ends up at the Content Router. The IP address of the Content Router is returned to the local DNS proxy server.
4. The local DNS proxy server returns the IP address of the Content Router to the client.
5. The client sends a request for the desired content to the Content Router.
6. The Content Router sends back a redirect response pointing to a DNS domain name and eventually resolves the DNS domain name to the closest Content Engine IP address.
7. The client sends a request for the content to the best-suited Content Engine.
8. The best-suited Content Engine sends the content to the client.

For Content Router request routing to work properly, the following information must be configured:

- Coverage zone

The Content Router chooses the Content Engine closest to the client end system based on the coverage zone. (See the next section, ["Coverage Zones and Coverage Zone Files."](#))

- Website assignment

A Content Engine can only serve content for websites to which the Content Engine is assigned. Content Engines are assigned to websites through channels. (See the [“Adding and Removing Content Engines from Channels”](#) section on page 5-16.)

- DNS server

DNS entries for all fully qualified domain names (FQDNs) must be delegated to the Content Router. In the DNS server’s zone database file, you must enter a name server (NS) record that delegates the FQDN to the Content Router, and then restart your DNS server.

In the following example, the FQDN `www.cisco.com` has been delegated to the Content Router named `bali-cr`:

```
www.cisco.com      IN NS    bali-cr
```

When you have configured an NS record on your DNS server, all requests for the FQDN and any of its subdomains are sent to the Content Router rather than to existing DNS servers. The Content Router automatically resolves DNS queries for the FQDN.

Coverage Zones and Coverage Zone Files

A *coverage zone* is a mapping of end system IP addresses to Content Engines. The Content Router uses the Content Engine IP addresses to create a static redirection table that maps end system IP addresses to Content Engines and provides information on the proximity of end systems to Content Engines. When content is requested by an end user, the Content Router checks the end user’s IP address to find the coverage zone that contains that IP address. The Content Router then selects the Content Engine that is serving this coverage zone. If a specific IP address is in multiple coverage zones, then the one with the more specific range is selected. For example, the IP address `172.1.1.1` uses the data for coverage zone `172.1.0.0/16` instead of `172.0.0.0/8`. If there is no match found in the coverage zone data, then the request is not redirected. The Content Router sends back an error response, and the content is not available.

A coverage zone can be associated with one or more than one Content Engine: each Content Engine can have its own unique coverage zone, or Content Engines can be associated with more than one coverage zone and have overlapping coverage zones. In ACNS 5.x software, coverage zones cannot be associated with Content Distribution Managers or Content Routers.

If a coverage zone is served by multiple Content Engines, the Content Engine with the lowest metric value is put in the routing table. The metric value (see [Table 4-10](#)) indicates the proximity of the Content Engine to the end user. When there are multiple Content Engines in a coverage zone that are on the same subnet and have the same metric value, the Content Router performs the redirect to the client in a round-robin fashion.

When a Content Engine is registered to the Content Distribution Manager, it is assigned a default coverage zone (by default) that is determined by the IP address assigned to the Content Engine and its subnet mask. The default coverage zone configuration can be unassigned and network administrators can create and assign their own custom coverage zones by defining the coverage zone in a coverage zone file.

A *coverage zone file* is an XML file used to specify a user-defined coverage zone. Each coverage zone entry in the file specifies the IP address range, the name of the Content Engine serving this subnet, metric value, and one or more optional agent fields, if the entry is designated for specific routing Content Engines. In addition to the coverage zone information, two optional elements are created for documentation purposes: a revision value to specify the version of the coverage zone file and a customer name.

Coverage zone files can be created using any ASCII text editing tool. You can use a single coverage zone text-format file to define all the coverage zones for your ACNS network.

Table 4-10 defines the coverage zone file elements.

Table 4-10 Coverage Zone File Elements

Element	Value	Description
Network	IP address	Coverage zone IP address range.
Content Engine	Content Engine name	Specifies the Content Engines serving the coverage zone specified in the network element. This can have one or more elements.
Metric	Integer	Value indicating the proximity of the Content Engine to the end user. The lower the value, the closer the Content Engine is to the end user.
Agent	Name of routing Content Engine	Optional field that specifies a Content Engine acting as the routing Content Engine. If not specified, the coverage zone entry is designated for all Content Routers registered in the system.

Registering Coverage Zone Files

The system administrator places a coverage zone file where the Content Distribution Manager or individual devices can access the URL. The administrator then registers the coverage zone file URL in the Content Distribution Manager GUI. Coverage zone files can be applied globally to the entire ACNS network, or locally to a specific Content Router or routing Content Engine in the ACNS network. If a coverage zone file is made global, then it is read and parsed by each Content Router or routing Content Engine that does not have a coverage zone file assigned. If the coverage zone is specified in an individual Content Router or routing Content Engine configuration, it is only applied to that particular Content Router or routing Content Engine.

Choosing the Coverage Zone

You have the choice of using two types of coverage zones:

- Default coverage zones
- User-defined coverage zones

A default coverage zone consists of all the Content Engines that reside in the same local network segment, or subnet. The Content Distribution Manager GUI provides a check box to specify whether the default coverage zone is to be used.

A user-defined coverage zone consists of all the Content Engines that are specified in a coverage zone file. This file defines the network segments to be covered in the routing process. The coverage zone file is registered with the Content Distribution Manager and then applied to a Content Router or routing Content Engine for routing definitions.



Note

The metric value of a default coverage zone is set to 20. If a particular Content Engine is preferred for a user-defined coverage zone, the metric value in the coverage zone file should be set to a value less than 20. If a default coverage zone is preferred, then the metric value in the coverage zone file should be set to a value greater than 20.

Using the Default Coverage Zone

To use the default coverage zone, follow these steps:

-
Step 1	From the Content Distribution Manager GUI, choose Devices > Devices .
Step 2	Click the Edit icon next to the name of the Content Engine that you want to modify. The Device Home window appears.
Step 3	To display the entire table of contents, click the Show All button above the Contents pane.
Step 4	In the Contents pane, choose Device Activation. The Device Activation window appears with fields for editing the properties of the selected Content Engine. (See Figure 4-11 .)
Step 5	Check the Set Default Coverage Zone File check box.
Step 6	Click Submit .
-

Registering and Applying a User-Defined Coverage Zone

To apply a custom coverage zone to a Content Router or routing Content Engine, you first need to register a coverage zone file URL in the Content Distribution Manager GUI. After you have registered the coverage zone file URL with the Content Distribution Manager, you can apply the coverage zone file in one of two ways:

- Globally—Deployed across the entire ACNS network
- Locally—Deployed on a specific Content Router or routing Content Engine



Note

If you apply a coverage zone file locally for a device, this file will overwrite the ACNS network-wide coverage zone file for that device.

To register a coverage zone file, follow these steps:

-
Step 1	From the Content Distribution Manager GUI, choose System > Configuration .
Step 2	In the Contents pane, choose Routing > Coverage Zone File Registration . The Coverage Zone File Registration window appears.
Step 3	Click the Create New Coverage Zone File icon. The Registering New Coverage Zone File window appears.
Step 4	Choose a file import method: • Upload—The upload method allows you to upload a coverage zone file from any location that is accessible from your PC by using the browse feature. • Import—The import method allows you to import the coverage zone file from an external HTTP, HTTPS, FTP, or CIFS server.

When you choose a method, the window refreshes and displays the configuration fields that are associated with the method that you chose. (For descriptions of the fields associated with the upload method, see [Table 4-11](#). For descriptions of the fields associated with the import method, see [Table 4-12](#).)

Table 4-11 Upload Method for Coverage Zone Files

Property	Description
File Import Method	Sets the file import method. Choose Upload to display the required fields to upload the coverage zone file.
Coverage Zone File Upload	Local directory path to the coverage zone file. Use the Browse button can be used to help you locate the file.
Destination Filename	Name of the coverage zone file. This field is filled in automatically with the filename from the local directory path.

Table 4-12 Import Method for Coverage Zone Files

Property	Description
File Import Method	Sets the file import method. Choose Import to display the required fields to import the coverage zone file.
Coverage Zone File URL	Full path of the origin URL (including the coverage zone filename) where the coverage zone file can be accessed.
Destination File Name	Name of the coverage zone file.
Update Interval (minutes)	Frequency with which the Content Distribution Manager looks for changes to the coverage zone file located at the coverage zone file URL. The default value is 10 minutes.
UserName	Name of the user to be authenticated to fetch the coverage zone file.
Password	User password for fetching the coverage zone file.
NTLM user Domain	NTLM user domain name for NTLM authentication.
Disable Basic Authentication	When checked, NTLM headers cannot be stripped off to allow fallback to the basic authentication method. If you leave this check box unchecked, NTLM authentication headers can be stripped to allow fallback to the basic authentication method, and the username and password information can be passed to the origin server in clear text with a basic authentication header.


Step 5 To save the settings, click **Submit**.

After you have registered the coverage zone file URL with the Content Distribution Manager, you can use this file as your global routing configuration.

To set a global coverage zone file, follow these steps:

Step 1 From the Content Distribution Manager GUI, choose **System > Configuration**.

Step 2 In the Contents pane, choose **Routing > Global Routing Config**. The Set Global Coverage Zone File window appears. (See [Figure 4-13](#).)

Figure 4-13 Set Global Coverage Zone File Window

- Step 3** From the Coverage Zone File drop-down list, choose a coverage zone file.
- Step 4** In the Keep Alive Port field, enter a port number that is used by the ACNS network-wide network to listen in on routing requests.
- Step 5** To configure the time period (in seconds) for caching DNS replies, enter a number (0–2147483647) in the DNS TTL field. The default is 60 seconds.
- Step 6** Click **Submit**.

To apply a coverage zone file to an individual Content Router to be used locally, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Router to which you want to assign a coverage zone file.
- Step 3** To display the entire table of contents, click the **Show All** button above the Contents pane.
- Step 4** In the Contents pane, choose **Device Activation**. The Modifying Content Router window appears.
- Step 5** Choose a coverage zone file from the Coverage Zone File drop-down list.
- Step 6** Click **Submit** to apply the coverage zone file to the Content Router.

Coverage Zone File Examples


The following sections show different coverage zone file examples in four different scenarios.

Scenario 1: No Network Address Translation Firewall

This is the simplest scenario. In this example, enterprise ent1.com is in the public address space with two Content Engines in different subnets, and each Content Engine backs up the other. There is no Network Address Translation (NAT) firewall. Both Content Engines use the default coverage zone. This configuration creates entry 3 and entry 4 of the coverage zone file (see the example) to be added by the Content Router.

To define a backup Content Engine in the coverage zone file, map the subnet to its backup Content Engine, and assign a metric value greater than 20 (entry 1 and entry 2).

Figure 4-14 Coverage Zone Scenario 1: No NAT



The following file is an example of a coverage zone file for scenario 1. (See [Figure 4-14](#).)

```
<?xml version="1.0" ?>
<CDNNetwork>
<revision>1.0</revision>
<customerName>ent.com</customerName>
<!-- entry 1 -->
<coverageZone>
<network>172.29.248.0/24</network>
<CE>ent_sanfrancisco</CE>
<metric>20</metric>
</coverageZone>


<!-- entry 2: backup CE for the San Francisco office -->
<coverageZone>
<network >172.29.249.0/24</network>
<CE>ent_sanjose</CE>
<metric>20</metric>
</coverageZone>

<!-- entry 3 -->
<coverageZone>
<network>172.29.248.0/24</network>
<CE>ent_sanjose</CE>
<metric>10</metric>
</coverageZone>

<!-- entry 4 -->
<coverageZone>
<network>172.29.249.0/24</network>
<CE>ent_sanfrancisco</CE>
<metric>10</metric>
</coverageZone>
</CDNNetwork>
```

Scenario 2: Content Engine Behind a NAT Firewall

Enterprise ent2.com is behind a NAT firewall with one Content Engine, and the Content Router is on the public Internet. All content requests from this public IP address are served by the Content Engine (ent2_sanfrancisco) behind the firewall. Because all content requests from the end systems behind the NAT have the same public IP address, there is one entry to map the public IP address to the Content Engine ent2_sanfrancisco in the coverage zone file. The Set Default Coverage Zone File check box should be unchecked in the General Configuration section in the Modifying Content Engine window (see [Figure 13-3](#)) because the subnet 10.1.1.0/24 is private and is not known to the Content Router.

Figure 4-15 Coverage Zone Scenario 2: Behind a NAT Firewall

The following coverage zone file applies to [Figure 4-15](#).


```
<?xml version="1.0" ?>
<!-- CZ for ent2.com -->

<CDNNetwork>
<coverageZone>
<network>171.29.250.1/32</network>
<CE>ent2_sanfrancisco</CE>
<metric> 10 </metric>
</coverageZone>
</CDNNetwork>
```

Scenario 3: Multiple Content Engines Behind a NAT Firewall

In this scenario, enterprise ent3.com is behind a NAT firewall with multiple subnets, and each subnet has one Content Engine. In this case, there must be at least one routing Content Engine behind the NAT firewall. See the [“Modifying Content Engine Properties”](#) section on page 13-9 for more information on how to reconfigure a Content Engine to act as a routing Content Engine. A routing Content Engine performs the routing process and is able to do content request routing. The Content Router redirects all content requests from behind the firewall to the routing Content Engine, and this routing Content Engine then performs the second redirect.

In this example, ent3_sanjose1 is the routing Content Engine. A coverage zone entry (CZ entry 1) must be defined to map all content requests from behind the NAT firewall to ent3_sanjose1. Coverage zone entries for routing Content Engine ent3_sanjose1 must also be defined to map all requests from 10.1.1.0/24 to itself (CZ entry 2) and requests from 10.1.2.0/24 to ent3_sanjose2 (CZ entry 3).

Figure 4-16 Scenario 3: Multiple Content Engines Behind a NAT Firewall

The following coverage zone file applies to [Figure 4-16](#).

```
<?xml version="1.0" ?>
<!-- CZ for ent3.com -->
<CDNNetwork>
<customerName>ent3.com</customerName>
<!-- CZ entry 1 -->
<coverageZone>
<network>171.29.1.1/32</network>
<CE>ent3_sanjose1</CE>
<metric> 10 </metric>
</coverageZone>


<!-- CZ for routing CE -->
<!-- CZ entry 2 -->
<coverageZone>
<network>10.1.1.0/24</network>
<CE>ent3_sanjose1</CE>
<metric> 10 </metric>
<agent> ent3_sanjose1</agent>
</coverageZone>

<!-- CZ entry 3 -->
<coverageZone>
<network>10.1.2.0/24</network>
<CE>
ent3_sanjose2
</CE>
<metric> 10 </metric>
<agent>ent3_sanjose1</agent>
</coverageZone>
```

Scenario 4: NAT Firewall with Multiple IP Addresses

Enterprise ent4.com is behind a NAT firewall with multiple public IP addresses. All content requests from behind the NAT firewall have one of these public IP addresses. Therefore, there must be one coverage zone data entry for each of these public IP addresses to map these IP addresses to the Content Engine behind the NAT firewall.

Figure 4-17 Scenario 4: NAT Firewall with Multiple IP Addresses



The following coverage zone file applies to [Figure 4-17](#).

```
<?xml version="1.0" ?>
<!-- CZ for ent4.com -->
<CDNNetwork>
  <coverageZone>
    <network>171.29.10.1/32</network>
    <CE>ent4_sanfrancisco</CE>
    <metric>10</metric>
  </coverageZone>

  <!-- CZ for ent4.com -->
  <coverageZone>
    <network>171.29.10.2/32</network>
    <CE>ent4_sanfrancisco</CE>
    <metric>10</metric>
  </coverageZone>
</CDNNetwork>
```

Configuring a Content Engine as a Routing Content Engine

You can enable routing capabilities on a Content Engine to allow this Content Engine to also redirect content requests. A routing Content Engine is required when multiple Content Engines are located behind a NAT firewall to be able to serve content requests from the clients behind the same NAT firewall. In this case, a Content Router first redirects the content requests to a routing Content Engine rather than to the Content Engine best-suited to deliver the content. The routing Content Engine then performs a

second redirect to find the best-suited Content Engine to serve the content. See the [“Scenario 3: Multiple Content Engines Behind a NAT Firewall”](#) section on page 4-49 for an example that illustrates the need for a routing Content Engine.

See the [“Modifying Content Engine Properties”](#) section on page 13-9 for information on how to configure a Content Engine to serve as a routing Content Engine.

Configuring Dynamic Content Routing Using a Content Router

In previous releases of ACNS software, Content Routers used a static coverage zone file to describe the preferred routing path between Content Engines and client end systems.

A coverage zone is a mapping of client end-system IP addresses to Content Engines. The Content Router uses the Content Engine IP addresses to create a static redirection table that maps end-system IP addresses to Content Engines and provides information on the proximity of end systems to Content Engines. When content is requested by a client, the Content Router checks the client IP address to find the coverage zone that contains that IP address. The Content Router then selects the Content Engine that serves this coverage zone.

In some ACNS network environments, Content Engine IP addresses keep changing, and coverage zones are dynamic instead of static. In such cases, the Content Router cannot create a static routing table, and it cannot successfully route the content.

When the following conditions are present, the content cannot be routed successfully by using static coverage zone tables in the Content Router:

- Multiple Content Engines are deployed in multiple locations.
- Each location contains a NAT firewall.
- One Content Router serves all locations.
- One root Content Engine serves all locations.
- Each location is configured with two uplink lines to the Internet for redundancy.
- Uplink lines for different locations can share an external public IP address pool so that the same IP address can be used by NAT firewalls in different locations at different times.

With multiple uplinks to the Internet, requests for content from clients and Content Engines that are in the same location can go out to the Content Router with different external IP addresses. The Content Router that is using static coverage zone files cannot accommodate sharing the same IP address pool among different locations.

In the ACNS 5.3.3 software release and later releases, the Content Router can detect changes in Content Engine coverage zones and can dynamically adjust its routing tables.

Enabling Dynamic Content Routing

For each Content Router that you want to configure for dynamic content routing, you must indicate that dynamic content routing will be used over static content routing.

To configure dynamic content routing, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Next to the name of the Content Router that you want to configure, click the **Edit** icon.
 - Step 3** In the Contents pane, choose **Device Activation**. The Device Activation window appears.

- Step 4** To indicate that dynamic content routing will be used over static content routing, check the **Dynamic Content Routing** check box.
- Step 5** To save the Content Router settings, click **Submit**.

Changing the Content Engine Metric Value for Dynamic Content Routing in the Coverage Zone File

In ACNS 5.3.3 software and later releases, a tag has been defined in the coverage zone file that allows you to change the metric that assigns a preference to a particular Content Engine in the routing table. This optional tag is the <dynamic> tag.

When a coverage zone file containing the <dynamic> tag and elements is assigned to a Content Router, the following applications are made:

- If dynamic content routing is enabled on a Content Router, the Content Router follows the specifications inside the <dynamic> tag, and it ignores the specifications in the <coverageZone> tag.
- If dynamic content routing is not enabled, the Content Router follows the specifications in the <coverageZone> tag, and it ignores the specifications in the <dynamic> tag.

The <dynamic> tag is a standalone tag; it is not defined as a subelement of the <coverageZone> tag. The subelements <CE> and <metric> must be defined within the <dynamic> tag. Table 4-13 describes the elements of the <dynamic> tag.

Table 4-13 Coverage Zone File Elements for Dynamic Content Routing

Tag Name	Elements	Value	Description
<dynamic>			
	<CE> ¹	Content Engine name	Specifies the Content Engine for which the metric value is to be applied. The <dynamic> tag can contain the names of multiple Content Engines.
	<metric> ¹	Number	Value indicates the proximity of the Content Engine to the end user. The lower the value, the greater the preference given to the Content Engine. If no metric is specified, the default value of 10 is applied.

1. This element is required.

The following example shows how the <dynamic> tag can be used in a coverage zone file:

```
<?xml version="1.0"?>
<CDNNetwork>
<revision>1.0</revision>
<dynamic>
<CE>hostname_of_ce</CE>-----> NOTE: The tag is <CE> and not <ce>.
<metric>number</metric>
</dynamic>
</CDNNetwork>
```

The following example shows an invalid coverage zone file where the <dynamic> tag is written as a subelement of the <coverageZone> tag:

```
<?xml version="1.0"?>
<CDNNetwork>
<revision>1.0</revision>
<coverageZone>
<dynamic>
<CE>hostname_of_ce</CE>
<metric>number</metric>
</dynamic>
</coverageZone>
</CDNNetwork>
```

The following example shows that the <dynamic> tag and the <coverageZone> tag can exist together in a coverage zone file:

```
<?xml version="1.0"?>
<CDNNetwork>
<revision>1.0</revision>
<coverageZone>
<network>a.b.c.d/mask</network>
<CE>hostname_of CE</CE>
<metric>0</metric>
</coverageZone>
<dynamic>
<CE>name_of_ce</CE>
<metric>number</metric>
</dynamic>
</CDNNetwork>
```

The following example shows that the <dynamic> tag can appear multiple times in a coverage zone file:

```
<?xml version="1.0"?>
<CDNNetwork>
<revision>1.0</revision>
<dynamic>
<CE>hostname_of_ce1</CE>
<metric>number</metric>
</dynamic>
<dynamic>
<CE>hostname_of_ce2</CE>
<metric>number</metric>
</dynamic>
</CDNNetwork>
```

The following example shows that the <dynamic> tag can contain the names of multiple Content Engines:

```
<?xml version="1.0"?>
<CDNNetwork>
<revision>1.0</revision>
<dynamic>
<CE>hostname_of_ce1</CE>
<CE>hostname_of_ce2</CE>
<CE>hostname_of_ce3</CE>
<metric>number</metric>
</dynamic>
</CDNNetwork>
```

You can use a coverage zone file with the `<dynamic>` tag, for example, when a preference order has to be given for a set of Content Engines that are serving the same client base. The following example shows that `ce1` is preferred over `ce2` because the metric value for `ce1` is lower than that of `ce2`:

```
<?xml version="1.0"?>
<CDNNetwork>
  <revision>1.0</revision>
  <dynamic>
    <CE>hostname_of_ce1</CE>
    <metric>20</metric>
  </dynamic>
  <dynamic>
    <CE>hostname_of_ce2</CE>
    <metric>30</metric>
  </dynamic>
</CDNNetwork>
```

Configuring Load-Based Content Routing Using a Content Router

When you enable load-based content routing, the Content Router redirects client requests to the Content Engine that reports the lowest average CPU load, given that each Content Engine in the routing table has the same metric value (or weight).

When you have multiple Content Engines that are defined with unequal weighting, you can configure threshold limits for CPU load, disk usage, and WMT stream count so that the Content Router redirects the client requests to the next preferred Content Engine that has not exceeded its threshold.

When a configured threshold is exceeded, messages are sent to the Content Router. Content Router algorithms compare the Content Engine assigned weight and current load to the configured threshold values when making routing decisions.

Enabling Load-Based Routing on the Content Router

To enable the Content Router for load-based routing, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon next to the name of the Content Router that you want to configure for load-based routing.
 - Step 3** In the Contents pane, choose **Device Activation**. The Modifying Content Router window appears.
 - Step 4** To enable load-based routing, check the **Least Loaded** check box.
 - Step 5** To save the settings, click **Submit**.
-

To enable load-based routing on the Content Router from the CLI, use the **contentrouting leastloaded** Content Router global configuration command.

Configuring Content Engine Threshold Limits for Load-Based Routing

To configure threshold limits on the Content Engine, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices** or **Devices > Device Groups**.
- Step 2** Click the **Edit** icon next to the Content Engine or device group that you want to configure.
- Step 3** In the Contents pane, choose **General Settings > Notification and Tracking > Service Monitor**. The Service Monitor Settings for Content Engine window appears. (See Figure 4-18.) Table 4-14 describes the fields in this window and lists the corresponding CLI commands.

Figure 4-18 Service Monitor Settings Window

The screenshot shows the 'Service Monitor Settings for Content Engine, CONTENTENGINE' window. The left pane shows the 'Contents' tree with 'Service Monitor' selected under 'Notification and Tracking'. The main pane displays the following settings:

Service Monitor Settings		
Current settings: None (Using Factory Defaults)		
CPU Settings		
Enable:	☐	
Threshold:	80	% (1-100)
Sample Period:	5	(seconds) (1-60)
Number of Samples:	60	(1-120)
DISK Settings		
Enable:	☐	
Sample Period:	5	(seconds) (1-60)
Number of Samples:	60	(1-120)
WMT Settings		
Enable:	☐	
Threshold:	90	% (1-100)
Sample Period:	5	(seconds) (1-60)

At the bottom right, there are 'Submit' and 'Cancel' buttons.

- Step 4** To monitor the Content Engine CPU load, configure the following service monitor settings:
- To enable collecting CPU load information from the Content Engine, check the **Enable** check box.
 - To configure a value that determines when the Content Engine is overloaded, enter a value (1–100) in the Threshold field. The threshold determines the extent of CPU usage allowable and is the percentage of total CPU capacity in use.
 - To set the time interval (in seconds) between two consecutive sample periods, enter a value (1–60) in the Sample Period field. The sample period is the time during which the Content Engine and the Content Router exchange Keep Alives that contain the Content Engine load information.
 - To configure the number of latest sampled values to be used by the Content Engine to calculate load information, enter a number (1–120) in the Number of Samples field.

- Step 5** To monitor the Content Engine disk load, configure the following service monitor settings:
- To enable collecting disk transaction information from the Content Engine, check the **Enable** check box. When the disk service monitor is enabled, the Content Engine samples the disk I/O transactions per second and converts the number to a value that is used in calculating the Content Engine load information.
 - To set the time interval (in seconds) between two consecutive sample periods, enter a value (1–60) in the Sample Period field.
 - To configure the number of latest sampled values that will be used by the Content Engine to calculate load information, enter a number (1–120) in the Number of Samples field.
- Step 6** To monitor the WMT stream load, configure the following service monitor settings:
- To enable collecting WMT stream count information from the Content Engine, check the **Enable** check box.
 - To configure the maximum allowable WMT streams, enter a value (1–100) in the Threshold field. This value is the percent of streams for which the Content Engine has been either configured or licensed.
 - To set the time interval (in seconds) between two consecutive sample periods, enter a value (1–60) in the Sample Period field.
 - To configure the number of latest sampled values to be used by the Content Engine to calculate load information, enter a number (1–120) in the Number of Samples field.
- Step 7** To save the settings, click **Submit**.

Table 4-14 Service Monitor Settings

GUI Parameter	Function	CLI Command
CPU Settings		
Enable	Allows the Content Router to collect CPU load information from the Content Engine.	contentrouting servicemonitor type cpu
Threshold	Value that determines when the Content Engine is overloaded. The range is 1–100. The default is 80 percent.	contentrouting servicemonitor threshold cpu percentage
Sample Period	Time interval (in seconds) between two consecutive samples. The range is 1–60. The default is 5 seconds.	contentrouting servicemonitor sampleperiod cpu seconds
Number of Samples	Number of most recent sampled values that will be used when calculating the average. The range is 1–120. The default is 60.	contentrouting servicemonitor numberofsamples cpu number
Disk Settings		
Enable	Allows the Content Router to collect disk transaction information from the Content Engine.	contentrouting servicemonitor type disk
Sample Period	Time interval (in seconds) between two consecutive samples. The range is 1–60. The default is 5 seconds.	contentrouting servicemonitor sampleperiod disk seconds

Table 4-14 Service Monitor Settings (continued)

GUI Parameter	Function	CLI Command
Number of Samples	Number of most recent sampled values that will be used when calculating the average. The range is 1–120. The default is 60.	contentrouting servicemonitor numberofsamples disk number
WMT Settings		
Enable	Allows the Content Router to collect stream count information from the Content Engine.	contentrouting servicemonitor type wmt
Threshold	Percentage of streams for which the Content Engine has been either configured or licensed. The range is 1–100. The default is 90 percent.	contentrouting servicemonitor threshold wmt percentage
Sample Period	Time interval (in seconds) between two consecutive samples. The range is 1–60. the default is 5 seconds.	contentrouting servicemonitor sampleperiod wmt seconds

Content Router Request Overflow Routing

In the ACNS 5.4 implementation of load-based content routing, when all Content Engines in a domain became overloaded and the configured resource utilization thresholds were exceeded, the Content Router began dropping requests.

In the ACNS 5.5 release, when the original domain is overloaded, client requests can be redirected to an overflow or alternate domain. This feature is applicable in the following two cases:

- Load-based routing is configured and all Content Engines in the domain are overloaded.
- All Content Engines in the domain are offline. (Load-based routing does not have to be configured in this instance.)

When request overflow routing is not configured for the domain requested, requests that exceed Content Engine load-based thresholds, or that are otherwise undeliverable, are redirected to an error notification message, and the requests are dropped.

Request overflow routing works dynamically when Content Engines are overloaded or deactivated. When the load of one or more Content Engines in the original host domain is reduced below threshold limits or Content Engines are reactivated, new requests are routed to the original host domain automatically.

Request overflow routing supports requests from RTSP, HTTP, and MMS clients.



Note

The request overflow routing feature must be configured on a per-domain basis.

Defining Alternate Domains for Request Overflow Routing

To configure request overflow routing on a Content Router or a routing-enabled Content Engine, you must create a list of domain and alternate domain pairs for the Content Router by using the **contentrouting lastresort** global configuration command. You can configure only one alternate domain for each Content Router FQDN.

In the following example, the original domain `cr.abc.com` will be redirected to `def.com` if all the Content Engines assigned to the `cr.abc.com` domain are overloaded and exceed their thresholds:

```
CR(config)# contentrouting lastresort domain cr.abc.com alternate def.com
```

The substitute URL will be based on the type of client issuing the request, as follows:

- An HTTP client requesting `http://cr.abc.com/test/test.html` will be redirected to `http://def.com/test/test.html`.
- An RTSP client requesting `rtsp://cr.abc.com/test/test.rm` will be redirected to `rtsp://def.com/test/test.rm`
- An MMS client requesting `http://cr.abc.com/test/test.asf.asx` will be redirected in the following order:
 - a. `rtsp://def.com/test/test.asf`
 - b. `mms://def.com/test/test.asf`
 - c. `http://def.com/test/test.asf`

To remove the overflow configuration and replace the alternate domain with another domain, use the **no** form of the command, and then reconfigure, as shown in the following example:

```
CR(Config)# no contentrouting lastresort domain cr.abc.com alternate def.com
```

```
CR(Config)# contentrouting lastresort domain cr.abc.com alternate cad.com
```

Viewing the Overflow Domain List

To view the overflow domain list configured on a Content Router or routing-enabled Content Engine, use the **show content-routing lastresort EXEC** command.

Troubleshooting Content Router Configurations

Because there are quite a few configuration steps required for the Content Router to redirect the request properly, you might see some content request errors from the Content Router when the configuration is not quite complete. Here are some areas to look at when troubleshooting:

- DNS delegation
 - Is the requested domain delegated to the Content Router on the DNS server that is authoritative for the parent domains? The Content Router's DNS name should be forward resolvable. Check with the system administrator to delegate a domain.
- Content Router routing properties
 - Is the Content Router activated? See the [“Activating Devices in the Content Distribution Manager GUT” section on page 2-22](#) to activate a Content Router.

- Is a default coverage zone set for a Content Engine, or is there an ACNS network-wide coverage zone file or a local coverage zone file set for the Content Router or routing Content Engine? See the [“Registering Coverage Zone Files”](#) section on page 4-43 to set a coverage zone file. See the [“Using the Default Coverage Zone”](#) section on page 4-44 to set a default coverage zone.
- If a coverage zone entry is for a routing Content Engine, does the “agent” field contain the name of the routing Content Engine? See the [“Coverage Zones and Coverage Zone Files”](#) section on page 4-42 for a description of agent fields in coverage zone files. See the [“Scenario 3: Multiple Content Engines Behind a NAT Firewall”](#) section on page 4-49 for an example of a coverage zone entry for a routing Content Engine.
- Is the content request from an end system covered by a Content Engine in a coverage zone based on the default coverage zone or the coverage zone file? This Content Engine is the “serving Content Engine.” See the [“Coverage Zones and Coverage Zone Files”](#) section on page 4-42 under [“Configuring Content Request Routing Using a Content Router or Routing Content Engine”](#) in Chapter 4 for information on how to select a serving Content Engine for a client.
- Is the serving Content Engine activated? See the [“Activating Devices in the Content Distribution Manager GUI”](#) section on page 2-22 to activate a Content Engine.
- Is there a channel created for the requested domain and a serving Content Engine assigned to this channel? See the [“Creating and Modifying Channels”](#) section on page 5-10.
- Is the serving Content Engine alive? Use the CLI **show statistics content-routing ce cename** command to show the status of a Content Engine. (Refer to the *Cisco ACNS Software Command Reference, Release 5.1* publication.)
- Content pre-positioned on a Content Engine
 - Is there a manifest file assigned to the channel associated with the serving Content Engine? See the [“Working with Manifest Files”](#) section on page A-6.
 - Is the manifest file accessible from the Content Distribution Manager? See the [“Creating and Modifying Channels”](#) section on page 5-10.
 - Is there any syntax error in the manifest file? See the [“Manifest File Structure and Syntax”](#) section on page A-33.
 - Is the requested content specified in the manifest file? See the [“Specifying a Single Content Item”](#) section on page A-6.
 - If the requested content is streaming media, is the application server enabled with the correct license key? See the [“Enabling RealProxy”](#) section on page 9-10, the [“Enabling RealSubscriber”](#) section on page 9-12, and the [“Enabling the Cisco Streaming Engine”](#) section on page 9-14.
 - Has disk space been configured for the mediafs to store streaming media objects or the cdnfs to store nonstreaming content? See the [“Configuring Disk Space”](#) section on page 2-19.
 - Has the requested content been successfully acquired by the Content Engine? See the [“Listing Website Content Using the Spider Script”](#) section on page A-80.

Where to Go Next

You have finished setting up the request routing infrastructure for your ACNS network so that client requests for content can be routed to the Content Engine that can best serve the content. You are now ready to configure your Content Engines to serve the content. Content Engines can be configured to serve content as proxy caching servers or they can be configured to distribute pre-positioned content. To configure the Content Engine for proxy caching, see [Chapter 8, “Configuring Caching Services.”](#)

Configuring your network for pre-positioned content distribution includes the following tasks:

- Building the location tree
- Assigning Content Engines to locations
- Creating a directory of content providers
- Defining websites for each content provider
- Creating channels
- Assigning Content Engines to each channel
- Assigning a root Content Engine to each channel

To accomplish these tasks, proceed to the next chapter, [Chapter 5, “Configuring the ACNS Network for Content Distribution.”](#)




Configuring the ACNS Network for Content Distribution

Once the hardware devices that make up your ACNS network are properly installed and configured, you are ready to use the information contained in this chapter to organize those devices into an ACNS network that is capable of serving pre-positioned content or live media to end users.

In the ACNS network, content is replicated through a channel distribution architecture. Content in channels can be transmitted by unicast pull or, if multicasting is enabled, by multicast push. This chapter provides information about configuring the network elements for channel distribution and multicasting. It contains the following sections:

- [Preparing for Channel Distribution of Content, page 5-3](#)
- [Configuring Channel Elements, page 5-6](#)
- [Configuring Channels for Channel Distribution, page 5-10](#)
- [Distributing Content Through Unicast Replication, page 5-20](#)
- [Distributing Content Through Multicast Replication, page 5-25](#)
- [Troubleshooting the Distribution Hierarchy, page 5-47](#)
- [Where to Go Next, page 5-48](#)

[Figure 5-1](#) shows the task flow for configuring channels for unicast and multicast content distribution.

Figure 5-1 *Distributing Content in the ACNS Network*

Preparing for Channel Distribution of Content

Channel distribution describes the manner in which content is distributed from the origin web server to all the Content Engines that are subscribed to the content or, more specifically, assigned to the channel that represents that group of content objects. (For an overview of the channel distribution architecture, see the [“Content Acquisition and Distribution Architecture”](#) section on page 1-11.)

In preparing for channel distribution of content, you should configure the following network elements in this order:

1. Locations—[Creating and Modifying Locations](#)
2. Content providers—[Creating and Modifying Content Providers](#)
3. Websites—[Creating and Modifying Websites](#)
4. Channels—[Creating and Modifying Channels](#)

Creating and Modifying Locations

Locations are set up in the Content Distribution Manager GUI, and they organize and group Content Engines into virtual networks for distribution of content through channels.

Locations need to be configured before you can activate Content Engines and Content Routers and bring them online in the ACNS network. A default location is created automatically when you use the Content Distribution Manager Quick Start tool to activate your devices. If you wish to create a new location or modify the default location, use the procedure in this section.

**Note**

For more information on working with locations, see the [“Working with Device Locations”](#) section on page 13-5.

To create a new location, follow these steps:

Step 1

From the Content Distribution Manager GUI, choose **Devices > Locations**. The Locations window appears. (See [Figure 5-2](#).)

Figure 5-2 **Locations Window**



Step 2 Click the **Create New Location** icon in the taskbar. The Creating New Location window appears. (See [Figure 5-3](#).)

Figure 5-3 *Creating New Location Window*

The screenshot shows the Cisco Application and Content Networking System (ACNS) web interface. The main navigation bar includes 'Devices', 'Services', and 'System'. The 'Locations' tab is selected. The 'Creating new Location' window is open, displaying the 'Location Information' section. The 'Name' field is marked as required with an asterisk. The 'Parent Location' dropdown is set to 'None'. The 'Level' is set to 1. Below the 'Location Information' section is a 'Comments' section with a text area. At the bottom of the window, there is a 'Note: * - Required Field' and 'Submit' and 'Cancel' buttons.

- Step 3** Enter a location name in the Name field.
- Step 4** Choose a parent location (or choose **None**) from the Parent Location drop-down list.
- A location with no parent location assigned is a level 1 location. A location with a level 1 parent becomes a level 2 location, and so forth. The location level is displayed after you choose a parent location (or choose None) and click **Submit** to save the configuration.
- Step 5** If you wish, enter comments about the location in the Comments field.
- Step 6** Click **Submit**.

To modify an existing location, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Locations**. The Locations window appears. (See Figure 5-2.)
- Step 2** Click the **Edit** icon next to the name of the location you want to modify.
- Step 3** Change the location name, parent location, or description as needed.
- Step 4** To save the changes, click **Submit**.
-

Configuring Channel Elements

This section describes how to create and modify content providers and websites. If you wish to associate a company name with its website, you must create a content provider directory listing. Websites must be set up in your ACNS network before you can create channels.

Creating and Modifying Content Providers

When you are creating a channel, you must choose the name of a content provider from the Content Provider drop-down list, or you can choose **None**. To populate this list, you must create a directory listing with the name, address, and contact information of a content provider. Content provider listings are created in the Content Distribution Manager GUI.

To create a content provider listing, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Content Provider**.
 - Step 2** Click the **Create New Content Provider** icon in the taskbar. The Creating New Content Provider window appears.
 - Step 3** Enter the company name. You also have the option to enter the address and contact information.
 - Step 4** Click **Submit**.
-

To modify a content provider listing, follow these steps:

-
- Step 1** Choose **Services > Web > Content Provider**, and click the **Edit** icon next to the content provider that you want to modify. The Modifying Content Provider window appears. (See [Figure 5-4](#).)

Figure 5-4 **Modifying Content Provider Window**

Modifying Content Provider, CP1

Content provider information

Name:

Street:

City: Postal Code:

State/Province: Country:

Primary Contact

Name:

Phone:

E-mail:

Secondary Contact

Name:

Phone:

E-mail:

Comments

Note: * - Required Field

- Step 2** Enter any changes you want to make to the company name and address, the primary contact information, or the optional secondary contact information.



Note Clicking **Cancel** returns you to the Modifying Channel window.

- Step 3** Click **Submit**.

Deleting Content Providers

You can delete content providers as needed.



Note When you delete a content provider, all websites and channels that have been added for the content provider are also deleted.

To delete a content provider, follow these steps:

- Step 1** Choose **Services > Web > Content Provider**, and click the **Edit** icon next to the content provider that you want to delete. The Modifying Content Provider window appears. (See [Figure 5-4](#).)

- Step 2** Click the **Trash** icon. You are prompted to confirm your decision.
- Step 3** Click **OK** to execute your request.

Creating and Modifying Websites

Websites must be defined in your ACNS network before you can create channels. Channels then map the website content to the Content Engines in your ACNS network. The Content Distribution Manager GUI allows you to define websites in your ACNS network. A website listing must provide the URL for the origin web server that stores the website content.

To create a website, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Websites**.
- Step 2** Click the **Create New Web Site** icon in the task bar. The Creating New Web Site window appears. (See [Figure 5-5](#).)

Figure 5-5 Creating New Web Site Window

- Step 3** Enter the name, content provider (or choose **None**), and origin server for the website. (See [Table 5-1](#) for a description of website properties.)

Table 5-1 Website Properties

Property	Description
Name* ¹	Unique name for the website.
Content Provider*	Name of the company that is providing the content. The default is None . Note The content provider must be predefined in the Creating New Content Provider window. (See the “ Creating and Modifying Content Providers ” section on page 5-6.)
Origin Server*	Fully qualified domain name of the origin server where the content resides.
Request Routed FQDN	Fully qualified domain name to route requests when using a Content Router.
WMT Authentication Type	Basic authentication or NTLM authentication. Note This setting <i>must</i> match the setting on the origin server for the stream to play correctly.
Pure DNS Routing	When checked, uses only the DNS server specified in the incoming packets for DNS requests. This causes the DNS caching name service to determine the original destination IP address. When unchecked (the default), allows the DNS caching name service to use the servers configured on the Content Engine and allows WCCP interception of DNS requests.
Ignore case of URLs during playback	When checked, allows case-insensitive URL matching during the playback of pre-positioned content.

1. * = required field. All other fields are optional.

Step 4 If you wish, choose a WMT authentication type from the **WMT Authentication Type** drop-down list.



Note This setting *must* match the setting on the origin server for the stream to play correctly.

Step 5 Check the **Pure DNS Routing** check box if you want to use only the DNS server specified in the incoming packets for DNS requests. This causes the DNS caching name service to determine the original destination IP address.

Step 6 Check the **Ignore case of URLs during playback** check box if you want to allow case-insensitive URL matches between the acquisition URL and the publishing URL.

Step 7 To save the configuration, click **Submit**.

To modify a website, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Websites**.
 - Step 2** Click the **Edit** icon of the website that you want to modify. The Modifying Web Site window appears.
 - Step 3** Enter changes to the website data as needed.
 - Step 4** To save the changes, click **Submit**.
-

Deleting Websites



Note

When you delete a website, all channels that have been added for the website are also deleted.

To delete a website, follow these steps:

-
- Step 1** To delete a website, go to the Modifying Web Site window.
 - Step 2** Click the **Trash** icon. You are prompted to confirm your decision.
 - Step 3** To execute your request, click **OK**.
-

Configuring Channels for Channel Distribution

Channel distribution describes the manner in which content is distributed from the origin web server to all the Content Engines that are subscribed to the content, or more specifically, assigned to the channel that represents that group of content objects. (For an overview of the channel distribution architecture, see the [“Content Acquisition and Distribution Architecture”](#) section on page 1-11.)

Creating and Modifying Channels

Channels map the content from a website to the devices in your ACNS network. When you create a channel, you first create a directory of your content providers and then provide URLs to their websites. (See the [“Creating and Modifying Content Providers”](#) section on page 5-6 and the [“Creating and Modifying Websites”](#) section on page 5-8.)

Using the Content Distribution Manager, you can create, modify, or remove channels from the ACNS network. Keep in mind that modifying channels can affect the availability of content on the ACNS network.

You create a channel by defining the following items:

- Name
- Content provider
- Website
- Channel quota
- Channel properties (channel priority, multicasting)

Creating a Channel

To create a channel, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**. The Channels window appears, listing all the channels in your ACNS network.
- Step 2** Click the **Create New Channel** icon in the taskbar. The Creating New Channel window appears. (See [Figure 5-6](#).)

Figure 5-6 Creating New Channel—Channel Properties

The screenshot shows the 'Creating New Channel' window in the Cisco Application and Content Networking System GUI. The window is titled 'Cisco Application and Content Networking System - Microsoft Internet Explorer'. The main navigation bar includes 'Devices', 'Services', and 'System'. The 'Services' tab is selected, and the 'Channels' sub-tab is active. The 'Creating New Channel' window is open, showing the 'Channel Information' and 'Acquisition and Distribution Properties' sections. The 'Channel Information' section includes fields for Name, Request Routed FQDN, Origin Server, Content Provider (set to None), Website (set to Please make a choice), Live Channel (checkbox), and Channel Quota (set to MB). The 'Acquisition and Distribution Properties' section includes fields for Distribution Priority (set to Medium), Use null cipher for Distribution (checkbox), Multicast only (radio button), Multicast/Unicast (radio button), Unicast only (radio button), Root CE failover/fallback grace period (set to 120 mins), Use system-wide settings for QoS for unicast data (checkbox), and QoS value for unicast data (set to Default). There is also a Comments section at the bottom.

- Step 3** Use the fields provided under the Channel Information heading to configure the channel name and website with which the channel is associated.
- Step 4** Use the fields provided under the Acquisition and Distribution Properties heading to configure channel priority, enable multicasting, or disable encryption. (See [Table 5-2](#) for a description of the channel properties fields. Required fields are indicated by an asterisk in the GUI and in [Table 5-2](#).)

Table 5-2 Channel Properties

Property	Description
Channel information	
Name* ¹	Unique name for the channel.
Request Routed FQDN	Fully qualified domain name to route requests. Read only. This field is configured from the Creating New Web Site window (see the “Creating and Modifying Websites” section on page 5-8) and is used if you have a Content Router in your network.
Origin Server	Fully qualified domain name of the origin server. Read only. This field is configured from the Creating New Web Site window. (See the “Creating and Modifying Websites” section on page 5-8 .)
Content Provider*	Name of the content provider associated with this channel. The default is None . To add a new content provider to this drop-down list, see the “Creating and Modifying Content Providers” section on page 5-6 .
WebSite*	Name of the website belonging to the content provider that is associated with this channel. Note The website must be predefined in the Creating New Web Site window.
Live Channel	When checked, creates a live channel to distribute live or scheduled programs to the Content Engines associated with this channel. This channel does not have a related manifest file and cannot be used to distribute file-based content as regular channels do. The live channel learns about a live program through a program file that describes the attributes of the program. Checking this check box disables the Channel Quota field and fields in the Acquisition and Distribution Properties section.
Channel Quota*	Maximum content storage size in megabytes for pre-positioning content for this channel. See the “Channel Quota” section on page 6-25 .
Acquisition and Distribution Properties	
Distribution Priority*	Content distribution priority setting. Options are High, Normal, and Low. The default is Normal. See the “Distribution Priority” section on page 6-24
Multicast only, Multicast/Unicast, or Unicast only*	Enables multicast-only, multicast transmission with failover to unicast, or unicast-only transmission. See the “About IP Multicasting” section on page 5-25 .

Table 5-2 *Channel Properties (continued)*

Property	Description
Use null cipher for Distribution	When checked, disables encryption for distribution.
Root CE failover/fallback grace period*	Number of minutes before a root Content Engine failover or a temporary root Content Engine fallback occurs. The range is 20 to 120 minutes. (See the “About the Root Content Engine Failover and Fallback Grace Period Setting” section on page 5-14.)
Never	When checked, Content Engine failover or fallback will never occur.
Use system-wide settings for QoS for unicast data	When checked, applies the system-wide QoS settings for unicast data to the channel. (See the “Configuring System-Wide Distribution QoS Settings Based on Channel Priority” section on page 20-16.) To override the system-wide QoS settings with channel-specific QoS values, leave this check box unchecked, and configure the channel-specific QoS values in the QoS value for unicast data field.
QoS value for unicast data*	Configures a Differentiated Services Code Point (DSCP) value for the QoS. (See Table 5-3 .) If you choose Other , enter a decimal value in the corresponding field. Note Cisco ACNS software allows you to set QoS settings on a per-channel basis and a system-wide global configuration basis. Channel settings take precedence over global settings.

1. * = required field. All other fields are optional.

- Step 5** If you wish to add comments regarding the channel for the benefit of other Content Distribution Manager users, enter them in the Comments field provided.
- Step 6** To save the channel configuration, click **Submit**.

[Table 5-3](#) lists and describes the DSCP values for unicast QoS.

Table 5-3 *DSCP Values*

DSCP value	Description
af11	Sets packets with AF11 dscp (001010)
af12	Sets packets with AF12 dscp (001100)
af13	Sets packets with AF13 dscp (001110)
af21	Sets packets with AF21 dscp (010010)
af22	Sets packets with AF22 dscp (010100)
af23	Sets packets with AF23 dscp (010110)

Table 5-3 DSCP Values

af31	Sets packets with AF31 dscp (011010)
af32	Sets packets with AF32 dscp (011100)
af33	Sets packets with AF33 dscp (011110)
af41	Sets packets with AF41 dscp (100010)
af42	Sets packets with AF42 dscp (100100)
af43	Sets packets with AF43 dscp (100110)
cs1	Sets packets with CS1(precedence 1) dscp (001000)
cs2	Sets packets with CS2(precedence 2) dscp (010000)
cs3	Sets packets with CS3(precedence 3) dscp (011000)
cs4	Sets packets with CS4(precedence 4) dscp (100000)
cs5	Sets packets with CS5(precedence 5) dscp (101000)
cs6	Sets packets with CS6(precedence 6) dscp (110000)
cs7	Sets packets with CS7(precedence 7) dscp (111000)
default	Sets packets with default dscp (000000)
ef	Sets packets with EF dscp (101110)

About the Root Content Engine Failover and Fallback Grace Period Setting

Cisco ACNS software supports failover of the root Content Engine to another Content Engine in your network when the root Content Engine is down or not functioning.

The secondary Content Engine polls the root Content Engine at preset intervals to make sure that the root Content Engine is functioning properly. If the polling of the root Content Engine does not return a positive response, the secondary Content Engine waits for the grace period and then takes over as the root Content Engine until the end of the grace period. The time interval before a Content Engine takes over as the root Content Engine is set by using the Root CE failover/fallback grace period option in the Create a New Channel window. You can choose a time interval between 20 and 120 minutes using increments of 10.

The acting root Content Engine keeps polling the original root Content Engine to see if the root Content Engine is back online. When the acting root Content Engine obtains a positive response from the original root Content Engine, the acting root Content Engine waits for the fallback grace period and then goes back to its original mode if the positive responses are consistent. When the acting root Content Engine goes back to its original mode, the original root Content Engine takes over the duties of the root Content Engine. Both failover and fallback grace period values are configured in the same field, and use the same value.



Note

To configure manifest file and manifest proxy information for a channel, see the [“Configure Manifest File and Proxy Information for the Channel”](#) section on page 6-26.

Modifying a Channel

To modify a channel, follow these steps:

- Step 1** Choose **Services > Web > Channels** to get to the Channels window.
- Step 2** Click the **Edit** icon next to the name of the channel that you want to change. The Modifying Channel window appears. (See [Figure 5-7](#).)

Figure 5-7 Modifying Channel Window

- Step 3** Make changes to the channel configuration as needed.
To modify the Content Provider or WebSite fields, click the **Edit Content Provider** or **Edit Website** button. (See the [“Configuring Channel Elements”](#) section on page 5-6.)
- Step 4** To effect any changes you have made to the channel configuration, click **Submit**. The Content Distribution Manager updates the channel information.

After you create a new channel, you need to designate a root Content Engine for the channel and then assign Content Engines or device groups to the channel. Although designating a root Content Engine and assigning devices to the channel can be done from the same GUI window, the following sections break out these steps into separate procedures.

Designating the Root Content Engine

The root Content Engine is the one Content Engine that is authorized to go directly to the origin web server for content. The root Content Engine then publishes the content to other Content Engines in the channel. You must designate a root Content Engine for content distribution to take place.

To designate one Content Engine to be the root Content Engine for a channel, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**.
- Step 2** Click the **Edit** icon next to the name of the channel to which you want to assign a root Content Engine. The Modifying Channel window appears. (See [Figure 5-7](#).)
- Step 3** You can designate a root Content Engine when assigning device groups to the channel or when assigning Content Engines to the channel.
- In the Contents pane, click **Assign Content Engines**. The Content Engine assignment to Channel window appears. Click the **Assign** button next to the name of the Content Engine, or click the **Assign All Content Engines** icon in the task bar and proceed to [Step 4](#).
 - Alternatively, click **Assign device groups**. The Device Group assignments for Channel window appears. Click the **Assign** button next to the name of the Content Engine, or click the **Assign All Device Groups** icon in the task bar and proceed to [Step 4](#).
- Step 4** Choose a Content Engine from the Root CE drop-down list. This list contains all of the Content Engines currently assigned to the channel in which you are working.
- 

Note If there are no Content Engines or device groups assigned to the channel, no choices appear in the Root CE drop-down list. To add Content Engines to a channel, see the next two sections, “[Adding and Removing Content Engines from Channels](#),” and “[Adding and Removing Device Groups from Channels](#).”
-
- Step 5** Click **Submit**.
-

Adding and Removing Content Engines from Channels

To add or remove Content Engines from a channel, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**.
- Step 2** Click the **Edit** icon next to the name of the channel that you want to change. The Modifying Channel window appears. (See [Figure 5-7](#).)
- Step 3** In the Contents pane, choose **Assign Content Engines**. The Content Engine assignments for Channel window appears.
- This window displays a list of the Content Engines in all locations to which they have been assigned, as well as the unreserved and total disk space in the cdnfs partition for each Content Engine. Content Engines can be added to the channel from more than one location.
- Step 4** Click the **Assign** icon (blue cross mark) next to the name of the Content Engine that you wish to associate with the channel. Alternatively, you can click the **Assign all Content Engines** icon in the taskbar to perform a bulk addition of all Content Engines in various locations to the channel.

- Step 5** To add the Content Engines to the channel, click **Submit**. Upon submission, a green tick mark appears next to the assigned Content Engine.
- Step 6** To remove Content Engines, click the **Unassign** icon (green tick mark) next to the name of the Content Engine that you want to remove from the channel. Alternatively, you can click the **Unassign all Content Engines** icon in the taskbar to perform a bulk removal of all Content Engines from the selected channel.
- Step 7** To remove the Content Engines from the channel, click **Submit**. The list of Content Engines in the channel shows a blue cross mark next to the name of the removed Content Engine.

If the Content Engine that you removed was the root Content Engine, and if there is at least one Content Engine still assigned to the channel, you must designate a new root Content Engine. (See the [“Designating the Root Content Engine”](#) section on page 5-16.)

Adding and Removing Device Groups from Channels

Device groups are assigned to channels through the Content Distribution Manager GUI. Whenever a channel is created and additional device groups are added, or when a channel assignment to the device group changes, devices in the group are notified of their assignment to the associated channel.

The relationship between device groups and channels is many-to-many. A channel can have multiple device groups and device groups, in turn, can belong to multiple channels.

To add a device group from a channel, follow these steps.

- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**.
- Step 2** Click the **Edit** icon next to the name of the channel to which you want to assign a device group or from which you want to remove a device group. The Modifying Channel window appears. (See [Figure 5-7](#).)
- Step 3** In the Contents pane, choose **Assign Device Groups**. The Device Group assignments for Channel window appears, displaying the minimum unreserved disk space in the cdnfs partition. An asterisk (*) under the Note column denotes that there is no free space available to be assigned to this channel. A double asterisk (**) under the Note column denotes that Content Engines associated with this device group have been already assigned to this channel.
- Step 4** Click the **Assign** icon (blue cross mark) next to the name of the device group that you wish to associate with the channel. A green tick mark appears next to the assigned device group after you click **Submit**. Alternatively, you can click the **Assign all Device Groups** icon in the taskbar to perform a bulk addition of all device groups in various locations to the channel.
- Step 5** To add the device groups to the channel, click **Submit**.



Note

When a device group is removed from a channel, the Content Engines that were part of the device group are also removed from the channel. When a device is removed from a device group containing the original channel assignment, this device is also unassigned from channels. Similarly, when a channel is removed from a device group, the associated devices are also unassigned.

To remove a device group from a channel, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**.
 - Step 2** Click the **Edit** icon next to the name of the channel to which you want to assign a device group or from which you want to remove a device group. The Modifying Channel window appears. (See [Figure 5-7](#).)
 - Step 3** In the Contents pane, choose **Assign Device Groups**. The Device Group assignments for Channel window appears, displaying the minimum unreserved disk space in the cdnfs partition. An asterisk (*) under the Note column denotes that there is no free space available to be assigned to this channel. A double asterisk (**) under the Note column denotes that Content Engines associated with this device group have been already assigned to this channel.
 - Step 4** Click the **Unassign** icon (green tick mark) next to the name of the device group that you want to remove from the channel. Alternatively, you can click the **Unassign all Device Groups** icon in the taskbar to perform a bulk removal of all device groups from the selected channel.
 - Step 5** To remove the device groups from the channel, click **Submit**. The browser window refreshes, listing the updated channels with a blue cross mark next to the unassigned device groups.
-

Priming a Live Channel Using the Channels Interface

The first client requesting a program often experiences the longest wait time for the program to begin playing. Users can experience long wait times because of the full RTSP negotiation that is required to pull the live stream from the source. Delays can also occur if the edge Content Engine has not buffered enough stream data to fill the media player's buffer at the time that the program is requested. When the buffer is not filled, some data to the client might be sent at the suboptimal line rate instead of at the Fast Start rate.

The ACNS 5.5 software release allows channels for unicast managed live programs to be primed for faster start up times. When a live channel is primed, a unicast-out stream is pulled from the origin server to a Content Engine before a client ever requests the stream. When the first request for the stream goes out, the stream is already in the channel.

To view the status of a primed channel, use the **show programs EXEC** command. Live unicast programs that have been configured to play forever will display current status as "Scheduled" when the channel is unprimed, whereas in previous releases that do not support channel priming, the status is displayed as "Playing." The status output has been changed to differentiate primed and unprimed channel programs in the ACNS 5.5 software release. (For more information, see the ["Priming a Live Channel Using the Programs Interface"](#) section on page 7-16.)

To configure channel priming, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**.
 - Step 2** Click the **Edit** icon next to the name of a live channel. The Modifying Channel window appears. (See [Figure 5-7](#).)
 - Step 3** In the Contents pane, choose **Assign Content Engines**. The Content Engine assignments for Channel window appears.
 - Step 4** Alternatively, choose **Assign Device Groups**. The Device Group assignments for Channel window appears.
 - Step 5** If you have not already done so, choose a root Content Engine from the Root CE drop-down list in the taskbar.
-

- Step 6** Click the **Assign** icon (blue cross mark) next to the name of the Content Engine or device group that you wish to associate with the channel.
- Step 7** Under the column heading Primed, check the check box. The check box is active only after a device or device group has been assigned to the channel.
- Step 8** To save the assignments, click **Submit**.

You can also find the check box to prime a channel by navigating to one of the following windows in the Content Distribution Manager GUI:

- **Devices > Devices > Assignments > Channels**
- **Devices > Device Groups > Assignments > Channels**

Deleting Channels

To delete a channel, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**.
- Step 2** Click the **Edit** icon next to the name of the channel that you want to delete. The Modifying Channel window appears. (See [Figure 5-7](#).)
- Step 3** Click the **Trash** icon. You are prompted to confirm your decision to delete the channel.
- Step 4** Click **OK** to confirm your decision. The channel is removed from the ACNS network.
-

Viewing Channels

To view all channels across your ACNS network, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**.
- The Channels window appears, listing all the channels in your ACNS network.
- The channel type and the website that is associated with each channel are displayed in the row for each channel.
- Step 2** To sort the channels by type, website, status, or state, click the respective column heading.
- Step 3** To print the channel display table, click the **Printer** icon.
-

Viewing All Content Engines Assigned to a Channel

In the Content Distribution Manager GUI, a distinction is made between devices that are assigned directly to a channel or devices that are assigned through a device group. The assignments are tracked separately. For tracking purposes, the View all Content Engines assigned to Channel window displays all Content Engines assigned to a channel, regardless of whether they are assigned directly or through device groups.

To view the list of all Content Engines assigned to a channel, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**.
- Step 2** Click the **Edit** icon next to the name of the channel that you want to view. The Modifying Channel window appears.
- Step 3** In the Contents pane, choose **List all assigned Content Engines**. The View all Content Engines assigned to Channel window appears and lists all Content Engines assigned to the specified channel. [Table 5-4](#) describes the information that is displayed in this window.

Table 5-4 Content Engine Channel Assignment Information

Column Heading	Description
Content Engine	Name of the Content Engine assigned to the channel.
Directly Assigned	Indicates whether the Content Engine was assigned directly to the channel (Yes or No).
Assigned through device group	Indicates whether or not the Content Engine was assigned through a device group. If yes, then this column displays the name of the device group. If no, then this column displays the value Not Applicable (N/A).

Distributing Content Through Unicast Replication

Once content is acquired from the origin server by the root Content Engine of a channel, it can be replicated through the channel either by unicast or multicast transmission. The basic channel distribution architecture provides for unicast content replication using a hop-by-hop, store-and-forward methodology with the forwarder Content Engines systematically selected on the basis of the manually configured location hierarchy.

To distribute content through unicast, the ACNS network automatically creates a unidirectional distribution tree for each channel. The root node of the tree is the root Content Engine of the channel, and each Content Engine subscribed to the channel is a node on the tree.

For each node, its parent node is also called its forwarder Content Engine. The algorithm for automatically designating the forwarder Content Engine is called the channel routing algorithm.

Three general rules in the current channel routing algorithm are as follows:

1. In each location for each channel, only one Content Engine fetches content from another location for that channel. We call this Content Engine the location leader of the channel. All other Content Engines in this location use the location leader as the forwarder for this channel. There can be only one location leader per channel per location. Note that within one location, different channels may have different location leaders.

The location leader is computed automatically by the channel routing algorithm. ACNS 5.5 software enables you to influence which Content Engines are chosen as location leaders by setting a priority level for individual Content Engines in the location. Content Engines are all set to the lowest priority by default. The Content Engine with the highest priority is more likely to be chosen as the location leader. (See the next section, [“Configuring Location Leader Preference and Forwarder Selection Probability”](#).)

Use the **show distribution channel EXEC** command to see which Content Engine is the current forwarder for a channel. The reason/status field in the command output shows why a Content Engine is unable find a forwarder. Use the **show distribution forwarder-list EXEC** command to see the forwarder selection order of a Content Engine for a channel.

2. The location leader finds a subscribed Content Engine from the closest location on the path toward the root Content Engine as its forwarder. If all the potential forwarders in a parent location are down (or unreachable) the location leader skips to the next location level in the hierarchy (towards the root Content Engine location) to find a forwarder.

ACNS 5.5 software allows you to set the number of levels that a location leader can search to find a forwarder. (See the [“Configuring the Forwarder Lookup Level”](#) section on page 5-22.)

You can also influence the selection of a forwarder by configuring the Content Engine with a distribution weight, by setting it to be used only as a last resort, or by taking it out of the forwarder list completely. The Content Engine configured with the heaviest weight has the highest probability of becoming the forwarder. (See the next section, [“Configuring Location Leader Preference and Forwarder Selection Probability”](#).)

3. If the location leader Content Engine fails for some time, another Content Engine in the location takes over as the location leader. If the root Content Engine fails, another Content Engine in the location takes over as the temporary root Content Engine.

When multicast transmission is desirable, channels and Content Engines can be enabled for multicasting. (See also the [“Configuring Weighted Load Balancing for Live Stream Splitting”](#) section on page 9-38.)

Configuring Location Leader Preference and Forwarder Selection Probability

To configure the location leader preference and forwarder probability settings, follow these steps:


-
Step 1	From the Content Distribution Manager GUI, choose Devices > Devices .
Step 2	Click the Edit icon next to the name of the Content Engine that you want to configure.
Step 3	In the Contents pane, choose Prepositioning > Location Leader and Forwarder . The Location Leader and Forwarder Settings for Content Engine window appears.
Step 4	To configure the location leader settings, click and slide the control button on the scale to the priority level (1–20) that you want. The lower the number, the lower the priority. The Content Engine with the highest priority is most likely to be chosen as the location leader.
Step 5	To configure the forwarder settings, choose an option (Weighted , Last Resort , Never Use) by clicking one of the Forwarder Probability radio buttons. If you choose Weighted , click and slide the control button on the scale to the weight (1–20) that you want. The Content Engine with the greatest weight is most likely to be chosen as the forwarder.
Step 6	To save the settings, click Submit .
-

Configuring the Forwarder Lookup Level

To configure the forwarder lookup level, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure.
- Step 3** In the Contents pane, choose **Prepositioning > Distribution**. The Distribution Settings for Content Engine window appears. (See [Figure 5-8](#).)

Figure 5-8 *Distribution Settings*



- Step 4** Click one of the Forwarder Lookup radio button options; chose **Unlimited** or chose **Maximum**, and enter a number to limit how many location levels the location leader can search to find a forwarder.
- Step 5** To save the settings, click **Submit**.

Configuring the Content Metadata Poll Interval Multiplier

Sender and receiver Content Engines exchange information before they transmit the actual content. This information is called *metadata*. Included in the metadata is the wait time between content updates (or the metadata polling interval), which the sender Content Engine sends to the receiver. The sender calculates this wait time based on a number of factors, such as the number of devices in the ACNS network or the number of receivers in the forwarder's queue. You can lengthen or shorten the wait time

by configuring the content metadata poll interval multiplier on receiver Content Engines. This multiplier enables you to control how often receiver Content Engines notice changes in content. The polling interval that is calculated by the sender Content Engine will be either shorter or longer by a factor of the content metadata poll interval multiplier that you configure.

When the receiver Content Engine obtains the wait time from the sender, it multiplies the wait time by the value you specify in the multiplier. You can scale the unicast receiver Content Engine's next metadata poll interval down to 10 percent of normal or up to 10 times normal. For example, if you set the multiplier to 10 percent of normal, and the wait time between update polls is 120 seconds, the effective polling interval is shortened to 12 seconds (10 percent of normal or 0.1×120 is 12 seconds).

The higher the multiplier, the longer it takes for the receiver to be notified of changes in content and the less the network traffic that results from polling. On the other hand, the lower the multiplier, the less time it takes for the receiver to be notified of content changes, and the greater the network traffic that results from polling.

**Note**

The content metadata poll interval multiplier applies only to a unicast receiver Content Engine, which regularly polls for new content metadata from its forwarder.

To configure the content metadata poll interval, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure.
 - Step 3** In the Contents pane, choose **Prepositioning > Distribution**. The Distribution Settings for Content Engine window appears. (See [Figure 5-8](#).)
 - Step 4** Under the Distribution Settings heading, slide the **Content Metadata Poll Interval Multiplier** control either down or up to scale the unicast receiver Content Engine's next metadata poll interval.
 - Step 5** To save the settings, click **Submit**.
-

Configuring Unicast Content Distribution Through a NAT Firewall

ACNS 5.5 software supports communication between two Content Engines that are located behind different NAT devices. Content Engines can be configured for unicast content distribution with a NAT address and port, as well as with a primary IP address and port.

For example, if a forwarder Content Engine is behind a NAT device, you can set its NAT address on the Content Distribution Manager GUI, and the Content Distribution Manager will propagate the forwarder Content Engine NAT address to all the devices in the ACNS network so that the receiver Content Engine can set up a connection with the forwarder using the NAT address.

NAT address configuration also supports the following types of communications between Content Engines and other devices in a centrally managed network:

- Receiver Content Engine polling for content metadata from a sender
- Unicast receiver polling for data files from a unicast sender
- Multicast receiver sending a negative acknowledgment (NACK) to a multicast sender
- Content Distribution Manager polling for configuration updates from Content Engines and Content Routers



Note ACNS 5.5 software does not support communication between a primary multicast sender and a backup multicast sender through a NAT firewall.

NAT configuration logic works in the following manner: A receiver Content Engine first attempts to contact a forwarder Content Engine using the Content Engine primary IP address. If the receiver cannot contact the forwarder Content Engine by using the primary IP address, the receiver attempts to contact the forwarder Content Engine using the Content Engine NAT IP address, if a NAT IP address has been configured. If the receiver successfully initiates communication with the forwarder using the NAT IP address, the receiver will retry the primary IP address periodically and fall back to using the primary IP address, if it becomes operational.

To configure a NAT IP address for a Content Engine, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine for which you want to configure a NAT IP address. The Device Home window appears.
- In the Contents pane, chose **Device Activation**. The Device Activation window appears with fields for editing the properties of the selected Content Engine.
- Step 3** Under the NAT Configuration heading, configure the NAT settings.
- To use the IP address configured on the primary interface of the Content Engine as the management IP address, check the **Use CE's primary IP Address** check box.
 - To allow the Content Distribution Manager to communicate with devices in the ACNS network that are behind the NAT firewall using an explicitly configured IP address, enter the NAT address of the Content Engine in the NAT Address field.
 - In the Port field, enter the port number for the NAT address.



Note If the Content Distribution Manager cannot contact a device using the primary IP address, it will attempt to communicate using NAT IP address.

- Step 4** To save the configuration, click **Submit**.
-

Commands Related to Distribution Settings

Table 5-5 lists and describes CLI commands that are related to distribution settings. For complete descriptions of these commands, including usage guidelines and examples, refer to the *Cisco ACNS Software Command Reference, Release 5.5* publication.

Table 5-5 Distribution-Related Commands

Command	Syntax	Description
show distribution	show distribution location-leader-preference channel-id <i>channel-id</i>	Shows the configured location-leader-priority of the Content Engines that are subscribed to the channel.
show distribution	show distribution forwarder-load-weight channel-id <i>channel-id</i>	Shows the configured forwarder-load-weight of the Content Engines that are subscribed to the channel.
show distribution	show distribution channel	Shows which Content Engine is the current forwarder for a channel.
show distribution	show distribution forwarder-list	Shows the forwarder selection order of a Content Engine for a channel.

Distributing Content Through Multicast Replication

Multicasting allows efficient distribution of content to multiple Content Engines and is useful when many end users are interested in the same content. ACNS software supports Pragmatic General Multicast (PGM)-based multicast replication using either satellite or multicast-enabled terrestrial infrastructures.

To prepare for multicast transmission of content, you should perform these tasks in the order presented in this section:

1. [Enabling Content Engines for Multicasting](#)
2. [Configuring Multicast Cloud Properties](#)
3. [Assigning Content Engines to a Multicast Cloud](#)
4. [Assigning and Removing Multicast Clouds from Channels](#)

About IP Multicasting

Multicast delivery enables the distribution of streaming media by allowing different receiving devices on the IP multicast to receive a single stream of media content from the Content Engine simultaneously. This type of delivery can save significant network bandwidth consumption because a single stream is sent to many devices, rather than sending a single stream to a single device every time that this stream is requested.

This multicast delivery feature is enabled by setting up a multicast address on the Content Engine to which different devices, configured to receive content from the same channel, can subscribe. The delivering device sends content to the multicast address set up at the Content Engine, from which it becomes available to all subscribed receiving devices.

To take advantage of multicasting, all devices, including Content Engines and routers, as well as clients, must be multicast-enabled. For this reason, multicasting is mostly used in local networks where routers can be configured for multicasting. Multicast delivery over the Internet can only be accomplished when all devices that participate in the multicast have been enabled for multicasting.

About Multicast Clouds

For multicast content replication, Content Engines are grouped into multicast clouds. A *multicast cloud* consists of one sender Content Engine, an optional backup sender Content Engine, and at least one receiver Content Engine in a hub and spoke topology. All Content Engines in one multicast cloud share a unique advertisement address that allows them to communicate multicast session information. The multicast cloud is then associated with one or more multicast-enabled channels.


In pull-based unicast content distribution, a unicast receiver pulls file data out of the proper forwarder (or root Content Engine) when a client requests the content. By contrast, in multicast content distribution, the sender Content Engine in a multicast cloud proactively pushes content into the cloud according to a preconfigured schedule.

The receiver Content Engines listen on the advertisement IP address for information on content to be replicated from the sender, and then the receiver Content Engines decide whether or not to accept an advertisement and receive the corresponding content.

Content *metadata* (machine-readable information that describes the characteristics of the content) must be distributed to a receiver first before the content itself can be replicated. Content metadata helps to define what content to retrieve, how content will be retrieved, how recently content has been updated, how content is to be pre-positioned (for example, expiration time), and so forth. Metadata is always distributed using unicast. Content, however, can be replicated using either multicast or unicast. A multicast receiver rejects the multicast sender's advertisement of a file if the proper content metadata has not yet arrived.

The Content Distribution Manager GUI allows you to configure multicast cloud parameters in the Creating New Multicast Cloud window. These parameters are discussed in the following sections.

[Figure 5-9](#) shows the task flow for configuring a multicast cloud.

Figure 5-9 Multicast Cloud Configuration Task Flow

About Configuring Carousel Passes

The carousel pass setting defines the maximum number of times that the sender Content Engine can retransmit the multicast for missing files. The primary sender sends the first carousel pass automatically. After the first round, multicast receiver Content Engines request missing content by sending a negative acknowledgment (NACK) to the sender that identifies the missing content. The multicast sender sends out the requested content when it receives the NACK from the receiver. After all receiver Content Engines have received all the multicast content or the sender has reached the maximum number of carousel passes, whichever comes first, the sender stops transmitting content.

If the primary sender fails and the backup sender becomes active, the backup sender takes charge of NACK processing. The backup sender's carousel passes are always triggered by a NACK. When the maximum number of carousel passes is reached for a file on the current active sender, you can configure file distribution to fall back to unicast, if this is desired. (See the [“About Configuring Channel Options for Content Replication”](#) section on page 5-32 for more information.)

To adjust the pacing of the multicast transmission, you can specify how much time must elapse before missing files are resent (the Delay between passes field) in the Content Distribution Manager GUI Modifying Multicast Cloud window (**Devices > Multicast Clouds > Modifying Multicast Cloud**).

About Multicast Logging Enhancements

The ACNS 5.3.5 software release includes enhanced multicast logging to identify the receiver Content Engine that is sending the retransmission requests (NACKs) and to identify why a file is scheduled for multicasting. These logging enhancements provide the following details:

- Any NACKs that are received by the multicast sender Content Engine are logged at the trace level in the dist-meta-sender error log on the multicast sender Content Engine.
- Any NACKs that are related to preparing a file for scheduling are logged in the transaction log.
- Any file that is scheduled for multicasting has the details about why it was scheduled for multicasting. You can obtain the details from the time-based queue or the priority-based queue.

About the NACK Interval Multiplier

To identify missing content and trigger a resend of a file, receiver Content Engines send a negative acknowledgement (NACK) message to the sender Content Engine. NACK messages generated by many receiver Content Engines could generate more traffic than the sender can handle. ACNS 5.x software allows you to adjust the average interval between NACKs by configuring a NACK interval multiplier for an individual receiver Content Engine. This value (an integer between 0.1 and 10) adjusts the default average NACK interval (the default is 20 minutes) by the value configured as the interval multiplier. For example, if you set the NACK interval multiplier to 3, the interval between NACKs becomes 20 minutes x 3, or 60 minutes.

This adjustment can be made as needed by choosing **Devices > Devices > Prepositioning > Distribution** in the Content Distribution Manager GUI.

To send an immediate NACK request rather than wait for the scheduled interval, you can issue the **distribution multicast send-nack-now** EXEC command on a multicast receiver Content Engine.

To configure the NACK interval multiplier, follow these steps:

-
- Step 1** In the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure.
 - Step 3** In the Contents pane, choose **Prepositioning > Distribution**. The Distribution Settings for Content Engine window appears. (See [Figure 5-8](#).)
 - Step 4** Click and drag the Content NACK Interval Multiplier slider control across the calibrated ruler to adjust the interval between NACK messages. The scale ranges from 10 percent of normal to 10 times normal. The center of the scale corresponding to “normal” denotes the default value of 20 minutes. The value corresponding to the position of the slider is displayed in the field next to it.
 - Step 5** To save the settings, click **Submit**.
-

About the File Resend or Stop Feature

To trigger resending of a particular file, or the content in a particular channel or all channels, you can issue the **distribution multicast resend** EXEC CLI command on a multicast sender. This command internally resets the number of carousel passes completed to zero for the desired target. Because the first resend of content always happens automatically without being triggered by a NACK, setting the carousel pass count to zero triggers a resend of the content. However, if the **on-demand-only** option is additionally specified, the first carousel pass made after the command has been issued can only be triggered by a NACK for the content involved.

To stop sending a particular file, a particular channel, or all channels, you can issue the **distribution multicast stop** EXEC command on a multicast sender. This command internally increases the number of carousels completed to the maximum integer, preventing further carousel passes from being triggered.

About Configuring Multicast Forward Error Correction

Forward error correction (FEC) is a type of data encoding that protects transmissions against errors, without requiring retransmission. The FEC number denotes the number of packets that will be encoded into one FEC transmission group. When the FEC number goes up, the transmission group becomes larger, so the multicast may be more error-resistant. However, there will be more computational overhead on the multicast sender and receivers. Note that there is no bandwidth overhead related to FEC.

In ACNS 5.0.1 software, the FEC default value is 8, whereas in ACNS 5.0.3 software and later, the FEC default is 16. If the multicast sender device is a high-end Content Engine model such as a CE-7325, you can set this number higher in order to improve multicast reliability, especially when your network connectivity has a high uniform loss rate. However, we do not recommend that you set this number beyond 64 because it may place too much of a load on all the receiver Content Engines.

Separately, you can also set proactive FEC using the PGM configuration file (a text file accessible from the Content Engine CLI in the /local1/multicast-expert-config/ directory). (See the [“Configuring PGM and File Transmission Parameters Using Multicast Expert Mode”](#) section on page 5-30.) Proactive FEC is the number of extra packets that the multicast sender proactively sends out for every FEC number of data packets. The proactive FEC default value is 0. You can set it higher for better multicast reliability, for example, 2 proactive packets for every 16 FEC packets, at the expense of 12.5 percent traffic overhead (2 divided by 16).

Proactive FEC is an additional reliability measure above and beyond that of normal FEC. Although normal FEC does not incur bandwidth overhead, proactive FEC does. Proactive FEC primarily protects the multicast from uniform losses. For example, if the network has a uniform loss rate of 15 percent, then a proactive FEC of 2 extra packets for every 16 FEC packets (a 12.5 percent bandwidth overhead) cuts the effective loss rate down to 2.5 percent. Most network losses are not completely uniform. Still, during bursts, proactive FEC similarly undercuts the effective burst loss rate. For example, if the burst loss rate is 20 percent while the average loss rate is 2 percent, with proactive FEC at 12.5 percent, the receiver Content Engines experience a burst loss rate of 7.5 percent and an average loss rate near 0 percent.

About Configuring the Multicast Sender Delay Interval

The multicast sender delay interval is the amount of time before each multicast transmission begins. A period of delay before the actual multicast transmission begins is required to allow content metadata time to propagate to the receiver Content Engine. Metadata contains the content file and configuration information that is necessary for the successful transmission of content files. The default sender delay interval is 16 minutes. The sender delay parameter is used to configure an extra delay before a multicast transmission can begin.

When configuring the sender delay interval, you must take into account that the content metadata must first be propagated to the receiver before the multicast transmission can commence. During a multicast session, a receiver Content Engine sends out periodic requests for files that it has not yet received. The sender retransmits files only as requested by the receiver Content Engine. A multicast receiver will reject a multicast sender's advertisement of a file if the associated content metadata has not arrived yet. The sender delay option allows you to configure enough time for the metadata to propagate to the receiver and avoid having the receiver reject the multicast sender's advertisement of a file.

To configure the sender delay interval, use the **multicast sender-delay** global configuration command on a sender Content Engine. You can configure the duration of the delay based on your expectation or best guess of the amount of time required for metadata to be propagated. The system takes this user-configured sender delay value and delays the multicast transmission by such a period.

**Note**

The sender delay interval cannot be configured using the Content Distribution Manager GUI. You must configure the sender delay interval using the CLI of the sender Content Engine.

Configuring PGM and File Transmission Parameters Using Multicast Expert Mode

PGM is a reliable multicast protocol defined in IETF RFC 3208. It is designed for applications that require ordered or unordered, duplicate-free, multicast data delivery from multiple sources to multiple receivers. Support for reliable multicasting and file transmission is provided by the TIBCO SmartPGM FX tool set, which is integrated with Cisco ACNS software, beginning with the ACNS 5.2 release. You can configure some PGM and file transmission (FX) parameters through the Content Distribution Manager GUI, such as the advertisement IP address, multicast-out bandwidth, TTL, FEC transmission group, and so forth.

In some cases, expert users might want to change other PGM and FX parameters to make the multicast file transfer more robust and efficient for their multicast environment. ACNS software allows you to change the configuration parameters of the TIBCO SmartPGM FX configuration file manually by using multicast expert mode on the Content Engine.

**Caution**

We do not recommend that you change the TIBCO SmartPGM FX configuration file unless you are an expert in PGM multicasting and know how to adjust the configuration parameters.

ACNS software contains default TIBCO SmartPGM FX configuration files for multicast sender and receiver Content Engines. The multicast sender and receiver Content Engines determine the medium being used for the multicast by checking the configuration of the multicast cloud, and they then read the configuration parameter values from the PGM configuration file that corresponds to the medium (terrestrial or satellite) being used for the multicast.

ACNS 5.5 software uses the following default TIBCO SmartPGM FX configuration files:

- fxd.conf.src
- fxd.conf.rcv

The Content Engine stores sample copies of the default TIBCO SmartPGM FX configuration files in the /local1/multicast-expert-config/ directory for reference. You can modify one of these sample configuration files, rename it, and save it back to the multicast sender or receiver Content Engine. When you create customized configuration files and copy them to the /local1/multicast-expert-config/ directory using the default configuration filenames, the Content Engine uses the customized configuration file instead of the default configuration file. The customized configuration file becomes effective after the Content Engine is restarted.

Content Engines running ACNS 5.5 software contain the following sample configuration files:

- `fxdSatellite.conf.src.sample`—Use for a sender Content Engine in a satellite network
- `fxdSatellite.conf.rcv.sample`—Use for a receiver Content Engine in a satellite network
- `fxdTerra.conf.src.sample`—Use for a sender Content Engine in a terrestrial network
- `fxdTerra.conf.rcv.sample`—Use for a receiver Content Engine in a terrestrial network

To change the configuration parameters of a default TIBCO SmartPGM FX configuration file, follow these steps:

Step 1 Log in to the multicast sender Content Engine using FTP.

- From a PC running Windows, choose **Start > Run**.
- In the Open field, enter **ftp ipaddress**, using the IP address of the multicast sender or receiver Content Engine.
- At the User prompt, enter your administrator-level username.
- At the password prompt, enter your password. The FTP prompt (`ftp>`) appears.

Step 2 At the FTP prompt, open the multicast-expert-config directory.

```
ftp> cd multicast-expert-config
250 CWD command is successful.
ftp>
```

Step 3 List the sample configuration files in the directory.

```
ftp> ls
200 PORT command successful.
150 Opening ASCII mode data connection for file list.
fxdSatellite.conf.rcv.sample
fxdSatellite.conf.src.sample
fxdTerra.conf.rcv.sample
fxdTerra.conf.src.sample
226 Transfer complete.
ftp: 297 bytes received in 0.01Seconds 29.70Kbytes/sec.
```

Determine which sample file you want to retrieve, based on whether you are configuring the multicast sender or the receiver and whether your network is using terrestrial or satellite media transmission.

Step 4 Return to binary mode from ASCII mode.

```
ftp> bin
200 Type set to I.
```

Step 5 Copy the configuration file to your desktop.

```
ftp> get fxdSatellite.conf.rcv.sample
200 PORT command successful.
Opening BINARY mode data connection for fxdSatellite.conf.rcv.sample (5607 bytes).
226 Transfer complete.
ftp: 5607 bytes received in 0.00Seconds 5607000.00Kbytes/sec.
ftp>
```

Step 6 End the FTP session.

```
ftp> quit
```

Step 7 Locate the configuration file on your PC and open it using any text editor.

Step 8 Edit the sample configuration file by making any needed changes.

Step 9 Save the file using **Save As**, and give it the same name as the default configuration file that you want to replace. For example, save the file named fxdSatellite.conf.rcv.sample as fxd.conf.rcv.

Step 10 Transfer the file back to the multicast sender or receiver Content Engine multicast-expert-config directory.

- a. Log in to the Content Engine using FTP.
- b. Open the multicast-expert-config directory.
- c. Enter binary mode.
- d. Transfer the file into the directory. For example:

```
C:\>ftp 128.19.220.79
Connected to 128.19.220.79.
220 CONTENTENGINE FTP server (Version wu-2.7.0(2) Tue Sep 7 17:20:20 P
DT 2004) ready.
User (128.19.220.79:(none)):admin
331 Password required for admin.
Password:
230 User admin logged in. Access restrictions apply.
ftp> cd multicast-expert-config
250 CWD command successful.
ftp> bin
200 Type set to I.
ftp> put fxd.conf.rcv
200 PORT command successful.
150 Opening BINARY mode data connection for fxd.conf.rcv.
226 Transfer complete.
ftp:5607 bytes sent in 0.01Seconds 560.70Kbytes/sec.
ftp> quit
```

Step 11 Restart the multicast sender or receiver Content Engine for the new configuration to take effect.

About Configuring Channel Options for Content Replication

Once a multicast cloud is configured, it must be associated with one or more multicast-enabled channels for multicast replication of content to take place within the channel. The channel configuration offers various transmission options for replicating content; a channel can be configured for multicast and unicast (multicast with failover to unicast), multicast-only, or unicast-only transmission.

Multicast and Unicast

When a channel is configured for multicast and unicast, the receiver Content Engine uses unicast to download content only after all carousel passes have been exhausted and after the preconfigured multicast transmission fails. In a multicast cloud configuration that uses a backup sender, when the channel is enabled for multicast and unicast, the failover to unicast occurs when the current active multicast sender has exhausted all the carousel passes for the file.

If the administrator wants the Content Engines to fall back to unicast (for example, with a multitier unicast deployment using a terrestrial multicast medium), the multicast cloud should be configured for a low number of carousel passes (such as 1, 2, or 3).

Multicast Only

However, if only multicast replication is desired (for example, with a hub and spoke or star topology deployment using a satellite multicast medium), the channels should be configured as multicast-only, with a high number of carousel passes configured in the multicast cloud (such as 10 or more).

When a channel is configured to be multicast only (that is, when the channels are associated with a multicast cloud and the subscribing receiver Content Engine has multicast service successfully licensed and enabled), content replication takes place only through multicasting. No retransmission takes place in unicast at all. This prevents background unicast polls from happening and taking up bandwidth. However, if a Content Engine in the multicast-only channel is not licensed and enabled for multicasting, it can continue to request all the content from a multicast-only channel through unicasting.

Enabling Content Engines for Multicasting

Before you can create a multicast cloud, you must have a multicast distribution license key (purchased from Cisco Systems) and Content Engines that are enabled for multicasting. These multicast-enabled Content Engines can then be assigned as sender and receiver Content Engines when you configure the multicast cloud.

Content Engines for multicasting must be assigned to the multicast cloud, which is assigned to multicast-enabled channels. Also, you need to assign individual Content Engine senders and receivers of the cloud to the particular multicast-enabled channel. You must do this additional step even though the multicast cloud is associated with the channel.



Note

You must assign the multicast cloud to a channel first, then assign the individual Content Engines to the channel.

To enable Content Engines for multicasting, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to enable for multicasting. The Content Engine Device Home window appears.
- Step 3** In the Contents pane, choose **Prepositioning > Multicast Distribution License**. The Multicast Distribution License Agreement window appears.
- The Multicast Distribution Agreement for Content Engine window appears only if you have not read and accepted the multicast distribution license for the device. To accept the license, follow these steps:
- Read the terms of this agreement.
 - Check the **Accept License** check box.
 - Click **Submit**.
- The window refreshes, and the Multicast Distribution License Settings for Content Engine window appears.
- Step 4** Check the **Enable** check box, and enter the license key in the License Key field.
- Alternatively, check the **Evaluate** check box to enable an evaluation license key for a set period.
- Step 5** Click **Submit**.
- Step 6** Return to the Modifying Content Engine window, and repeat [Step 1](#) through [Step 5](#) to enable multicasting on other Content Engines in your ACNS network.
- Step 7** To modify or delete the multicast distribution license settings, you have the following options:
- To delete the configured settings for the device, click the **Remove Device Settings** icon in the taskbar.
 - To apply the factory default settings for the device, click the **Apply Defaults** icon in the taskbar.

- To override the device group settings that have been applied from device groups to which the device is associated, click the **Override Group Settings** icon in the taskbar.

When a device is associated with one or many device groups that have been configured with the settings displayed in this window, you can choose the device group name from the drop-down list whose settings you want to apply to the device.

Step 8 Proceed to the next section, “[Configuring Multicast Cloud Properties](#),” to create a multicast cloud.

Configuring Multicast Cloud Properties

A multicast cloud is configured by specifying an IP multicast address for advertising the data being transferred, an IP multicast address range for transferring the data, a primary multicast sender Content Engine and an optional backup sender, a set of receiver Content Engines, and a maximum rate at which to send the data.



Note

We highly recommend that you avoid using multicast addresses of the form x.0.0.y (for example, 238.0.0.1). These addresses hash to the same Ethernet address space as 224.0.0.x, which is used frequently by routers and switches for local multicasts. Additional traffic on these addresses adds to the workload of these network elements.

To configure a multicast cloud, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Multicast Clouds**. The Multicast Clouds window appears.
- Step 2** Click the **Create New Multicast Cloud** icon in the taskbar. The Creating New Multicast Cloud window appears. (See [Figure 5-10](#).) See [Table 5-6](#) for a description of multicast cloud parameters.

Figure 5-10 *Creating New Multicast Cloud Window*

Creating New Multicast Cloud

Multicast Cloud Information

Name:

Advertisement IP address: Port:

Multicast Address Settings

Start IP address: End IP address:

Default Multicast-out Bandwidth: Kbps ☒ Synchronize Primary and Backup CE Multicast-out Bandwidth

Backup CE Default Multicast-out Bandwidth: Kbps

Advanced Settings

Multicast medium: # of Carousel passes:

TTL: hop Delay between passes: mins

FEC transmission group: PGM Router-assist: ☐

Primary-to-backup failover grace period: mins

Backup-to-primary fallback grace period: mins

Comments

Step 3 Configure the basic multicast cloud information:

- a. Enter a name for the multicast cloud in the Name field.
- b. Enter an advertisement address in the Advertisement IP address field.

When creating the multicast cloud, you must assign a unique advertisement address. Assigning a unique advertisement address provides all the Content Engines in one cloud with the same advertisement address and allows them to communicate multicast session information.

- c. Enter a port number in the Port field (the port used by file transfer addresses). The default port is 7777.

Step 4 Configure the multicast address settings:

- a. Enter a start IP address and an end IP address in the fields provided.

Once the multicast cloud is configured and you associate the multicast cloud with one or more multicast-enabled channels, this multicast address range is used to provide each channel associated with it a unique data channel multicast address. When you assign a channel to the multicast cloud, an unused IP address is automatically selected from this range to ensure that the address is used by only one channel and by only one multicast cloud.

- b. Enter a maximum multicast rate in kilobits per second (kbps) in the Default Multicast-out Bandwidth field. This value applies to the sender Content Engine.

**Tip**

The multicast-out bandwidth should be a realistic value, no greater than 102400 kbps (100 mbps). If your network has both Gigabit Ethernet and Fast Ethernet connections and you have low-end Content Engine models in the multicast cloud, set the default multicast-out bandwidth no higher than 10240 kbps (10 mbps). If all of the Content Engines in your network are high-end models (CE-732x) and your network contains all Gigabit Ethernet connections, you can set the default multicast-out bandwidth as high as 102400 kbps (100 mbps). If you set the value higher than this, the receiver might fail to receive the entire file.

- c. If you have a backup sender configured, enter a maximum multicast rate in kilobits per second (kbps) in the Backup CE Default Multicast-out Bandwidth field.

Step 5 Configure the multicast cloud advanced settings as necessary:

- a. Based on your means of transmitting the multicast, choose either **Satellite** or **Terrestrial** from the Multicast medium drop-down list.
- b. Enter the number of hops for Time To Live in the TTL field.

The TTL field dictates how many hops a packet travels before it is discarded, regardless of whether or not the packet has reached its destination.

- c. Enter the number of carousel passes in the # of Carousel passes field.

This number of carousel passes is the maximum number of times that a multicast sender can send the content of the channel to which it is assigned. The multicast sender can be set to fall back to unicast after the maximum number of carousel passes has been reached.

If you want the Content Engines to fall back to unicast (for example, with a multitier unicast deployment using a terrestrial multicast medium), the multicast cloud should be configured for a low number of carousel passes (such as 1, 2, or 3).

However, if multicast replication is preferred (for example, with a hub and spoke or star topology deployment using a satellite multicast medium), use a high number of carousel passes, such as 10 or more.

**Note**

If a channel is set to multicast only and the receiver is multicast-ready, the maximum number of configurable carousel passes is 1,000,000,000. If the number of carousel passes configured is used up, the syslog displays a warning message as an alert.

- d. Choose a value for the FEC transmission group from the drop-down list.

FEC data encoding protects transmissions against errors, without requiring retransmission. The FEC number denotes the number of packets that will be encoded into one FEC transmission group. The higher the number, the bigger the transmission group size, and hence, the more error-resistant the multicast may be, but also the more computational overhead there will be on the multicast sender and receivers. Note that there is no bandwidth overhead related to FEC.

- e. Enter a value, in minutes, for the Delay between passes field. This sets the interval between file transmissions.
- f. To enable the IP router alert option for Pragmatic Group Multicasting (PGM) packets, check the **PGM Router-assist** check box.

This option specifies whether IP routers are to be used to assist in distribution of content.

- g. Enter a value, in minutes, for the Primary-to-backup failover grace period. The range is 5 to 7200 minutes. The default is 30 minutes.

This is the amount of time allotted for the backup sender to detect whether the primary sender is active.

- h. Enter a value, in minutes, for the Backup-to-primary fallback grace period. The range is 5 to 7200 minutes. The default is 30 minutes.

This is the amount of time allotted for the primary sender to detect whether the backup is active.

Step 6 If you wish, enter any comments in the Comments field.

Step 7 To save the configuration, click **Submit**.

Your multicast cloud configuration is saved, and you can proceed to [Step 3](#) in the “Assigning Content Engines to a Multicast Cloud” section on page 5-40.



Note If you have not assigned any receiver Content Engines, a message appears, stating that your transaction is not complete until you assign receiver Content Engines to the multicast cloud.

[Table 5-6](#) describes the parameters that you need to define in the Content Distribution Manager GUI when configuring a multicast cloud. Required fields are marked with an asterisk in the GUI and in the table.

Table 5-6 Multicast Cloud Parameters

Field	Description
Multicast Cloud Information	
Name ^{*1}	Identifier for the multicast cloud. The name must be unique across the system.
Advertisement IP address [*]	Multicast address collectively known and used by a multicast cloud to implement a communication channel. The advertisement IP address must conform to these guidelines: • It must be unique across the system. • Each multicast cloud must have one advertisement address. • It cannot be used in another multicast cloud. • It must be within the RFC multicast range (224.0.0.0–239.255.255.255). • It must not be within the start and end range specified by this cloud. • It must not be any user-specified multicast address for programs. • It must not be within the program multicast pool start and end range (as defined in System > Program Addresses). See Appendix B, “IP Multicast Addressing,” for information on multicast address restrictions.
Port [*]	Port used for file addresses. The default port is 7777.

Table 5-6 Multicast Cloud Parameters (continued)

Field	Description
Multicast Address Settings	
Start IP address*	Start of the IP address range, which must be within the range 224.0.0.0 to 239.255.255.255. The IP address range must conform to the following: - The IP address range cannot overlap with program multicast pool addresses. - The IP address range must contain all multicast addresses used by this cloud with its associated channel. Note The IP address range in one multicast cloud can overlap that of another multicast cloud. A message alerts you that there is an overlap, but allows the operation. See Appendix B, “IP Multicast Addressing,” for information on multicast address restrictions.
End IP address*	End of the IP address range, which must be higher than the start IP address. The end IP address must be within the range 224.0.0.0 to 239.255.255.255. See Appendix B, “IP Multicast Addressing,” for information on multicast address restrictions.
Default Multicast-out Bandwidth*	Maximum multicast rate in kilobits per second (kbps). The minimum rate is 10 kbps. This is the default multicast-out bandwidth, which is applied 24 hours a day, 7 days a week. The multicast-out bandwidth is associated with the sender Content Engine only. To customize bandwidth rates for different days, use the A&D Bandwidth Settings window. These settings override the default for the period that you specify. (See the “Configuring Acquisition and Distribution Default Bandwidth Settings” section on page 6-54.)
Synchronize Primary and Backup CE Multicast-out Bandwidth	When checked, default multicast-out bandwidth settings are used by both primary and backup senders, if a backup sender is configured.
Backup CE Default Multicast-out Bandwidth	Maximum multicast rate in kilobits per second (kbps) for the backup sender. The minimum rate is 10 kbps.
Advanced Settings	
Multicast medium	Means of transmitting the multicast, either satellite or terrestrial. The default is satellite.
TTL	Time To Live expressed as the number of hops (1–255). The default is 255.
FEC transmission group	Size of an FEC (forward error correction) block in packets. (Refer to RFC 3208 <i>PGM Reliable Transport Protocol Specification</i> for more information.) The allowable inputs are 2, 4, 8, 16, 32, 64, and 128. The default is 16.

Table 5-6 Multicast Cloud Parameters (continued)

Field	Description
# of Carousel passes	Maximum number of times (1–1000000000) that a multicast sender will attempt to send missing content for on-demand carousels. The default is 5. Note If the multicast sender finishes the last carousel on an object at time t and the multicast sender receives a NACK within $t + \text{carousel\_delay}$, the multicast sender starts the next carousel of this object at time $t + \text{carousel\_delay}$. That is, the multicast carousel is not triggered immediately upon receipt of the NACK if the carousel delay is greater than zero (0).
Delay between passes	Delay between file transmissions, in minutes (0–10080) (7 days * 24 hours * 60 minutes). The default is 0 minutes.
PGM Router-assist	Specifies whether IP routers are to be used to assist in distribution of content. Checking this check box enables the IP router alert option for PGM packets. Note PGM packet routing is supported in Cisco IOS 12.0(5)T and later 12.0 T releases.
Primary-to-backup failover grace period	Amount of time allotted for the backup sender to detect whether the primary sender is active. If the backup sender does not hear a heartbeat from the primary sender within this grace period, the backup sender assumes the active role. The range is 5–7200 minutes. The default is 30 minutes.
Backup-to-primary fallback grace period	Amount of time allotted for the primary sender to detect whether the backup is active. If the backup sender does not respond within this grace period, the primary sender assumes the active role. The range is 5–7200 minutes. The default is 30 minutes.
Comments	Comments about the multicast cloud.

1. * = required field. All other fields are optional.

Modifying a Multicast Cloud

To modify a multicast cloud, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Multicast Clouds**. The Multicast Clouds window appears.
- Step 2** Click the **Edit** icon next to the name of the multicast cloud that you want to change. The Modifying Multicast Cloud window appears.

- Step 3** Edit the information or settings as needed.
- Step 4** To save the changes, click **Submit**.

Deleting Multicast Clouds

To delete a multicast cloud, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Multicast Clouds**. The Multicast Clouds window appears.
- Step 2** Click the **Edit** icon next to the name of the multicast cloud that you want to delete. The Modifying Multicast Cloud window appears.
- Step 3** Click the **Trash** icon in the taskbar. You are prompted to confirm your action.
- Step 4** Click **OK**. The multicast cloud is deleted.

Viewing Multicast Clouds

To view a list of all the multicast clouds configured in your network, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Multicast Clouds**. The Multicast Clouds window appears, listing all the multicast clouds in your ACNS network.
- The table also displays the sender Content Engine for each multicast cloud.
- Step 2** To print the table data, click the **Printer** icon.

Assigning Content Engines to a Multicast Cloud



Note

Content Engines must be multicast-enabled with the multicast distribution license before you can add them to a multicast cloud. (See the [“Enabling Content Engines for Multicasting”](#) section on page 5-33.)

To add Content Engines to a multicast cloud, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Multicast Clouds**. The Multicast Clouds window appears.
- Step 2** Click the **Edit** icon next to the name of the multicast cloud to which you want to add receiver Content Engines. The Modifying Multicast Cloud window appears.
- Step 3** In the Contents pane, click **Assign Content Engines**. The CE assignments for Multicast Cloud window appears.

Step 4 Assign multicast receiver, sender, and backup sender Content Engines. [Table 5-7](#) describes the Content Engine role assignments for a multicast cloud.

- a. Choose **Receiver** from the Role drop-down list and click the **Assign** icon (blue cross mark) next to the name of the Content Engine that you want to assign to the receiver role. Click **Submit**. Repeat this step if you wish to assign multiple receiver Content Engines to the multicast cloud.

Alternatively, click the **Assign all Content Engines** icon in the taskbar to perform a bulk addition of all Content Engines in various locations to the multicast cloud and click **Submit**.

- b. Choose **Primary Sender** from the Role drop-down list and click the **Assign** icon next to the name of the Content Engine that you want to assign to the primary sender role and click **Submit**.
- c. If you wish, choose **Backup Sender** from the Role drop-down list and click the **Assign** icon next to the name of the Content Engine that you want to assign to the backup sender role and click **Submit**.

**Note**

The Content Distribution Manager GUI allows you to assign Content Engines that are not multicast enabled. However, you must ensure that any Content Engine that you assign to a multicast cloud is multicast enabled with a multicast license. (See the [“Enabling Content Engines for Multicasting”](#) section on page 5-33.)

Table 5-7 **Content Engine Role Assignments for Multicast Clouds**

Role Assignment	Description
Receiver	Content Engines that receive content from the primary sender. To create a functional multicast cloud, you must add at least one receiver Content Engine. Use the following guidelines when adding receiver Content Engines: - You must choose at least one Content Engine to be a receiver. - The maximum number of receivers you can add is the total number of Content Engines in the system (excluding the sender Content Engine). - The receiver cannot be a receiver in another multicast cloud. - The receiver cannot be a sender in the same multicast cloud. Only Content Engines that are <i>not</i> assigned to another multicast cloud are displayed as selections in the Content Distribution Manager GUI. - Only a fully configured multicast cloud (with at least one receiver Content Engine) can be associated with a channel to enable multicast capability. - The list of Content Engines for selection is sorted by location. You can choose individual Content Engines from one or more locations to add to this multicast cloud. Note Only Content Engines that are not assigned as receiver Content Engines to another multicast cloud or as sender Content Engines to this cloud are displayed as selections in the Content Distribution Manager GUI.
Primary sender	Primary Content Engine forwarder that pushes content to a set of Content Engine receivers using multicast. Primary senders must conform to the following rules: - The primary sender cannot be the backup sender for the same cloud. - The primary sender cannot be the receiver for the same cloud. - The primary sender cannot be the primary sender for a different cloud. - The primary sender cannot be the backup sender for a different cloud. A sender Content Engine cannot be deleted from the network. Before deleting a sender Content Engine, you must choose another Content Engine as the sender for the multicast cloud.
Backup sender	Content Engine configured for backup in the event of failure of the primary sender Content Engine. Backup senders must conform to the following rules: - The backup sender cannot be the primary sender for the same cloud. - The backup sender cannot be the receiver for the same cloud. - The backup sender cannot be the primary sender for a different cloud. - The backup sender cannot be the backup sender for a different cloud. The primary and backup senders of a multicast cloud should subscribe to the same set of multicast-enabled channels.

To remove a Content Engine from the multicast cloud, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Multicast Clouds**. The Multicast Clouds window appears.
 - Step 2** Click the **Edit** icon next to the name of the multicast cloud from which you want to remove receiver Content Engines. The Modifying Multicast Cloud window appears.
 - Step 3** In the Contents pane, click **Assign Content Engines**. The CE assignments for Multicast Cloud window appears.
 - Step 4** Click the **Unassign** icon next to the name of the Content Engine that you want to remove. Alternatively, you can click the **Unassign all Content Engines** icon in the taskbar to perform a bulk removal of all Content Engines in various locations from the multicast cloud.
 - Step 5** To remove the Content Engines from the multicast cloud, click **Submit**. After you click **Submit**, a blue cross mark appears next to the unassigned Content Engine.
-

Assigning and Removing Multicast Clouds from Channels

You can assign or remove multicast clouds only from channels that are multicast-enabled. A channel is multicast-enabled when you click the Multicast/Unicast radio button or the Multicast-only radio button when you create or modify the channel. (To create a new channel, see the [“Creating and Modifying Channels”](#) section on page 5-10.)




Note

When a multicast cloud is assigned to a channel, the Content Engines that are part of the multicast cloud must also be individually assigned to the channel for multicasting. To assign Content Engines to a channel, follow the procedure in the [“Adding and Removing Content Engines from Channels”](#) section on page 5-16.


To assign a multicast cloud from a multicast-enabled channel, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**. The Channels window appears, listing all the channels in your ACNS network.
 - Step 2** Click the **Edit** icon next to the name of the channel to which you want to assign a multicast cloud a multicast cloud. The Modifying Channel window appears. (See [Figure 5-7](#).)
 - Step 3** Enable multicasting on the channel, if you have not already done so.
 - a. Click the **Multicast only** or **Multicast/Unicast** radio button.
 - b. Click **Submit**.
 - Step 4** In the Contents pane, choose **Assign Multicast Cloud**. The Assigning Multicast Clouds to Channel window appears. (See [Figure 5-11](#).)

Figure 5-11 Assigning Multicast Clouds to Channel Window

- Step 5** Click the **Assign Multicast Cloud to Channel** icon in the taskbar. The Assign Multicast Cloud to Channel Assignment window appears.
- Step 6** Choose a multicast cloud from the drop-down list. The window refreshes, showing the IP address range for that multicast cloud. In the IP address to use for this Channel field, enter any available IP address from this range to assign to this channel. (The first available address automatically populates this field; however, you can change the address.)

In [Figure 5-12](#) a multicast cloud has already been chosen and the window has refreshed, showing the IP address range for the cloud.

Figure 5-12 Assign Multicast Cloud to Channel Assignment Window

Step 7 To save the settings, click **Submit**.

To remove a multicast cloud from a channel, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**. The Channels window appears, listing all the channels in your ACNS network.
- Step 2** Click the **Edit** icon next to the name of the channel from which you want to remove a multicast cloud. The Modifying Channel window appears. (See [Figure 5-7](#).)
- Step 3** Click the **Edit** icon next to the name of the multicast cloud that you want to remove. The Assign Multicast Cloud to Channel Assignment window appears.
- Step 4** Click the **Trash** icon. You are prompted to confirm the action.
- Step 5** Click **OK**. The channel is removed.
-

Verifying the Multicast Configuration

Use the following checklist to verify your multicast configuration:

1. Verify that all Content Engines being added to the multicast cloud are online and enabled for multicasting by using the **show multicast** EXEC command.
2. After multicast connectivity has been confirmed, configure the multicast cloud parameters: sender, advertisement address, transfer address, range, multicast-out bandwidth, multicast medium (satellite or terrestrial), number of carousel passes, and receiver Content Engines.
3. Verify the properties of the multicast cloud by using the **show distribution mcast-data-sender** EXEC command for the sender Content Engine and **show distribution mcast-data-receiver** EXEC command for receiver Content Engines.
4. Configure a channel for multicasting and assign a multicast cloud to it.
5. Assign multicast senders and the desired receiver Content Engines from the assigned multicast cloud to the channel.
6. After a multicast channel has been configured on the Content Distribution Manager, use the **show statistics distribution mcast-sender** EXEC command and the **show statistics distribution mcast-receiver** EXEC command to show the progress of a data transfer.

Commands Related to Multicasting

[Table 5-8](#) lists and describes CLI commands that are related to multicasting. For complete descriptions of these commands, including usage guidelines and examples, refer to the *Cisco ACNS Software Command Reference, Release 5.5* publication.

Table 5-8 Multicast-Related Commands

Command	Syntax	Description
distribution multicast	distribution multicast send-nack-now	Generates a NACK for missing objects and sends them to the multicast receiver.
	distribution multicast resend	Resends an individual object (object option) or specifies that the first round of carousel passes be triggered by a NACK only (on-demand-only option).
	distribution multicast stop	Sets the number of carousel passes completed to the maximum value so that no more multicast carousels can be triggered.
multicast	multicast {accept-license-agreement enable evaluate license-key <i>key</i> sender-delay <i>delay</i>}	Controls enabling and disabling of multicasting, including reading and accepting the license and entering the key. Note We recommend that you use the Content Distribution Manager GUI rather than the CLI to configure multicasting.
	multicast connectivity-test	Verifies multicast connectivity between a sender and one or more receivers.
	multicast fixed-carousel enable	Enables fixed carousel sending, regardless of receiver feedback; NACK triggered carousel passes are disabled. Only applies to the primary sender and is disabled if the Content Engine becomes a backup sender.
	multicast priority-weight	Sets the percentage of multicast bandwidth for priority-based scheduling. The remaining bandwidth is given to time-based scheduling of queued jobs. The default is 50%.
	multicast back-version-compatibility	Makes the multicast sender backward compatible with the receivers of the specified software version.
show multicast	show multicast [license-agreement]	Displays the current status of the multicast client. If the license-agreement option is included in the command string, the full text of the multicast license agreement is displayed.
show distribution	show distribution mcast-data-receiver show distribution mcast-data-sender show distribution location	Shows information about the multicast receiver or the sender. Displays information about the multicast cloud.

Table 5-8 Multicast-Related Commands (continued)

Command	Syntax	Description
show statistics distribution	show statistics distribution mcast-data-receiver show statistics distribution mcast-data-sender	Displays the content distribution statistics of the multicast data receiver or sender.
clear statistics distribution	clear statistics distribution mcast-data-receiver clear statistics distribution mcast-data-sender	Clears multicast statistics.

Troubleshooting the Distribution Hierarchy

Because distribution-related problems are design-dependant, your initial strategy is to discover whether or not the correct Content Engine is sending content in the correct distribution path.

- To determine which Content Engines are in the distribution path of a particular Content Engine, use the **show distribution remote traceroute EXEC** command, as shown in the following example:

```
cel# show distribution remote traceroute ?
  forwarder-next-hop  next forwarder along the path
  unicast-sender       check status for unicast sender

cel# show distribution remote traceroute forwarder-next-hop ?
  channel-id          Channel-id of a Channel

cel# show distribution remote traceroute forwarder-next-hop channel-id 133 ?
  max-hop             Trace route till specified number of hops is reached
  trace-till-good     traceroute till probe is good or the object is found
  trace-till-root     traceroute till the root ce

cel# show distribution remote traceroute forwarder-next-hop channel-id 133
trace-till-root

Hop NextHop_CElD  NextHop_CEName NextHop_CElp  GenID Status/Reason
---
  1    1100          ce3    10.255.0.43    1      LOC-LEAD
  1    1100          ce3  128.107.193.183  1      LOC-LEAD (Reached RootCE)
```

- To verify that the Content Engine is reachable and that it is in the distribution hierarchy, use **show distribution remote traceroute EXEC** command, as shown in the following example:

```
cel# show distribution remote traceroute unicast-sender channel-id 133 ?
  cdn-url             check the object on remote CE using cdn-url
  probe              probe the remote unicast sender
  relative-cdn-url    check the object on remote CE using relative-cdn-url

cel# show distribution remote traceroute unicast-sender channel-id 133 probe ?
  max-hop             Max-hop to traceroute to
  trace-till-good     traceroute till probe is good or the object is found
  trace-till-root     traceroute till the root ce

cel# show distribution remote traceroute unicast-sender channel-id 133 probe
trace-till-root
Polling .... ce3 [10.255.0.43]  Fwdr_Id:1100
Polling .... ce3 [128.107.193.183]  Fwdr_Id:1100
(Reached RootCE)
```

Where to Go Next

In this chapter you have accomplished the following tasks:

- Configured your distribution topology by creating and modifying locations
- Defined websites, providing URLs for the origin web servers that store the website content
- Created channels to map content from the origin server to the Content Engines in your ACNS network
- Enabled channels for unicast or multicast transmission, or both
- Assigned Content Engines to channels
- Created multicast clouds for multicasting

Now you are ready to configure additional device settings for acquiring pre-positioned content and to define the content that you want to acquire. Proceed to the next chapter, [Chapter 6, “Configuring the ACNS Network for Content Acquisition.”](#)



Configuring the ACNS Network for Content Acquisition

Content acquisition is an important part of content pre-positioning in the ACNS network. ACNS software uses the concept of a *channel* to map a set of content objects to a set of Content Engines. Before pre-positioned content can be acquired or distributed in the ACNS network, users must create a new channel, subscribe Content Engines to that channel, and designate one of the Content Engines as the root Content Engine. The source of the content might be stored on various file servers or web servers. The root Content Engine of the channel fetches the content objects from these *origin* servers and then replicates this content to all the Content Engines in the channel.



Note

Pre-positioned content is served only on ports that are standard for the protocol. If the incoming URL contains a port number other than the protocol's standard port (for example, HTTP uses port 80), then the Content Engine does not attempt to serve the content from the pre-positioned file system (cdnfs). Instead, the Content Engine tries to serve the content from the cache file system (cfs) or tries to fetch the content from the origin server, depending on the existing configuration of the Content Engine. To server pre-positioned content from a nonstandard port, set the *ignoreOriginPort* attribute to true in the manifest file. For a description of this attribute, see the “[item](#)” section on [page A-44](#).

This chapter outlines the procedures necessary for acquiring pre-positioned content in the ACNS network and contains information on the following topics:

- [Preparing to Acquire Pre-positioned Content](#), page 6-2
- [Choosing a Content Acquisition Method for a Channel](#), page 6-6
- [Acquiring Pre-Positioned Content Using the Content Distribution Manager GUI](#), page 6-9
- [Acquiring Pre-Positioned Content Using an Externally Hosted Manifest File](#), page 6-23
- [Acquiring Content through a Proxy Server](#), page 6-35
- [Acquiring Content Using Indexed Web Servers](#), page 6-40
- [Authentication Support](#), page 6-42
- [Verifying the Results](#), page 6-48
- [Troubleshooting Content Acquisition](#), page 6-48
- [Retry and Refresh Mechanisms](#), page 6-52
- [Bandwidth Control](#), page 6-54
- [Error Codes](#), page 6-59
- [Where to Go Next](#), page 6-62

Preparing to Acquire Pre-positioned Content

To prepare your ACNS network to acquire pre-positioned content, you must complete the following tasks:

1. [Set Up the Origin Server](#)
2. [Create a Channel](#)
3. [Assign a Root Content Engine](#)

About the Acquirer

The root Content Engine uses a software agent, referred to as the acquirer, that gathers channel content before it is distributed to the receiver Content Engines in the ACNS network. The acquirer maintains a task list, which it updates after receiving a notification of changes in its channel configuration. The following items in this section are supported or not supported by the acquirer:

File Acquisition from SMB Servers

ACNS 5.5 software supports file acquisition from Windows file servers with shared folders and Unix servers running the SMB protocol. The acquirer first mounts the share folder to the root Content Engine. This mount point then acts as the origin server from which the content is to be fetched. The acquirer fetches the content and pre-positions it in the cdnfs storage space. (Note that files greater than 2 GB cannot be acquired using SMB.)



Note

Symbolic links within exported filesystems (SMB or NFS) must contain a relative path to the target file, or the target file should be copied into the exported volume. If the symbolic link points to an absolute path outside of the volume that the file server is exporting to the Content Engine, the file path is inaccessible to the Content Engine, SMB acquisition fails, and the ACNS 5.5 software issues a 404 error message.

SMB is supported in the Content Distribution Manager GUI for fetching an externally hosted manifest file. The Content Distribution Manager GUI content import feature also supports SMB file acquisition for content items and crawl tasks.

You can pre-position content items and crawl tasks from SMB servers by using the following URL format: `\\SMBserver\sharefolder\filepath`. The full URL can be specified in the manifest file `<item>` or `<crawler>` tags, as shown in this example:

```
<CdnManifest>
<server name="myserver1">
<host name="\\<SMB-Server>" />
</server>
<item src = "<Share Folder>\<File>" server="myserver1"/>
<item src="\\<SMB-Server>\<Share Folder>\File" />
</CdnManifest>
```

File acquisition from an SMB server for crawl jobs is similar to FTP acquisition in which the crawler crawls the folder hierarchy rather than parsing the HTML file.

You can specify the port to be used for the SMB protocol in the URL or in the *port* attribute. If the port is not specified, the default port is 139.

Authentication for SMB server file acquisition is supported by means of the *user*, *password*, and *userDomainName* attributes in the <host> and <item> tags. (See the “[item](#)” section on page A-44 for a description of this attribute.) These attributes must be specified separately rather than as part of the full URL. Proxy and proxy authentication support are not available for this feature.

**Note**

If you have a Windows XP server share that is using the anonymous method of authentication, file acquisition fails. (Refer to the *Release Notes for ACNS Software, Release 5.5* document for more information.)

EMBED and OBJECT ID Tag Support

If an HTML file includes an EMBED tag or an OBJECT ID tag, the acquirer parses the file, obtains the file link from the *src* and *filename* attributes, and fetches these files.

Crawl .asx File

The acquirer can parse an .asx file and fetch its contents.

expires Attribute Specified in the Manifest File

The *expires* attribute, specified as part of the <item> tag or <crawler> tag in the manifest file, designates a time in yyyy-mm-dd hh:mm:ss format for content to be removed.

If you specify the *expires* attribute in the manifest file, the root Content Engine deletes the content at the time specified. The root Content Engine then replicates the deletion event to receiver Content Engines, which also delete the content from the channel.

Parsing Scripts in an HTML File

If an HTML file contains a JavaScript or a VBScript (Microsoft’s Visual Basic Scripting Edition is a scaled-down version of Visual Basic), and the script codes have links, the acquirer does not parse the script code to get the links.

Cookie Support

Cookies are a general mechanism that server-side connections (such as CGI scripts) use to both store and retrieve information on the client side of the connection. The addition of a simple, persistent, client-side state significantly extends the capabilities of web-based client/server applications. Although the acquirer is also a client, it does *not* support cookies.

Set Up the Origin Server

Content is usually stored on file servers or web servers that are not part of the ACNS network of managed devices. To pre-position this content for ACNS network acquisition and distribution, the origin file server or web server must support at least one of the following protocols that are used by ACNS software to acquire content:

- HTTP
- HTTPS
- FTP

**Note**

Beginning with the ACNS 5.5 release, the software no longer supports the MMS protocol or the MMS-over-HTTP protocol for content acquisition.

Using HTTP and HTTPS

Any standard web server supports the HTTP and HTTPS protocols. You can set up your web server as an origin server for pre-positioned content intended for the ACNS network by moving the content over to the web server or by configuring the web server to access the desired content. The following two web servers are the most popular:

- Apache—Supported on UNIX, Linux, and Microsoft NT platforms
- Microsoft IIS—Only supported on Microsoft platforms

For the HTTP and HTTPS protocols, content can be fetched as single content items by using the `<item>` tag in the manifest file, or content can be fetched by using the crawling feature to crawl web server directories. The crawler crawls the folder hierarchy rather than parsing the HTML file. Therefore, if you want to use the crawl feature, you must enable directory indexing and make sure that the directory does not contain `index.html`, `default.html`, or `home.html` files.



Tip

You might need to install SSL certificates in order to set up the web server for HTTPS content acquisition. If your server is using an expired certificate, or a self-signed certificate, you should set `sslAuthType` to “weak” in the manifest file `<host>` tag.

Redirection Support

For a single-item HTTP redirection, the acquirer follows the redirect to the new link to obtain the file. If the HTTP redirect is a crawl job, the acquirer checks the new link against the crawl parameters. For example, if a start-url `http://www.ibm.com/` is redirected to `http://www.cnn.com/`, the acquirer checks `www.cnn.com` to see if the start-url matches the crawl parameters. Unless the *externalPrefixes* attribute is specified, the redirection is usually rejected because it does not match the default prefix.

HTTP Response no-cache Directive

The no-cache directive is related to the HTTP protocol. The HTTP protocol uses headers to communicate between the client and the server. The client sends a request header, and the server responds with a response header. Headers specify the various attributes for the resource being requested. For example, the client sends the request:

```
GET /index.html HTTP/1.0
Host:www.google.com
```

The server might send a response, as follows:

```
HTTP/1.0 200 OK
Content-Length:10
pragma:no-cache
Cache-control:no-cache
```

The no-cache directive in the response header tells the client that the content being requested is not cacheable. When an HTTP server responds to a request with a no-cache directive in the header, the acquirer behaves as follows:

- If the content to be acquired is specified in an `<item>` tag, the acquirer ignores the *no-cache* directive and fetches the content anyway.
- If the content to be acquired is specified in a `<crawler>` tag in the manifest file, and the response sent out by the server contains a *no-cache* directive in the response header, the acquirer honors the directive and does not fetch the content or pre-position it.

HTTP Response expires Directive

An expires directive in the HTTP response header indicates that the content being served is not valid after the time specified in the response. For example, if the response header contains

```
HTTP/1.0 200 OK
expires:<some time>
```

the client (the acquirer in this case) is directed not to use the content after the time specified. The expires directive is honored only for content that is specified in the <crawler> tag. Single items configured using the <item> tag in the manifest file do not honor the expires directive sent by the server.

Using FTP

The root Content Engine acquirer supports acquiring files from FTP servers. When you use FTP, content can be acquired as single content items by using the <item> tag in the manifest file, or content can be fetched by using the crawling feature to crawl the FTP server directories. In FTP acquisition, the crawler crawls the folder hierarchy rather than parsing the HTML file. The following popular FTP servers are supported:

- Microsoft IIS 4.0, 5.0, 6.0—For Windows platforms
- Wu-2.6.1-18—For Linux platforms
- FTP Server—In SunOS (Version 5.6)
- proFTPD—For Linux platforms

Other supported Windows FTP servers are:

- WS_FTP server
- Bulletproof FTP server
- SurgeFTP
- SlimFTPd

You can use other FTP servers, as long as the following FTP commands are supported:

- USER, PASS
- [SIZE, MDTM] [or] [LIST -a]
- PASV [or] PORT
- CWD ~ [or] CWD <SPACE> [or] CWD /
- RETR

Secure FTP

The acquirer currently does not support secure FTP.

Create a Channel

Channels map a set of content objects to a set of Content Engines. You must have channels configured before you can acquire pre-positioned content or distribute it. To create channels, see [Chapter 5, “Configuring the ACNS Network for Content Distribution.”](#) Chapter 5 discusses how to create channels for content acquisition and distribution by configuring the following network elements:

1. Locations—[Creating and Modifying Locations](#)
2. Content providers—[Creating and Modifying Content Providers](#)

3. Websites—[Creating and Modifying Websites](#)
4. Channels—[Creating and Modifying Channels](#)

Assign a Root Content Engine

A root Content Engine is used to acquire content for a channel. A channel can have only one root Content Engine. We recommend that you choose a root Content Engine that has enough bandwidth to access the content at the origin server.

For information on how to assign a root Content Engine, see the [“Designating the Root Content Engine” section on page 5-16](#).

Assigning a Backup Root Content Engine

You do not need to specify a backup root Content Engine; however, for backup purposes, you must have a second Content Engine in the same location as the primary Content Engine that is subscribed to the channel. If the designated root Content Engine becomes inactive, the other Content Engine in the location automatically becomes a temporary root Content Engine. If the designated root Content Engine comes back online, it takes over as the root and the temporary root Content Engine becomes a regular Content Engine.

During content acquisition, the acquirer uses a lot of CPU power. We recommend that you set up a dedicated high-end root Content Engine for content acquisition. The root Content Engine can be used for streaming video files or serving pre-positioned content; however, streaming video quality might suffer during periods of heavy content acquisition.

Choosing a Content Acquisition Method for a Channel

When you configure a channel to acquire content, you must choose one of the following content acquisition methods:

- Specify the content using an externally hosted manifest file.

Manifest files contain the XML tags, subtags, and attributes that you wish to use to define the parameters for content acquisition. You must be familiar with the structure of the XML-based manifest file and be sure that XML tags are properly formatted and syntactically correct before you can effectively create and use manifest files. (See the [“Acquiring Pre-Positioned Content Using an Externally Hosted Manifest File” section on page 6-23](#) and [Appendix A, “Creating Manifest Files.”](#))

- Specify the content using the Content Distribution Manager GUI.

The Content Distribution Manager GUI provides a user-friendly interface that you can use to easily add content items and specify crawl tasks without having to create and update a manifest file. The Content Distribution Manager GUI automatically validates all user input and generates an XML-formatted manifest file in the background that is free of syntax errors. (See the [“Acquiring Pre-Positioned Content Using the Content Distribution Manager GUI” section on page 6-9](#).)

Only one manifest file is generated per channel for all content items. You can save your GUI-generated manifest file to any desired location.

To configure the content acquisition method for a channel, follow these steps:


- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**. The Channels window appears, listing all channels in your ACNS network.
- Step 2** Click the **Edit** icon next to the channel for which you want to configure the content acquisition method. The Modifying Channel window appears.



Note Channels for pre-positioning content are listed as type “Content” in the Channels window.

- Step 3** In the Contents pane, choose **Channel Content**. The Content Acquisition Method for Channel window appears. By default, the content acquisition method for a newly created channel is set to **Use GUI to specify content acquisition**. (See [Figure 6-1](#).)

Figure 6-1 Content Acquisition Method for Channel—Use Content Distribution Manager GUI



Use this window to add content items and specify crawl tasks for the channel. (See the [“Acquiring Pre-Positioned Content Using the Content Distribution Manager GUI”](#) section on page 6-9.)

- Step 4** To use an externally-hosted manifest file:
- Change the content acquisition method by clicking the **Change Method** button.
 - Choose **Specify external manifest file** from the drop-down list that appears.

**Note**

When you change the content acquisition method from using the Content Distribution Manager GUI to using an external manifest file, any content items that you added using the GUI are removed from the GUI. You can save the content items in XML format by clicking the **Save settings locally** icon in the taskbar. A new window displays the GUI-generated manifest file. Save the file to your local disk.

- c. Click the **Save** button next to the drop-down list. You are prompted to confirm your decision.
- d. Click **OK** to confirm. The window refreshes and displays fields for defining manifest file settings and manifest proxy information. (See [Figure 6-2](#).)

Figure 6-2 Content Acquisition Method for Channel—Use External Manifest File

Use this window to define basic manifest settings and proxy information. (See the “[Configure Manifest File and Proxy Information for the Channel](#)” section on page 6-26.)


Acquiring Pre-Positioned Content Using the Content Distribution Manager GUI

ACNS 5.5 software allows you to pre-position content items and crawl tasks directly in the Content Distribution Manager GUI without having to create a manifest file. Use this method to set up simple demonstrations or systems, or to add a few items that do not require advanced acquisition and distribution features.

To pre-position content using the Content Distribution Manager GUI, follow these steps:

-
- Step 1** In the Content Distribution Manager GUI, choose **Services > Web > Channels**.
- Step 2** Click the **Edit** icon next to the name of the channel that you want to modify. The Modifying Channel window appears.
- Step 3** In the Contents pane, choose **Channel Content**. The Content Acquisition Method for Channel window appears.
- Step 4** Specify the method for content acquisition.
- a. Click the **Change Method** button, if necessary.
 - b. Choose **Use GUI to specify content acquisition** from the drop-down list.
 - c. Click **Save**.
- See the [“Choosing a Content Acquisition Method for a Channel”](#) section on page 6-6.
- Step 5** To add content items and crawl tasks, click the **Add Content** icon in the taskbar. The Content Manager window for the channel appears. (See [Figure 6-3](#).)

Figure 6-3 Content Manager Window

**Step 6** Specify the content source.

- a. Choose the protocol to be used for content requests from the Source URL drop-down list. Then enter the origin server domain name or IP address in the field that follows.
- b. If you wish to acquire the source URL as a single content item, check the **Single Item** check box.
- c. If you wish to acquire content by crawling the website or if you wish to add individual content items, leave the **Single Item** check box unchecked. For crawl tasks, enter the number of levels to which the website is to be crawled in the Link Depth field.

Step 7 Set up the parameters for acquiring the content.

- a. If you click the **Define a crawl task** radio button, you can further define and apply rules to the crawl task. Click the **Show optional content acquisition rules** drop down arrow. (See the [“Adding a Crawl Task to a Channel”](#) section on page 6-13 and the [“Configuring Content Acquisition Rules for a Crawl Task”](#) section on page 6-15.)
- b. If you click the **Select individual items** radio button, you can further define content item parameters by clicking the **Launch Quick Crawl** button. This button launches the Quick Crawl Filter window. (See the [“Adding Content Items to a Channel Using the Quick Crawl Utility”](#) section on page 6-11.)

Step 8 If you wish to configure advanced settings, such as content serving time, authentication, URL settings and content settings, click the **Show advanced settings** drop-down arrow. The window expands, allowing you to configure advanced settings. (See the [“Configuring Advanced Content Serving Settings”](#) section on page 6-17.)

- Step 9** To process the content request, click **Submit**. When you click **Submit**, the local manifest file is automatically reparsed, changes are detected, and the corresponding items are acquired or removed. This action, however, does not trigger a recheck of all the content in the channel.

Adding Content Items to a Channel Using the Quick Crawl Utility

Quick Crawl is a utility that automatically crawls websites starting from the specified source URL. You can use this utility when you know only the domain name and not the exact location of the content item. You can set a rule-based quick crawl filter so that only content with predefined characteristics will be matched for crawling. The results of the filtering are displayed, and you can select the content items that you want to be acquired. Quick Crawl supports crawling only for websites that use HTTP and HTTPS.

To add individual content items using the Quick Crawl utility, follow these steps:

- Step 1** Choose **Services > Web > Channels**. The Channels window appears, listing all channels in your ACNS network.
- Step 2** Click the **Edit** icon next to the channel for which you want to add content items. The Modifying Channel window appears.
- Step 3** In the Contents pane, choose **Channel Content**. The Content Acquisition Method for Channel—Use GUI to Specify Content Acquisition window appears. (See [Figure 6-1](#).)
- Step 4** Click the **Add Content** icon in the taskbar. The Content Manager window for the channel appears. (See [Figure 6-3](#).)
- Step 5** From the Source URL drop-down list, choose the communication protocol (HTTP or HTTPS) to be used. The default is HTTP.
- Step 6** In the Source URL field, enter the URL of the host server (origin server) to specify the location from which to start Quick Crawl. You can copy and paste URLs from a web browser to the Source URL field.



Note The values specified for Quick Crawl are not directly associated with manifest file attributes. Quick Crawl simply provides a list of content URLs that you can use as a list of individual content items to be added to the channel.

- Step 7** Leave the **Single Item** check box unchecked.
- Step 8** In the Link Depth field, specify how many levels of a website are to be crawled or how many directory levels of an FTP server are to be crawled. This is optional. The range is -1 to 2147483636.
- If depth = -1, there is no depth constraint.
 If depth = 0, acquire only the starting URL.
 If depth = 1, acquire the starting URL and all the content it references.
- Alternatively, you can choose to specify the link depth after you launch Quick Crawl.
- Step 9** Click the **Select individual items** radio button to select content items that you want to be fetched, and click the **Launch Quick Crawl** button to start the Quick Crawl utility. The Quick Crawl Filter window pops up.

**Note**

Items fetched using Quick Crawl are not associated with the <crawl> element of the manifest file because Quick Crawl is not a manifest generator. Quick Crawl merely lists a number of content items or content objects found on the origin server, such as a graphic file, MPEG video, or RealAudio sound file. When added to the channel, these items are associated with the <item> tag in the GUI-generated manifest file.

Step 10 Configure the Quick Crawl filter settings. Note that all fields in this window are optional.

- a. In the MIME type field, specify the MIME type of files to be crawled. An object is acquired only if its MIME type matches this MIME type.
- b. In the Extension field, specify the extension of files to be crawled. An object is acquired only if its extension matches this extension.
- c. In the Modified After field, enter the date that must be matched if content is to be acquired. An object is acquired only if it was modified after this date.
- d. In the Modified Before field, enter the date that must be matched if content is to be acquired. An object is acquired only if it was modified before this date.


Alternatively, click the **Calendar** icon next to the Modified After and Modified Before fields to display the Date Time Picker popup window. In the Date Time Picker popup window, use the left or right arrow icons to choose previous and following years if required. Choose a month from the drop-down list. Click a day of the month. The chosen date is highlighted in blue. Click **Apply**. Alternatively, click **Set Today** to revert to the current day. The chosen date will be displayed in the Modified After or Modified Before field.

- e. In the Minimum Size field, specify the minimum size of content to be acquired. Content equal to or larger than this value is acquired. Choose **MB**, **KB**, or **Bytes** as the unit of measure. The range is 0 to 2147483636.
- f. In the Max Size field, specify the maximum size of content to be acquired. Content equal to or less than this value is acquired. Choose **MB**, **KB**, or **Bytes** as the unit of measure. The range is 0 to 2147483636.
- g. In the Link depth field, specify how many levels of a website are to be crawled or how many directory levels of an FTP server are to be crawled. The range is -1 to 2147483636.

If you have already entered the depth level in the Link Depth field of the Content Manager window, it is displayed here.
- h. The Domain field displays only the *host.domain* portion of the Source URL that you entered in [Step 6](#). Edit this field if you want to limit the search to a specific host on a domain.
- i. In the Username field, enter the name of the user who requires secure login to the host server (for websites requiring authentication).
- j. In the Password field, enter the password for the user account that is required to access the host server (for websites requiring authentication).

Step 11 When you have finished adding Quick Crawl filter rules, click the **Start Quick Crawl** button. The Searching for Content window appears. A progress bar shows the number of items that are being crawled. Before the specified number of items have been crawled or the progress bar reaches 100 percent, you can click **Show Results** to display the list of content item URLs that have been crawled so far. Also, you can click **Refresh Results** to refresh the progress bar display.

Once the specified number of items has been crawled, the window refreshes itself, and the Content Importer window appears. (See [Figure 6-4](#).)

Figure 6-4 Content Importer Window

This window displays the associated MIME type, size, date modified, and the URL.

- Step 12** Check the check box next to the content items that you want to add to the list of URLs to be acquired and distributed.
- Step 13** To add all the listed content items in the current window to the Content Manager, click **Select All** at the bottom of the window. If the number of content items exceeds ten, you might need to select the items from multiple pages. To deselect all selected items and choose different items, click **Select None**.
- Step 14** Click the **Add Selected** button at the bottom of the Content Importer window to add all selected content items to be imported. The Content Importer window closes, and the Content Acquisition Method for Channel–Use GUI to Specify Content Acquisition window appears with all the content items added.

Alternatively, click the **Show Filter** button to return to the Quick Crawl Filter window to change the filter settings.

Adding a Crawl Task to a Channel

The crawler application methodically and automatically searches websites and FTP sites that have met the specified crawl criteria and makes a copy of the visited pages for later processing. The crawler application starts with a list of URLs to visit and identifies every link on the page associated with that URL and adds these links to the list of URLs to visit.

The process ends after one or more of the following conditions are met:

- Links have been followed to a specified depth.
- The maximum number of objects has been acquired.
- The maximum content size has been acquired.

If you use a manifest file, you can specify both the *max-number* and *max-size* attributes in the <crawler> tag as the criteria to stop a crawler job, whichever condition occurs first stops the job. For example, if the crawler has acquired the maximum number of objects specified in the manifest file, the crawler stops, even if it has not yet acquired the maximum content size.

If you use the Content Distribution Manager GUI, you cannot specify both attributes for crawl tasks. However, after you save the GUI-generated manifest file (click the **Save settings locally** icon), you can modify it and add additional XML tags manually using any text or XML editor.

**Note**

If you change the manifest file that you saved, and you want to use that manifest file instead of the content that you defined in the Content Distribution Manager GUI, or if you want to use the manifest file for another channel, then you must use the **Specify external manifest file** setting (see [“Choosing a Content Acquisition Method for a Channel”](#) section on page 6-6) and point to the manifest file. When you change the content acquisition method from using the Content Distribution Manager GUI to using an external manifest file, any content items that you added using the GUI are removed from the GUI.

To add a crawl task to a channel using the Content Distribution Manager GUI, follow these steps:

- Step 1** Choose **Services > Web > Channels**. The Channels window appears, listing all channels in your ACNS network.
- Step 2** Click the **Edit** icon next to the channel for which you want to configure crawl jobs. The Modifying Channel window appears.
- Step 3** In the Contents pane, choose **Channel Content**. The Content acquisition method for channel—Use GUI to specify content acquisition window appears.
- Step 4** Click the **Add Content** icon in the taskbar to specify a crawl job. The Content Manager window for the channel appears.
- Step 5** From the Source URL drop-down list, choose the communication protocol (HTTP, HTTPS, FTP) that will be used to crawl the content from the origin server. The default is HTTP.
- Step 6** In the Source URL field, enter the URL of the host server (origin server) to specify the location from which to start crawling the website or FTP server. You can right-click to copy and paste URLs from a web browser to the Source URL field.

The value entered here is associated with the full URL of the *start-url* attribute of the <crawler> element. This value is also associated with the *proto* attribute of the <host> element. The *proto* attribute can be empty if the *name* attribute of the <host> element is a fully qualified domain name (FQDN).
- Step 7** In the Link Depth field, specify how many levels of a website are to be crawled or how many directory levels of an FTP server are to be crawled. The default depth is 10 levels. The range is –1 to 2147483636.

If depth = –1, there is no depth constraint.
If depth = 0, acquire only the starting URL.
If depth = 1, acquire the starting URL and all the content it references.

This value is associated with the *depth* attribute of the <crawler> element.
- Step 8** Click the **Define a crawl task** radio button.
- Step 9** To add the crawl task to the channel, click **Submit**. When you click **Submit**, the local manifest file is automatically reparsed, changes are detected, and the corresponding items are acquired or removed.

To define optional content acquisition rules for the crawl task, see the [“Configuring Content Acquisition Rules for a Crawl Task”](#) section on page 6-15.

Configuring Content Acquisition Rules for a Crawl Task

To configure content acquisition rules, follow these steps:

- Step 1** Navigate to the content acquisition rules configuration area:
- Choose **Services > Web > Channels**. The Channels window appears, listing all channels in your ACNS network.
 - Click the **Edit** icon next to the channel for which you want to configure crawl jobs. The Modifying Channel window appears.
 - In the Contents pane, choose **Channel Content**. The Content Acquisition Method for Channel—Use GUI to Specify Content Acquisition window appears, listing the content items and crawl tasks defined for this channel.
 - Click the **Edit** icon next to the crawl task for which you want to configure acquisition rules. The Content Manager window appears.
 - Click the **Show optional content acquisition rules** drop-down arrow. The fields for configuring content acquisition rules appear. (See [Figure 6-5](#).) [Table 6-1](#) describes the fields in this window and provides the corresponding manifest file attributes.

Figure 6-5 Configuring Content Acquisition Rules for a Crawl Task

The screenshot shows the Cisco Application and Content Networking System GUI. The main window is titled "Cisco Application and Content Networking System - Microsoft Internet Explorer". The left sidebar contains a "Contents" pane with a tree view showing "Definition", "Channel Content", "Assign Multicast Cloud", "Assign Content Engines", "Assign Device Groups", "List all assigned Content Engines", "Replication Status", and "Tools". The main area is titled "Channel: Product Information - Content Manager". It contains a "Content Source" section with a "Source URL" field set to "http://www.cisco.com" and a "Link depth" field set to "10". Below this is a table for "Content Acquisition Rules" with columns for "MIME Type", "Extension", "Time Before", "Time After", "Min Size", and "Max Size". The table shows one rule for "jpeg" files. Below the table are fields for "MIME type", "Extension", "Time before", "Time after", "Minimum size", and "Maximum size". At the bottom, there are "Add" and "Cancel" buttons, and a "Show advanced settings" link. A note at the bottom states "Note: * - Required Field".

MIME Type	Extension	Time Before	Time After	Min Size	Max Size
jpeg	jpg	8/12/2005 14:07:42	8/2/2005 14:15:00		

Step 2 In the MIME type field, specify the MIME type of the content to be acquired.

Step 3 In the Extension field, specify the extension of the files to be acquired.

Step 4 In the Time before field, enter the time that must be matched if content is to be acquired. Files that were modified before this time will be acquired. Use the format dd-mm-yyyy hh:mm:ss [TMZ] format, where TMZ (the time zone) is optional. UTC is the default. For example, you can specify 5/27/2004 14:02:07 UTC, which is a valid date format.

Alternatively, click the **Calendar** icon next to the Time before and Time after fields to display the Date Time Picker popup window.

In the Date Time Picker popup window, use the left or right arrow icons to choose previous and following years, if required. Choose a month from the drop-down list. Click a day of the month. The chosen date is highlighted in blue. Alternatively, click **Set Today** to revert to the current day. The Time field displays the current system time. Edit the fields to configure a different time. Click **Apply**.

Step 5 In the Time after field, enter the time that must be matched for content to be acquired. Files that were modified after this time will be acquired. Use the format dd-mm-yyyy hh:mm:ss [TMZ] format, where TMZ (the time zone) is optional. UTC is the default.

Step 6 In the Minimum size field, specify the minimum size of content to be acquired. Content equal to or larger than this value is acquired. Choose **MB**, **KB**, or **Bytes** as the unit of measure.

Step 7 In the Max size field, specify the maximum size of content to be acquired. Content equal to or smaller than this value is acquired. Choose **MB**, **KB**, or **Bytes** as the unit of measure.

Step 8 Click **Add** to add the rule to the list of rules. An entry is added, showing the values under each column head.



Note A maximum of 10 rules can be configured for each crawl task.

Step 9 To modify a content acquisition rule, click the **Edit Rule** icon next to the rule whose settings you want to change. Once you have finished, click the small **Update** button in the content acquisition rules section to save the edits.

Step 10 To delete a content acquisition rule, click the **Edit Rule** icon next to the rule. Click the **Delete** button in the content acquisition rules section. The rule is removed from the rules listing.

Step 11 When you have finished adding and modifying content acquisition rules, click **Update** in the Content Manager window to save your crawl task configurations and return to the list of content items.

Table 6-1 Content Acquisition Rules for a Crawl Task

GUI Parameter	Function	Corresponding Manifest Attribute
MIME Type	MIME type of content to be acquired. The object is acquired only if its MIME type matches the MIME type specified here.	<match> <i>mime-type</i>
Extension	File extension of content to be acquired. The object is acquired only if its extension matches the extension specified here.	<match> <i>extension</i>

Table 6-1 **Content Acquisition Rules for a Crawl Task (continued)**

GUI Parameter	Function	Corresponding Manifest Attribute
Time before	Date and time limit before which modified content is to be acquired. The object is acquired only if its last modified date and time are before the value specified here. The format of this field is dd-mm-yyyy hh:mm:ss [TMZ], where the time zone (TMZ) is optional. UTC is the default.	<match> <i>time-before</i>
Time after	Date and time limit after which modified content is to be acquired. The object is acquired only if its last modified date and time are after the value specified here. The format of this field is dd-mm-yyyy hh:mm:ss [TMZ], where the time zone (TMZ) is optional. UTC is the default.	<match> <i>time-after</i>
Minimum Size	Minimum size for content; acquired content size must be larger than or equal to this number, in bytes, kilobytes, or megabytes. The range for this field is 0 to 2147483636.	<match> <i>minFileSizeInB/KB/MB</i>
Maximum Size	Maximum size for content; acquired content size must be smaller than or equal to this number, in bytes, kilobytes, or megabytes. The range for this field is 0 to 2147483636.	<match> <i>maxFileSizeInB/KB/MB</i>

Configuring Advanced Content Serving Settings

In the Content Distribution Manager GUI under the Show advanced settings heading in the Content Manager window, you can configure the following settings:

- Content Serving Time
- Authentication
- URL Settings
- Content Settings

By configuring these advanced settings, you can control the manner in which content is served by the Content Engines. These settings correspond to attributes in the manifest file and are associated with the <item> and <crawler> tags. These same attributes can also be specified in the <item-group> or <options> tags in a manifest file, so they can be shared by their <item> and <crawler> subtags.

To configure advanced settings for serving content, follow these steps:

- Step 1** Navigate to the advanced settings configuration area:
- Choose **Services > Web > Channels**.
 - Click the **Edit** icon next to the name of the channel for which you wish to configure content serving settings.
 - In the Contents pane, choose **Channel Content**.
 - Click the **Edit** icon next to the content item or crawl task that you want to configure.
Alternatively, click the **Add Content** icon and configure the optional advanced settings along with the required settings for a new content item or crawl task.
 - In the Content Manager window, click the **Show advanced settings** drop-down arrow. (See [Figure 6-5](#).) The fields for configuring advanced settings are displayed, and the arrow becomes the Hide advanced settings arrow. (See [Figure 6-6](#) and [Figure 6-7](#).) [Table 6-2](#) describes the fields in these windows and provides the corresponding manifest file attributes.


Figure 6-6 Advanced Settings for Serving Content—Top of Window

The screenshot displays the Cisco Application and Content Networking System GUI in a Microsoft Internet Explorer browser window. The main navigation pane on the left shows a tree structure with 'Contents' expanded, and 'Channel Content' selected. The main content area is titled 'Channel: Product Information - Content Manager'. It contains several configuration sections:

- Content Source:** Includes a 'Source URL' field set to 'http://www.cisco.com' and a 'Link depth' field set to '10'. A note specifies: 'Specifies the link depth to which a website is to be crawled. If 0 is specified, only the start URL will be acquired.' Below this is a checkbox for 'Show optional content acquisition rules (0 rules defined)'.
- Content Serving Time:** Includes a 'High priority content' checkbox with a note: 'If selected, acquisition of this content will take precedence.' Below this are 'Start serving time' and 'Stop serving time' fields.
- Authentication:** Includes a 'Use weak SSL certificate' checkbox with a note: 'If selected, allow https protocol to accept expired or self-signed certificate.' Below this is a 'Disable basic authentication' checkbox with a note: 'If selected, acquirer will not use basic authentication while fetching content.' Below these are 'Playback authentication' (set to 'As acquired'), 'User name', 'Password', and 'User Domain Name' fields.
- URL Settings:** A section at the bottom of the configuration area.

At the bottom right of the configuration area are 'Update' and 'Cancel' buttons. The status bar at the bottom right shows the number '134683'.

Figure 6-7 Advanced Settings for Serving Content—Bottom of Window



Step 2 Under the Content Serving Time section, configure the following:

- a. Check the **High priority content** check box to specify the order of importance and hence the processing order of the item acquisition or crawl job. The higher the priority, the earlier the content is acquired and distributed.
- b. In the Start serving time field, specify a time in dd-mm-yyyy hh:mm:ss [TMZ] format at which the Content Engine is allowed to start serving the content. The time zone (TMZ) is optional; UTC is the default. For example, you can specify 5/27/2004 14:02:07 UTC, which is a valid date format. If you do not specify a serving start time, content is ready to serve once it is distributed to the Content Engine.
- c. In the Stop serving time field, specify a time in dd-mm-yyyy hh:mm:ss [TMZ] format at which the ACNS network temporarily stops serving the content. The time zone (TMZ) is optional; UTC is the default. If you do not specify a serving stop time, the ACNS network serves the content to the Content Engine until it is removed by modifying the manifest file or renaming the channel.

Alternatively, click the **Calendar** icon next to the Start serving time and Stop serving time fields to display the Date Time Picker popup window. In the Date Time Picker popup window the current date is highlighted in yellow by default. Use the left or right arrow icons to choose the previous or following years if required. Choose a month from the drop-down list. Click a day of the month. The chosen date is highlighted in blue. The Time field displays the current system time. Edit the fields to configure a different time. Click **Apply**. Alternatively, click **Set Today** to revert to the current day. The date and time chosen are displayed in the Start serving time or Stop serving time fields.

Step 3 Under the Authentication section, configure the following:

- a. Check the **Use weak SSL certificate** check box if you want to allow the HTTP protocol to accept expired or self-signed certificates.
- b. Check the **Disable basic authentication** check box if you do not want NTLM headers to be stripped off to allow fallback to the basic authentication method while acquiring content.
If you leave this check box unchecked, NTLM authentication headers can be stripped off and fallback to the basic authentication method occurs. The username and password information are passed to the origin server in clear text with a basic authentication header.
- c. In the Username field, specify the name of the user for proxy authentication.
- d. In the Password field, specify the password of the user for proxy authentication.
- e. In the User Domain Name field, specify the NTLM user domain name for the NTLM authentication scheme configured on the proxy.

Step 4 Under the URL Settings section, configure the following:

- a. Check the **No redirect to origin server** check box if you do not want to redirect content requests to the origin server. If unchecked, the Content Engine is allowed to redirect content requests to the origin server when it does not have the content in its cache.
- b. Check the **Ignore query string** check box for the ACNS software to ignore any string after the question mark (?) character in the requested URL for playback.
- c. In the Alternate URL field, specify an alternative URL to which the ACNS network can redirect the request if the content requested by the user has not yet been replicated to the Content Engine (that is, it is not ready in the ACNS network).

Step 5 Under the Content Settings section, configure the following:

- a. In the TTL field, specify a time period for revalidation of content, and choose a unit of measure from the drop-down list. If no TTL is specified, the content is fetched only once, and its freshness is never checked again.
- b. In the Retry interval field, specify a time period after which the ACNS software attempts to acquire the content again if acquisition fails, and choose a unit of measure from the drop-down list.
- c. In the Play duration field, specify the playtime duration for a video file, and choose a unit of measure from the drop-down list.
- d. In the Bit rate field, specify the bit rate of the content for download and playback, and choose a unit of measure from the drop-down list.

Table 6-2 Advanced Settings for Serving Content

GUI Parameter	Function	Corresponding Manifest Attribute
Content Serving Time		
High priority content	Processing order of the content item or crawl job.	<i>priority</i>
Start serving time	Time at which the Content Engine is allowed to start serving the content.	<i>serveStartTime</i>
Stop serving time	Time at which the ACNS network temporarily stops serving the content.	<i>serveStopTime</i>

Table 6-2 **Advanced Settings for Serving Content (continued)**

GUI Parameter	Function	Corresponding Manifest Attribute
Authentication		
Use weak SSL certificate	Allows acceptance of expired or self-signed certificates during authentication.	<i>sslAuthType</i>
Disable basic authentication	Permits NTLM headers to be stripped off to allow fallback to the basic authentication method while acquiring content.	<i>disableBasicAuth</i>
User name	Name of the user for proxy authentication.	<i>user</i>
Password	Password of the user for proxy authentication.	<i>password</i>
User Domain Name	NTLM user domain name for the NTLM authentication scheme configured on the proxy.	<i>ntlmUserDomain</i>
URL Settings		
No redirect to origin server	Disallows redirecting content requests to the origin server. This attribute is a per content object attribute, meaning that if the content has been removed, the redirection settings do not apply. However, if the software fails to acquire the content, then the settings apply.	<i>noRedirectToOrigin</i>
Ignore query string	Software ignores any string after the question mark (?) character in the requested URL for playback.	<i>ignoreQueryString</i>
Alternate URL	Alternative URL to which the ACNS network can redirect the request if the content requested by the user has not yet been replicated to the Content Engine. This attribute supports only the full URL and applies to the content being requested. If the content is deleted, the attribute does not apply. If content acquisition fails, then this attribute should apply.	<i>alternateUrl</i>
Content Settings		
TTL	Time period for revalidation of content.	<i>ttl</i>
Retry interval	Time period after which the ACNS software attempts to acquire the content again if acquisition fails.	<i>failRetryInterval</i>

Table 6-2 **Advanced Settings for Serving Content (continued)**

GUI Parameter	Function	Corresponding Manifest Attribute
Play duration	Playtime duration of a video file.	<i>playDuration</i>
Bit rate	Bitrate of the content for download and playback.	<i>bitrate</i>

Modifying Multiple Content Items

If multiple content items (single items and crawl tasks) need to be configured with the same content acquisition and distribution properties, you can select them and modify their link depth and advanced content serving settings in one instance. However, source URLs cannot be modified while you edit multiple content items and crawl tasks. Similarly, content acquisition rules, applicable for crawl tasks, cannot be modified for multiple crawl tasks. If you have selected both single items and crawl tasks for modification, even if you specify a common link depth for all crawl tasks, this value is not applied to single items.

To modify multiple content items, follow these steps:

- Step 1** Choose **Services > Web > Channels**. The Channels window appears, listing all channels in your ACNS network.
- Step 2** Click the **Edit** icon next to the channel which has content items that you want to modify. The Modifying Channel window appears.
- Step 3** In the Contents pane, choose **Channel Content**. The Content Acquisition Method for Channel–Use GUI to Specify Content Acquisition window appears.
- Step 4** Check the check box next to each content item that you want to modify.
 - To choose all listed content items in the current window, click the **All** button at the bottom of the window. (You can choose to display 10, 20, 40, or ALL content items from the **Rows** drop-down list.)
 - To deselect chosen items, click the **None** button and make a different choice.
- Step 5** Click the **Edit Selected Items** button at the bottom of the window to modify selected content items. The Channel–Content Manager window appears. By default, all fields are deactivated in this window.
- Step 6** Click the **Click to override individual settings** icon next to each field that you want to modify. The fields become available for editing. Alternatively, click the **Click to revert initial settings** icon next to the field to preserve the original settings set during configuration.
- Step 7** Click the **Show advanced settings** drop-down arrow and configure advanced settings.
- Step 8** Click **Update** to save and apply modified settings to the chosen content items. The Content Acquisition Method for Channel–Use GUI to Specify Content Acquisition window appears.

Deleting Multiple Content Items

To delete multiple content items, follow these steps:

-
Step 1	Choose Services > Web > Channels . The Channels window appears, listing all channels in your ACNS network.
Step 2	Click the Edit icon next to the channel which has content items that you want to delete. The Modifying Channel window appears.
Step 3	In the Contents pane, choose Channel Content . The Content Acquisition Method for Channel–Use GUI to Specify Content Acquisition window appears.
Step 4	Check the check box next to each content item that you want to delete. • To choose all listed content items in the current window, click the All button at the bottom of the window. (You can choose to display 10, 20, 40, or ALL content items from the Rows drop-down list.) • To deselect chosen items, click the None button, and make a different choice.
Step 5	Click the Delete Selected Items button in the taskbar to remove the selected content items. You are prompted to confirm your decision. The Content Acquisition Method for Channel–Use GUI to Specify Content Acquisition window refreshes itself with an updated list of content items.
-

Acquiring Pre-Positioned Content Using an Externally Hosted Manifest File

To configure your ACNS network to acquire content using a manifest file, you must complete the following tasks:

1. [Create a Manifest File](#)
2. [Specify Channel Acquisition and Distribution Properties](#)
3. [Configure Manifest File and Proxy Information for the Channel](#)
4. [Generate the Publishing URL](#)
5. [Assign a Playserver](#)

**Caution**

If you are using an application that dynamically generates and publishes manifest files, you must take precautions before you upgrade this application software in your ACNS network. During the software upgrade of the manifest publishing software, the manifest file could become invalid, causing the distributed content to be inadvertently deleted.

To avoid losing content in this situation, we recommend that before you upgrade the manifest publishing software, you disable acquisition and distribution processes on the root Content Engine and on any potential temporary root Content Engines in the root location. Any Content Engine that is in the same channel and same location as the root Content Engine is a potential temporary root. After the application software upgrade is complete, you can restart the acquisition and distribution process on each Content Engine.

To disable the acquisition and distribution process, use the **acquisition-distribution stop** EXEC command. To restart the acquisition and distribution process, use the **acquisition-distribution start** EXEC command.

Create a Manifest File

See [Appendix A, “Creating Manifest Files,”](#) for details on creating manifest files.

After you create the manifest file, use the Manifest Validator utility to verify the syntax. (See the [“Manifest Validator Utility”](#) section on page A-28.) Next, specify the manifest URL in the Content Distribution Manager GUI. If authentication is required to fetch the manifest file, specify a username and password as well. (See the [“Configure Manifest File and Proxy Information for the Channel”](#) section on page 6-26.)

Specify Channel Acquisition and Distribution Properties

This section discusses four important content acquisition and distribution properties that you need to define:

- Distribution priority
- Item priority
- Channel quota
- Update interval

Some of these properties are configured in the **Services > Web > Channels** section of the Content Distribution Manager GUI, and some are defined in the manifest file.

Distribution Priority

The distribution priority setting determines the priority of content acquisition and distribution. You configure this setting from the Distribution Priority drop-down list in the Content Distribution Manager GUI. The distribution priority values are High (750), Normal (500), or Low (250). [Figure 6-8](#) shows the acquisition and distribution properties and the manifest properties that you can configure in the Content Distribution Manager GUI. See the [“Creating a Channel”](#) section on page 5-11 for channel configuration procedures and field descriptions.

Figure 6-8 Acquisition and Distribution Properties for a Channel

Channel information

Name:

Request Routed FQDN:

Origin Server:

Content Provider: [New Content Provider](#)

WebSite: [New Website](#)

Live Channel: ☐

Channel Quota: MB

Acquisition and Distribution Properties

Distribution Priority: Use null cipher for Distribution: ☐

* ☒ Multicast only ☐ Multicast/Unicast ☐ Unicast only

Root CE failover/fallback grace period: mins ☐ Never ☐ If there is only one CE in root location or if 'Never' is selected, failover/fallback will not occur.

Use system-wide settings for GoS for unicast data: ☒ System-wide GoS settings page is under the Systems Tab.

GoS value for unicast data: or Select DSCP from the drop-down or enter decimal value in the text box.

Comments

The priority of content acquisition also depends on the origin server. Requests from different origin servers are processed in parallel. Requests from the same origin server are processed sequentially by their overall priority.

Item Priority

The item priority is determined by the manifest file. If the *priority* attribute is specified in the manifest file under the <item> tag, it is used as the item priority; if it is not specified, the order in which the item is listed in the manifest file (index number) is used as the item priority. (All crawled pages have the same item priority; therefore, you do not need to specify the *priority* attribute under the <crawler> tag.) See the [“Specifying Content Priority” section on page A-13](#) for more information.

Channel Quota

The channel quota is the disk space allowed for the channel. You configure the channel quota in the Content Distribution Manager GUI. (See [Figure 6-8](#).) See the [“Configure Manifest File and Proxy Information for the Channel” section on page 6-26](#) for configuration information.

When configuring the channel quota, keep in mind the following:

- The total of channel quota in all subscribed channels should not exceed the cdnfs disk space allocation of the Content Engine. See the [“Updating Storage Capacity on Your Content Engines” section in the Cisco ACNS Software Update and Maintenance Guide](#).

- The total of used disk space in a channel should not exceed the amount of disk space that you allocated for the channel in the Content Distribution Manager GUI Channel Quota field (**Services > Web > Channels > Definition**).

Because of overhead, the amount of disk space used by a file is always larger than the size of the file itself. To figure the amount of disk space needed for a file, follow these steps:

- a. Divide the actual file size in kilobytes (KB) by the file system block size, which is a fixed 4-KB (4096-byte) unit, and then round up the result to the nearest integer. This provides the number of filled and partially filled 4-KB blocks used by a file.

(File size in KB / 4096) rounded up to the next integer value = Total number of blocks per file

- b. Multiply the total number of file system blocks used by 4 KB (4096 bytes) to calculate the actual disk space consumed in bytes.

Total blocks per file * 4096 = Total disk usage in bytes

- c. Multiply 4 KB by 4 and add the product to the total disk space consumed. (The integer 4 represents disk space that is reserved for internal system usage.)

Total disk usage in bytes + (4096 bytes * 4) = Disk usage per file

Also, because the software attempts to reserve enough space for other minor internal system functions, it is helpful to configure your channel quotas (and pre-positioned disk space) with a modest amount (perhaps 10 percent) of extra space beyond the total disk space consumed.

Channel quota in kilobytes = (Total disk usage in kilobytes) + (0.1 * Total disk usage in kilobytes)

Manifest File Update Interval

The update interval is the interval for the root Content Engine to check the manifest file itself. (This is *not* the update interval for checking content.) You configure the update interval in the Content Distribution Manager GUI. (See the Check Manifest field in [Figure 6-9](#).) For configuration information, see the next section.

Configure Manifest File and Proxy Information for the Channel

The manifest file provides information about the content to be pre-positioned through the channel or information about the live and video-on-demand (VoD) content served through the channel.

To configure manifest file and proxy information for the channel, follow these steps:

-
- Step 1** In the Content Distribution Manager GUI, choose **Services > Web > Channels**.
 - Step 2** Click the **Edit** icon next to the name of the channel that you want to modify. The Modifying Channel window appears.
 - Step 3** In the Contents pane, choose **Channel Content**.
 - Step 4** In the Content Acquisition Method for Channel title bar, make sure that the content acquisition method for the channel is to specify an external manifest file. If this is not the case, click the **Change Method** button and choose **Specify external manifest file** from the drop-down list. Then click **Save**. The window refreshes, displaying the manifest settings and proxy information fields.

**Note**

If you are changing the method of content acquisition from using the GUI to using an externally hosted manifest file, all content that was defined in the GUI is removed. If you want to save the content that you defined using the GUI, you must click the **Save settings locally** icon in the taskbar. A window pops up showing the GUI-generated manifest file. Choose **File > Save As** to save the file.

- Step 5** Use the fields provided under the Define Basic Manifest Settings heading to configure manifest file information for the channel.
- Step 6** Use the fields provided under the Define Basic Manifest Proxy Information heading to configure the proxy information for the acquirer to fetch the manifest file. If a proxy server is configured, requests to fetch the manifest file from the origin server will go through the proxy server.

(See [Figure 6-9](#).) [Table 6-3](#) describes the manifest file fields in this window. Required fields are indicated by an asterisk in the GUI and in the table.

(See also the “[Configuring a Proxy Server and Port for the Manifest File](#)” section on page 6-40 and the “[Configuring Authentication for a Proxy to Fetch the Manifest File](#)” section on page 6-43.)

Figure 6-9 Manifest File and Proxy Information for a Channel

The screenshot shows the Cisco Application and Content Networking System GUI in a Microsoft Internet Explorer browser window. The main navigation bar includes 'Devices', 'Services', and 'System'. The 'Services' tab is active, showing a list of services: 'Web', 'Video', 'Platform', 'Channels', 'Content Provider', 'Websites', 'Multicast Clouds', and 'Baseline Settings'. The 'Channels' service is selected, and the 'Content acquisition method for channel, chan3' is set to 'Specify external manifest file'. The 'Define Basic Manifest Settings' section includes fields for 'Manifest URL', 'Check Manifest Every' (with a 'Fetch Manifest Now' button), 'Weak Certificate Verification', 'Disable basic authentication', 'Manifest Username', 'Manifest Password', 'Confirm Password', and 'NTLM user domain name'. The 'Define Manifest Proxy Information' section includes fields for 'Disable All Proxy', 'Proxy Hostname', 'Proxy Port', 'Proxy Username', 'Proxy Password', 'Confirm Password', 'Disable proxy basic authentication', and 'Proxy NTLM user domain name'. A 'Note' at the bottom indicates that fields with an asterisk are required. The 'Submit' and 'Cancel' buttons are at the bottom right.

- Step 7** To save the manifest file settings, click **Submit**. If you have not specified the manifest file URL, the system prompts you to confirm your decision not to supply one. You can specify the manifest file URL at a later time after configuring other settings. Click **OK** to confirm.

Table 6-3 **Manifest Properties**

Property	Description
Define Basic Manifest Settings	
Manifest URL	Address of the manifest file for the channel. The manifest URL must be a well-formed URL. If the protocol (FTP, HTTP, or HTTPS) for the URL is not specified, HTTP is used.
Check Manifest Every*	Frequency in minutes (0 to 52560000) with which the Content Engines assigned to the channel check for updates to the manifest file. (This field is required if the manifest file is specified.)
Weak Certificate Verification	When checked, enables weak certificate verification for the manifest file. This is applicable when the manifest file is fetched using the HTTPS protocol. Note To use weak certification for channel content, you need to specify weak certification within the manifest file.
Disable basic authentication	When checked, NTLM headers cannot be stripped off to allow fallback to the basic authentication method. If you leave this check box unchecked, NTLM authentication headers can be stripped to allow fallback to the basic authentication method, and the username and password information can be passed to the origin server in clear text with a basic authentication header.
Manifest Username	Username to fetch the manifest. The manifest username must be a valid ID. If the server allows anonymous login, the user ID can be null. Note The Username and Password fields allow you to enter any secure login information needed to access the manifest file at its remote location.
Manifest Password	Password for the user.
Confirm Password	Password confirmation.
NTLM user domain name	NTLM user domain name to be allowed access by the NTLM authentication scheme configured on the origin server.
Define Manifest Proxy Information	
Disable All Proxy	Disables the outgoing proxy server for fetching the manifest file. Any outgoing proxy server configured on the root Content Engine will be bypassed, and the acquirer will contact the origin server directly.
Proxy Hostname	Host name or IP address of the proxy server used by the acquirer to retrieve the manifest file.
Proxy Port	Port number of the proxy on which the acquirer fetches the manifest file. The range is from 1 to 65535.

Table 6-3 **Manifest Properties (continued)**

Property	Description
Proxy Username	Name of the user to be authenticated to fetch the manifest file.
Proxy Password	Password of the user to pass authentication on the proxy.
Confirm Password	Reentry of the same password for confirmation to pass authentication on the proxy.

Table 6-3 *Manifest Properties (continued)*

Property	Description
Disable proxy basic authentication	When checked, NTLM headers will not be stripped off to allow fallback to the basic authentication method against Microsoft Internet Information Services (IIS) servers. If you leave this check box unchecked, NTLM authentication headers can be stripped to allow fallback to the basic authentication method; the username and password information can be passed to the proxy server in clear text with a basic authentication header.
Proxy NTLM user domain name	NTLM user domain name to be allowed access by the NTLM authentication scheme configured on the proxy.

Generate the Publishing URL

A publishing URL is the URL that plays back pre-positioned content in the ACNS network. A complete publishing URL consists of three parts:

- Scheme
- Domain name
- Path

The path includes both the file directory path and the file name. The playserver list determines the publishing URL for the ACNS network. The playserver list is generated directly through the manifest file, through the <playServerTable> tag in the manifest file, or through the default playserver table.

Scheme

The scheme of the publishing URL is the protocol used to play the content type. For example, if an .asf video file can be played by both an HTTP and a RealMedia playserver, two URL schemes can be used to access this content: HTTP and RTSP.

The *scheme* is determined by the type of playserver and has the following direct mapping between playserver and scheme:

Playserver	Scheme
HTTP	HTTP
RealMedia	RTSP
QTSS	RTSP

Domain Name

The domain name of the publishing URL is determined by the configuration of the ACNS network. If WCCP is used to redirect content requests to a Content Engine, the domain name in the request URL is the origin server FQDN (fully qualified domain name) in the website or channel. If content routing is used, the content routing FQDN (the FQDN of the website) becomes the domain name.

All content acquired through the manifest file is published under the domain name that is entered in the Content Distribution Manager GUI (**Services > Web > Channels > Edit channel > Edit Website**) in the website Origin Server field (in the case of WCCP routing) or in the Request Routed FQDN field (in the case of a Content Router).

The content *must* be accessible from the website origin server FQDN, even if it is acquired from a different server.

**Note**

The manifest file contains a <server> tag. The “server” in the case of this manifest tag refers to “acquisition server,” and may or may not be the same as the website origin server FQDN that is published during content serving. Note the distinction between the acquisition server and the website origin server FQDN.

Manifest file attributes such as *requireAuth* and *noRedirectToOrigin* refer to the website origin server. Attributes such as *ttl* and *user* and *password* are related to the acquisition server. Thus, if you specify the *requireAuth* attribute for any content item, make sure that the origin server FQDN you entered in the Content Distribution Manager GUI Create New Web Site window or Modifying Web Site window (accessed through **Services > Web > Channels > Edit channel > Edit Website**) is accurate and can do the following:

1. Accept requests for the path as stated in the manifest.
2. Accept authentication requests from end users for this URL.

This reminder is applicable even when you have multiple acquisition servers specified in the manifest file. Because the “real” origin server is still the website origin server FQDN, you need to make sure that content is accessible from the website origin server FQDN *and* that the website origin server can accept authentication requests.

Path

In most cases, the path of the publishing URL is the relative source URL, or the *src* attribute in the <item> tags. For content crawling, it is a relative URL, relative to the host name of the origin server.

Certain attributes in the manifest file allow you to alter the publishing URL path. These attributes are *cdn-url* in the <item> tag, and *srcPrefix* or *cdnPrefix* in the <crawler> and <item-group> tags. These attributes convert a relative source URL into a completely new relative ACNS network URL.

For the content in the following example, the path uses default.html instead of index.html.

```
<item src="index.html" cdn-url="default.html" />
```

The relative URL is always relative to the host name. In the following example, the relative URL is index.html, not sport/index.html.

```
<server>
  <host name="http://www.cnn.com/sport/" />
</server>

<item src="index.html" />
```

In the following example, the *srcPrefix* and *cdnPrefix* attributes convert the prefix of every crawled content object from NBA/ to ABC/. The relative *cdn-url* is ABC/*. The path for the *start-url* attribute is ABC/index.html.

```
<crawler
  start-url="NBA/index.html"
  srcPrefix="NBA/"
  cdnPrefix="ABC/"
/>
```

Special Characters in URLs

Certain special characters are not allowed in URLs, according to the RFC 2396 standard. If the URL of the content to be acquired contains illegal characters, you must rewrite the URL using the ASCII code equivalent escape characters to represent the illegal special characters.

For example, the escape characters for a blank space are %20. An illegal URL that uses blank spaces in the filename, such as:

`http://www.cnn.com/file with space`

Can be rewritten as:

`http://www.cnn.com/file%20with%20space`

You might have to rewrite the URL using the ASCII equivalent escape characters when you specify a URL in the Manifest URL field in the Channel definition window of the Content Distribution Manager GUI (**Services > Web > Channels > Definition**) or when you specify a URL in the manifest file within the following tags:

`<host name=>`

`<item src=>`

`<crawler start-url=>`

`<crawler externalPrefixes=>`

`<match prefix=>`

ACNS software enforces the following rules for escaping a URL:

- The special characters ! # \$ & ' () , ; = ? serve special functions in a URL. If your filename or folder name on the origin server uses these special characters, you must rewrite the URL.

In addition to the above special characters, you must also avoid using the following characters in filenames and folder names and substitute the ASCII equivalent shown:

space ==> %20

“ ==> %22

% ==> %25

< ==> %3c

> ==> %3e

[==> %5b

\ ==> %5c

] ==> %5d

^ ==> %5e

{ ==> %7b

| ==> %7c

} ==> %7d

~ ==> %7e

- If a URL in the manifest file or in an HTML file uses a “?” character, the URL is rejected. If a URL uses a “#” character, the portion of the URL after the “#” is discarded.
- If a URL in a manifest file or in an HTML file uses other special characters, follow the requirements in RFC 2396.

Assign a Playserver

The playserver is assigned in the manifest file. See the [“Generating a Playserver List”](#) section on page A-13 for more details.

The <playServer> tag is very important for playing back pre-positioned content; it contains a list of playservers to play back the content. All ACNS pre-positioned content needs to use one of the ACNS software-supported playservers to play back content to end users.

ACNS 5.x software supports playservers that play back the following pre-positioned content types on the ACNS network: HTTP, HTTPS, WMT, and RTSP (RealMedia and QuickTime Streaming Server [QTSS]).

You can use any protocol to request content. Actually, the protocol information implies which playserver is needed to play the content. The ACNS software checks whether the requested protocol matches the list in the playserver table. If it matches, the request is delivered. If it does not match, the request is rejected.

You can generate a playserver list through:

- The manifest file, by configuring *playServer* attributes in an <item> tag
- The <playServerTable> tag, by configuring playserver MIME-type extension names

To create the playserver list directly through the manifest file, configure *playServer* attributes of the playserver list in an <item> tag. If an <item> tag does not have a *playServer* attribute, its playserver list is generated through the <playServerTable> tag. If the <playServerTable> tag is omitted in the manifest file, a built-in default <playServerTable> tag is used to generate the playserver list. Multiple servers are separated by commas, as shown in the following example:

```
<item src="video.mpg" playServer="real,wmt" />
```

You can also generate the playserver list that supports these streaming media types through the <playServerTable> tag. The <playServerTable> tag maps content into a playserver list based on the MIME-type extension name. If there is a <playServerTable> tag in the manifest file, use it to generate the playserver list.

To generate the playserver list through the <playServerTable> tag, use MIME-type extension names to configure which playserver can play the particular pre-positioned content, as shown in the following example:

```
<playServerTable>
<playServer name="real">
  <contentType name="application/x-pn-realaudio" />
  <contentType name="application/vnd.rn-rmadriver" />
  <extension name="rm" />
  <extension name="ra" />
  <extension name="rp" />
  <extension name="rt" />
  <extension name="smi" />
</playServer>
<playServer name="wmt">
  <extension name="wmv" />
  <extension name="wma" />
  <extension name="wmx" />
  <extension name="asx" />
  <extension name="asf" />
  <extension name="avi" />
</playServer>
<playServer name="http">
  <contentType name="application/pdf" />
  <contentType name="application/postscript" />
```

```

        <extension name="pdf" />
        <extension name="ps" />
    </playServer>
</playServerTable>

```

The <playServerTable> tag is used to generate a playserver list for each content type. Note that in the preceding example, any file with a PDF or a PostScript extension uses HTTP to play the content.

HTTP and HTTPS are the default playservers. In other words, if you did not use the customized <playServerTable> tag or the playServer name field to specify playservers for content, HTTP and HTTPS are added to the playServer name fields to make sure the content can always be played back through HTTP and HTTPS.

If you do use a <playServerTable> tag or *playServer* attribute in the manifest file, HTTP and HTTPS are not automatically allowed for playback. In this case, if you want to use HTTP or HTTPS to play certain content, you must specify the protocol, as shown in this example:

```

<CdnManifest>

<playServerTable>
  <playServer name="wmt">
    <extension name="asf" />
  </playServer>
  <playServer name="tvout">
    <extension name="mpg" />
  </playServer>
  <playServer name="https">
    <contentType name="text/html" />
  </playServer>
</playServerTable>

<item src="http://www.aaa.com/a.asf" />
<item src="http://www.aaa.com/b.html" />
<item src="http://www.aaa.com/c.mpg" />
<item src="http://www.aaa.com/d.jpg" playServer="http,https" />
<item src="http://www.aaa.com/e.rm" />

</CdnManifest>

```

In this example, the playserver list is generated as follows:

a.asf:	wmt:	from <playServerTable> by extension
b.html:	https:	from <playServerTable> by contentType
c.mpg:	tvout:	from <playServerTable> by extension.
d.jpg:	http + https:	from "playServer" attribute
e.rm:	real + http + https:	from build-in <playServerTable>

Because there is no playserver specified in the manifest file for the .rm extension, the built-in rule matches it to the RealMedia playserver. Also HTTP and HTTPS playservers are automatically added if the built-in default <playServerTable> tag is used.

The TV-out playserver is a special playserver because it allows all pre-positioned content to be played back. If you want to exclude content from being played back by other playservers, specify the TV-out playserver. ACNS software does not allow other playservers to play content once a playserver has been specified.

Acquiring Content through a Proxy Server

Acquisition through a proxy server can be configured when the root Content Engine cannot directly access the origin server because the origin server is set up to allow access only by a specified proxy server. When a proxy server is configured for root Content Engine content acquisition, the acquirer contacts the proxy server instead of the origin server, and all requests to that origin server go through the proxy server.

**Note**

Content acquisition through a proxy server is only supported for HTTP requests. It is not supported for HTTPS or FTP.

If a transparent WCCP proxy is used, you do not need to configure a proxy for content acquisition. In such cases, HTTP requests from the acquirer are redirected to the WCCP proxy by a router. However, when you do not have a router to redirect HTTP requests to the proxy server, you must configure the Content Engine to use the proxy server.

There are three ways to configure the proxy server: through the Content Distribution Manager GUI, through the Content Engine CLI, or through the manifest file. If you need to configure the Content Engine to use the proxy for both caching and pre-positioned content, use the CLI to configure the proxy. The CLI command is a global configuration command that configures the entire Content Engine to use the proxy. If only the acquirer portion of the Content Engine needs to use the proxy for acquiring pre-positioned content, use the manifest file or to specify the outgoing proxy. When you configure the proxy server in the manifest file, you are configuring the acquirer to use the proxy to fetch content for a particular channel.

**Note**

Proxy configurations in the manifest file take precedence over proxy configurations in the CLI. Furthermore, a *noProxy* configuration in the manifest file takes precedence over the other proxy server configurations in the manifest file.

You can also configure a proxy for fetching the manifest file by using the Content Distribution Manager GUI (the Creating New Channel or Modifying Channel window). When you configure a proxy server in the Content Distribution Manager GUI, the proxy configuration is valid only for acquiring the manifest file itself and not for acquiring the channel content. Requests for the manifest file go through the proxy server, whereas requests for content go directly to the origin server.

**Tip**

Before configuring a proxy server, verify that the root Content Engine is able to ping the proxy server. If the proxy is not servicing the configured port, you will get the message: “failed: Connection refused.”

Configuring Host and Proxy Server Settings Using the Content Distribution Manager GUI

To configure proxy server settings for content items defined using the Content Distribution Manager GUI, follow these steps:

Step 1

Choose **Services > Web > Channels**. The Channels window appears, listing all channels in your ACNS network.

- Step 2** Click the **Edit** icon next to the channel for which you want to configure proxy server settings for the manifest file items. The Modifying Channel window appears.
- Step 3** In the Contents pane, choose **Channel Content**. The Content Acquisition Method for Channel–Use GUI to Specify Content Acquisition window appears.
- Step 4** Click the **Manage Host and Proxy Settings** icon in the taskbar to configure proxy server information. The Content Hosts window appears, listing all previously created host URLs, the number of content items for each host, and proxy server (if configured).
- Step 5** Check the check box next to the host URL for which you want to configure proxy server settings. You can also select multiple host URLs and configure proxy servers.
- Step 6** Click the **Manage Proxy for Selected Hosts** button. A new window appears. Under the Defining proxy server for the following hosts heading, a bulleted list of host servers is displayed for which proxy servers are being configured.
- To return to the Content Acquisition Method for Channel–Use GUI to Specify Content Acquisition window, click the **Return to Content Listing** button.
- Step 7** Use the fields provided under the Proxy Server specifications section on the right to configure the proxy information for the acquirer to fetch the manifest file. If the proxy server is configured, requests to fetch the manifest file from the origin server will go through the proxy server. (See [Table 6-4](#) for a description of the fields for configuring proxy server properties.)
- You can configure the proxy server for the acquirer by using the fields in this section. These fields are associated with the attributes specified in the <proxyServer> tag in the manifest file. The proxy server configured in the manifest file is valid for all the hosts named in the manifest file. If the specified proxy fails, the acquirer, by default, contacts the origin server directly and tries to fetch the object.
- Step 8** Click **Add** to add the proxy server to the Select a proxy server section on the left. If you need to change any of the fields, click **Cancel** to erase the values and start over.
- Step 9** To modify a proxy host, click **Edit** at the bottom of the Select a proxy server section. The configured values are displayed in the fields provided under the Proxy Server specifications section on the right and can be modified. Once you have finished, click **Update** to save the modified settings.
- Step 10** To delete a proxy host, click **Delete** at the bottom of the Select a proxy server section.
- Step 11** In the Select a Proxy Server section, a list of already defined proxy servers (listed by proxy host name) is displayed.
- To define a proxy server for a host URL, choose the proxy host name (the selected proxy host is highlighted) and click the **Save Assignment** button. The Content Hosts window appears.
 - To specify that all the requests for a particular host should go directly to the origin server and should not use the proxy, choose **Do not use proxy server** and click the **Save Assignment** button. The Content Hosts window appears.
- This option is associated with the *noProxy* attribute specified in the <host>, <item>, or <crawler> tags. The noProxy= “1” designation specifies that the item and the crawl task data are to be fetched directly from the origin server. (See the [“About the noProxy Attribute”](#) section on page 6-39 for more information.)
- If you have selected multiple host URLs (from the Content Hosts window) for which you want to define and assign a proxy server, a “Hosts selected use different proxy servers.” warning message is displayed on top of the window only if the selected host URLs have been previously assigned to different proxy servers.

To assign a common proxy server for all selected host URLs, choose the proxy host name (the selected proxy host is highlighted) and click the **Save Assignment** button. Otherwise, choose **Do not change assignments** to retain the previously defined proxy host assignments for the host URLs and click the **Save Assignment** button.

- If you have changed the proxy server for host URLs and want to leave the assignment unchanged, click the **Cancel Assignment** button. The Content Hosts window appears.

Table 6-4 Proxy Server Settings

GUI Parameter	Function	Manifest Attribute
Proxy Hostname	Host name or IP address of the proxy server used by the acquirer for content acquisition. When you use a domain name instead of an IP address, make sure that the domain name can be resolved by the DNS server configured.	<i>serverName</i>
Proxy Port	Port number of the proxy server on which the acquirer fetches content. The range is 1 to 65535.	<i>port</i>
Disable Basic Authentication	When checked, NTLM headers cannot be stripped off to allow fallback to the basic authentication method. If you leave this check box unchecked, NTLM authentication headers can be stripped to allow fallback to the basic authentication method and the username and password information can be passed to the origin server in clear text with a basic authentication header.	<i>disableBasicAuth</i>
User Name	Name of the user to be authenticated to fetch the manifest file.	<i>user</i>
Password	Password of the user to pass authentication from the proxy.	<i>password</i>

Configuring an HTTP Proxy Server Using the CLI

To configure a proxy server using the CLI, use the **http proxy outgoing** global configuration command. This is a global proxy configuration for caching and pre-positioned content acquisition. You can use this command to configure a list of outgoing proxy servers through the CLI, and when configured, all the requests (for any host) will go through the proxy server. The global proxy configuration can also be reset using the **no** form of the command, as shown in this example:

```
ContentEngine(config)# http proxy outgoing ?
connection-timeout Timeout period used for probing outgoing proxy servers in microseconds
host               Use Outgoing HTTP Proxy
monitor            Interval at which to monitor the outgoing proxy servers
origin-server      Use Origin Server if all outgoing proxies failed.
```



```

preserve-407      Preserve 407 HTTP authentication header

ContentEngine(config)# http proxy outgoing host ?
Hostname or A.B.C.D  Hostname or IP address of outgoing proxy

ContentEngine(config)# http proxy outgoing host 128.107.192.24 ?
<1-65535>  Port of outgoing proxy

ContentEngine(config)# http proxy outgoing host 128.107.192.24 80

```

If the acquirer is unable to contact a proxy server, the proxy is considered disabled. The acquirer retries a disabled proxy after the period specified in the **monitor** option of the **http proxy outgoing** command. If no monitor interval is configured, the retry interval defaults to every 30 minutes.

If all outgoing proxies fail and you want the acquirer to contact the origin server directly, use the **origin-server** key word, as shown in this example:

```
ContentEngine(config)# http proxy outgoing origin-server
```

If a list of outgoing proxy servers is configured in the CLI and the first proxy server fails, the next proxy server in the list will be tried, and so on. If all the proxies fail, the acquirer will fail over to the origin server, depending on the configuration specified in the **http proxy outgoing origin-server** command.

To verify that the outgoing proxy is configured in the CLI, use the **show http proxy EXEC** command, as shown in this example:

```

CE-507# show http proxy
      Incoming Proxy-Mode:
        Not servicing incoming proxy mode connections.
      Outgoing Proxy-Mode:
        Primary Proxy Server:          128.107.192.190 port    9090

        Monitor Interval for Outgoing Proxy Servers is 60 seconds

        Timeout period for probing Outgoing Proxy Servers is 300000 microseconds

        Use of Origin Server upon Proxy Failures is disabled.

```

To confirm whether the content is being acquired through a proxy, use the acquisition and distribution transaction log, and check the proxyUsed line. This line shows the proxy server used for the request.

Configuring a Proxy Server Using the Manifest File

You can configure the proxy server for the acquirer by using the <proxyServer> tag in the manifest file. The proxy server configured in the manifest file is valid for all the hosts named in the manifest file. If the specified proxy fails, the acquirer, by default, contacts the origin server directly and tries to fetch the object.

If you have multiple proxy servers in the manifest file, and if you want to associate different proxy servers for different hosts, you must include the *proxyServer* attribute in the <host> tag. In this first example, the <proxyServer> tag is associated with the <host> tag through its *proxyServer* attribute (proxyServer="172.19.226.242"). All items that use this <host> tag will use this proxy server.



Note

The <proxyServer> tag must be located at the top level of the manifest file, directly under the <CdnManifest> tag; it cannot be used as a subtag of any other tags. It is associated with the <host>, <item>, and <crawler> tags through the *proxyServer* attribute.

```

<CdnManifest>
  <proxyServer
    serverName="172.19.226.242"
    port="8080"
  />
  <server name="my-devbox">
    <host name="http://vista2.cisco.com"
      proxyServer="172.19.226.242"
      proto="http" />
    </server>
    <item src="HR/salary.html" />
    <item src="HR/Holidays.html" />
  </CdnManifest>

```

In the following example, the *proxyServer* attribute is associated with the `<item>` tag and the `<crawler>` tag. The proxy server is applied to the second `<item>` tag and the `<crawler>` tag, but not the first `<item>` tag.

```

<CdnManifest>
  <item src="http://www.msnbc.com/computer/jokes.html" />
  <proxyServer serverName="128.107.192.24" port="80" />
  <item src="http://www.cnn.com/War/World-war-two.jpg" />
  <crawler start-url="http://www.abc.com/War/World-war-two/index.html" />
</CdnManifest>

```

Instead of using an IP address for the *serverName* attribute in the `<proxyServer>` tag, you can also use a domain name, as shown in this example:

```

<proxyServer
  serverName="spachiap-unil"
  port="8080"
/>
<server name="my-devbox">
  <host name="http://128.107.150.26/nfs-obsidian/Unicorn"
    proxyServer="spachiap-unil"
    proto="http" />
  </server>

```

When you use a domain name instead of an IP address, make sure that the domain name can be resolved by the DNS server configured.

About the *noProxy* Attribute

You can configure the *noProxy* attribute in the manifest file to specify that all the requests for a particular host should go directly to the origin server and should not use the proxy specified in the **http proxy outgoing host** command, if any proxies were configured in the CLI.

The *noProxy* attribute can be specified in the `<host>`, `<item>`, or `<crawler>` tags. In the following example, if the root Content Engine has a proxy configured through the **http proxy outgoing host** command, the first item in the manifest file uses that proxy, but the rest of the items or crawl tasks in this manifest file do not. The *noProxy="1"* designation specifies that the second item and the crawl task data are to be fetched directly from the origin server.

```

<CdnManifest>
  <item src="http://www.msnbc.com/computer/jokes.html" />
  <item src="http://www.cnn.com/War/World-war-two.jpg"
    noProxy="1" />
  <crawler start-url="http://www.abc.com/War/World-war-two/index.html"
    noProxy="1" />
  <server name="cisco">
    <host name="http://www.cisco.com" noProxy="1" />
  </server>

```

```

    </server>
    <crawler start-url="product/routers/" />
  </CdnManifest>

```

Configuring a Proxy Server and Port for the Manifest File

A proxy server and port for fetching the manifest file can be set in the Channel Content window of the Content Distribution Manager GUI. If a proxy server is configured, requests for the manifest file go through the proxy server that is specified in the Content Distribution Manager GUI.



Note

The proxy configuration from the Channel Content window only applies to fetching the manifest file. It does not apply to content acquisition.

If the manifest file resides on an origin server that requires a proxy, you need to specify the proxy name or IP address and the proxy port in the fields provided in the Define Manifest Proxy Information section of the Channel Content window in the Content Distribution Manager GUI. (See [Figure 6-9](#). For a description of the fields in this window, see [Table 6-3](#).)

If the manifest file resides on an origin server that cannot use the proxy configured from the **http proxy outgoing host** CLI command, you must check the **Disable All Proxy** check box.

The **Disable All Proxy** check box, when checked, overrides the values in the Proxy Hostname and Proxy Port fields.



Note

To activate the Manifest Proxy Information fields in the Channel Content window, you must first enter a URL in the Manifest URL field.

Acquiring Content Using Indexed Web Servers

A root Content Engine using HTTP that crawls an origin server to acquire content first makes a request for the starting URL that was specified when you configured the Channel Content window settings in the Content Distribution Manager GUI (**Services > Web > Channels > Channel Content**). The root Content Engine then parses the HTML of that location for links to other files. Those other files are then requested. If those files are HTML files, they are also parsed for links to additional files. In this manner, the Content Engine crawls through the origin server much like a client browses through a website.

A website that has indexing enabled and the default document feature disabled generates HTML that contains a directory listing whenever a directory URL is given. That HTML contains links to the files contained in that directory. This indexing feature makes it very easy for the root Content Engine to get a full listing of all the content in that directory so that it can be acquired.



Note

When you specify a starting URL for acquisition, if that URL does not end with “/” then the acquirer parses the “..” link that is returned with a directory index, and the acquirer crawls the parent directory of that URL. This behavior is not usually desired. When you enter an indexed directory as the starting URL for a crawl, always end that URL with “/” so that crawling is limited to that subdirectory.

Certain limitations to the acquirer's crawling functionality might prevent the complete set of desired content from being acquired by the Content Engine. For example, the acquirer does not crawl within script elements. So, if a client browser runs a script that provides links to additional content, those links and that additional content cannot be seen or acquired by the Content Engine.

One way to ensure that all the desired content is acquired is to configure the Channel Content window settings or manifest file to acquire individual items rather than perform a crawl. However, that method is not practical for large web sites or for web sites that are constantly changing. A better workaround would be to create a multi-homed web server, or virtual website, to allow indexed access to the same directory structure. This virtual website can be secured so that only the root Content Engine is allowed access.

The steps for creating a virtual website and performing acquisition will vary based on the web server software you use, how your website is configured, where the content exists, and many other factors. For this discussion we use the Microsoft IIS server as an example, but other web servers can be configured similarly.

To create a virtual website and configure an IIS server to allow indexed access, follow these steps:

-
- Step 1** Login to the IIS server and create a virtual website.
 - Step 2** In the Properties dialog box for the virtual website that you created, assign an IP address or TCP port to the virtual website that is different from your original website IP address or TCP port setting.
 - Step 3** Choose the Document tab and uncheck the **Enable Default Document** check box.
When you clear this check box, it prevents the default document from interfering with directory indexing.
 - Step 4** Set the Directory Security settings with the same security settings as your original website, or restrict access so that only the root Content Engine and any backup root Content Engines have access to the site.
 - Step 5** Choose the Home Directory tab, and in the Local Path field, enter the same path as your original website directory.
 - Step 6** Check the **Direct browsing** check box.
 - Step 7** Create the manifest file making sure that you enter the hostname of the virtual website and include the correct port.

The first line of your manifest file must contain the XML version tag `<?xml version="1.0"?>` so that the manifest can be validated without error, as shown in the following example:

```
<?xml version="1.0" ?>
<CdnManifest>
  <server name="Horizon">
    <host name="http://10.86.46.42:81" />
  </server>
  <crawler start-url="/intro/" depth="5" />
</CdnManifest>
```

The root Content Engine is now able to acquire all the content within the subdirectories.

Authentication Support

The acquirer supports two types of authentication schemes: basic and NTLM authentication. Authentication information is configured either in the Content Distribution Manager GUI or in the manifest file, depending on what is being acquired.

If authentication is required to fetch the manifest file from the origin server or from a proxy server, you can specify the authentication information in the Content Distribution Manager GUI. From the **Services > Web > Channels** Channel Content window, you can disable basic authentication, enter a username and password, and specify an NTLM user domain name. (Figure 6-9 shows the fields for entering authentication information, and Table 6-3 describes these fields.)

**Note**

The authentication configuration from the Channel Content window only applies to fetching the manifest file. It does not apply to content acquisition.

If authentication is required to fetch pre-positioned content from the origin server or from a proxy server, you must specify the authentication information in the manifest file. (See the next section, “[Acquiring Content that Requires NTLM Authentication](#).”)

Acquiring Content that Requires NTLM Authentication

In ACNS 5.x software, a root Content Engine can acquire content from HTTP origin servers or proxy servers after performing basic authentication. However, origin servers often support NTLM authentication only. In ACNS 5.5 software, the acquirer can act as an NTLM client to talk to NTLM-capable origin servers.

**Note**

The Content Engine acquirer only supports NTLM Version 1.

If authentication is required to fetch content, you must specify the authentication information in the manifest file by using following manifest file attributes:

- *user*—User account
- *password*—Password

If NTLM authentication is required, you can specify the following two additional manifest file attributes:

- *ntlmUserDomain*—NTLM user domain name. This attribute is required for an NTLM authentication scheme.
- *disableBasicAuth*—(Optional attribute) Disables basic authentication, if needed. If you always want your username and password to be used for NTLM authentication and not basic authentication, set this attribute to true. If this attribute is omitted, the default is false.

These attributes can be specified in the <host>, <item>, <item-group>, and <crawler> tags. See [Appendix A, “Creating Manifest Files,”](#) for more information.

If an origin server or proxy server supports both basic authentication and NTLM authentication, the acquirer chooses the first supported authentication scheme from the response challenge list. The authentication scheme used by the acquirer can also depend on the specified authentication information. For example, an IIS server responds with NTLM as the first scheme and basic authentication as the second. In this case the acquirer chooses NTLM. However, if the *ntlmUserDomain* attribute is not specified in the manifest file, the acquirer chooses basic authentication. If the *disableBasicAuth* attribute is set to true, the acquirer chooses NTLM authentication.

You can view the acquisition and distribution transaction log from the root Content Engine CLI to determine which authentication scheme the acquirer has used.

Proxy Authentication Support

If the proxy for the root Content Engine was configured in the root Content Engine CLI (**http proxy outgoing host** command), you can set the authentication information for the proxy from either the Content Distribution Manager GUI (see the [“Configuring Authentication for an HTTP Proxy Using the Content Distribution Manager GUI”](#) section on page 6-44) or from the root Content Engine CLI (see the next section, [“Configuring Authentication for a Proxy Using the CLI”](#)).

Configuring Authentication for a Proxy to Fetch the Manifest File

To configure the authentication information for fetching the manifest file from a proxy server, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**.
 - Step 2** Click the **Edit** icon next to the name of the channel that you want to modify.
 - Step 3** In the Contents pane, choose **Channel Content**.
 - Step 4** Enter the authentication information in the appropriate fields under the Define Manifest Proxy Information heading. (Figure 6-9 shows the fields for entering authentication information, and Table 6-3 describes these fields.)
 - Step 5** To save the settings, click **Submit**.
-

Configuring Authentication for a Proxy Using the Manifest File

The following example specifies authentication information for the proxy server. See [Appendix A, “Creating Manifest Files,”](#) for more information.

```
<CdnManifest>
  <!-- specify a proxy server/port and its authentication info
    this proxy supports Basic auth so only user/password is needed -->
  <proxyServer serverName="128.107.192.24" port = "80"
    user="johnz" password="xxx123yyy"/>

  <!-- this item is below the proxyServer tag; it is using the above proxy -->

  <item src="http://www.cnn.com/War/World-war-two.jpg" />

  <!-- specify a proxy server/port and its authentication info
    this proxy requires NTLM auth, so ntlmDomainName is needed
    for extra password security, users want to disable basic auth
    so their user/password is not sent over the wire -->

  <proxyServer serverName="company-proxy" port = "80"
    user="johnz" password="xxx123yyy"
    ntlmUserDomain="cisco-eng"
    disableBasicAuth="1" />

  <!-- this crawler is below the proxyServer tag; it is using the proxy -->

  <crawler start-url="http://www.abc.com/War/World-war-two/index.html"/>
</CdnManifest>
```

Configuring Authentication for an HTTP Proxy Using the Content Distribution Manager GUI

If a root Content Engine is configured to receive content through a proxy server, the acquirer running on the root Content Engine must be authenticated by the proxy server before it can obtain content from the origin server.

To configure authentication information for a proxy that was specified using the **http proxy outgoing host** command, you can use the Acquirer Outgoing Proxy Authentication section in the Content Distribution Manager GUI (**Devices > Devices > Applications > Web > HTTP > HTTP Connections**) to set the authentication information for the proxy. ACNS software supports multiple proxies; therefore, you can set the proxy authentication information for multiple proxies as well.

**Note**

If you specify the `<proxyServer>` tag in the manifest file, or if you enter the proxy host IP address in the Manifest URL field in the Channel Content window of the Content Distribution Manager GUI (**Services > Web > Channels**), the **http proxy outgoing host** command is ignored.

To configure the acquirer outgoing proxy authentication information, follow these steps:

-
- Step 1** In the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the root Content Engine.
- Step 3** In the Contents pane, choose **Applications > Web > HTTP > HTTP Connections**. The HTTP Connection Settings for Content Engine window appears.
- Step 4** Under the Acquirer Outgoing Proxy Authentication heading, enter the following information for the outgoing proxy that is listed in each row:
- To acquire content from the origin server, enter the name of the user to be authenticated in the Username field. This username will be used for both NTLM and basic authentication.
 - Enter the password of the user in the Password field. Reenter the same password in the Confirm Password field for confirmation. The password details appear as asterisks.
 - In the NTLM User Domain field, enter the NTLM server domain name to be used to authenticate user access.
 - To disallow the removal of NTLM headers and fall back to the basic authentication method, check the **Disable basic authentication** check box.
- If you leave this check box unchecked, NTLM authentication headers can be stripped to allow fallback to the basic authentication method against Microsoft Internet Information Services (IIS) servers; the username and password information can be passed to the origin server in clear text with a basic authentication header.
- Step 5** To save the authentication settings for the outgoing proxy, click **Submit**.
-

Configuring Authentication for a Proxy Using the CLI

To configure the authentication information for a proxy that was specified using the **http proxy outgoing host** command, you can also use the **acquirer proxy authentication outgoing** command from the root Content Engine CLI. The following example shows the authentication configuration for a nontransparent proxy server (IP address 192.168.1.1, port 8080) with NTLM authentication:

```
CE(config)# acquirer proxy authentication outgoing 192.168.1.1 8080 myname password
password ntlm mydomain basic-auth-disable
```

**Note**

If there is a transparent proxy between the root Content Engine and the origin server, you can use the **acquirer proxy authentication transparent** command from the root Content Engine CLI to specify the authentication information for the transparent proxy. (See the [“Configuring Authentication for a WCCP Proxy Using the CLI”](#) section on page 6-47.)

To verify that the authentication information has been set up correctly, use the **show acquirer proxy authentication** EXEC command.

The acquirer supports proxy chaining as long as there is only one proxy in the chain that requires authentication. The acquirer will fail if more than one proxy in the chain requires authentication.

Configuring Authentication for a WCCP Proxy Using the Content Distribution Manager GUI

In a transparent caching environment using a WCCP proxy, requests for content are redirected to a Content Engine by a WCCP-enabled router. When an HTTP proxy-style request from the root Content Engine to the origin server is intercepted by a WCCP proxy that requires authentication, you can configure authentication settings for the WCCP proxy.

To configure acquirer WCCP proxy authentication settings in a transparent proxy environment, follow these steps:

- Step 1** Choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to specify acquirer WCCP proxy authentication settings. The Device Home window appears.
- Step 3** In the Contents pane, choose **Prepositioning > Acquirer WCCP Proxy Authentication**. The Acquirer WCCP Proxy Authentication for Content Engine window appears. (See [Figure 6-10](#).)

Figure 6-10 Acquirer WCCP Proxy Authentication

The screenshot shows the Cisco Application and Content Networking System GUI in a Microsoft Internet Explorer browser window. The main navigation pane on the left is expanded to 'Prepositioning' > 'Acquirer WCCP Proxy Authentication'. The main content area displays the 'Acquirer WCCP Proxy Authentication Settings' for Content Engine, CONTENTENGINE. The settings include:

- Enable:** ☒ (checked)
- Username:** (required field, indicated by a red asterisk)
- Password:**
- Confirm Password:**
- NTLM user domain:**
- Disable basic authentication:** ☐

A note at the bottom states: 'Note: * - Required Field'. The bottom right corner of the window has 'Submit' and 'Cancel' buttons. The system status bar at the top right indicates 'Devices: 3 Devices, Critical' and 'Content: 17 Channels, Minor'.

- Step 4** Check the **Enable** check box to enable transparent proxy authentication for the acquirer.
- Step 5** In the Username field, enter the name of the user to be authenticated to acquire content from the origin server. This username is used for both NTLM and basic authentication.

- Step 6** Enter the password of the user in the Password field. Reenter the same password in the Confirm Password field for confirmation. The password details appear as asterisks.
- Step 7** In the NTLM user domain field, enter the NTLM server domain name to be used to authenticate user access.
- Step 8** Check the **Disable basic authentication** check box to disallow the removal of NTLM headers. When checked, NTLM headers will not be stripped off, and therefore fallback to the basic authentication method is prevented.
- If you leave this check box unchecked, NTLM authentication headers can be stripped to allow fallback to the basic authentication method against Microsoft Internet Information Services (IIS) servers; the username and password information can be passed to the origin server in clear text with a basic authentication header.
- Step 9** To save your settings, click **Submit**.
- A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The Reset button is visible only when you apply default or device group settings to change the current device settings but the settings have not yet been submitted.

Modifying WCCP Proxy Authentication Settings

You can modify the following WCCP proxy authentication settings in any order:

- To delete the configured settings for the Content Engine, click the **Remove Device Settings** icon in the taskbar to delete the settings. This icon appears only if you have configured the settings for the Content Engine.
- To restore the factory default settings to the Content Engine, click the **Apply Defaults** icon in the taskbar.
- To override the device group settings applied to the Content Engine with the factory default settings, click the **Override Group Settings with Defaults** icon in the taskbar. This icon appears only if you have applied the device group settings to the Content Engine.
- When settings have been applied from device groups with which the Content Engine is associated, click the **Override Group Settings** icon in the taskbar to override the device group settings and configure the device settings. This icon appears only if you have applied the device group settings to the Content Engine.
- When a Content Engine is associated with one or many device groups that have been configured with acquirer WCCP proxy authentication settings, a drop-down list appears in the taskbar. You can choose a device group name from this list if you want to apply settings from a different device group to this Content Engine.

Configuring Authentication for a WCCP Proxy Using the CLI

For a transparent WCCP proxy, if authentication is required, you can specify the authentication information by using the **acquirer proxy authentication transparent** global configuration command in the root Content Engine CLI. The following example shows the authentication configuration for a transparent WCCP proxy server with basic authentication. The username is *admin*, and the password is *default*.

```

CE(config)# acquirer proxy authentication transparent
CE(config)# acquirer proxy authentication transparent admin
CE(config)# acquirer proxy authentication transparent admin password
CE(config)# acquirer proxy authentication transparent admin password default

```

Verifying the Results

Use the following **show** CLI commands to verify acquisition results.

Table 6-5 *show Commands for Content Acquisition*

Command	Description
show acquirer channel	Shows how many channels use the Content Engine as the root Content Engine.
show acquirer progress	Shows the progress of the acquisition.
show statistics acquirer	Shows the result of content acquisition: how many items have been acquired and how much disk space has been used.
show statistics acquirer job-list	Displays the details of all the single items and crawler jobs for the channel specified by ID or name.
show statistics acquirer error	Shows the detailed error message for an acquisition failure. Note For a crawl job, only the first 100 errors encountered are displayed.
show statistics replication	Shows the replication status.

Troubleshooting Content Acquisition

To monitor acquisition progress and to troubleshoot, use the following commands from the root Content Engine CLI:

- Use **show acquirer channel EXEC** command to obtain channel information, such as the channel-id and channel-name, that you need to enter in other show acquirer commands, such as the **show acquirer progress** command. In the following example, the channel-id is 793 and the channel-name is group01-cifs.

```

CE# show acquirer channels
Querying Database.....
Acquirer information for all channels:
-----
Channel-id       : 793
Channel-Name     : group01-cifs
WebSite-Name     : group01-cifs
Root-CE-Type    : Configured
State           : Enabled
Disk Quota      : 200 MB
Origin FQDN     : cdn.allcisco.com
Channel Priority : 500
Manifestfile-TTL : 5
Manifestfile-URL : ftp://10.1.1.1/cifs.xml

```

- Use the **show acquirer EXEC** command to make sure that the acquirer process on the root Content Engine is working correctly, and that the device is using the expected amount of bandwidth for acquisition. The following example shows that the acquirer is running properly and that the device is configured with unlimited bandwidth for acquisition of content.

```
Content Engine# show acquirer
Acquirer is running OK
Current Acquisition Bandwidth:Not Limited
```

- Use the **show acquirer progress EXEC** command to check how far the acquisition of content has progressed. A specific channel ID or channel name can be specified to obtain the progress for a specific channel. In the example below, the acquirer has already acquired 2237 items.

```
ContentEngine# show acquirer progress channel-id 793
Querying Database.....

Acquirer progress information for channel ID:793 Channel-Name:group01-cifs
-----

Acquired Single Items      :          0 / 0
Acquired Crawl Items       :    2237 / 2500    -- start-url=www.mtv.com//
```

- Use the **show statistics acquirer channel-id** or **show statistics acquirer channel-name EXEC** command to obtain the detailed acquisition statistics for a given channel. In the example below, there was an error acquiring two items.

```
ContentEngine# show statistics acquirer channel-id 793
Querying Database.....

Statistics for Channel Channel-id :793 Channel-Name :group01-cifs
-----

Manifest:
-----
      Fetch Errors      :0
      Parsing Errors    :0
      Parsing Warnings  :0

Acquisition:
-----
      Total Number of Acquired Objects      :2237
      Total Disk Used for Acquired Objects  :981511280 Bytes
      Total Number of Failed Objects        :2
      Total Number of Re-Check Failed Objects :0
```

- Use the **show statistics acquirer errors channel-id** or **show statistics acquirer errors channel-name EXEC** command to see the reasons why the errors occurred. In the example below, one error occurred because there was a problem acquiring the URL. The other error occurred because the disk quota for the channel configured in the Content Distribution Manager GUI would have been exceeded if the specified URL had been acquired. You can increase the channel disk quota to correct this error.

```
Content Engine# show statistics acquirer errors channel-id 793
Querying Database.....

Acquisition Errors for the Channel ID:793
-----
Crawl job:start-url http://www.mtv.com//
Crawl Errors
-----
Internal Server Error(500):http://cgi.cnn.com/entries/intl-emailsubs-confirm
Exceeded Disk Quota(703):http://www.cdt.org/copyright/backgroundchart.pdf
```

- If more detailed troubleshooting of content acquisition is required, you can increase the debug level of the acquirer using the **debug acquirer trace EXEC** command. The logs are written to `local1/errorlog/acquirer-errorlog.current`.
- To verify that an expected object has been pre-positioned on the Content Engine, use the **show distribution object-status EXEC** command, as shown in the following example:

```
CE# show distribution object-status
http://172.18.81.168/Videos/SM-final%20Innebandy%202003.wmv

===== Website Information =====
Name:          RTPServer5
Origin Server FQDN:      172.18.81.168
Request Routed FQDN:    N/A
Content UNS Reference #: 1
===== Channels Information =====
*** Channel 1903 (name = A_Multicast) ***

Object Replication
-----
Replication:          Done
File State:           Ready for distribution
Multicast for Channel: Not Enabled
Replication Lock:     Received by Unicast-Receiver/Acquirer
Reference Count:      1
Total Size:           2756437
Transferred Size:     2756437
MD5 of MD5:          tjs#DxqE5oUc024Z8XtFDw..
Source Url:           http://172.18.81.168/Videos/SM-final%20Innebandy%202003.wmv
Source Last Modified Time: Wed Jan  7 19:03:48 2004

Object Properties
-----
Redirect To Origin:    Yes
Requires Authentication: No
Alternative URL:
Serve Start Time:     N/A
Serve End Time:       N/A
Play servers:         HTTP HTTPS WMT
Content Metadata:      None
Content uns_id:       NgcJTCU#JaY4ZGIPbsrONw..
Content gen-id:       1768:1136512329:2

===== CDNFS Information =====
Internal File Name:
/disk00-04/d/http-172.18.81.168-k5bsm1o+y14jgiqsvwaohq/19/19f6d5cec7266c33f419709dc28c
8d9b.0.data.wmv
Actual File Size:      2756437 bytes
MD5 of MD5 (Re-calculated): tjs#DxqE5oUc024Z8XtFDw..
Content metadata:      None
Metadata match with:   Channel 1903
Number of Source-urls: 1

Source-url to CDN-object mapping:
Source-url:           http://172.18.81.168/Videos/SM-final%20Innebandy%202003.wmv
Used by CDN object:   ---- Yes ----
Internal File Name:
/disk00-04/d/http-172.18.81.168-k5bsm1o+y14jgiqsvwaohq/19/19f6d5cec7266c33f419709dc28c
8d9b.0.data.wmv
Actual File Size:      2756437 bytes
```

```

===== CDNFS lookup output =====
CDNFS File Attributes:
  Status           3   (Ready)
  File Size        2756437 Bytes
  Start Time       null
  End Time         null
  Allowed Playback via HTTP WMT HTTPS
  Last-modified Time Wed Jan 7 19:03:48 2004
  cache-control    max-age=864000
  cdn_uns_id       NgcJTCU#JaY4ZGIPbsrONw..
  content-type     video/x-ms-wmv
  etag             "042e6fa50d5c31:b39"
  file_duration    65
  last-modified    Wed, 07 Jan 2004 19:03:48 GMT
  server           Microsoft-IIS/6.0
  x-powered-by     ASP.NET
Internal path to data file:
/disk00-04/d/http-172.18.81.168-k5bsml0+y14jgigsvwaohq/19/19f6d5cec7266c33f419709dc28c
8d9b.0.data.wmv

```

By comparing fields, such as Total Size, Transferred Size, and Source URL in the Object Replication output and Actual File Size and Source URL in the Source-URL-to-CDN-Object Mapping output, you can determine whether or not the object that is stored is the same as the object that was requested.

- To view the file directory structure on the Content Engine and verify the physical file on the disk, you can use the **cdnfs browse** EXEC command, as shown in the following example:

```

CE# cdnfs browse

----- CDNFS interactive browsing -----
dir, ls:   list directory contents
cd, chdir: change current working directory
info:      display attributes of a file
more:      page through a file
cat:       display a file
exit,quit: quit CDNFS browse shell

/>ls

srv2-static.cisco.com/
http-172.18.81.151-glfc4h-b9gywnf5rlnfweg/
http-172.18.81.163-og5o2lu178nrhlwlmctgtiq/
172.18.81.163/
file--xrnfwxifgu62jtiwtyixvg/
172.18.81.168/
http-172.18.81.168-k5bsml0+y14jgigsvwaohq/
172.18.81.155/

/>cd 172.18.81.168/
/172.18.81.168/>ls
376 Bytes      manifest-Channel_1903.xml-1EYmrfnjt2o5GbUNwLCAPA
                Videos/
/172.18.81.168/>cd Videos
/172.18.81.168/Videos/>ls
2756437 Bytes  SM-final Innebandy 2003.wmv <=====Physical file on disk
/172.18.81.168/Videos/>quit
CE#

```

Retry and Refresh Mechanisms

ACNS 5.x software supports various retry and refresh mechanisms in the manifest file. Following are some things to note.

Using the <schedule> and <repeat> Subelements

ACNS 5.5 software supports two subelements in the manifest file: <schedule> and <repeat>. These tags allow you to specify a time to begin a recrawl or refetch and are more powerful than the *ttr* schema, which only allows you to specify the interval between refetch or recrawl occurrences. The <schedule> element can be a subtag of <options>, <item-group>, <item> or <crawler>. If both the <schedule> tag and the *ttr* attribute are specified in the manifest file, the <schedule> designation takes precedence over *ttr*. (See the “item” section on page A-44 for a detailed description of these subelements.)

Using the *failRetryInterval* Attribute

When the acquirer tries to acquire a content item specified in the <item> or <crawler> tags, and the acquisition fails (for example, because of some intermittent error or for some other reason), the acquisition task (corresponding to the <item> or <crawler> tag) is retried after the interval specified in the *failRetryInterval* attribute.

If the *failRetryInterval* attribute is not specified, the default interval for retrying the task is 5 minutes. (For a crawl job, the task is considered to be a failure only when crawling fails to occur at all. If only some of the crawled pages fail to be acquired, the task is still considered a success.) The following rules apply to the retry mechanism:

- Single item tasks specified in the <item> tag are retried for all errors except the EXCEED_DISK_QUOTA error.
- Crawl tasks are not considered a failure if the error status is between 300 and 500. The EXCEED_DISK_QUOTA error does not cause a retry either.
- When you change the disk quota using the Content Distribution Manager GUI, the acquirer is notified automatically and retries all status error nodes containing the EXCEED_DISK_QUOTA error.

Using the *ttr* Attribute or the Fetch Manifest Now Button

When an item is acquired successfully, and if a positive value is specified in the *ttr* attribute, the acquirer rechecks the content for freshness at the interval specified by the *ttr* attribute. A *ttr* value of 0 (zero) means that the acquirer will not recheck the item unless the manifest file is updated. A negative *ttr* value means that the acquirer will never recheck the item. The following rules apply to the refresh mechanism:

- When *ttr* > 0: Recheck every *ttr* minutes. The acquirer also rechecks the content when the manifest file is reparsed and the content specification in the manifest file has changed, or when you click the **Fetch Manifest Now** button in the Content Distribution Manager GUI. (See the next section, “Updating Channel Content.”)
- When *ttr* = 0: Only acquire once. The acquirer only rechecks when the manifest file is reparsed and the content specification in the manifest file has changed, or when you click the **Fetch Manifest Now** button in the Content Distribution Manager GUI.
- When *ttr* < 0: Only acquire once. The acquirer will *not* recheck even if manifest file is reparsed, or when you click the **Fetch Manifest Now** button in the Content Distribution Manager GUI.

Reparsing the Local Manifest File From Content Distribution Manager GUI

When you add content to a channel using the Content Distribution Manager GUI to define the content, a manifest file is automatically generated and stored locally in the Content Distribution Manager. When you make any changes such as adding, removing, or modifying content definitions in the GUI and click the **Submit** or **Update** button, the local manifest file is automatically reparsed, changes are detected, and the corresponding items are acquired or removed. Note that when you click the **Submit** or **Update** button, it does not trigger a recheck of all the content in the channel.

Updating Channel Content

At any point after you have replicated content to the Content Engines that are associated with your channel, you can update that content using the fetch manifest feature. For example, if you modify your manifest file to point to new content or remove references to content that you want to make obsolete, you must fetch the manifest file to begin replication of any new channel content, and to sever connections to content that you want to make obsolete.



Note

Content that is removed from the manifest file is made unavailable as soon as that updated manifest file is fetched. Obsolete content is not immediately deleted from the channel cache but is eventually removed to make room for new channel content.

To fetch a new or updated manifest file, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**.
- Step 2** Click the **Edit** icon next to the name of the channel to open your channel for editing.
- Step 3** In the Contents pane choose **Channel Content**.
- Step 4** Verify that the Manifest URL field points to the correct manifest file for the channel.
- Step 5** Click the **Fetch Manifest Now** button. You are prompted to confirm your decision.
When you click this button, the software checks to see if the manifest file has been updated, and the updated manifest file is downloaded and reparsed. Also, regardless whether the manifest file has been updated, all content in the channel is rechecked and the updated content is downloaded.
- Step 6** Click **OK** to execute your request.

To force the replication of channel content and refresh the information, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**.
- Step 2** Click the **Edit** icon next to the name of the channel that you want to modify to open your channel for editing.
- Step 3** In the Contents pane, click **Replication Status**. The Replication Status for Channel window appears.
- Step 4** Click the **Force Replication information refresh** icon in the taskbar. You are prompted to confirm your decision.
- Step 5** Click **OK**. You are notified that your request has been sent and prompted to check back in a few minutes.

Step 6 Click **OK**. After a time, the Replication Status for Channel window refreshes.

For a description of the channel replication status data, see [Chapter 11, “Viewing Content Replication Status.”](#)

Bandwidth Control

The bandwidth control feature allows you to specify how much bandwidth in the network is consumed by data replication from the root Content Engine to the edge Content Engines. Bandwidth controls allow you to specify the amount of bandwidth to be used at various times during the day and also to set up a weekly schedule that repeats week after week. This section describes the bandwidth controls as they relate to content acquisition and distribution. Bandwidth controls are available for the acquisition process, replication process, and multicast sender.

Configuring Acquisition and Distribution Default Bandwidth Settings

Default bandwidth settings can be configured for acquisition and distribution of content. Default bandwidth is the amount of bandwidth allocated for content acquisition and distribution when there is no scheduled bandwidth.

If a Content Engine is assigned to a device group and no default bandwidth has been set for the device, the device group default bandwidth settings are applied. If the Content Engine is part of multiple device groups, the most recently updated default bandwidth settings are applied.

However, if default bandwidth is specified for a device, it will override the settings at the device group level. This occurs when a Content Engine is a member of a device group.

To configure the default bandwidth settings for acquisition and distribution for the Content Engine, follow these steps:

-
- Step 1** In the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon next to the desired Content Engine. The Device Home window for the Content Engine appears.
 - Step 3** In the Contents pane, choose **Prepositioning > Default Bandwidth**. The Acquisition and Distribution Default Bandwidth for Content Engine window appears. (See [Figure 6-11](#).)

Figure 6-11 Acquisition and Distribution Default Bandwidth Window

- Step 4** In the Acquisition-in Bandwidth field, enter the bandwidth value in kbps for incoming content acquisition traffic from origin servers. The default is 1024 kbps.
- Step 5** In the Distribution-in Bandwidth field, enter the bandwidth value in kbps for incoming unicast content distribution traffic from Content Engines. The default is 56 kbps.
- Step 6** In the Minimum Nonstreaming Acquisition-in Bandwidth field, enter the bandwidth value in kbps for incoming nonstreaming content acquisition traffic from the origin server. The default is 50 kbps.
- Minimum nonstreaming bandwidth is the bandwidth allocated for acquiring content using nonstreaming protocols such as HTTP, HTTPS, and FTP. This setting is useful when streaming protocols, such as RTSP, take up all the available bandwidth in certain cases, yet bandwidth needs to be reserved for acquiring manifest files and other content items from origin servers using nonstreaming protocols. We recommend that you retain the minimum nonstreaming bandwidth default setting, and adjust it only when it might be necessary to increase the value, such as when certain large files fail to be acquired because of many stream acquisitions taking place simultaneously.
- Step 7** In the Distribution-out Bandwidth field, enter a bandwidth value in kbps for outgoing unicast content distribution traffic to various Content Engines in the ACNS network. The default is 128 kbps.



Note The Multicast-out Bandwidth field is read-only. This field displays the multicast-out bandwidth setting for sender Content Engines. This value is configured in the Default Multicast-out Bandwidth field in the Creating New Multicast Cloud window. (See the [“Configuring Multicast Cloud Properties”](#) section on page 5-34.) For receiver Content Engines, this field is blank.

- Step 8** To revert to the previously configured window settings, click **Reset**. The Reset button is visible only when you apply default or group settings to change the current device settings, but you have not yet clicked **Submit**.
- Step 9** To save your settings, click **Submit**.

Displaying a Graphical Representation of the Acquisition and Distribution Bandwidth Settings

You can view a graphical representation of the bandwidth settings configured on a Content Engine for the acquisition and distribution of files. The vertical axis of the graph represents the amount of bandwidth in kbps and the horizontal axis represents the days of the week. The scale shown on the vertical axis is determined dynamically based on the bandwidth rate for a particular type of bandwidth and is incremented appropriately. The scale shown on the horizontal axis for each day is incremented for each hour. Each type of bandwidth is represented by a unique color. A legend at the bottom of the graph maps the colors to the corresponding bandwidths.

To view the graph that displays the acquisition and distribution bandwidth, follow these steps:

- Step 1** In the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to view the acquisition and distribution bandwidth. The Device Home window appears.
- Step 3** In the Contents pane, choose **Prepositioning > Default Bandwidth**. The Acquisition and Distribution Default Bandwidth for Content Engine window appears. (See [Figure 6-11](#).)
- Step 4** Click the **Display Graph** icon in the taskbar. A new Acquisition and Distribution bandwidth for Content Engine popup window appears, displaying the bandwidth graph.
- You can choose the view that you wish to apply to the bandwidth graph by clicking the different viewing options. (See [Table 6-6](#) for a description of the viewing options and their descriptions.)
- Step 5** Click **Close** once you have finished viewing the settings. Alternatively, you can click **Refresh** to view the most currently applied bandwidth settings.

Table 6-6 Viewing Options in Acquisition and Distribution Bandwidth Graph

Item	Description
View specific servers	Displays the bandwidth settings for the corresponding bandwidth type selected.
Distribution In	Displays the bandwidth settings for incoming content distribution traffic.
Distribution Out	Displays the bandwidth settings for outgoing content distribution traffic.
Acquisition In	Displays the bandwidth settings for incoming content acquisition traffic.
All Servers	Displays a consolidated view of all configured bandwidth types. This is the default view combined with the Full Week view.

Table 6-6 Viewing Options in Acquisition and Distribution Bandwidth Graph (continued)

Item	Description
View mode	Displays detailed and composite bandwidth settings.
Show Detailed Bandwidth	Toggles with the Show Effective Bandwidth option. Displays the detailed bandwidth settings for the device and its associated device groups. The bandwidth settings of the device and device groups are shown in different colors for easy identification.
Show Effective Bandwidth	Toggles with the Show Detailed Bandwidth option. Displays the consolidated or composite bandwidth settings for the device and its associated device groups.
Show Aggregate View	Toggles with the Show Non-Aggregate View option. Displays the bandwidth settings configured for the corresponding device groups.
Show Non-Aggregate View	Toggles with the Show Aggregate View option. Hides the bandwidth settings configured for the corresponding device groups.
View by day	Displays the bandwidth settings for a particular day or all days of the week.
Sun, Mon, Tues, Wed, Thurs, Fri, Sat	Displays the bandwidth settings for the corresponding day of the week.
Full Week	Displays the bandwidth settings for the entire week. This is the default view combined with the All Servers view.

Configuring Acquisition and Distribution Bandwidth Settings for Scheduled Times

In addition to being able to set the default bandwidth limits for incoming and outgoing traffic on the Content Engine, you can use the Content Distribution Manager GUI to configure different limits for different time segments that form a week-long cycle. For example, you can configure the acquisition-in limit at a maximum of 100 kbps from 8:00 a.m. to 8:00 p.m., Monday through Friday, and extend that limit to as high as 10 Mbps during the nights from 8:00 p.m. to 8:00 a.m., Monday through Friday and all day Saturday and Sunday. These settings override the default value for the time and period that you specify.


Note

For a schedule from 8:00 p.m. to 8:00 a.m., the administrator must configure two schedules in order to span the two days: one from 8:00 p.m. to 11:59 p.m. (2000 to 2359) and another from 12:00 a.m. to 8:00 a.m. (0000 to 0800).


Note

Distribution bandwidth settings apply only to unicast distribution.

**Note**

To disable acquisition and distribution activity during unscheduled times of the day, you must enter zeros (0) for the default bandwidth settings in the Acquisition and Distribution Default Bandwidth window. (See [Figure 6-11](#).) If you do not set the default bandwidth fields to 0, acquisition and distribution traffic continues during the unscheduled times of the day. If you leave the default bandwidth fields blank, acquisition and distribution traffic may use the maximum bandwidth allowance for the unscheduled periods.

To configure distribution bandwidth settings for specific days and times, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to view. The Device Home window appears.
 - Step 3** In the Contents pane, choose **Prepositioning > Scheduled Bandwidth**. The Acquisition and Distribution Bandwidth Schedule for Content Engine window appears.
 - Step 4** The Aggregate Settings **Yes** radio button is selected by default. This means that the acquisition and distribution bandwidth schedule is displayed for the Content Engine as well as for any associated device groups. The settings for the acquisition and distribution bandwidth schedule for device groups cannot be modified or deleted; they are read-only. If you click the Aggregate Settings **No** radio button, you can view and modify the settings for the acquisition and distribution bandwidth schedule for the Content Engine only. The the acquisition and distribution bandwidth schedule settings for any associated device groups are not even displayed.
 - Step 5** Click the **Create New Bandwidth Setting** icon in the taskbar. The Creating New Acquisition and Distribution Bandwidth Schedule for Content Engine window appears. (See [Figure 6-12](#).)

Figure 6-12 Creating New Acquisition and Distribution Bandwidth Schedule Window

The screenshot shows the Cisco Application and Content Networking System web interface. The main title is "Cisco Application and Content Networking System - Microsoft Internet Explorer". The interface includes a navigation menu on the left with options like "Contents", "Device Home", "Device Activation", "Assignments", "Prepositioning", "Default Bandwidth", "Scheduled Bandwidth", "Distribution", "Location Leader and Forwarder", "Multicast Distribution License", "Acquirer WCCP Proxy Authentication", "Replication Status", "Request Routing", "Request Processing", "Applications", "General Settings", and "Device Monitoring". The main content area is titled "Creating New Acquisition and Distribution Bandwidth Schedule for Content Engine, CONTENTENGINE". It contains two main sections: "Bandwidth Configuration" and "Bandwidth Time Settings".

Bandwidth Configuration

- Bandwidth Type*: Please make a choice (dropdown menu)
- Bandwidth Rate*: [text input] Kbps

Bandwidth Time Settings

The local time on the CDM is Tue Aug 2 22:12:28 UTC 2005.
Note: All times use a 24-hour clock (hh:mm) and are interpreted as local time on the CE(s).

- Time Selection*: Start Time [text input] End Time [text input]
- Day Selection*: ☐ Full Week ☐ Sun ☐ Mon ☐ Tue ☐ Wed ☐ Thu ☐ Fri ☐ Sat

Note: * - Required Field

At the bottom right, there are "Submit" and "Cancel" buttons.

Step 6 Choose a bandwidth type from the drop-down list. (See [Table 6-7](#) for a description of each field in this window. All fields are required.)

Step 7 Enter the bandwidth rate, start time, end time, and day of the week in the appropriate fields.

Table 6-7 Configuring Acquisition and Distribution Bandwidth Settings

Field	Description
Bandwidth Type	Distribution-in—For incoming unicast content distribution traffic from Content Engines. Distribution-out—For outgoing unicast content distribution traffic to Content Engines. Acquisition-in—For incoming content acquisition traffic from origin servers. Multicast-out—For outgoing multicast content distribution traffic to Content Engines.
Bandwidth Rate	Maximum amount of bandwidth that you want to allow (in kbps).
Start Time	Time of day for the bandwidth setting to begin, using a 24-hour clock in local time (hh:mm).

Table 6-7 *Configuring Acquisition and Distribution Bandwidth Settings (continued)*

Field	Description
End Time	Time of day for the bandwidth setting to end (hh:mm).
Day Selection	Days on which bandwidth settings apply. - Full Week—Specifies that the allowable bandwidth settings are applied for an entire week. - Sun, Mon, Tue, Wed, Thu, Fri, and Sat—Specifies individual days of the week on which the allowable bandwidth settings take effect.

Step 8 Click **Submit**.

Scheduling Bandwidth for Streaming Acquisition

When streaming acquisition is configured, the streaming content consumes a large amount of bandwidth over the playtime duration of the file. You must ensure that sufficient bandwidth is available for sufficiently large blocks of time. For multi-bit-rate streams (MBR), the required bandwidth is the sum total of all bit rates encoded in the stream.

Error Codes

This section describes the error codes for content replication. The following error codes are shown in the Content Distribution Manager GUI Replication Status window:

ACNS Unified Name Space Errors

- 1401: Bad magic number in unified name space (UNS) meta file
- 1402: Unknown version in UNS meta file
- 1403: Bad checksum on UNS meta file
- 1404: Internal error - URL mismatch between caller specified URL and URL in meta file
- 1405: Invalid URL syntax
- 1406: Attempt to create URL already in UNS
- 1407: Insufficient space to store requested object
- 1408: Internal logic error in UNS server code
- 1409: Requested object not found in UNS
- 1410: Requested UNS operation not implemented
- 1411: Failure in underlying RPC transport
- 1412: Destination URL already exists
- 1413: Channel does not exist
- 1414: Writes failed to all metadata files

1415: Object is not servable
1416: Object is out of presentation time
1417: Object playback not allowed by this playback server
1418: Live object, but attributes invalid
1419: Alternate media attributes invalid
1420: Alternate media is not servable
1421: Channel would be over disk quota
1422: Cannot change channel ID
1423: Object already in specified channel
1424: Object not in specified channel
1425: Metadata operations disabled (no file systems)
1426: Channel operations disabled (no file systems)
1427: Nonobject metadata services not initialized
1428: Out of handles for nonobject metadata files
1429: Internal error loading metadata file
1430: (No error text available)
1431: Too many CDN files (cannot add to URL for file system map [UFM])
1432: Specified legacy ECDNFS file not found
1433: Bad MD5 checksum passed by caller
1434: Bad tags present in supplied attributes (OBSOLETE)
1435: Cannot resize content file migrated from ECDN
1436: UNS symlink references nonexistent URL
1437: URL is not a UNS symlink
1438: Too many levels of UNS symlinks
1439: UNS entry metafile truncated
1440: Output to be returned is too big in size
1441: URL request was made on a nondefault service port
1442: Specified legacy file not on any UNS filesystem
1443: Specified legacy file is not in the 'cache' directory
1444: Specified legacy file is referenced by a UNS entry
1445: Specified legacy file is not a .data file
1446: File open failed during ASX/SMIL rewrite operation
1447: Parse error during ASX/SMIL rewrite operation
1448: Initialization error during ASX/SMIL rewrite operation
1449: Actual Rewriting failed during ASX/SMIL rewrite operation
1450: No more file system slots
1451: Specified file system is already in use
1452: Specified file system not known to UNS

- 1453: Specified file system bytes in use exceeds target
- 1454: Cannot unuse the file system containing the symlink tree
- 1455: Cannot add file system because there is no local disk-based CDNFS storage

Acquirer Error Codes

The following error codes are used for content acquisition errors:

- 700: Acquirer internal error
- 701: Manifest parser error
- 702: Manifest parse warning
- 703: Exceeded disk quota
- 704: No space in UNS
- 705: It is a folder
- 706: ACCEPT failed
- 707: Connection refused
- 708: Listen failed
- 709: Mismatched in crawling
- 710: No-cache instructed from server
- 711: Disabled by the user
- 712: Downloaded size mismatched with content length
- 713: Invalid response received from server
- 714: Database access error
- 715: Not to acquire URL with ?
- 716: Invalid redirect foo to foo/
- 717: Illegal folder for file import
- 718: Unable to connect to proxy server
- 719: URL is too long
- 720: HTTP metadata is too long
- 721: Connection closed by peer
- 722: Invalid content length
- 730: MmsOverHttp is not supported
- 905: Socket timeout
- 906: No host
- 907: Zero bandwidth
- 908: File download aborted
- 909: Content expired before fetch
- 912: Request timed out
- 1000: UNS error

Where to Go Next

In this chapter you accomplished the following tasks:

- Defined content for pre-positioning either directly through the Content Distribution Manager GUI or using a manifest file as described in [Appendix A, “Creating Manifest Files”](#)
- Configured the basic manifest file settings in the Content Distribution Manager GUI so that the root Content Engine could fetch the manifest file and update the channel content at specified intervals
- Configured bandwidth settings for the root Content Engine to optimize bandwidth usage in your network
- Configured a proxy to fetch the manifest file, if needed
- Configured authentication as needed to fetch the manifest file, acquire content, or both

Now you need to consider how the content is going to be played back to the end user. If the content is to be scheduled or rebroadcast, proceed to [Chapter 7, “Creating and Managing Programs.”](#)



Creating and Managing Programs

This chapter explains how to create and manage live, rebroadcast, TV-out, and export programs (for export to set top boxes). It contains the following sections:

- [About Programs, page 7-1](#)
- [Viewing Programs, page 7-8](#)
- [Defining a Program, page 7-11](#)
- [Creating a Live Program, page 7-13](#)
- [Creating a Scheduled Rebroadcast Program, page 7-25](#)
- [Creating a TV-Out Program, page 7-31](#)
- [Creating an Export Program, page 7-40](#)
- [Copying a Program, page 7-45](#)
- [Managing Multicast Addresses for Programs, page 7-46](#)

About Programs

A *program* in ACNS software is defined as a scheduled event in which the content is presented to the end user. Attributes of the program include the start and end time of the program (whether the content is pre-positioned or live) the set of Content Engines assigned to the program, and the streaming server to be used for content delivery.

Program types determine the hardware or software component involved in delivering content to the user. Different program types are:

- Cisco Streaming Engine live
- Cisco Streaming Engine rebroadcast
- Windows Media live
- Windows Media rebroadcast
- TV-out
- Export (meant for a third-party device such as a set top box)

The Content Distribution Manager GUI allows you to create a program, schedule it, and deliver it over your multicast-enabled ACNS network. Programs can consist of prerecorded video or live events. In the case of live events, ACNS software interoperates with third-party encoders and IP/TV Broadcast Servers to transmit the live stream across the ACNS network.

Retention of m4e File Extensions for Envivio-Based Programs

ACNS software supports streaming media that is generated by Envivio-based programs and IP/TV programs. For IP/TV-ACNS programs, the program description is generated as a Session Description Protocol (SDP) file by the IP/TV Program Manager. ACNS software uses this SDP file to distribute the program specifications to the appropriate Content Engines.

Envivio generates the program description as an .m4e file. In ACNS 5.3.3 software and later releases, m4e extensions in the reference URLs are retained. This new functionality allows you to use the Envivio TV plugin to render streams that are encoded by the Envivio encoder using specific codecs, such as H.264. For CLI-based programs, you can now publish broadcast_id.m4e files if the source is an .m4e file.

Content Engine Support for Publishing of Multicast SDP Files for Cisco Streaming Engine Live Programs

In previous releases of ACNS software, multicast SDP files were not published for Cisco Streaming Engine programs that were created through the Content Distribution Manager or the Content Engine CLI. In previous releases, you could publish the SDP files from any HTTP or FTP server.

ACNS 5.3.3 software and later releases support multicast reference URLs (Announce URLs) for programs that are created through the Content Distribution Manager or the CLI. The multicast reference URL, which is in the form of `http//Content Engine-IP address/ProgramID.sdp`, is resolved by Content Engines that are serving the live program.

Program Prerequisites

When creating a program in the Content Distribution Manager GUI, the administrator specifies the following information:

- Streaming server—How the program will be streamed or delivered to the end user
 - Cisco Streaming Engine—The event is played from a PC using the Apple QuickTime application or using the IP/TV Viewer application.
 - WMT server—The event is played using a Windows Media Technologies player.
 - TV-out device—The event is displayed over a television monitor that is directly connected to the Content Engine.
 - Exported—The program is exported over HTTP to a set top box.
- Media source
 - Live stream transmitted from a third-party encoder or an IP/TV Broadcast Server.
 - Pre-positioned content stored on Content Engines.
- Destination
 - A set of Content Engines or device groups.
 - A set of Content Engines or device groups assigned to a channel (live or export).
- Delivery method
 - Unicast or multicast live events.
 - Multicast scheduled rebroadcasts.
 - TV-out or export playlists.
- Schedule
 - Start time and duration.

- Looping forever (enabled or disabled).
- Multiple playback times.
- Repeating playback times.
- Publishing URL—The URL used to view the program

Playing Hinted and Nonhinted IP/TV Programs over an ACNS Network



Note

Hinted files contain hint tracks, which store packetization information that tells the streaming server how to package the media data. The streaming server uses the packetization information in the hint tracks to stream the media data to the network.

The Cisco Streaming Engine supports only hinted files (MOV and hinted MP4 files) for streaming. If you are creating a file-based IP/TV program for streaming over an ACNS network, make sure that you use only hinted files such as those with .mov or .mp4 extensions.

However, you can pre-position IP/TV on-demand programs based on nonhinted files such as .mpg files on Content Engines when IP/TV is integrated with an ACNS network by using the TV-out feature. (See the [“Creating a TV-Out Program” section on page 7-31.](#))

Configuring Multicast IP Addresses for IP/TV Programs

If an IP/TV program that you want to deliver over an ACNS network uses live multicast mode, you must use the same multicast IP address for the audio, video, and SlideCast streams.

This requirement does not apply if the Content Engine used for live splitting is running ACNS 5.1.5 or later software. However, this requirement still applies if the Content Engine is running ACNS 5.1.1 software, even if you are running IP/TV 5.1.5 or later software on IP/TV Program Manager.

About Live Events

Live events are streamed from third-party encoders (such as the Windows Media encoder Version 9 and the QuickTime encoder) or from streaming servers (such as Windows Media Server). The live stream is acquired by the root Content Engine and transmitted to edge Content Engines using unicast. From the edge Content Engines, the live stream is transmitted to end users using either multicast or multicast and unicast live splitting. The live program is available to viewers only during its scheduled times.

Before setting up a live event in the Content Distribution Manager GUI, the administrator must complete the following tasks:

1. Set up the live event on the encoder or streaming server. This task is done externally to the Content Distribution Manager GUI.
2. In the Content Distribution Manager GUI, accept the streaming server’s license agreement (WMT or RTSP) on the Content Engines designated to acquire and distribute the stream. (See the [“Licensing and Enabling Streaming Servers” section on page 9-6.](#))
3. Create a live channel (or designate an existing channel as a live channel), assign Content Engines to the live channel, and designate a root Content Engine to acquire the stream. (See the [“Creating and Modifying Channels” section on page 5-10.](#))

After you have completed the preceding tasks, you are ready to define the program and configure the program settings in the Content Distribution Manager GUI. See the [“Defining a Program” section on page 7-11](#). The Content Distribution Manager allows you to define two types of live events based on the streaming server being used:

- Windows Media live
- Cisco Streaming Engine live

See the [“Creating a Live Program” section on page 7-13](#).

Configuration Requirements for Managed Live Events

If you have configured channels for live programs in your ACNS network, make sure that there are no external proxy servers physically located between your receiver Content Engines and your root Content Engine that require proxy authentication. Also, make sure that proxy authentication is not enabled on any receiver Content Engines that might be in the logical, hierarchical path between the root Content Engine and the receiver Content Engine that is going to serve the live stream to the requesting clients. If a live stream encounters any device that requires proxy authentication, the stream will be dropped before it reaches its destination.

If your network is set up with intermediary devices that require proxy authentication, you can work around the problem by configuring rules to bypass authentication on these devices.

For example, to enable the formation of the unicast splitting tree and, in turn, enable live broadcasting from all receiver Content Engines, you can specify the following rule on all of the parent Content Engines in the channel:

```
ContentEngine(config)# rule pattern-list 1 downstream-CE-ipaddress
ContentEngine(config)# rule no-auth pattern-list 1
```

In ACNS 5.4 and releases prior to ACNS 5.5, you must make sure that MMST is not disabled on the root Content Engine or on any receiver Content Engines that might be in the logical, hierarchical path between the root Content Engine and the receiver Content Engine that is going to serve the Windows Media managed live program stream to the requesting clients. If a Windows Media live stream encounters any device that has MMST disabled, the stream fails on that Content Engine and all its down-level devices.

In ACNS 5.5 software, the MMS protocol is no longer supported. Therefore, when you configure a Windows Media managed live program using ACNS 5.5, all Content Engines in the channel must be using ACNS 5.5. You cannot deliver a Windows Media managed live program in a mixed-version environment.



Note

Furthermore, all Content Engines in the live channel must have RTSPT enabled because Content Engines must use the RTSPT protocol to communicate with each other.

Live Stream Interruptions

During a Windows Media live broadcast, any interruption of the live stream that lasts five minutes or longer causes the multicast broadcast to cease for the duration of the currently scheduled period. If the live stream is interrupted for less than five minutes, the broadcast resumes.

Live stream interruptions can be caused by unexpected encoder failures or by an operational restart. If you have early trials scheduled prior to the live event, we recommend that you configure multiple schedules, one for each trial. We also recommend that you start the encoder prior to the scheduled time.

If the live stream stops for more than five minutes and resumes later while the program is still scheduled, you can modify the schedule or any other attribute of the program (such as the description) to trigger a restart of the multicast broadcast. Restarting might take up to five minutes under these circumstances.

This note does not apply to unicast delivery of a Windows Media live event or to Cisco Streaming Engine live programs.

About Root Content Engine Failover and Fallback Support for Cisco Streaming Engine Live Programs

Cisco IP/TV Program Manager Release 5.x and ACNS software Release 5.5 support failover and fallback mechanisms for the root Content Engine to ensure seamless streaming of a Cisco Streaming Engine live program in an IP/TV-integrated ACNS network. In the event of a primary root Content Engine failure, the failover mechanism supports the election of a backup root Content Engine. Fallback occurs when the primary root Content Engine regains its primary position. (See also the [“About the Root Content Engine Failover and Fallback Grace Period Setting”](#) section on page 5-14.)

Root Content Engine failover requires that both the primary and the backup Content Engine be located in the root location and be multicast-enabled to receive multicast streams from the Cisco IP/TV Broadcast Server. When the primary root Content Engine fails, the backup root Content Engine assumes control of the live program streaming, and the live program continues without interruption. The Content Engine that is elected as the location leader in the child location continues to receive the live stream from the backup root Content Engine. This process is transparent to the person viewing the program on IP/TV Viewer or on Apple QuickTime Player.



Note

If the root Content Engine receives the program as a unicast stream from the IP/TV Broadcast Server, the failover mechanism is not supported. If the primary root Content Engine fails while a program is playing, the person viewing the program must re-request the program.

Root Content Engines can fail for any of the following reasons:

- The Content Engine is switched off.
- The Content Engine is rebooted.
- The Cisco Streaming Engine is accidentally disabled in the Content Engine.
- A process running on the Content Engine fails.
- The Content Engine is disconnected from the rest of the network.



Note

For failover to work correctly, make sure that the Cisco IP/TV Broadcast Server, the root Content Engine, and at least one additional backup Content Engine in the root location are in the same multicast cloud.

Fallback occurs when the primary root Content Engine regains its primary position. When the primary root Content Engine resumes functioning after a failure, it receives the live stream from the active secondary root Content Engine and does not try to regain its primary position. The fallback mechanism occurs in the following circumstances:

- No other Content Engine in the location has been elected as a location leader.
- A program running on the Content Engine is restarted.
- The current program schedule ends and a new program schedule begins.

Configuring Root Content Engine Failover for Incoming Multicast Streams

To configure root Content Engine failover using the Content Distribution Manager GUI, choose the multicast mode of delivery to the root Content Engine by checking the **Enable Multicast Delivery to Client** check box when you configure the live stream settings for a new live program. Alternatively, you can check the **Enable Multicast Delivery to Client** check box when you edit an existing Cisco Streaming Engine live program.

In addition, when you create or edit a live streaming program, you must enter **0** (zero) for the root Content Engine port. (For more information about configuring live stream settings, see [“Configuring Live Stream Settings for a Cisco Streaming Engine Live Program”](#) section on page 7-20.)



Note

If the program is created using the IP/TV Program Manager, and Multicast is chosen as the stream delivery mode, the IP/TV program Manager automatically assigns a zero to the port setting. (For more information, see the *Cisco IP/TV Program Manager Guide, Release 5.x* publication.)

To disable failover to the root Content Engine, choose the unicast mode of stream delivery to the root Content Engine.

About Root Content Engine Failover and Fallback Support for a Windows Media Live Program

ACNS 5.5 software supports root Content Engine failover and fallback for a Windows Media live program. This feature is configured in the same manner as described for a Cisco Streaming Engine live program except that in a Windows Media setup, an encoder replaces the IP/TV Broadcast Server. (See the [“About Root Content Engine Failover and Fallback Support for Cisco Streaming Engine Live Programs”](#) section on page 7-5.)



Note

For failover to work correctly during a Windows Media live program, the root location must contain a primary root Content Engine and at least one additional backup Content Engine. The encoder is not required to be in the root location, and the encoder is not required to be in the multicast cloud.

About ACNS 5.5 Support for Windows Media Server-Side Playlists

The ACNS 5.5 software release supports server-side playlists for Windows Media live streams. A playlist is an ordered list of media items, such as media files, live streams, and so on, that the server can stream to a client. Cisco ACNS software Release 5.5 supports two types of server-side playlists for Windows Media files: broadcast server-side playlists and on-demand server-side playlists. The only protocol supported for ACNS 5.5 server-side playlists is RTSP.

For broadcast server-side playlists, ACNS 5.5 software supports live splitting of the broadcast stream so that multiple clients can view the program simultaneously, while the Content Engine maintains just a single connection to the Windows Media server fetching the data.

For on-demand server-side playlists, ACNS 5.5 software supports pass through delivery of the program, which can be viewed in the same manner as video on demand (VoD).

ACNS software supports the following features of Windows Media server-side playlist:

- Multiple files and Windows Media encoders
- Looping of playlists
- Managed live events (multicast and unicast)
- Stream splitting for unicast streams
- Encoder redundancy
- Custom scripts



Note For custom scripts, you must use the player as a plug-in.

However, ACNS software does not support the following protocols and features for Windows Media server-side playlists:

- MMS (live, on-demand playlist, and proxy caching)
- Text and URL events
- Multiple bit rate files

ACNS software supports only those server-side playlists that are created using the publishing point feature of Windows Media Server. From the Windows Media Server GUI, you can create a broadcast publishing point that contains all media items that you want to be part of the playlist. If you want to use the playlist for multicast, you must create an .nsc file that contains all media items that are included in the playlist. (For information about creating a broadcast publishing point and .nsc file, refer to the Microsoft Windows Media Services documentation.)

An .nsc file is a Windows Media metafile that contains multicast information, such as the multicast IP address, multicast port, time-to-live value, stream formats of media content, and so forth. You must make sure that all media items mentioned in the playlist are also included in the .nsc file. If you made any changes to the publishing point, you should re-create your .nsc file, as well.

You must enter an HTTP URL for the .nsc file in the Streaming Settings window of the Windows Media live program if you want that program to use the server-side playlist for multicast. (See the [“Configuring Live Stream Settings for a Windows Media Live Program”](#) section on page 7-18 and the [“About NSC Files for Multicasting Windows Media Server-Side Playlists”](#) section on page 7-23.)

About Scheduled Rebroadcasts

In a scheduled rebroadcast, pre-positioned content is scheduled to be streamed from edge Content Engines using multicast. To define a program for a rebroadcast event, the administrator must complete the following tasks:

1. Create a content acquisition channel and pre-position the content.
2. Assign media files to the program.
3. Specify the multicast address and port or define a program address pool.
4. Specify the days and times for the rebroadcast.



Note

For rebroadcast programs, media can only be selected from one channel. The Content Engines and device groups assigned to the channel are selected automatically when you choose the media files for the program.

The Content Distribution Manager allows you to define a program as a Windows Media rebroadcast or as a Cisco Streaming Engine rebroadcast. (See the [“Creating a Scheduled Rebroadcast Program”](#) section on page 7-25.)

About TV-Out Programs

A TV-out program is a scheduled program that is based on a playlist. A *playlist* is list of media files and their associated attributes. When you define a TV-out program, you set up a schedule called a *playtime*. A playlist can have between one and ten different playtimes.

To define a program for TV-out, the administrator must complete the following tasks:

1. Create a content acquisition channel and pre-position the content.
2. Select a content channel to be associated with the program.
3. Specify the days and times for the program to play.

(See the [“Creating a TV-Out Program”](#) section on page 7-31.)



Note

For TV-out programs, Content Engines are assigned directly to the program. Content can be acquired from one channel and distributed through a different channel. In other words, the program distribution channel does not have to be the same as the content acquisition channel or channels. To ensure a successful playback, however, the administrator must verify that the media files are pre-positioned on all Content Engines assigned to the program distribution channel.

(For information about Content Engine software driver support for the TV-out feature, see the *Cisco ACNS Software Upgrade and Maintenance Guide, Release 5.x*.)

About Export Programs

An export program is a program that is played across one or more set top boxes (STBs). The administrator selects the media files and image overlays, selects an export channel, and assigns devices to the export channel. The schedule for the playtimes is controlled by the set top box.

(See the [“Creating an Export Program”](#) section on page 7-40.)


Viewing Programs

The Programs window of the Content Distribution Manager GUI lists all of the programs defined in your ACNS network. Programs can be defined through the Content Distribution Manager GUI or through an API.

The Programs window allows you to view scheduled programs by day, week, month, or year. You can sort and filter programs by name, type, or schedule. You can also preview live programs while they are playing.

To view all the programs defined in your ACNS network, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Video > Programs**. The Programs window appears listing all the programs that have been defined through either the Content Distribution Manager GUI or the Program API. (See [Figure 7-1](#).)

Figure 7-1 Viewing Programs Defined in the Network

Step 2 Choose the **Day**, **Week**, **Month**, or **Year** tab to view the playback schedules. Scheduled playlists are listed by start time (initial start time plus any repeat intervals). Times begin with the current device time (current system time plus device time zone offset).

The **Unscheduled** tab displays all unscheduled programs defined in your ACNS network. The **All** tab displays all the programs defined in your ACNS network. The Programs window opens to this view by default.

Step 3 Sort columns by clicking the column heading. You can also combine filtering conditions. For example, you can filter only Windows Media live programs and then choose the **Week** tab to view the week of November 23 to November 29, 2003. [Table 7-1](#) describes the information that is displayed in this window.

Table 7-1 Programs Window Information

Item	Description
Tabs	
Day/Week/Month/Year	Lists programs based on their schedule. The current day, week, month, or year is displayed by default. You can navigate to the next or previous day, week, month, or year by clicking the back or forward arrows in the date field.
Unscheduled	Lists only programs with no schedule defined.
All	Lists all programs. This is the default view.

Table 7-1 Programs Window Information

Item	Description
Program Listing Table	
Program	Program name, which must be unique to the Content Distribution Manager.
Type	Program type. Program types are: • Cisco Streaming Engine live • Cisco Streaming Engine rebroadcast • Export • TV-out • Windows Media live • Windows Media rebroadcast
Schedule	Describes the schedule. Options are: • None (the program has no schedule) • Loop continuously • Number of playtimes (the number of times that the program is scheduled to be shown)
Start Time	Program start time in a scheduled view (Day, Week, Month, or Year tab). Lists up to 3 start times if repeat broadcasts are configured.
Duration	Duration of the program or the looping time in a scheduled view (Day, Week, Month, or Year tab).

Viewing and Modifying API Programs

Programs created through APIs are based on a program file. A *program file* contains the elements that define the schedule, content, and presentation parameters. It is a text file written in XML format, similar to the manifest file. The program file contains most of the program settings and resides on an external server. The Content Distribution Manager obtains the program file, parses it, and saves the program file to the database. The program is automatically updated at intervals by refetching the program file and reparsing it. ACNS 5.5 software supports RTSP in the program file.

In contrast, programs defined using the Content Distribution Manager GUI are not based on a program file; the settings entered in the Content Distribution Manager GUI are saved directly to the database.

Programs created using an API can be viewed in the Content Distribution Manager as read only, and modifications to API programs can be done through the API. You can also edit the API program using the Content Distribution Manager GUI; however, if you choose this option, the information about the API program file is deleted and the program can no longer be modified through the API. A third option is to copy the API program using the Content Distribution Manager copy program feature. The new copy will not contain the program file information and will be treated as a Content Distribution Manager GUI-generated program for the purposes of editing. (See the [“Copying a Program”](#) section on page 7-45.)

You can delete any program from the list (whether created through the Content Distribution Manager GUI or through an API) in the Programs window of the Content Distribution Manager GUI.

Previewing a Program

You can preview live programs by live split or by joining a multicast broadcast. Live programs can only be viewed during the scheduled playtime. You can preview a rebroadcast program by joining the multicast broadcast during the scheduled playtime.

To preview a live Cisco Streaming Engine or Windows Media program or scheduled rebroadcast, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Video > Programs**. The Programs window appears.
- Step 2** Click the **Play** icon next to the name of a program. A program preview window pops up, displaying the program information with links to view the program.



Note The **Play** icon only appears while the live program is playing. If a program is not currently playing, you cannot view it.

- Step 3** Click the URL reference link for the program. You have the option to choose a multicast or unicast URL reference, if such are defined for the program. A new window with the URL reference opens.

To successfully view the program, you must meet these conditions:

- You must be able to access the client network.
- You must have a WMT plug-in installed to view Windows Media live programs.
- You must have a QuickTime or IP/TV plug-in installed to view Cisco Streaming Engine live programs.

Defining a Program

When you define a new program, you must choose one of the following program types:

- Cisco Streaming Engine live
- Cisco Streaming Engine rebroadcast
- Export
- TV-out
- Windows Media live
- Windows Media rebroadcast

After you have completed the following initial configuration, proceed to the section corresponding to the specific program type you are defining.

To define a program, follow these steps:

- Step 1** In the Content Distribution Manager GUI, choose **Services > Video > Programs**. The Programs window appears.
- Step 2** Click the **Create a new Program** icon in the taskbar. The Adding New Program window appears. (See [Figure 7-2](#).)

Figure 7-2 Adding New Program Window

The screenshot shows the 'Adding New Program' window in the Cisco Application and Content Networking System GUI. The window is titled 'Adding New Program' and contains a 'Program Information' section. The fields are as follows:

- Name:** A required text field, currently empty.
- Type:** A dropdown menu with 'Please make a choice' selected.
- Auto Deletion:** An unchecked checkbox with the text '(automatically deletes the program about 24 hours after its end-time)'.
- Description:** A text area.
- Shuffle:** An unchecked checkbox with the text '(plays media in a random order)'.
- Graceful Exit:** An unchecked checkbox with the text '(avoids schedule exit in the middle of a media file)'.

A note at the bottom of the form states: '* - Required Field'. The window has 'Submit' and 'Cancel' buttons at the bottom right.

- Step 3** Enter a unique name for the program in the Name field.
- Step 4** Choose the type of program you are creating from the Type drop-down list.
- Step 5** If you wish the program to be automatically deleted from the Content Manager GUI 24 hours after its scheduled end time, check the **Auto Deletion** check box. This option only applies to live programs.
- Step 6** If you wish, enter a description of the program in the Description field.
- Step 7** For TV-out programs, check the **Shuffle** check box if you want media files to be played in random order. If left unchecked, media files are played in the order listed.
- Step 8** For TV-out programs, check the **Graceful Exit** check box if you want the program to continue playing to the end of the current media file when the program schedule ends. This avoids ending the program in the middle of a media file.
- Step 9** To save your program definition, click **Submit**.

You have defined the type of program that you want to configure. Proceed to the section for configuring that type of program:

- To configure Cisco Streaming Engine live and Windows Media live programs, see the next section, [“Creating a Live Program.”](#)
 - To configure Windows Media rebroadcast and Cisco Streaming Engine rebroadcast programs, see the [“Creating a Scheduled Rebroadcast Program”](#) section on page 7-25.
 - To configure TV-out programs, see the [“Creating a TV-Out Program”](#) section on page 7-31.
 - To configure set top box programs, see the [“Creating an Export Program”](#) section on page 7-40.
-

Creating a Live Program

A live event can be defined as a Windows Media live program or it can be defined as a Cisco Streaming Engine live program. The Cisco Streaming Engine is an open standards-based streaming server that sends hinted MP1, MP2, MP4, and streaming MOV files to clients across the Internet using the industry standard RTP and RTSP protocols.


To program a live event, define the program by following the steps in the previous section ([“Defining a Program”](#)) and choose **Cisco Streaming Engine Live** or **Windows Media Live** from the Type drop-down list. Now you must choose a live channel. Proceed to the next section, [“Choosing a Live Channel.”](#)

Choosing a Live Channel

To choose a live channel, follow these steps:

-
Step 1	From the Content Distribution Manager GUI, choose Services > Video > Programs .
Step 2	Click the Edit icon of the program for which you want to select a live channel. The Modifying Program window appears.
Step 3	In the Contents pane, choose Select Live Channel . The Selecting Live Channel for Program window appears. (See Figure 7-3 .)

Figure 7-3 Selecting Live Channel for Program Window



This window lists all the live channels in the system.

- Step 4** Click the radio button next to the name of the live channel that you want to associate with the program. Only one live channel can be associated per program.
- Step 5** To save this setting, click **Submit**.

Adding a New Live Channel

If you wish to add a new live channel to associate with your live programs, follow these steps:

- Step 1** From the Selecting Live Channel for Program window (see [Figure 7-3](#)), click the **Create a new Live Channel** icon in the taskbar. The Adding New Live Channel window appears. (See [Figure 7-4](#).)

Figure 7-4 Adding New Live Channel Window

The screenshot shows the Cisco Application and Content Networking System web interface. The top navigation bar includes 'Devices', 'Services', and 'System'. The left sidebar shows 'Contents' with sub-items: Definition, Select Live Channel, Live Channel Assignments, Live Streaming, and Schedule. The main content area is titled 'Adding New Live Channel' and contains a 'Live Channel Information' form. The form has three fields: 'Name' with the value 'liveChannel_3', 'Content Provider' with a dropdown set to 'None', and 'Web Site' with a dropdown set to 'Please make a choice'. A note at the bottom of the form states 'Note: * - Required Field'. At the bottom right of the form are 'Submit' and 'Cancel' buttons.

- Step 2** The Name field is automatically populated with a unique channel name. If you wish to change the name given by default, enter a unique name for the channel in this field.
- Step 3** Choose a predefined content provider from the Content Provider drop-down list or choose **None**, which is the default. (See the [“Creating and Modifying Content Providers”](#) section on page 5-6.)
- Step 4** Choose a predefined website from the Web Site drop-down list. (See the [“Creating and Modifying Websites”](#) section on page 5-8.)



Note For Windows Media live programs, the origin server for the live channel’s website should be the root Content Engine DNS name or IP address.

- Step 5** To save the live channel settings, click **Submit**.
- Step 6** To deselect a channel, click the **Clear** button.

After you have associated a live channel with the program, you need to assign Content Engines, device groups, or both to the associated live channel and choose a root Content Engine. Proceed to the next section, [“Assigning Content Engines and Device Groups.”](#)

Assigning Content Engines and Device Groups

The Content Engine and device group assignment windows are only available from the Contents pane *after* you have associated a live channel with a program.

To assign Content Engines to the live channel, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI choose **Services > Video > Programs**.
 - Step 2** Click the **Edit** icon next to the name of a live Windows Media or Cisco Streaming Engine program. If this program is associated with a live channel, the **Live Channels Assignments** option appears in the Contents pane.
 - Step 3** Choose **Live Channels Assignments > Assign Content Engine**. The Content Engine Assignments for Live Channel window appears.
 - Step 4** To make batch Content Engine assignments, click the **Assign all Content Engines** icon in the taskbar, or click the **Remove all Content Engines** icon.
 - Step 5** To make individual assignments, click the **Assign** icon (blue cross mark) next to the name of an individual Content Engine.
 - Step 6** Choose a root Content Engine from the Root CE drop-down list in the task bar.
 - Step 7** To save the assignments, click **Submit**.
-

To assign device groups to the live channel, follow these steps:

-
- Step 1** From Content Distribution Manager GUI, choose **Services > Video > Programs**.
 - Step 2** Click the **Edit** icon next to the name of a live Windows Media or Cisco Streaming Engine program. If this program is associated with a live channel, the **Live Channels Assignments** option appears in the Contents pane.
 - Step 3** Choose **Live Channels Assignments > Assign Device Groups**. The Device Group Assignments for Live Channel window appears.
 - Step 4** To make batch device group assignments, click the **Assign all Device Groups** icon in the taskbar, or click the **Remove all Device Groups**.
 - Step 5** To make individual assignments, click the **Assign** icon (blue cross mark) next to the name of an individual device group.



Note For information about creating device groups, see the [“Creating a Device Group”](#) section on page 13-26.

- Step 6** Choose a root Content Engine from the Root CE drop-down list in the taskbar.
- Step 7** To save the assignments, click **Submit**.

You are now ready to configure the source URL for the live event as well as the unicast and multicast publishing URLs for the event. Proceed to the next section, [“Configuring Live Stream Settings.”](#)

Priming a Live Channel Using the Programs Interface

The first client requesting a program often experiences the longest wait time for the program to begin playing. Users can experience long wait times because of the full RTSP negotiation that is required to pull the live stream from the source. Delays can also occur if the edge Content Engine has not buffered

enough stream data to fill the media player's buffer at the time that the program is requested. When the buffer is not filled, some data to the client might be sent at the suboptimal line rate instead of at the Fast Start rate. (For information about the Fast Start feature, see the [“About Fast Start” section on page 9-24.](#))

The ACNS 5.5 software release allows channels for unicast managed live programs to be primed for faster startup times. When a live channel is primed, a unicast-out stream is pulled from the origin server to a Content Engine before a client ever requests the stream. When the first request for the stream goes out, the stream is already in the channel.

To view the status of a primed channel, use the **show programs EXEC** command. Live unicast programs that have been configured to play forever will display current status as “Scheduled” when the channel is unprimed, whereas in previous releases that do not support channel priming, the status is displayed as “Playing.” The status output has been changed to differentiate primed and unprimed channel programs in the ACNS 5.5 software release.

You can enable priming for each Content Engine or each device group that is assigned to a live channel on a per-channel basis. A particular program is primed on a Content Engine when the following conditions exist:

1. The Content Engine is assigned to the program's corresponding channel either as a standalone device or as part of a device group.
2. Either the Content Engine itself or any of the device groups in (1) are marked as primed for the program's corresponding channel.

Because Content Engines and device groups are assigned to channels (and not programs) with a many-to-one relationship between programs and channels, enabling primed devices and device groups on a per-channel basis follows the same structure. Given that a live channel exists to which Content Engines or device groups are assigned, the following two cases are possible:

1. A program is assigned to the channel, whereby none of the Content Engines in the channel should prime the program.
2. A program is assigned to the channel, and some subset of the channel's Content Engines should prime the program. This subset can be different for different programs assigned to the channel.

Priming for these cases can be satisfied if the following rules are observed:

- If you want to prime a program on a subset of the Content Engines assigned to an existing channel, where the desired subset differs from the subset that is priming the existing channel, you must create a new channel with the same membership, publishing point, and website information as the existing channel and assign the program to new channel.
- If the desired subset can be formed by using priming granularity at the device group level, that is what you should do. If the granularity is not sufficient because some Content Engines in the device group are to be primed while others are not, then do one of the following:
 - Create a new device group which results in the desired granularity.
 - Assign the Content Engines on which priming should be enabled directly to the channel; they will remain assigned as part of the device group. Priming can then be enabled on the individual Content Engines and disabled at the device group level.

To configure channel priming, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Video > Programs**.
 - Step 2** Click the **Edit** icon next to the name of a live Windows Media program. If this program is associated with a live channel, the Live Channels Assignments option appears in the Contents pane. (To assign a live channel, see the [“Choosing a Live Channel” section on page 7-13.](#))

- Step 3** Choose **Live Channels Assignments > Assign Content Engines**. The Content Engine Assignments for Live Channel window appears.
- Step 4** Alternatively, choose **Live Channels Assignments > Assign Device Groups**. The Device Group Assignments for Live Channel window appears.
- Step 5** If you have not already done so, choose a root Content Engine from the Root CE drop-down list in the taskbar.
- Step 6** Under the column heading Primed, check the check box. The check box is only active after a device or device group has been assigned.
- Step 7** To save the assignments, click **Submit**.

To complete the configuration, configure the live stream settings for a Windows Media program and schedule the event. (See the [“Configuring Live Stream Settings for a Windows Media Live Program”](#) section on page 7-18 and the [“Scheduling Live Programs”](#) section on page 7-23.)

**Note**

You can also configure this feature using the Channels interface. See the [Priming a Live Channel Using the Channels Interface](#), page 5-18.

Configuring Live Stream Settings

In this section, you will configure the source URL as well as the unicast and multicast publishing URLs for the live event.

**Note**

The monitoring interval for a live program is 30 seconds. The retry interval for a live program is 60 seconds. Retries will be done forever. These values are applicable for both CSE and WMT live programs.

Configuring Live Stream Settings for a Windows Media Live Program

To configure live stream settings for a Windows Media live program, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Video > Programs**.
 - Step 2** Click the **Edit** icon next to the name of a live Windows Media program.
 - Step 3** In the Contents pane, choose **Live Streaming**. The Streaming Settings for Program window appears. (See [Figure 7-5](#).)

Figure 7-5 Streaming Settings for a Windows Media Live Program

The screenshot shows the Cisco Application and Content Networking System GUI. The 'Services' tab is selected, and the 'Live Streaming Settings' section is expanded. The 'Live Source URL' field contains the text 'http_or_rtsp://encoder#host_or_sourceServer#host.port#fileName'. Below this, there are checkboxes for 'Enable Unicast Delivery to Client' and 'Enable Multicast Delivery to Client'. The 'Unicast URL Reference' field is set to 'Please make a choice', and the 'Multicast URL Reference' field is set to 'Customized URL Reference'. The 'Multicast Address' field is set to '192.168.1.1' and the 'Port' field is set to '8080'. The 'Multicast TTL' field is set to '15'. The 'Auto Select' button is visible next to the 'Port' field.

Step 4 In the Live Source URL field, enter the URL of the origin Windows Media encoder or Windows Media Server using the following format:

- **http:// WMTencoder:port**
- **rtsp://wmStreamingServer:port/publishPoint**

For encoder failover, you can specify more than one encoder. Separate live source URLs in the list by using a semicolon (;), as in this example:

http://encoder_1:8080;http://encoder_2:8080

Step 5 If the program requires unicast transmission, check the **Enable Unicast Delivery to Client** check box. If a live channel has been associated with the program, a unicast URL is automatically generated and appears in the Unicast URL Reference field.

**Note**

The Unicast URL Reference field is auto-populated from the Origin Server field and the Request Routed FQDN field that you configured when you created the website for the live channel. For more information about creating websites and channels, see the [“Creating and Modifying Websites”](#) section on page 5-8 and the [“Creating and Modifying Channels”](#) section on page 5-10.

**Note**

For Windows Media managed live streaming, HTTP unicast delivery (Unicast in- Unicast out) to the client is not supported. Only RTSP URLs are supported. In the Live Source URL field in the Content Distribution Manager GUI, you can have an HTTP or RTSP URL, but in the Unicast URL Reference field only RTSP URLs are allowed.

Alternatively, if you want to enter a new URL, choose the **Customized Reference URL** option from the Unicast URL Reference drop-down list, and enter a unicast publishing URL in the Customized URL field. Use one of the following formats for a customized URL for unicast:

rtsp://livechannelOriginFQDN/customized-name

http://livechannelOriginFQDN/customized-name



Note Customized URLs can use HTTP; however, they must use an origin server FQDN and not a Content Router FQDN.



Note For Windows Media live programs, URLs of the format *rtsp://rootContentEngine/programName* (used to enable unicast delivery to end clients) must be intercepted transparently using WCCP or explicitly configured to use an RTSP proxy before they can be served by a Content Engine.

Step 6 To enable multicast delivery, check the **Enable Multicast Delivery to Client** check box. If a live channel has been associated with the program, a multicast URL is automatically generated and appears in the Multicast URL Reference field.

If there is no URL listed, or if you want to enter a new URL, choose the **Customized Reference URL** option from the Multicast URL Reference drop-down list, and enter a multicast publishing URL in the Customized URL field. Use the following format for a customized URL for multicast:

http://anyCEinLiveChannel/customized-name.nsc

Step 7 If you want to use a server-side playlist as the media source for a multicast program, enter the URL for the NSC file in the NSC Reference for Multicast field. (See the [“About NSC Files for Multicasting Windows Media Server-Side Playlists”](#) section on page 7-23.)



Note ACNS software supports the HTTP protocol for fetching the NSC file only. Make sure that the NSC file is saved in a location from where it can be accessed using the HTTP protocol. You must use a URL with a format similar to the following: *http://<hostIP>/<filename>.nsc*.

Step 8 If you enabled multicasting, you must enter the multicast address and the port number. The address range is 224.0.0.0 to 239.255.255.255. The port number range is 1 to 65535. These values must be unique within the system.

Alternatively, the multicast address and port can be automatically selected from the program address pool, if a pool has been defined. (See the [“Managing Multicast Addresses for Programs”](#) section on page 7-46.) To automatically choose the multicast address and port, click the **Auto Select** button.

Step 9 In the Multicast TTL field, specify the multicast Time To Live (number of hops). The default is 15 hops.

Step 10 To save the settings, click **Submit**.

Configuring Live Stream Settings for a Cisco Streaming Engine Live Program

There are two major differences between Cisco Streaming Engine live program streaming settings and Windows Media live program streaming settings:

- Cisco Streaming Engine live programs can have more than one live stream (audio, video, and slides). Each stream requires a unique multicast address and port.

- Publishing multicast and unicast URLs for Cisco Streaming Engine Live programs cannot be autogenerated. The URLs point to meta files (sdp) that are generated and reside on an external server controlled by the administrator.

To configure live stream settings for a Cisco Streaming Engine Live program, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Video > Programs**.
- Step 2** Click the **Edit** icon next to the name of a live Cisco Streaming Engine program.
- Step 3** In the Contents pane, choose **Live Streaming**. The Streaming Settings for Program window appears. (See [Figure 7-6](#).)

Figure 7-6 Streaming Settings for a Cisco Streaming Engine Live Program

The screenshot displays the 'Streaming Settings for Program, CSE-Live (CSE Live)' configuration window. The 'Contents' pane on the left shows 'Live Streaming' selected. The main area contains the following settings:

- CSE Live Streaming Settings**
 - Origin Server SDP File URL:
 - Unicast URL Reference: (Note: Please select one of the suggested URLs, or select "Customized".)
 - Customized URL:
 - Enable Multicast Delivery to Client: ☒
 - Multicast URL Reference: (Note: http://sourceHost_or_Fqdn/path/file)
 - Multicast TTL: (1-255)
- Live Streams**
 -
 - Table with 4 columns: Source Server, Root CE Port, Multicast Address, Multicast Port.

Source Server	Root CE Port	Multicast Address	Multicast Port
1 * streamserver	8080	224.1.1.1	1096
 -

Note: * - Required Field

- Step 4** In the Origin Server SDP File URL field, enter the URL for the SDP file generated on the encoder.
- Step 5** If the program requires unicast transmission, check the **Enable Unicast Delivery to Client** check box, and choose one of the suggested publishing URLs from the Unicast URL Reference drop-down list.



Note

The list of suggested URLs is auto-populated from the Origin Server field and the Request Routed FQDN field that you configured when you created the website for the live channel. For more information about creating websites and channels, see the [“Creating and Modifying Websites”](#) section on page 5-8 and the [“Creating and Modifying Channels”](#) section on page 5-10.

Alternatively, choose **Customized Reference URL** from the Unicast URL Reference drop-down list, and enter a unicast publishing URL in the Customized URL field using one of the formats described in the [“About Unicast Publishing URLs” section on page 7-22](#).

- Step 6** To enable multicast delivery, check the **Enable Multicast Delivery to Client** check box. If a live channel has been associated with the program, a multicast URL is automatically generated and appears in the Multicast URL Reference field.



Note If you wish to enable support for root Content Engine failover, you must check this check box. Root Content Engine failover for a live program works only when the incoming stream is a multicast stream. (See the [“About Root Content Engine Failover and Fallback Support for Cisco Streaming Engine Live Programs” section on page 7-5](#).)

- Step 7** In the Multicast TTL field, specify the multicast Time To Live (number of hops). The default is 15 hops.

- Step 8** Define at least one live stream.

- a. Enter the stream source server IP address in the Source Server field.
- b. Enter the root Content Engine port number in the Root CE Port field.



Note To activate failover support for the root Content Engine, you must enter **0** (zero) in the Root CE Port field. (See the [“About Root Content Engine Failover and Fallback Support for Cisco Streaming Engine Live Programs” section on page 7-5](#).)

You can define up to 10 different live streams for each program.

- Step 9** To add a new live stream, click the **Add New Live Stream** button.

- Step 10** Enter the multicast address and port for each live stream defined. The address range is 224.0.0.0 to 239.255.255.255. The port number range is 1 to 65535. These values must be unique within the system.

Alternatively, the multicast address and port can be automatically chosen for each stream from the program address pool, if one has been defined. (See the [“Managing Multicast Addresses for Programs” section on page 7-46](#).) To automatically choose the multicast address and port, click the **Auto Select** button.

- Step 11** To save the settings, click **Submit**.

You have completed live stream configuration. You must now set the playtimes for the program. Proceed to the next section, [“Scheduling Live Programs.”](#)

About Unicast Publishing URLs

The unicast publishing URL can use one of the following formats:

- For proxy or WCCP, use one of these formats:
 - **rtsp://liveChannelOriginFqdn/programName**
 - **http://liveChannelOriginFqdn/programName**
- For Content Router routing, use one of these formats:
 - **rtsp://liveChannelCrFqdn/programName)**
 - **http://liveChannelCrFqdn/programName.asx**

- For Global Site Selector (GSS), use one of these formats:
 - `rtsp://liveChannelCrFqdn/programName`
 - `http://liveChannelCrFqdn/programName`

About Multicast Publishing URLs

The multicast publishing URL can use one of the following formats:

- For proxy or WCCP, use one of these formats:
 - `http://anyCEinLiveChannel/programName.nsc`
- For Content Router routing, use one of these formats:
 - `http://liveChannelCrFqdn/programName.nsc.asx`
- For GSS, use one of these formats:
 - `http://liveChannelCrFqdn/programName.nsc`



Note

ACNS 5.2 software supported the multicast publishing URL format, `http://liveChannelOsFqdn/programName.nsc`. However, the live channel origin server FQDN had to be the live channel's root Content Engine. This requirement no longer applies to ACNS 5.3 software and later releases.

About NSC Files for Multicasting Windows Media Server-Side Playlists

Windows Media Station (NSC) files are essentially playlists that contain information to allow Windows Media Player to connect to and play streaming media. Windows Media Player uses NSC (.nsc file extension) files to get the information it needs to receive multicast content over the Internet. These files can contain information such as stream location and rollover URL, as well as descriptive information about the station. Whereas standard streaming multimedia sends a single media stream to a single recipient, multicasting allows a single media stream to be received by more than one person, much like a television or radio broadcast. NSC files contain the information necessary to allow multicast streams to be processed correctly by Windows Media.

For proper multicast delivery of a Windows Media live program, provide the URL of the web server that is hosting the program NSC file in the Stream Settings window for Windows Media live program. (See the [“Configuring Live Stream Settings for a Windows Media Live Program”](#) section on page 7-18.)

Scheduling Live Programs

Each live program can have up to 10 different playtimes scheduled. The program start time can be specified in Universal Coordinated Time (UTC), which is the same as Greenwich mean time (GMT), or in the Content Engine local time. The program is broadcast from all Content Engines simultaneously; local Content Engine time may differ based on the time zone settings of the device.

To schedule a live program, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Video > Programs**.
- Step 2** Click the **Edit** icon next to the name of a live program.
- Step 3** In the Contents pane, choose **Schedule**. The Scheduling Program window appears. (See Figure 7-7.)
If you have not previously scheduled or added a playtime, you must choose one of the following two options:
 - To schedule a playtime, add a new playtime, or edit a playtime, choose the **Schedule Playtime** radio button. The window refreshes and displays the playtime editor.
 - Alternatively, you can choose the **Play Forever** radio button. If you chose this option, the schedule is preset to play continuously, and the playtime editor not displayed.

Figure 7-7 Scheduling Program Window—Adding a New Playtime

The screenshot shows the Cisco Application and Content Networking System GUI. The main window is titled "Scheduling Program, CSE-Live". It features a "Contents" pane on the left with options like Definition, Select Live Channel, Live Channel Assignments, Live Streaming, and Schedule. The main area displays a "Playlist Schedule" table with two entries:

	Initial Start Time	Duration	Repeat
1	Tue, 2 Aug 2005 01:02 (UTC)	00:00:00	No
2	Wed, 1 Jun 2005 22:15 (UTC)	01:00:00	No

Below the table is the "Editing Playtime 1" section. It includes a "Start Playback on:" field set to August 2, 2005, at 1:02 (24-hour clock) in UTC. The "Duration:" field is set to 1 hour. There are three radio button options for repetition: "Do Not Repeat" (selected), "Repeat Every 1 Days", and "Repeat Weekly at 1:02 on Selected Days". The weekly option includes checkboxes for Sun, Mon, Tue, Wed, Thu, Fri, and Sat.

- Step 4** In the playtime editor, enter a start date and time for the program. The playtime editor displays the current time and a duration of one hour as the default. You can edit these fields as you want.
- Step 5** From the drop-down list, choose either **UTC** or **CE(s) Local Time**. The default is UTC.
In general, a live event should be based on UTC time for successful delivery across time zones. However, if all Content Engines are located in the same time zone, the program can use local Content Engine time.
- Step 6** To configure the duration of the program, in the Duration field, enter a number, and from the drop-down list, choose minutes, hours, or days as the unit of time.

- Step 7** Choose whether or not to repeat the program. If you choose to repeat the program, choose how often to repeat it.
- Do Not Repeat—Configures a one-time live event.
 - Repeat Every—Configures the program to repeat every so many days, hours, or minutes.
 - Repeat Weekly—Configures the program to repeat at the same hour on the days that you choose.
- Step 8** To save the settings, click **Submit**.
- Step 9** To add a new playtime, click the **Add Playtime** button and repeat the above steps. You can have up to 10 different playtimes scheduled for a program.
-

Creating a Scheduled Rebroadcast Program

You can define a scheduled rebroadcast program in the Content Distribution Manager GUI by choosing **Window Media Rebroadcast** or **Cisco Streaming Engine Rebroadcast** as the program type. (See the [“Defining a Program”](#) section on page 7-11.) After you define the program as a rebroadcast program, you are ready to choose the media files that you want to include in the program. Proceed to the next section, [“Choosing the Media Files”](#).

Choosing the Media Files


In this section you can define a series of media files to be played during your scheduled rebroadcast program. You can define the media file playlist by performing the following tasks in the Content Distribution Manager GUI:

1. Choose the content channel.
2. Perform a content query based on search criteria.
3. Select media that has replicated on the root Content Engine.
4. Define an ordered list of the selected media files.

To choose media files for playback, follow these steps:


-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Video > Programs**.
- Step 2** Click the **Edit** icon next to the name of a rebroadcast program.
- Step 3** In the Contents pane, choose **Select Media**. The Selecting Media Files for Program window appears. (See [Figure 7-8](#).)

Figure 7-8 Selecting Media Files for Rebroadcast Program



- Step 4** To choose a content channel from the list, click the radio button next to the name of the channel.
- Step 5** Click the **Show Media in Selected Channel** button. The Media File Selection pane appears. (See [Figure 7-9](#).)

Figure 7-9 Media File Selection Pane



- Step 6** You can search the channel for the media files that you want to add to the playlist by entering the search criteria in the Criteria field. Use an asterisk (*) to match zero or more characters, or a question mark (?) to match exactly one character. For example, use “*.mpg” for all files with the suffix “mpg,” and “file?.mpg” to match file1.mpg, file2.mpg, and so on. Click **Use Criteria**. All media files in the selected channel that match the criteria appear.
- Step 7** Check the **Pick** check box next to each media file that you want to add to the playlist and click the **Add Media** button. The media file is added to the bottom of the list in the Media Files in Program pane on the right.
- Step 8** In the right pane, use the **Up Arrow** and **Down Arrow** keys to move a media file up or down one position in the list. Unless you check the **Shuffle** check box in the Playlist Information window, files are played in the order in which they are listed.
- Note that the Up Arrow and Down Arrow are only displayed if the list of media files in the playlist is sorted by position. If you sort the media files alphabetically, for example, then the arrows do not appear.
- Step 9** To remove a media file from the list, check the **Pick** check box next to the file, and click the **Remove Media** button.



Note Media for a Windows Media rebroadcast program can only be chosen from a single channel. If you choose media files from different channels, a warning message appears, and all previously chosen media files are erased.

- Step 10** To save the playlist, click **Submit**.

You have chosen a content channel for the program, and you have also chosen the media files from that channel that you want to add to the program playlist. You are now ready to assign Content Engines or device groups from the channel to the program. Proceed to the next section, “[Assigning Content Engines](#).”

Assigning Content Engines

For rebroadcast programs, media can only be selected from one channel. The Content Engines and device groups assigned to that channel are selected automatically when you choose the media files for the program. If at a later time you add new Content Engines or device groups to the channel, you must manually add the new Content Engines (individually, or as a device group, or both) to the program.

To assign new Content Engines to the rebroadcast program that you are creating, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Video > Programs**.
 - Step 2** Click the **Edit** icon for the rebroadcast program that you want to modify. The Modifying Program window for the rebroadcast program that you are working on appears.
 - Step 3** In the Contents pane, choose **Assign Content Engines**. The Assignment of Content Engine of Channel to Program window appears.



Note If you have not chosen a channel and added media files to the program from the channel, this window does not appear. Instead, you see the Content Engine Assignments for Program window, which directs you choose your media files first.

- Step 4** Click the **Assign** icon (blue cross mark) next to as many Content Engines as you want to assign to the program, and click **Submit**. A green tick mark appears next to the assigned Content Engine.
 - Step 5** To unassign a Content Engine that has been assigned to the program, click the **Unassign** icon (green tick mark) next to the assigned Content Engine, and click **Submit**.
-

Assigning Device Groups

To assign a new device group to the rebroadcast program that you are creating, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Video > Programs**.
 - Step 2** Click the **Edit** icon for the rebroadcast program that you want to modify. The Modifying Program window for the rebroadcast program that you are working on appears.
 - Step 3** In the Contents pane, choose **Assign Device Groups**. The Assignment of Device Groups of Channel to Program window appears.



Note If you have not chosen a channel and added media files to the program from the channel, this window does not appear. Instead, you see the Device Group Assignments for Program window, which directs you choose your media files first.

- Step 4** Click the **Assign** icon (blue cross mark) next to as many device groups as you want to assign to the program, and click **Submit**. A green tick mark appears next to the assigned device group.
-

- Step 5** To unassign a device group that has been assigned to the program, click the **Unassign** icon (green tick mark) next to the assigned device group, and click **Submit**.

Configuring Streaming Settings

To configure streaming settings for a scheduled rebroadcast program, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Video > Programs**.
- Step 2** Click the **Edit** icon for the rebroadcast program that you want to modify. The Modifying Program window for the rebroadcast program that you are working on appears.
- Step 3** In the Contents pane, choose **Streaming**. The Streaming Settings for Program window appears.
- Step 4** Choose one of the suggested publishing URLs from the Multicast URL Reference drop-down list.



Note The list of suggested URLs that appears in the drop-down menu is based on the pre-positioned channel that you chose when making your media file selection. URLs are auto-generated according to a pre-defined pattern that is derived from the Origin Server field and the Request Routed FQDN field that you configured when you created the website for the pre-positioned channel. For more information about creating websites and channels, see the [“Creating and Modifying Websites”](#) section on page 5-8 and the [“Creating and Modifying Channels”](#) section on page 5-10.

Alternatively, choose **Customized URL Reference**, and enter a publishing URL in the field provided using the following format:

- For Windows Media: **http://ce_ipaddress/programName.nsc**
- For Cisco Streaming Engine: **http://OS-FQDN/programName.sdp** or **http://CR-FQDN/programName.sdp**

- Step 5** In the Multicast TTL field, specify the multicast Time To Live (number of hops). The default is 15 hops.
- Step 6** Enter up to three multicast addresses and port numbers. The address range is 224.0.0.0 to 239.255.255.255. The port number range is 1 to 65535. These values must be unique within the system.
- Alternatively, the multicast address and port can be automatically chosen from the program address pool, if one has been defined. (See the [“Managing Multicast Addresses for Programs”](#) section on page 7-46.) To automatically choose the multicast address and port, click the **Auto Select** button.



Note Because Cisco Streaming Engine rebroadcast files can contain multiple tracks (1–3), you can define up to three multicast addresses and ports for each track in the file. At least one multicast address and port must be specified, even if a program address pool is defined; otherwise, program creation fails. Click the **Add Multicast Address/Port** button to define a new track. The GUI displays a track number, which corresponds to the tracks in the stream.

- Step 7** To save the settings, click **Submit**.

Scheduling Rebroadcast Programs

To schedule a rebroadcast program, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Video > Programs**.
- Step 2** Click the **Edit** icon for the rebroadcast program that you want to schedule. The Modifying Program window for the rebroadcast program that you are working on appears.
- Step 3** In the Contents pane, choose **Schedule**. The Scheduling Program window appears.
- Step 4** If you want the playlist to play in a continuous loop, click the **Loop Playback Continuously** radio button, and click **Submit**.



Note The **Loop Playback Continuously** radio button toggles with the **Schedule Playback** radio button. These options are mutually exclusive; you can choose one option or the other.

- Step 5** To schedule playback times for the playlist, click the **Schedule Playback** radio button. The window refreshes, and the Playtimes control window appears. You can schedule up to ten playback times per playlist.



Note No two playback times for a single playlist can have the same start time.

- Step 6** To add a playtime (that is, the time a playback is scheduled to occur), click the **Add a New Playtime** icon, which is in the Playtimes header. (See [Figure 7-7](#).) A new playtime selection is added to the list, and the Playtime Editor pane appears. Configure the playtime by entering the necessary information in the appropriate fields. To save the playtime, click **Submit**.
- Step 7** To remove an existing playtime, click the **Delete** icon that is next to the playtime selection you want to remove. The playtime is removed from the schedule.
- Step 8** To edit an existing playtime, click the **Edit** icon next to the scheduled playtime that you want to change. Change the schedule and playback options as necessary. [Table 7-2](#) explains the fields under the Editing Playtime heading.



Note The dates and times that are entered and displayed are observed according to the local times on the Content Engine.

Table 7-2 *Playlist Schedule Settings*

Schedule Setting	Description
Start Playback on	Local time that the playlist starts playing. This field is required.
Playback Once and Stop	Playlist is played only once and is stopped.
Loop for	Playlist repeats for the specified duration.
Do Not Repeat	Playback occurs only once and does not repeat.
Repeat Every	Playback repeats every so many minutes, hours, or days according to the interval specified (for example, every 10 minutes).

Table 7-2 *Playlist Schedule Settings (continued)*

Schedule Setting	Description
Repeat Weekly at <start time> on selected days	Playback repeats weekly at the local time indicated in the Start Playback on fields.
Repeat Forever ¹	Playback schedule runs indefinitely.
Repeat Until ¹	Playback schedule runs until the date and time that you set.

1. This option only appears in the window after you choose the Repeat Every or Repeat Weekly option.

Step 9 Click **Submit**.

Creating a TV-Out Program

A TV-out program uses a Content Engine equipped with an MPEG decoder card to play back the media files. You can combine media from multiple channels in the playlist and add image overlays to media files in TV-out programs. Before you create a TV-out program, you must first enable the TV-out feature on the Content Engine and designate the format of the signal output.

Enabling the TV-Out Feature

Content Engines that are equipped with an integrated Moving Picture Experts Group (MPEG) decoder can convert digital media into analog TV signals. These Content Engines play media files using National Television Standards Committee (NTSC) or Phase Alternation Line (PAL) video signals and can thus play video directly to a TV monitor for applications such as kiosks, cable TV systems, and video walls.

Before you create a TV-out program, you must first enable the TV-out feature on the Content Engine that delivers the program. Enabling the TV-out feature will succeed only if supported audio and video hardware can be detected on the Content Engine. Autodetection of audio and video supported hardware occurs when you enable TV-out. You also need to designate the TV signal output format that is used when delivering content from a TV-out enabled Content Engine to a TV monitor.

To enable the TV-out feature on a Content Engine and designate the format of the signal output, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to enable for TV-out.
 - Step 3** In the Contents pane, choose **Applications > Set Top Box > TV-out**. The TV-out Settings for Content Engine window appears. [Table 7-3](#) describes the corresponding global configuration CLI commands.
 - Step 4** To enable the TV-out feature, check the **TV-out Enable** check box.
 - Step 5** Choose a TV signal output format (**ntsc** or **pal**) from the TV-out Signal drop-down list.
 - Step 6** To save the settings, click **Submit**.
 - Step 7** To disable TV-out, uncheck the **TV-out Enable** check box and click **Submit**. Disabling TV-out unloads the TV-out driver and stops all associated software.

Table 7-3 GUI Fields and Corresponding CLI Commands

GUI Field	Command
TV-out Enable (checked)	tvout enable
TV-out Enable (unchecked)	no tvout enable
TV-out Signal	tvout signal { ntsc pal } Where ntsc is the NTSC standard and pal is the PAL standard. NTSC is the default.

Supported MPEG File Specifications

Table 7-4 lists the supported MPEG file specifications for TV-out enabled devices.

Table 7-4 MPEG File Specifications

Specification	Description
MPEG stream type	MPEG-1 system stream MPEG-2 program stream MPEG-2 transport stream
MPEG video	MPEG-1—4:2:0 Constrained Parameter Bitstream MPEG-2—4:2:0 Main Profile (MP) at Main Level (ML)
MPEG audio	Layers 1 and 2 at 32 kHz, 44.1 kHz, 49 kHz
Data rate	Up to 15 Mbps
Video resolution	NTSC 352 x 240 at 30 frames per second (fps) 352 x 480 at 30 fps 480 x 480 at 30 fps 544 x 480 at 30 fps 704 x 480 at 30 fps 720 x 480 at 30 fps PAL 352 x 288 at 24 fps 352 x 576 at 24 fps 480 x 576 at 24 fps 544 x 576 at 24 fps 704 x 576 at 24 fps 720 x 576 at 24 fps

Selecting Media Files and Overlay Images for Playlists


The Content Distribution Manager GUI lets you assemble a series of media files from different channels to create a playlist. The media is then played on any video-out enabled device according to the defined order in the playlist. You can associate one overlay image with each playlist. The user sees this image displayed in a determinable screen location when the playlist it is associated with is playing. Overlay images are useful for displaying logos.

Adding Media Files to a Playlist

To add media files to a playlist, follow these steps:

- Step 1** Choose **Services > Video > Programs**. The Programs window appears with existing playlists listed.
- Step 2** Click the **Edit** icon next to the TV-out program to which you want to add media files.
- Step 3** In the Contents pane, choose **Media and Overlays**. The Selecting Media Files for program window appears. (See Figure 7-10.) The left pane displays channels and the right pane displays media files currently in the program playlist. You can add more media files to the playlist.

Figure 7-10 Selecting Media Files for Program Window



- Step 4** In the left pane, choose the channel whose media files you want to display. Click **Show Media in Selected Channel**.

- Step 5** Search for the channel's media files by entering the search criteria in the Criteria field. Use an asterisk (*) to match zero or more characters, or a question mark (?) to match exactly one character. For example, use "*.mpg" for all files with the suffix "mpg," and "file?.mpg" to match file1.mpg, file2.mpg, and so on. Click **Use Criteria**. All media files in the selected channel that match the criteria appear.
- Step 6** Check the check box next to each media file that you want to add to the playlist. Click **Add**. The media file appears in the right pane at the bottom of the list.
- Step 7** In the right pane, you can use the **Up Arrow** and **Down Arrow** keys to move a media file up or down one position in the list. Unless you check the **Shuffle** check box in the Playlist Information window, files are played in the order in which they are listed.

Note that the Up Arrow and Down Arrow are only displayed if the list of media files in the playlist is sorted by position. If you sort the media files alphabetically, for example, then the arrows do not appear.

Adding an Overlay Image to a Playlist



Note

For CE-507AV and CE-560AV models, overlay images must be Windows bitmap files with 16 colors. They must have 4 bits per pixel and must be less than 64 KB in size. The width of the image (in pixels) must be divisible by 8. Before associating an overlay image, open the image file in an image editor, and view its properties to make sure that it meets the above requirements.



Note

For CE-510 and CE-565 models, overlay images must be bitmap files with either 16 or 256 colors. They must have 4 or 16 bits per pixel, and must be less than 128 KB in size. The width of the image (in pixels) must be divisible by 8. Before associating an overlay image, open the image file in an image editor, and view its properties to make sure that it meets the above requirements.

To add an overlay image to a playlist, follow these steps:

- Step 1** Choose **Services > Video > Programs**. The Programs window appears with existing playlists listed.
- Step 2** Click the **Edit** icon next to the TV-out program to which you want to add an overlay image. The Modifying Program window appears.
- Step 3** In the Contents pane, choose **Media and Overlays**. The Selecting Media Files for program window appears. (See [Figure 7-10](#).) The left pane displays channels and the right pane displays media files currently in the program playlist.
- Step 4** In the Channels pane (on the left), choose the channel whose media files you want to display. Click **Show Media in Selected Channel**.
- Step 5** Search for the channel's media files by entering the search criteria in the Criteria field. Use an asterisk (*) to match zero or more characters, or a question mark (?) to match exactly one character. For example, use "*.mpg" for all files with the suffix "mpg," and "file?.mpg" to match file1.mpg, file2.mpg, and so on. Click **Use Criteria**. All media files in the selected channel that match the criteria appear.
- Step 6** In the left pane, check the **Pick** check box next to the media file that you want to use as an overlay image, and click **Add Overlays**.

The Selecting Media Files for program window becomes the Selecting Overlays for program window. The right pane refreshes and becomes the Overlay Images in Program pane. The media file that you added as an overlay appears in the right pane at the bottom of the list of overlay images.

Step 7 Click **Submit**.

Editing a Overlay Image

To edit an overlay image, follow these steps:

- Step 1** Choose **Services > Video > Programs**. The Programs window appears with existing playlists listed.
- Step 2** Click the **Edit** icon next to the TV-out program for which you want to edit an overlay image. The Modifying Program window appears.
- Step 3** In the Contents pane, choose **Media and Overlays**. The Selecting Media Files for program window appears. (See [Figure 7-10](#).) The left pane displays channels and the right pane displays media files currently in the program playlist.
- Step 4** In the Selecting Media Files for program window, click **Edit Program Overlays**. A list of overlay images for the selected media file appears in the Overlay Images in Program pane (on the right).
- Step 5** In the Overlay Images for Program pane, click the **Edit** icon next to the overlay image that you want to edit. The Overlay Settings for Playlist pane appears and contains the image name and the channel from which it was obtained.
- Step 6** You can edit the parameters of the overlay image by specifying and adjusting its location on the TV-out or set top box monitor, specifying the overall opacity of the image, and choosing whether you want one color to be transparent. If desired, you can remove the entire overlay by clicking **Remove Overlay**. (For a description of overlay image parameters, see [Table 7-5](#).)

Table 7-5 *Overlay Image Parameters*

Image Parameter	Description
Name	Name of the overlay. A default name is given to each overlay. You can change the default name, but you cannot remove the name. The name must be a valid XML name; that is, it can contain alphanumeric characters and underscores, and it must be unique.
Set as Default Overlay	Specifies whether this overlay is the default overlay for this playlist. Each playlist can have only one default overlay.
Placement	Specifies where the overlay image should appear on the TV-out or STB monitor. Your options are upper left, upper right, lower left, lower right, and center.
Height Offset	Adjusts placement along the vertical axis. The valid range is from –40 to 40 pixels. –40 = Moves the image to the lowest possible position on the display. 40 = Moves the image to the uppermost position on the display.
Width Offset	Adjusts placement along the horizontal axis. The valid range is from –60 to 60 pixels. –60 = Moves the image to the left-most possible position on the display. 60 = Moves the image to the right-most possible position on the display.

Table 7-5 Overlay Image Parameters (continued)

Image Parameter	Description
Opacity	Adjusts the opacity of the overlay image. The valid range is from 0 to 15. 0 = Renders the image completely opaque. 15 = Renders the image nearly transparent.
Transparent Color	Enables or disables the use of transparent color. When checked, the color of the pixel in the lower left corner of the image is transparent over the entire image. When a transparency-enabled image is placed over a video, the background color of the overlay image is not visible against the video playing behind it. This feature is useful if you want to use a nonrectangular image as your overlay.

Step 7 Click **Submit**.

Assigning an Overlay Image to a Media File

To assign an overlay image to a media file, follow these steps:

-
- Step 1** Choose **Services > Video > Programs**. The Programs window appears with existing playlists listed.
- Step 2** Click the **Edit** icon next to the TV-out program that you want to edit. The Modifying Program window appears.
- Step 3** From the Modifying Program window, choose **Assign overlays** from the Contents pane. The Assigning Overlay to Media window appears.
- The left pane displays a list of overlay images in the playlist, and the right pane displays media files in the program. The default overlay in the left pane is highlighted. Note that you cannot remove the default overlay.
- Step 4** Click the radio button next to the overlay that you want to associate with one or more media files.
- Step 5** In the Media Files in Program pane, check the check box next to the media file or files to which you want to associate the overlay.
- Step 6** Click **Submit**.
-

Scheduling TV-Out Programs



Note

You can only set playback schedules for playlists that are played over TV-out enabled devices.

To set automatic playback schedules for TV-out playlists, follow these steps:

-
- Step 1** Choose **Services > Video > Programs**. The Programs window appears with existing playlists listed.
- Step 2** Click the **Edit** icon next to the TV-out program that you want to schedule. The Modifying Program window appears.

- Step 3** From the Modifying Program window, follow Step 3 through Step 9 in the “[Scheduling Rebroadcast Programs](#)” section on page 7-30.

Assigning Content Engines to Playlists




Note

You can only assign TV-out enabled Content Engines to playlists by using the Content Distribution Manager.

To assign a TV-out enabled Content Engine or a group of Content Engines to a playlist, follow these steps:

- Step 1** Choose **Services > Video > Programs**. The Programs window appears with existing programs listed.
- Step 2** Click the **Edit** icon next to the TV-out program to which you want to assign a Content Engine. The Modifying Program window appears.
- Step 3** From the Modifying Program window, choose **Assign Content Engines** in the Contents pane. The Content Engine Assignments for Playlist window appears. (See [Figure 7-11](#).)

Figure 7-11 Content Engine Assignments for Playlist Window



- Step 4** Click the **Assign** icon (blue cross mark) next to as many Content Engines as you want to assign to the playlist, and click **Submit**. A green tick mark appears next to the assigned Content Engine.

- Step 5** To unassign a Content Engine that has been assigned to the playlist, click the **Unassign** icon (green tick mark) next to the assigned Content Engine, and click **Submit**.
-


Assigning Device Groups to Playlists

You can assign one or more device groups to a playlist. A device group is a set of devices that share common qualities and capabilities and can be assigned to one or more channels. Some common qualities might include disk capacity, distribution minimum bandwidth, or routing properties. One channel can have many device groups assigned to it.

Device groups are assigned through the Content Distribution Manager GUI. When a channel is created, when additional groups are added, or when a device group assignment changes, the devices in the group are notified of their assignment to the associated channel. Each device can subscribe to more than one channel.

To assign a device group to a playlist, follow these steps:

-
- Step 1** Choose **Services > Video > Programs**. The Programs window appears with existing programs listed.
- Step 2** Click the **Edit** icon next to the program to which you want to assign a device group. The Modifying Program window appears.
- Step 3** From the Modifying Program window, choose **Assign Device Groups** in the Contents pane. The Device Group assignments for Playlist window appears. (See [Figure 7-12](#).)

Figure 7-12 Assigning Device Groups to Playlists Window

- Step 4** Click the **Assign** icon next to as many device groups as you want to assign to the current playlist, and click **Submit**. A green tick mark replaces the blue cross mark.
- Step 5** To remove device groups from the playlist, click the **Unassign** icon (green tick mark) next to the device groups that you want to remove, and click **Submit**.

Viewing the TV-Out Playback Guide

You can view a playback guide that displays playlists by date for a selected Content Engine. Playback schedules can be viewed by day, week, month, or year.

To view the TV-out playback guide, follow these steps:

- Step 1** Choose **Devices > Devices**. The Content Engines window appears with a current list of Content Engines.
- Step 2** Click the **Edit** icon next to the Content Engine with the playback guide you want to view.
- Step 3** Choose **Applications > Set Top Box > TV-out Playback Guide** from the Contents pane. The TV-out Playback Guide for Content Engine window appears. (See [Figure 7-13](#).)

Figure 7-13 Playback Guide for Content Engine Window



- Step 4** Choose the **Day**, **Week**, **Month**, or **Year** tab to view the playback schedules. Scheduled playlists are listed by start time (initial start time plus any repeat intervals). Times begin with the current device time (current system time plus device time zone offset).

Creating an Export Program

You create an export program when you want to play the program using a set top box. A set top box is a device capable of taking information from a network interface using IP and displaying it on a monitor attached to a video interface (digital-to-analog conversion).

To play video over a set top box, you must complete the following tasks:

1. Define a program as an export program—[Defining a Program, page 7-11](#)
2. Create a playlist by selecting media files and overlay images—[Selecting Media Files and Overlay Images for Playlists, page 7-33](#)
3. Configure an export channel to export the playlist to the set top box—[Selecting Export Channels for Set Top Box Playback, page 7-41](#)

Selecting Export Channels for Set Top Box Playback


Note

You can create and use export channels only for playlists that are played over set top boxes.

To create an export channel for a playlist to be played over an set top box, follow these steps:

Step 1 Choose **Services > Video > Programs**. The Programs window appears with existing programs listed.


Step 2 Click the **Edit** icon next to the export program for which you want to create an export channel. The Modifying Programs window appears.


Note


You must choose a playlist listed as Export in the Type column for the **Select Export Channel** link to appear as a choice in the Contents pane.

Step 3 Choose **Select Export Channel** from the Contents pane. The Selecting Export Channel for program window appears. All current export channels already in the system are listed. (See Figure 7-14.) If a channel has already been associated with the current playlist, it appears at the top of the channel column, and its radio button is checked.

Figure 7-14 Selecting Export Channel for Program Window



Step 4 To add a new export channel, click the **Create a new Export Channel** icon. The Adding New Export Channel window appears. (See Figure 7-15.)

Figure 7-15 Adding New Export Channel Window

- Step 5** Enter the export channel name in the Name field, or use the default channel that is displayed. The export name must follow this naming convention: `export_playlist name`.
- Step 6** Choose a content provider from the Content Provider drop-down list. All previously configured content providers in the system are displayed.
- Step 7** From the Website drop-down list, choose a website that contains the playlist file you want to export. All websites mapped to the content provider chosen in the previous step are displayed.
- The playlist file URL should have the following format:
`http://<website_fqdn>/playlists/<playlistName>.xml`
- The playlist file does not exist on the origin server and only maps to the routed FQDN.
- Step 8** In the Channel Quota field, enter the amount of disk space (in megabytes) that the export channel will use. The default is 5 megabytes.
- Step 9** Click **Submit**. The Selecting Export Channel for program window appears.
- Step 10** Check the radio button next to the export channel that you wish to associate with the current playlist.
- Step 11** Click **Submit**. The Export Channel Assignments link appears in the Contents pane.
- Choose **Export Channel Assignments > Assign content engines** to assign Content Engines to the export channel. (See the “[Assigning Content Engines to an Export Channel](#)” section on page 7-43.)
 - Choose **Export Channel Assignments > Assign device groups** to assign device groups to the export channel. (See the “[Assigning Device Groups to an Export Channel](#)” section on page 7-44.)

- Step 12** To disassociate an export channel from the current playlist, check the radio button next to an associated channel, and click the **Clear** button. Click **Submit** to ensure that the current playlist is no longer distributed with the selected export channel.

Assigning Content Engines to an Export Channel

You can assign Content Engines to export channels by using the Content Distribution Manager GUI.

To assign Content Engines to an export channel (which has been associated with export-type playlists), follow these steps:

- Step 1** Complete Step 1 through Step 11 in the previous section. (See the [“Selecting Export Channels for Set Top Box Playback”](#) section on page 7-41.)

If you have exited the Content tab, you can return to the Content Engine assignment window as follows:

- a. From the Content Distribution Manager GUI, choose **Services > Video > Programs**. The Programs window appears.
- b. Click the **Edit** icon next to the name of the export program to which you want to assign a Content Engine. The Modifying Program window appears.
- c. In the Contents pane, choose **Export Channel Assignments > Assign content engines**. The Content Engine Assignments for Export Channel window appears, displaying the unreserved and total disk space in the cdnfs partition. Content Engines can be added to the channel from more than one location. A list of the Content Engines in all locations to which they have been assigned appears.



Note If you have not selected an export channel, the **Export Channel Assignments** link will not appear in the Contents pane.

- Step 2** Click the **Assign** icon (blue cross mark) next to the name of the Content Engine that you wish to associate with the export channel. Alternatively, you can click the **Assign all Content Engines** icon in the taskbar to perform a bulk addition of all Content Engines in various locations to the export channel.
- Step 3** To add the Content Engines to the export channel, click **Submit**. A green tick mark appears next to the assigned Content Engine when you click submit.
- Step 4** You can perform the following additional tasks while in this window (Content Engine Assignments for Export Channel):

- Remove Content Engines

To remove Content Engines, click the **Unassign** icon (green tick mark) next to the name of the Content Engine that you want to remove from the export channel. Alternatively, you can click the **Unassign all Content Engines** icon in the taskbar to perform a bulk removal of all Content Engines from the selected export channel.

To remove the Content Engines from the export channel, click **Submit**. The list of Content Engines in the export channel shows a blue cross mark next to the name of the removed Content Engine.

- Designate a root Content Engine for the export channel

To designate a root Content Engine for an export channel, choose a Content Engine from the Root CE drop-down list in the taskbar. If there are no Content Engines assigned to the channel, no choices appear in the Root CE drop-down list.

Assigning Device Groups to an Export Channel

You can assign device groups to an export channel with which export-type programs have been previously associated. To assign device groups to an export channel, follow these steps:

- Step 1** Complete Step 1 through Step 11 in the [“Selecting Export Channels for Set Top Box Playback” section on page 7-41](#).

If you have exited the Content tab, you can return to the Device Group assignment window as follows:

- From the Content Distribution Manager GUI, choose **Services > Video > Programs**. The Programs window appears.
- Click the **Edit** icon next to the name of the export program to which you want to assign a device group. The Modifying Program window appears.
- In the Contents pane, choose **Export Channel Assignments > Assign device groups**. The Device Group Assignments for Export Channel window appears, displaying a list of the device groups with their minimum unreserved cdnfs disk space.

An asterisk (*) under the Note column denotes that there is no free space available to be assigned to this channel. A double asterisk (**) under the Note column denotes that Content Engines associated with this device group have been already assigned to this channel.



Note If you have not selected an export channel, the **Export Channel Assignments** link will not appear in the Contents pane.

- Step 2** Click the **Assign** icon (blue cross mark) next to the name of the device group that you wish to associate with the export channel. A green tick mark appears next to the assigned device group when you click **Submit**. Alternatively, you can click the **Assign all Device Groups** icon in the taskbar to perform a bulk addition of all device groups in various locations to the export channel.

- Step 3** To add the device groups to the export channel, click **Submit**.

- Step 4** You can perform the following additional tasks while in this window (Device Group Assignments for Export Channel):

- Remove device groups

To remove device groups, click the **Unassign** icon (green tick mark) next to the name of the device group that you want to remove from the export channel. Alternatively, you can click the **Unassign all Device Groups** icon in the taskbar to perform a bulk removal of all device groups from the selected export channel.

To remove the device groups from the export channel, click **Submit**. The list of device groups in the export channel shows a blue cross mark next to the name of the removed device group.

- Designate a root Content Engine for the export channel

To designate a root Content Engine for an export channel, choose a Content Engine from the Root CE drop-down list in the taskbar. If there are no Content Engines assigned to the channel, no choices appear in the Root CE drop-down list.

Viewing Export Playlist Settings

You can view the export-type playlists configured on a Content Engine. Set top boxes (STBs) directly connect to these Content Engines to download the playlists and play the media files added to the playlist. For the STBs to function properly, both the playlist file and the playlist content (media files and overlay images) must be available on the particular Content Engine. The STBs download the playlist file using the URL `http://<website_fqdn>playlists/<playlistName>.xml`.

Export playlists are displayed for a Content Engine only if Content Engines have been assigned to the export channel associated with the export playlist.

To view the export playlists configured on a Content Engine, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to view the associated export playlists. The Modifying Content Engines window appears.
- Step 3** In the Contents pane, choose **Applications > Set Top Box > Export Playlists**. The Export Playlists for Content Engine window appears.

The playlist name, number of media files, number of overlay images in the playlist, length of the playlists, and playlist file URLs are displayed. [Table 7-6](#) describes the fields in this window.

Table 7-6 Export Playlists for Content Engine—Field Descriptions

Field Name	Description
Playlist	Name of the export playlist associated with the export channel associated with this Content Engine
Titles	Number of media files added to the playlist
Length	Number of overlay images assigned to the media files in the playlist
Overlays	Total playtime of the playlist in hh:mm:ss
Playlist File URL	URL of the playlist file, which is of the syntax <code>http://<website_fqdn>/playlists/<playlistName>.xml</code>

Copying a Program

The program copy feature allows you to create a copy of an existing program and then modify a subset of attributes, which eliminates the need to reenter all the program settings each time you create programs with similar characteristics.

When you copy a program, a duplicate of the program is created and saved to the database. Any changes that you make to the new copy of the program do not affect the original program and vice versa. Note, however, that if multicast is configured, the multicast address and port cannot be copied. These parameters must be unique across the system. If a program address pool is configured, these parameters can be automatically selected by the system. A notification message alerts you when a new multicast address and port have been selected for a new copy of an original program.

To create a copy of an existing program, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Video > Programs**.
 - Step 2** Click the **Edit** icon next to the name of the program that you want to copy. The Modifying Program window appears.
 - Step 3** In the taskbar, click the **Copy Program** icon. You are prompted to confirm your decision. Click **OK**. The window refreshes, displaying *ProgramName_dup* in the Name field.
 - Step 4** Edit any program information that you want to change. (See the [“Defining a Program” section on page 7-11](#).)



Note You cannot change the program type.

- Step 5** To save your program definition, click **Submit**.
 - Step 6** Edit any of the other program properties found in the Contents pane, such as the program schedule, playlist, or device assignments. (See the appropriate sections in this chapter for configuration information.)
-

Managing Multicast Addresses for Programs

The Content Distribution Manager manages multicast addresses used for programs. Each Content Engine assigned to a program uses a multicast address for broadcasting. The Content Engine determines which multicast address to use based on the program data. A set of multicast addresses can be specified either in the Program API or by using the Content Distribution Manager. When a program requires a multicast address, the Content Distribution Manager associates one of the addresses with the program. Addresses are allocated for the life of a program. Programs can be configured with an autodelete feature that frees up program addresses automatically about 24 hours after a program schedule is complete.



Note For more information about the Program API, refer to the *Cisco ACNS Software API Guide, Release 5.x* publication.

Configuring Program Multicast Addresses

When you configure a pool of addresses for programs, ACNS software ensures that none of these addresses are currently in use by any multicast clouds. When the user requests a specific address or a set of addresses to be used for a program, ACNS software issues only those addresses that have not been

configured for multicast clouds and are not used by any existing programs. Users receive an error message if no multicast address is associated with the imported program file, if no addresses are available to be configured from the pool, or if the multicast pool has not been configured.

**Note**

If you create a program and specify a set of multicast addresses in the program file, it is not necessary to configure a pool of addresses using the Content Distribution Manager GUI. The pool of multicast addresses is required only if you have not specified the addresses to be used for multicasting in the program file.

To configure a pool of multicast addresses to be used for programs, follow these steps:

- Step 1** In the Content Distribution Manager GUI, choose **Services > Video > Program Addresses**. The Program Multicast Addresses window appears. (See [Figure 7-16](#).)

Figure 7-16 Program Multicast Addresses Window

The screenshot shows the Cisco Application and Content Networking System GUI. The breadcrumb navigation is **Services > Video > Program Addresses**. The window title is "Program Multicast Addresses". It features a "Multicast Address Pool" section with input fields for "Start Address:", "End Address:", and "TTL:" (with a unit of "hops"). A note indicates that asterisks (*) denote required fields. Below these fields are "Reset" and "Submit" buttons. A table lists existing multicast addresses with columns for "Multicast address", "User specified", and "true". The table shows two entries: "224.1.1.1" and "224.224.0.1". The bottom of the window shows pagination: "<< Page 1 >>" and "Showing 1-2 of 2 Multicast Addresses".

- Step 2** In the Start Address field, specify the first multicast IP address in the pool.
- Step 3** In the End Address field, specify the last multicast IP address in the pool.
- Step 4** In the TTL field, specify the Time To Live (number of hops) for all addresses configured in the pool.
- Step 5** To save your settings, click **Submit**. The multicast addresses are included in the database record for the program when it requests an address.

The list of multicast addresses that have been currently configured for specific programs is displayed in the Multicast Address column. The User Specified column displays *true* if the user has already specified the particular address for a program.



PART 3

Configuring ACNS Software Applications





Configuring Caching Services

This chapter describes how to configure conventional caching services (HTTP, FTP [FTP-over-HTTP caching and native FTP caching], HTTPS, and DNS caching) for centrally managed Content Engines. It also describes how to configure the TFTP gateway, and the Internet Cache Protocol (ICP) for the Content Engine and device groups. This chapter includes the following sections:

- [Content Engine Proxy Caching Overview, page 8-1](#)
- [Configuring HTTP Settings, page 8-2](#)
- [Configuring HTTPS Settings, page 8-27](#)
- [Configuring HTTP and HTTPS Outgoing Proxy Exclusion Settings, page 8-48](#)
- [Configuring FTP Settings, page 8-50](#)
- [Configuring TFTP Settings, page 8-65](#)
- [Configuring DNS Caching Name Service for WCCP Transparent Caching, page 8-71](#)
- [Configuring the DNS Server for HTTP Proxy Caching, page 8-78](#)
- [Configuring ICP Settings, page 8-80](#)

Content Engine Proxy Caching Overview

Caching is the ability to store web objects for later retrieval. Caching services are generally associated with the type of routing method being used, such as direct proxy routing or transparent redirection.

Caching services that are associated with direct proxy routing are said to be deployed in *nontransparent mode*. Nontransparent mode caching is described as *nontransparent forward proxy caching*. Proxy caching in nontransparent mode in an ACNS network is handled by the Content Engine, which acts as a forward proxy server. In deployments that use direct proxy routing to route content requests to the Content Engine, the Content Engine acts as a network gateway device that retrieves content on behalf of web clients.

Caching services that are associated with transparent redirection routing are said to be deployed in *transparent mode*. Proxy caching in transparent mode in an ACNS network is handled by Content Engines that are configured for WCCP caching services in conjunction with a WCCP-enabled router or Layer 4 switch. Transparent mode caching is described as *transparent forward proxy caching* or *transparent reverse proxy caching*, depending on the placement of the Content Engine in the network. The Content Engine can serve as a transparent proxy server to web clients; in which case, the Content Engine is acting as a transparent *forward* proxy server. Or, the Content Engine can serve as a proxy for a web server or group of web servers (for example web servers in a server farm at a company headquarters); in which case, the Content Engine is acting as a transparent *reverse* proxy server.

Network Protocols and Caching

The interaction between a web browser (web client) and a web server uses the existing standard application-layer Internet protocols, such as HTTP, Hypertext Protocol Secure (HTTPS), and File Transfer Protocol (FTP). The Content Engine must support these standard protocols to receive, cache, and send web content. Support for HTTP, FTP, TFTP, and HTTPS protocols is included as part of the ACNS software product.

**Note**

Protocols commonly used for streaming content, such as RealNetworks RTSP protocol, are discussed in Chapter 9, “[Configuring Streaming Media Services](#).”

The Content Engine services any protocols for which it has been configured. If the Content Engine is not configured to support a certain protocol, the Content Engine returns an error. For example, if port 8080 is configured for HTTP and HTTPS, an FTP request coming to this port is rejected.

The Content Engine supports up to eight incoming ports each for FTP, HTTPS, HTTP, and RTSP nontransparent mode caching services. The incoming proxy ports can be the same ports that are used by transparent mode services. The incoming proxy ports can be changed without stopping any WCCP services operating on the Content Engine or on other Content Engines in the network.

TCP ports reserved for system or network services should not be used for proxy caching services (for example, DNS or FTP). In transparent mode, if more than eight ports for a protocol are required, the administrator can configure multiple dynamic WCCP services. Intercepted FTP, HTTP, and HTTPS requests addressed to other proxy servers (received on transparent mode ports) are serviced according to the **proxy-protocols transparent** command parameters.

Configuring HTTP Settings

You can configure, modify, and view HTTP settings for Content Engines and device groups by completing the following tasks:

- [Configuring HTTP Connection Settings, page 8-2](#)
- [Configuring HTTP Cache Settings, page 8-6](#)
- [Configuring HTTP Cache Freshness Settings, page 8-10](#)
- [Configuring Authenticated HTTP Cache Settings, page 8-13](#)
- [Configuring Advanced HTTP Cache Settings, page 8-19](#)
- [Adding or Modifying Additional HTTP Request Methods for the Content Engine, page 8-23](#)
- [Configuring HTTP Cache Vary User Agent, page 8-25](#)

Configuring HTTP Connection Settings

To configure HTTP connection settings, follow these steps:

Step 1	From the Content Distribution Manager GUI, choose Devices > Devices . The Devices window appears.
Step 2	Click the Edit icon next to the name of the Content Engine that you want to configure.

- Step 3** From the Contents pane, choose **Applications > Web > HTTP > HTTP Connections**. The HTTP Connection Settings for Content Engine window appears. (See [Figure 8-1](#).) [Table 8-1](#) describes the fields in this window.

Figure 8-1 HTTP Connection Settings Window

The screenshot shows the 'HTTP Connection Settings for Content Engine, CONTENTENGINE' window. The 'Contents' pane on the left is expanded to 'HTTP Connections'. The main area shows the following settings:

- Enable Incoming Proxy:** ☒ Space separated list of incoming HTTP ports:
- Preserve 407 headers:** ☒
- Enable Outgoing Proxy:** ☒
- Use Origin Server:** ☒
- Outgoing Connection Timeout:** (microseconds) (200-5000000)
- Outgoing Monitor Period:** (seconds) (10-300)

Below these settings is a table for 'Outgoing Proxy' configuration:

	Hostname	Port	Username	Password	Confirm Password	HTLM User Domain	Disable basic authentication
1 *	10.77.155.221	8090					☐
2	10.77.155.221	8080	admin	*****	*****	my_domain	☒
3							☐
4							☐
5							☐
6							☐

At the bottom right are 'Submit' and 'Cancel' buttons.

- Step 4** To accept incoming requests on ports in addition to port 80, check the **Enable Incoming Proxy** check box.
- Step 5** In the List of Incoming Ports field, enter the port numbers used by the proxy server or Content Engine to receive requests. Separate port numbers with a space.
- Step 6** To allow an outgoing proxy server or another Content Engine to receive HTTP cache miss request traffic, check the **Enable Outgoing Proxy** check box.
- Step 7** To preserve HTTP authentication header 407 by sending header 407 to the client when asking the client for Internet proxy authentication credentials, check the **Preserve 407 headers** check box.
- If you are connected to the Content Engine in transparent mode, and you preserve 407 error codes by checking this check box, your browser will not recognize the 407 error codes. By default, the Content Engine strips the hop-to-hop 407 (Proxy Authentication Required) error code sent by the Internet proxy.
- Step 8** To direct requests to the origin server specified in the user request if all outgoing proxy servers fail, check the **Use Origin Server** check box.
- Step 9** In the Outgoing Connection Timeout field, enter a value in microseconds for the timeout period to use when probing for outgoing proxy servers.

Step 10 In the Outgoing Monitor Period field, enter an interval (in seconds) for monitoring an outgoing proxy server. This is the time period over which the outgoing proxy servers are queried. If one of the outgoing proxy servers is unavailable, the polling mechanism waits for the connection timeout before polling the next server.

Step 11 Configure the outgoing proxy:

- a. Enter the host name or IP address for the outgoing proxy host names in the Hostname field under the Outgoing Proxy heading.

The first host name or IP address entered designates that outgoing proxy server as the primary server. You can configure up to eight backup proxy servers for the HTTP proxy failover feature. If the primary proxy server fails to respond to the HTTP connection request, the request is redirected to the next proxy server in the list until one of the proxy servers services the request.

- b. In the Port field, enter a port number that corresponds to the outgoing proxy host names from the previous step. (A port is required for the primary server entered in the Hostname field.)

Step 12 If a root Content Engine is configured to receive content through a proxy server, the acquirer operating on the root Content Engine must be authenticated by the proxy server before it can obtain content from the origin server. To configure outgoing proxy authentication for the acquirer, follow these steps:

- a. In the Username field, enter the name of the user to be authenticated. This username will be used for both NTLM and basic authentication.
- b. In the Password field, enter the password of the user. Reenter the same password in the Confirm Password field for confirmation. The password details will be encrypted on display.
- c. In the NTLM User Domain field, enter the NTLM server domain name to be used to authenticate user access.
- d. To disallow removal of NTLM headers for fallback to basic authentication, check the **Disable basic authentication** check box.

ACNS software supports multiple proxies; therefore, you can set the proxy authentication information for multiple proxies as well.

Step 13 To save the settings, click **Submit**.

A “Click Submit to Save” message appears in red next to the current settings when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The **Reset** button is visible only when you apply default or group settings to change the current device settings but have not yet submitted the changes.

Table 8-1 HTTP Connection Settings

GUI Parameter	Function	CLI Command or Manifest Attribute
Enable Incoming Proxy	Configures the Content Engine to accept incoming requests on ports in addition to port 80.	http proxy incoming
Space separated list of incoming HTTP ports	Port numbers used by the proxy server or Content Engine to receive requests. This number ranges from 1 to 65535. You can specify up to 8 ports.	http proxy incoming ports
Enable Outgoing Proxy	Configures an outgoing proxy server or another Content Engine to receive HTTP cache miss request traffic without using ICP or WCCP.	http proxy outgoing

Table 8-1 HTTP Connection Settings (continued)

GUI Parameter	Function	CLI Command or Manifest Attribute
Preserve 407 headers	Preserves 407 HTTP authentication headers. These headers indicate that the client must first authenticate itself with the proxy server to request cache miss traffic.	http proxy outgoing preserve-407
Use Origin Server	When enabled, handles requests using the origin server if all outgoing proxy servers fail. If disabled, the user receives an error message if all outgoing proxy servers fail.	http proxy outgoing origin-server
Outgoing Connection Timeout	Timeout period (in microseconds) to use when probing for outgoing proxy servers. The range is from 200 to 5000000. The default value is 300 microseconds.	http proxy outgoing connection-timeout <i>microsecs</i>
Outgoing Monitor Period	Specifies the interval in seconds at which to monitor one outgoing proxy server. If multiple outgoing proxy servers are configured, they are monitored sequentially. The default value is 200 seconds.	http proxy outgoing monitor <i>seconds</i>

Outgoing Proxy

Hostname	Configures multiple outgoing proxy servers using the host name or IP address for the outgoing proxy server.	http proxy outgoing host { <i>hostname</i> <i>ipaddress</i> }
Port	Port number corresponding to the outgoing proxy host names.	http proxy outgoing host { <i>hostname</i> <i>ipaddress</i> } <i>port</i> 1–65535

Acquirer Outgoing Proxy Authentication

Note These settings can also be defined in the manifest file. (See [Chapter 6, “Configuring the ACNS Network for Content Acquisition.”](#))

Username	Configures the name of the user to be authenticated.	Manifest file attribute: <i>user</i>
Password	Configures the user password.	<i>password</i>
Confirm Password	Confirms the user password.	
NTLM User Domain	Configures the NTLM server domain name to be used to authenticate user access.	<i>ntlmUserDomain</i>
Disable basic authentication	Disallows the removal of NTLM headers for fallback to the basic authentication method.	<i>disableBasicAuth</i>



Note For other operations that you can perform using the icons in the taskbar, see [Table 3-3, “Content Distribution Manager GUI Icons.”](#)

Configuring HTTP Cache Settings

The HTTP Cache Settings window allows you to configure caching parameters for the caching of HTTP requests.

To configure HTTP caching parameters, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure.
- Step 3** Choose **Applications > Web > HTTP > HTTP Caching**. The HTTP Caching for Content Engine window appears. (See Figure 8-2.) Table 8-2 describes the fields in this window.

Figure 8-2 HTTP Caching for Content Engine Window

The screenshot shows the Cisco Application and Content Networking System GUI. The left sidebar contains a tree view with categories like Contents, Applications, and Web. The main area displays the 'HTTP Caching' configuration for Content Engine 'CONTENTENGINE'. The configuration includes various checkboxes and input fields for settings such as L4 Switch Support, Cache Binaries with Cookies, and Max Object Size.

HTTP Caching	
Current applied settings from Content Engine, CONTENTENGINE	
L4 Switch Support:	☒
Cache Binaries with Cookies:	☐
Enable Cache Vary User Agent:	☐
L4 Switch Spoof Support:	☐
Enforce Max Object Size:	☐
Max Object Size:	204800 (KB) (1-204800)
URL Validation:	☒
Client No Cache:	Do not set
Enable Smart Range:	☒
Max Start:	16384 (Bytes) (0-2147483647)
Max Interval:	16384 (Bytes) (0-2147483647)
Default Bitrate:	10 (Kbps) (0-2000000)
Always Resolve Host:	☐

- Step 4** To enable redirection using a Layer 4 switch, such as the Cisco Content Services (CSS) switches, check the **L4 Switch Support** check box.
- Step 5** To enable stripping of the requesting IP address, check the **Anonymize Requests** check box. This option sets the HTTP anonymizer feature, which hides the identity of the sender.
- Step 6** To enable caching of objects that are served with cookies attached and no expiration information, check the **Cache Binaries with Cookies** check box. These requests are for binary content with Set-cookie headers attached.
- Step 7** To disable Nagle's algorithm for HTTP requests, check the **Enable Fast Response** check box. Use this option only if particular server applications require an immediate, timely response to frequent, small bursts of information from the client. Disabling Nagle's algorithm results in performance degradation, because it disables buffering of data.

Step 8 To enable the Content Engine to cache responses with the Vary: User-Agent header, check the **Enable Cache Vary User Agent** check box. (To configure cache vary user agent settings, see the “[Configuring HTTP Cache Vary User Agent](#)” section on page 8-25.)

Step 9 To have the Content Engine check for required content length in HTTP requests, check the **Strict Length Checking** check box.

Step 10 To enable IP spoofing on the Content Engine for Layer 4 redirected traffic from Layer 4-enabled switches, check the **L4 Switch Spoof Support** check box.

Enabling this feature makes the Content Engine spoof the client’s IP address as the source IP address for requests originating from the Content Engine to the origin server. Requests to the origin server from the Content Engine generally happen on a cache miss.

Step 11 If the request contains a host name in the HTTP request header that is not a fully qualified domain name (FQDN), and you do not want the Content Engine to rewrite the request header by using a FQDN instead of the host name, check the **Do not modify request host header** check box.

When a client request is intercepted by a Content Engine, the default behavior of the Content Engine is to query the DNS server for the IP address of the origin server and use the FQDN of the origin server when forwarding the HTTP request. If the origin server is configured with a modified host name and is not using a FQDN in its host header, the origin server will not be able to process the Content Engine request.

If you want the Content Engine to use a FQDN in the request host header, leave this check box unchecked.

Step 12 To have the Content Engine not cache objects that exceed the specified maximum size, check the **Enforce Max Object Size** check box. An object with a size above the configurable upper limit is not stored by the Content Engine.



Note When this check box is unchecked, the cache file system (cfs) size limit remains at 2096128 KB (2 GB) for all cacheable objects.

Step 13 In the Max Object Size field, specify the maximum object size in kilobytes.

Step 14 To have an object that is in the Content Engine cache validated each time the object is requested, check the **URL Validation** check box.

When this option is enabled, cached objects are tested with respect to the validity of the host header and the server IP address. If the host header and server IP address are still valid, the content is served from the cache. If these parameters fail the validity test, the object is not served from the cache and a cache miss occurs.

Step 15 To configure how no-cache requests should be managed, choose an option from the **Client No Cache** drop-down list.

- If you want the Content Engine to ignore the no-cache request, choose the **ignore** option.
- Choose the **revalidate** option if you want to revalidate the request with the origin server before serving a no-cache client request.
- If you do not want to set this option, choose the default **Do not set** option.

Step 16 To cache an entire HTTP response even if you issue a Range request (that is, a request for only portions of a requested content object) and the object is not in the cache, check the **Enable Smart Range** check box. The Max Start and Max Interval fields become available if you check this check box.

Step 17 In the Max Start field, specify the maximum starting offset (in bytes) in the client’s Range request of the object to be cached. The range is 0 to 2147483647. The default is 16384.

- Step 18** In the Max Interval field, specify the maximum interval (in bytes) between any two consecutive ranges in the Range request. The range is 0 to 2147483647. The default is 16384.
- Only if a Range request satisfies both conditions (offset and interval), and a cache miss occurs, does the proxy issue a normal request to the origin server, cache the response (if cacheable), and send a Range response to the client. If the response is not cacheable, then a full response is sent to the client.
- Step 19** To limit the maximum bit rate per session for large files delivered using the HTTP protocol, specify the default bit rate in kilobits per second (kbps) in the Default Bitrate field.
- Step 20** To always resolve the Host header to determine the origin server IP address, check the Always Resolve Host check box.
- Step 21** To save the settings, click **Submit**.
- A “Click Submit to Save” message appears in red next to the current settings when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The **Reset** button is visible only when you apply default or group settings to change the current device settings but have not yet submitted the changes.

Table 8-2 HTTP Cache Settings

GUI Parameter	Function	CLI Command
L4 Switch Support	Enables redirection using a Layer 4 switch, such as the Cisco Content Services (CSS) switches.	http l4-switch enable
Anonymize Requests	Enables stripping of the requesting IP address. This option sets the HTTP anonymizer feature, which hides the identity of the sender.	http anonymizer enable
Cache Binaries with Cookies	Enables caching of objects that are served with cookies attached and no expiration information. These requests are for binary content with Set-cookie headers attached.	http cache-cookies
Enable Fast Response	Disables Nagle’s algorithm for HTTP requests. Use this option only if particular server applications require an immediate, timely response to frequent, small bursts of information from the client. Disabling Nagle’s algorithm results in performance degradation, because it disables buffering of data.	http fast-response enable
Enable Cache Vary User Agent	Configures parameters to enable the Content Engine to cache responses with the Vary: User-Agent header.	http cache-vary-user-agent enable
Strict Length Checking	Content Engine checks for required content length in HTTP requests.	http strict-request-content-length-checking enable
L4 Switch Spoof Support	Enables IP spoofing on the Content Engine for Layer 4 redirected traffic from Layer 4-enabled switches.	http l4-switch spoof-client-ip enable

Table 8-2 HTTP Cache Settings (continued)

GUI Parameter	Function	CLI Command
Do not modify request host header	Stops the Content Engine from rewriting the host name if the request contains a host name in the HTTP request header that is not a fully qualified domain name (FQDN).	http request-header host unmodified
Enforce Max Object Size	Stops the Content Engine from caching objects that exceed the specified maximum size. Note When this check box is unchecked, the cache file system (cfs) size limit remains at 2096128 KB (2 GB) for all cacheable objects.	http object max-size <i>kbytes</i>
Max Object size	Maximum object size in kilobytes.	
URL Validation	Revalidate every object requested in the cache with the origin server.	http object url-validation enable
Client No Cache	Configures the manner in which no-cache requests should be managed.	http client-no-cache-request {ignore revalidate}
Enable Smart Range	Caches an entire HTTP response even if you issue a Range request (that is, a request for only portions of a requested content object) and the object is not in the cache.	http smart-range enable
Max Start	Maximum starting offset (in bytes) in the client's Range request of the object to be cached. The range is 0 to 2147483647. The default is 16384.	http smart-range max-start <i>offset</i>
Max Interval	Maximum interval (in bytes) between any two consecutive ranges in the Range request. The range is 0 to 2147483647. The default is 16384.	http smart-range max-start <i>offset</i> max-interval <i>interval</i>
Default Bitrate	Maximum pacing bitrate (in kbps) for large pre-positioned HTTP files. The range is 0 to 2000000. The default is 1500 kbps.	bitrate http default <i>kbps</i>
Always Resolve Host	Resolves the Host header to determine the origin server IP address.	http always-resolve-host

**Note**

For other operations that you can perform using the icons in the taskbar, see [Table 3-3, “Content Distribution Manager GUI Icons.”](#)

Configuring HTTP Cache Freshness Settings

You can configure HTTP cache object freshness parameters for Content Engines on the ACNS network using the HTTP Cache Freshness Settings window from the Content Distribution Manager GUI. These parameters can be configured for either directory listings (the files and subdirectories listed under a directory) or particular objects in the cache.

To configure HTTP freshness parameters, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure.
- Step 3** From the Contents pane, choose **Applications > Web > HTTP > HTTP Cache Freshness**. The HTTP Cache Freshness Settings window appears. (See [Figure 8-3](#).) [Table 8-4](#) describes the fields in this window.

Figure 8-3 HTTP Cache Freshness Settings Window

The screenshot shows the 'HTTP Cache Freshness Settings' window for Content Engine, CONTENTENGINE. The window is titled 'Cisco Application and Content Networking System - Microsoft Internet Explorer'. The left pane shows the 'Contents' tree with 'Web' > 'HTTP' > 'HTTP Cache Freshness' selected. The main pane displays the following settings:

HTTP Cache Freshness Settings	
Current applied settings from Content Engine, CONTENTENGINE	
Enable:	☒
Text Object Age Multiplier:	30
Binary Object Age Multiplier:	60
Max TTL Scale:	days
Max Text Object TTL:	5
Max Binary Object TTL:	3
Minimum TTL:	5 (Minutes)
Re-evaluate Request:	none
Enable If-Modified-Since:	☐
IMS Text Age Modifier:	50 (%)
IMS Binary Age Modifier:	50 (%)

Note: * - Required Field

Buttons: Submit, Cancel

- Step 4** To enable the HTTP cache freshness settings, check the **Enable** check box. This check box is checked by default.
- Step 5** Specify a value in the Text Object Age Multiplier field.
The age multiplier value enables the Content Engine to guess the life of a text object by multiplying the time since the object was last modified by a percentage to obtain an approximate expiration date. After this date, the object is considered stale, and subsequent results cause a fresh retrieval by the Content Engine. Valid values range from 0 to 100 percent. The default value is 30 percent.
- Step 6** Specify a value in the Binary Object Age Multiplier field.

The age multiplier value enables the Content Engine to guess the life of a binary object by multiplying the time since the object was last modified by a percentage to obtain an approximate expiration date. After this date, the object is considered stale, and subsequent results cause a fresh retrieval by the Content Engine. Valid values range from 0 to 100 percent. The default value is 60 percent.

- Step 7** Choose a scale (**seconds**, **hours**, **minutes**, or **days**) from the Maximum Time-to-Live (TTL) Scale drop-down list.

The TTL sets a ceiling on estimated expiration dates. If an object has an explicit expiration date, this takes precedence over the configured TTL.

- Step 8** Specify a value in the Max Text Object TTL field. [Table 8-3](#) provides valid values for the TTL field.

Table 8-3 Time To Live Range of Values for HTTP Freshness

Scale	Range
Days	1–1825
Hours	1–43800
Minutes	1–2628000
Seconds	1–157680000

- Step 9** Specify a value in the Max Binary Object TTL field.
See [Table 8-3](#) for a list of maximum values depending on the scale used.
- Step 10** Specify a value in the Minimum TTL field.
A minimum TTL is the minimum Time To Live for objects in the cache. The range of values is from 0 to 86400 minutes. The default value is 30 minutes.
- Step 11** Choose how reevaluation requests are to be handled from the Re-evaluate Request drop-down list.
- To apply these parameters to the directory listing (the files and subdirectories listed under a directory), choose **text**.
 - To apply these parameters to both objects and directory listing, choose **all**.
 - To ensure that a TTL is not applied and objects in the cache have no expiration dates, choose **none**.
- Step 12** To specify text and age modifiers in the next two steps, check the **Enable If-Modified-Since** check box.
- Step 13** Specify a percentage in the IMS Text Age Modifier (%) field. This value specifies the age percentage to serve a text object without revalidation.
- Step 14** Specify a percentage in the IMS Binary Age Modifier (%) field. This value specifies the age percentage to serve a binary object without revalidation.
- Step 15** To save the settings, click **Submit**.

A “Click Submit to Save” message appears in red next to the current settings when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The **Reset** button is visible only when you apply default or group settings to change the current device settings but have not yet submitted the changes.

Table 8-4 HTTP Cache Freshness Settings

GUI Parameter	Function	CLI Command
Enable	Enables the HTTP cache freshness settings.	—
Text Object Age Multiplier	Enables the Content Engine to guess the life of a text object by multiplying the time since the object was last modified by a percentage to obtain an approximate expiration date. After this date, the object is considered stale, and subsequent results cause a fresh retrieval by the Content Engine. Valid values range from 0 to 100 percent. The default value is 30 percent.	http age-multiplier text num
Binary Object Age Multiplier	Enables the Content Engine to guess the life of a binary object by multiplying the time since the object was last modified by a percentage to obtain an approximate expiration date. After this date, the object is considered stale, and subsequent results cause a fresh retrieval by the Content Engine. Valid values range from 0 to 100 percent. The default value is 60 percent.	http age-multiplier text num binary num
Max TTL Scale	Sets ceiling on estimated expiration dates. If an object has an explicit expiration date, this takes precedence over the configured TTL.	http max-ttl { days text textdays binary bindays hours text texthours binary binhours minutes text textminutes binary binminutes seconds text textseconds binary binseconds }
Max Text Object TTL	Value.	
Max Binary Object TTL	Value.	
Minimum TTL	Minimum Time To Live for objects in the cache. The range of values is from 0 to 86400 minutes. The default value is 30 minutes.	http min-ttl minutes
Re-evaluate Request	Configures how reevaluation requests are to be handled.	http reval-each-request { all none text }
Enable If-Modified-Since (IMS)	Enables IMS modifiers.	http serve-ims
IMS Text Age Modifier	Age percentage to serve a text object without revalidation.	http serve-ims text percentage
IMS Binary Age Modifier	age percentage to serve a binary object without revalidation.	http serve-ims text percentage binary percentage

**Note**

For other operations that you can perform using the icons in the taskbar, see [Table 3-3, “Content Distribution Manager GUI Icons.”](#)


Configuring Authenticated HTTP Cache Settings

The authenticated HTTP caching feature allows basic and NT LAN Manager (NTLM) authenticated content to be cached and served to more than one user, while maintaining security. If an authenticated object is cached, then subsequent requests for that object (from new users) require authentication. The cached object is revalidated with the origin server through the authorization header for the new user. If the user is not authorized, the server sends a 401 (Unauthorized) response. If the user is authorized and the object is not modified, the cached object is served to the client.

To configure authenticated HTTP caching parameters, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure.
- Step 3** From the Contents pane, choose **Applications > Web > HTTP > Authenticated HTTP Caching**. The Authenticated HTTP Caching for Content Engine window appears. (See [Figure 8-4.](#)) [Table 8-5](#) describes the fields in this window.

Figure 8-4 Authenticated HTTP Caching for Content Engine Window—Top



- Step 4** To enable use of authenticated HTTP cache settings for the Content Engine, check the **Enable** check box.
- Step 5** From the Cache Authenticated Content drop-down list, choose an authentication method to cache authenticated content.
- To cache web objects authenticated with any scheme, choose **all**.
 - To cache web objects authenticated with basic or NTLM, choose **basic** or **ntlm** from the drop-down list.
 - To remove the authentication method, choose **Do Not Set**. This choice is the equivalent of the **no http cache-authenticated basic** command and the **no http cache-authenticated ntlm** command.
- Step 6** To strip NTLM headers from the authenticated content, check the **Strip NTLM Headers** check box.
- When checked, NTLM headers are stripped from the authenticated content being sent from the Content Engine to the client browser. This feature enables browsers to fall back to basic authentication. Browsers can authenticate against a basic-style authentication challenge posed by Microsoft Internet Information Service (IIS) servers.
- When left unchecked, the NTLM header in the response is preserved, and the browser will authenticate against the NTLM challenge.
- Step 7** To enable the maximum number of entries retained in the authentication cache to be configured on the Content Engine, check the **Set Max Cache Authentication Entries** check box.
- Step 8** In the Max Cached Authenticated Entries field, enter a value for the maximum number of entries that is maintained in the cache. The default value is 32000 entries.
- Step 9** To enable the maximum number of entries retained in the authentication group cache to be configured on the Content Engine, check the **Set Max Group Cached Authenticated Entries** check box.

- Step 10** In the Max Group Cached Authenticated Entries field, specify the maximum number of entries retained in the authentication group cache on the Content Engine. The range is 500 to 12000. This maximum is subject to physical resources on the Content Engine.
- Step 11** In the Time Authenticated Entries are cached field, enter the length of time (in minutes) between the last access and the removal of the entry from the cache on the Content Engine. The range is from 1 to 1440 minutes, and the default value is 480 minutes.
- The authentication record entry is stored in the cache for the specified period. Once this period is exceeded, subsequent access to restricted content requires reauthentication through NTLM, RADIUS, or LDAP servers. (See the [“Specifying a Reauthentication Interval”](#) section on page 8-19.)
- Step 12** To enable the absolute timeout for entries in the authentication cache to be configured on the Content Engine to help reduce the possibility of individuals gaining access by using previously authenticated browsers, check the **Set Authentication Cache TTL** check box. By default, this check box is unchecked, which means that there is no TTL timeout in effect. This means that there will be no check to time out an authentication cache entry based on its creation time relative to a TTL value.
- Step 13** In the Authentication Cache TTL (max valid time) field, specify an absolute value for the maximum time an authentication cache entry is considered valid after its creation. The TTL timeout is specified in minutes. The minimum is 1 minute; the maximum is 1440 minutes (24 hours).
- When the absolute TTL period has expired, the client browser is forced to re-authenticate itself and the user must enter valid credentials. This ability to specify an absolute TTL timeout, adds security to the inactivity timeout mechanism for content that is served through the Content Engine.
- The timeout value in the Time Authenticated Entries are cached field is not affected by the absolute TTL; they are independent of each other. If both the Time Authenticated Entries are cached timeout and the absolute TTL timeout are configured on the Content Engine, the authentication cache entry times out depending on which timeout occurs first for that entry.
- Step 14** To send an authentication message in the header, choose either the **401** (Unauthorized) or **407** (Proxy Authentication Required) header type message from the **Re-Authenticate Header type** drop-down list. If you do not want to send a message, choose **Do not set**. (See [Figure 8-5](#).) [Table 8-5](#) describes the fields in this window.
- These authentication headers are used to query users for credentials when an entry is not found in the authentication cache, requiring server lookup.

Figure 8-5 Authenticated HTTP Caching for Content Engine Window—Bottom

Authenticated HTTP Caching for Content Engine, CONTENTENGINE

Max Group Cached Authenticated Entries: 512 (500-12000)

Time Authenticated Entries are cached: 2 (mins) (1-1440)

Set Authentication Cache TTL: ☒

Authentication Cache TTL (max valid time): 2 (mins) (1-1440)

Re-Authenticate Header type: 401

Basic Realm String: "Cache"

Append Headers: ☒

Proxy Header Server:

Use Via Headers: ☐

Use Custom String Via Headers: "cust"

WWW Authentication Server:

Use X-FWD-Headers: ☒

Use X-FWD-Headers with Multiple Ip Address: ☐

Use Host Header: ☒

Note: * - Required Field

Submit Cancel

- Step 15** To modify the HTTP authentication realm string, enter the desired string in the Basic Realm String field. Cisco ACNS software allows you to configure the realm that is displayed in the authentication popup window when you connect to the Content Engine. The default realm is "Cisco Content Engine." This setting allows you to reconfigure the realm to reflect the nature of the device.
- Step 16** If you want HTTP headers to be appended by downstream Content Engines, check the **Append Headers** check box. The fields that follow are available only when you have checked the **Append Headers** check box.
- Step 17** In the Proxy Header Server field, enter the host name or IP address of the server that is configured to receive the proxy authorization header. When this setting is configured, the Content Engine appends the proxy authorization header to the request when it sends the request to an upstream Content Engine proxy or server that is configured to receive the proxy authorization header. Proxy authorization headers are hop-by-hop headers, and so, the default behavior is to not append proxy authorization headers to upstream proxies or servers. Configuring this setting allows a proxy authorization header to be sent to a specific server that is configured to receive this type of header.
- Step 18** To include a "Via" header in responses and replies, check the **Use Via Headers** check box.
- Step 19** To add a custom string in the Via header description, enter up to 99 characters in the Use Custom String Via Headers field.
- Step 20** In the WWW Authentication Server field, enter the host name or IP address of the server that is configured to receive the WWW authorization header.

For example, you might have the following set of conditions:

- WCCP is enabled for the Content Engine.
- Proxy authentication (RADIUS, LDAP, TACAS, or NTLM) is enabled for the Content Engine.
- You configured **base2** as the host name in the WWW Authentication Server field.

When a client sends a request for **http://base2/** the Content Engine authenticates the request, and while sending the request to the **base2** server, the Content Engine includes the WWW authorization header received from the client. (The **base2** server has been configured to receive the WWW authorization header.) If the request from the client is for a server other than **base2**, the Content Engine authenticates the request and sends the request to the server without including the WWW authorization header received from the client.

This setting is similar to the append proxy authorization header setting (configured in the Proxy Header Server field).

Step 21 To notify the web server of the client IP address through the X-Forwarded-For header, check the **Use X-FWD-Headers** check box.

Step 22 To enable support for appending multiple IP addresses to the X-Forwarded-For header, check the **Use X-FWD-Headers with Multiple IP Address** check box.

If you specify this command, and an X-Forwarded-For header already exists, then the Content Engine IP address is appended to the existing client IP address in the header, separated by a comma.

Step 23 To append host headers to HTTP/1.0 requests that do not possess these headers, check the **Use Host Header** check box.

Step 24 To save the settings, click **Submit**.

A “Click Submit to Save” message appears in red next to the Current Settings when there are pending changes to be saved after you have applied the default or device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The **Reset** button is visible only when you apply default or group settings to change the current device settings but the settings have not yet been submitted.

Table 8-5 Authenticated HTTP Cache Settings

GUI Parameter	Function	CLI Command
Enable	Enables use of authenticated HTTP cache GUI settings for the Content Engine.	—
Cache Authenticated Content	Authentication method (all, basic, ntlm, Do Not Set) to cache authenticated content.	[no] http cache-authenticated {all basic ntlm}
Strip NTLM Headers	Strips NTLM headers from the authenticated content.	http authenticate-strip-ntlm
Set Max Cached Authenticated Entries	Enables the maximum number of entries retained in the authentication cache to be configured on the Content Engine.	—
Max Cached Authenticated Entries	Maximum number of entries that is maintained in the cache. The default value is 32000 entries.	http authentication cache max-entries entries
Set Max Group Cached Authenticated Entries	Enables the maximum number of entries retained in the authentication group cache to be configured on the Content Engine.	—

Table 8-5 Authenticated HTTP Cache Settings (continued)

GUI Parameter	Function	CLI Command
Max Group Cached Authenticated Entries	Maximum number of entries retained in the authentication group cache on the Content Engine. The range is 500 to 12000. This value is subject to physical resources on the Content Engine.	http authentication cache max-group-entries <i>entries</i>
Time Authenticated Entries are cached	Length of time (in minutes) between the last access and the removal of the entry from the cache on the Content Engine. The range is 1 to 1440. The default is 480 minutes.	http authentication cache timeout <i>minutes</i>
Set Authentication Cache TTL	Enables the absolute timeout for entries in the authentication cache to be configured on the Content Engine.	—
Authentication Cache TTL (max valid time)	Maximum time an authentication cache entry is considered valid after its creation. The TTL timeout is specified in minutes. The minimum is 1 minute; the maximum is 1440 minutes (24 hours).	http authentication cache ttl <i>minutes</i>
Re-Authenticate Header type	Sends an authentication message in the header.	http authentication header {401 407}
Append Response Headers	Directs downstream Content Engines to append HTTP headers.	—
Proxy Header Server	Host name or IP address of the server that is configured to receive the proxy authorization header.	http append proxy-auth-header <i>{ipaddress hostname}</i>
Use Via Headers	Includes a “Via” header in responses and replies.	http append via-header
Use Custom String Via Headers	Adds a custom string in the Via header description.	http append via-header custom <i>string</i>
WWW Authentication Server	Host name or IP address of the server that is configured to receive the WWW authorization header.	http append www-auth-header <i>{hostname ipaddress}</i>
Use X-FWD-Headers	Notifies the web server of the client’s IP address through the X-Forwarded-For header.	http append x-forwarded-for-header
Use X-FWD-Headers with Multiple IP Address	Appends multiple IP addresses to the X-Forwarded-For header.	http append x-forwarded-for-header multiple-ip-address
Use Host Header	Appends host headers to HTTP/1.0 requests.	http append host-header



Note For other operations that you can perform using the icons in the taskbar, see [Table 3-3, “Content Distribution Manager GUI Icons.”](#)

Specifying a Reauthentication Interval

In the ACNS 5.1 software, an inactivity timer was used to determine when the client browser was prompted to reenter authentication credentials after being initially authenticated. This inactivity timer is configured with the `http authentication cache timeout` global configuration command. By default, the inactivity timeout period was 8 hours. This meant that as long as someone continued to use the client browser after the legitimate authenticated user was initially authenticated, the client was not forced to reenter that person's authentication credentials.

In the ACNS 5.2.1 software and later releases, you can specify an absolute TTL for HTTP authentication cache entries for increased security in a shared workstation environment. This ability to specify an absolute TTL timeout adds security to the inactivity timeout mechanism for content that is served through the Content Engine. When the absolute TTL period has expired, the client browser is forced to reauthenticate itself, and the user must enter valid credentials.


A security vulnerability exists in a shared workstation environment that is using WCCP-enabled router redirection mode with any authentication method, or proxy redirection mode and the NTLM authentication method. In these cases, the Content Engine uses the client's IP address as the index into the authentication record kept in the authentication cache, and the Content Engine can therefore authenticate users who have not presented valid credentials of their own if a different user using the same workstation has previously presented valid credentials that are cached in the authentication cache. To provide additional security in this situation, you can configure the absolute TTL for HTTP authentication cache entries; this will specify the absolute time during which an authentication cache entry is valid in the cache. If a cache lookup occurs on an entry and its configured TTL time is exceeded, then the entry is deleted and the Content Engine will query the user for credentials.

To support this feature, the `ttl` option was added to the `http authentication cache` global configuration command.

Configuring Advanced HTTP Cache Settings

To configure advanced HTTP cache settings, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
 - Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure.
 - Step 3** From the Contents pane, choose **Applications > Web > HTTP > Advanced HTTP Caching**. The Advanced HTTP Caching for Content Engine window appears. (See [Figure 8-6](#) and [Figure 8-7](#).) [Table 8-6](#) describes the fields in these windows.

Figure 8-6 Advanced HTTP Caching for Content Engine Window—Top

- Step 4** Under the Advanced HTTP Caching heading, check the **Enable Cache on Abort** check box to continue to cache an object even though the client has aborted the request.
- Step 5** To cache an object if the number of kilobytes remaining to download from the server is less than the maximum threshold value, check the **Use Max Threshold** check box.
- Step 6** In the Abort Max Threshold field, enter a value in kilobytes for the maximum threshold. The default value is 256 kilobytes.
- Step 7** To cache an object if the number of kilobytes remaining to download from the server is greater than the minimum threshold value, check the **Use Min Threshold** check box.
- Step 8** In the Abort Min Threshold field, enter a value in kilobytes for the minimum threshold. The default value is 32 kilobytes.
- Step 9** To cache an object if the percentage of the object already downloaded is greater than the percentage threshold value entered, check the **Use Percent Threshold** check box.
- Step 10** In the Abort Percent Threshold field, enter a value for the percentage threshold. The default value is 80 percent.
- Step 11** Under the Cluster Settings heading, check the **Enable Clustering** check box to allow a Content Engine in a Content Engine farm to query and obtain cache objects from other Content Engines in the cluster.
- Step 12** In the Cluster Heal Port field, enter the port number over which requests from the healing Content Engine are sent to other Content Engines in the cluster, if you are not using the default HTTP port (14333).

The Content Engine that responds to queries from another Content Engine in a Content Engine cluster is called a healing server. The Content Engine that requests a cache object from the cluster is called a healing client.

- Step 13** In the Cluster HTTP Port field, enter the HTTP port number over which requests from the healing Content Engine are sent to other Content Engines in the cluster. The default port is 80.
- Step 14** In the Cluster Max Delay field, enter the maximum time in seconds that a healing Content Engine waits for a healing Content Engine response.
- Step 15** In the Cluster Misses field, enter the maximum number of misses that the healing Content Engine can receive from the cluster after the last healing mode hit response.
- Step 16** Under the Persistent Connections Settings heading, check the **Enable Persistent Connections** check box to allow persistent connections on the Content Engine. (See [Figure 8-7](#).)

Figure 8-7 Advanced HTTP Caching for Content Engine Window—Bottom

The screenshot shows the 'Advanced HTTP Caching for Content Engine, CONTENTENGINE' configuration window. The left sidebar contains a tree view with categories like Contents, Applications, Web, HTTP, and others. The main area is divided into several sections:

- Thresholds:** Includes fields for 'Use Max Threshold', 'Use Min Threshold', and 'Use Percent Threshold', each with an 'Abort' counterpart and a unit of (kB) (1-99999).
- Cluster Settings:** Includes 'Enable Clustering' (checkbox), 'Cluster Heal Port' (14999), 'Cluster HTTP Port' (80), 'Cluster Max Delay' (0 seconds), and 'Cluster Misses' (0).
- Persistent Connection Settings:** Includes 'Enable Persistent Connections' (checkbox), 'Persistent Connection Type' (dropdown set to 'all'), and 'Persistent Connection Timeout' (600 seconds).
- TCP KeepAlive Settings:** Includes 'Enable TCP Keepalive' (checkbox, checked).

A note at the bottom states: 'Note: * - Required Field'. The window has 'Submit' and 'Cancel' buttons at the bottom right.

- Step 17** From the Persistent Connections Type drop-down list, choose a persistent connection type. A persistent connection can be set for client-only, server-only, or all connections on the Content Engine.
- Step 18** In the Persistent Connection Timeout field, enter the number of seconds that the Content Engine should keep an idle persistent connection open before it closes the connection. The default value is 600 seconds.
- Step 19** Under the TCP KeepAlive Settings heading, check the **Enable TCP Keepalive** check box to force the Content Engine to send keepalive probes, which allows persistent connections on the Content Engine.

For HTTP connections, the default timeout value is 5 minutes. When no keepalive messages are sent by the Content Engine to the clients and to the edge Content Engines, the connection is closed. When the HTTP TCP keepalive feature is enabled, the Content Engine continues to send TCP keepalives on idle TCP connections using keepalive configuration parameters, such as TCP keepalive timeout, TCP keepalive probe count, and TCP keepalive probe interval.

Step 20 To save the settings, click **Submit**.

A “Click Submit to Save” message appears in red next to the current settings when there are pending changes to be saved. You can also revert to the previously configured window settings by clicking **Reset**. The **Reset** button is visible only when you apply default or group settings to change the current device settings but the settings have not yet been submitted.

Table 8-6 Advanced HTTP Cache Settings

GUI Parameter	Function	CLI Command
Advanced HTTP Caching		
Enable Cache On Abort	Continues to cache an object even though the client has aborted the request.	http cache-on-abort enable
Use Max Threshold	Caches an object if the number of kilobytes remaining to download from the server is less than the maximum threshold value.	http cache-on-abort max-threshold max_thresh
Abort Max Threshold	Maximum threshold in kilobytes.	
Use Min Threshold	Caches an object if the number of kilobytes remaining to download from the server is greater than the minimum threshold value.	http cache-on-abort min-threshold min_thresh
Abort Min Threshold	Minimum threshold in kilobytes.	
Use Percent Threshold	Caches an object if the percentage of the object already downloaded is greater than the percentage threshold value entered.	http cache-on-abort percent percentthresh
Abort Percent Threshold	Percentage threshold value.	
Cluster Settings		
Enable Clustering	Allows a Content Engine in a Content Engine farm to query and obtain cache objects from other Content Engines in the cluster.	—
Cluster Heal Port	Port number over which requests from the healing Content Engine are sent to other Content Engines in the cluster, if you are not using the default HTTP port (14333).	http cluster heal-port num
Cluster HTTP Port	HTTP port number over which requests from the healing Content Engine are sent to other Content Engines in the cluster.	http cluster http-port num
Cluster Max Delay	Maximum time in seconds that a healing Content Engine waits for a healing Content Engine response.	http cluster max-delay seconds
Cluster Misses	Maximum number of misses that the healing Content Engine can receive from the cluster after the last healing mode hit response.	http cluster misses num

Table 8-6 *Advanced HTTP Cache Settings (continued)*

GUI Parameter	Function	CLI Command
Persistent Connection Settings		
Enable Persistent Connections	Allows persistent connections on the Content Engine.	—
Persistent Connection Type	Sets persistent connection type (client-only, server-only, or all) for connections on the Content Engine.	http persistent-connections { all server-only client-only }
Persistent Connection Timeout	Number of seconds that the Content Engine should keep an idle persistent connection open before it closes the connection.	http persistent-connections timeout seconds
TCP KeepAlive Settings		
Enable TCP Keepalive	Forces the Content Engine to send keepalive probes.	http tcp-keepalive enable



Note For other operations that you can perform using the icons in the taskbar, see [Table 3-3, “Content Distribution Manager GUI Icons.”](#)

Adding or Modifying Additional HTTP Request Methods for the Content Engine

Content Engines accept or reject an HTTP request depending on whether the HTTP request method is supported. HTTP 1.1 request methods (for example, GET, HEAD, or POST) are supported by default. Nonstandard request methods, such as Web-Based Distributed Authoring and Versioning (WebDAV) are not.

When the Content Engine receives an HTTP request from a client using an unsupported method, ACNS software adds the method to the list of unsupported methods and sends an error to the client. You can use the Creating New HTTP Method for Content Engine window to add a method to the list of methods already supported by the Content Engine.

To add a new HTTP request method, follow these steps:

-
- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
 - Step 2** Click the **Edit** icon next to the name of the Content Engine for which you want to add an HTTP request method. The Device Home for Content Engine window appears.
 - Step 3** In the Contents pane, choose **Applications > Web > HTTP > HTTP Method**. The HTTP Methods for Content Engine window appears, displaying the method names and device group sources. The Device Group Source column specifies the name of the device group whose settings have been applied to the Content Engine.
 - Step 4** The Aggregate Settings **Yes** radio button is selected. This means that the HTTP methods are displayed for the Content Engine as well as for any associated device groups. Click the **No** radio button to apply the HTTP methods to the Content Engine only.



Note When the Aggregate Settings **Yes** radio button is selected, the HTTP methods that have been configured for device groups to which the Content Engine belongs cannot be modified or deleted; the settings are read-only. You can, however, modify the HTTP methods configured for the Content Engine. If the Aggregate Settings **No** radio button is selected, you can view or modify the HTTP methods for the Content Engine only, and not for the device groups with which it is associated.

- Step 5** Click the **Create New HTTP Method** icon in the taskbar. The Creating new HTTP method for Content Engine window appears.
 - Step 6** In the Method Name field, enter the name of the HTTP request method to be added to the list of supported methods.
 - Step 7** To save your settings, click **Submit**. The HTTP Methods for Content Engine window appears, displaying the newly added HTTP method.
-


To add a new HTTP request method using the CLI, use the **http add-method** global configuration command.

Configuring HTTP Cache Vary User Agent

To configure the parameters for caching responses with the Vary header, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices** (or **Devices > Device Groups**).
- Step 2** Click the **Edit** icon next to the Content Engine or Device Group that you want to configure.
- Step 3** In the Contents pane, choose **Applications > Web > HTTP > HTTP Cache User Agent**. The HTTP User Agents for Content Engine (or Device Group) window appears.
- Step 4** Click the **Create New HTTP User Agent** icon in the task bar. The Creating New HTTP User Agent window appears. (See Figure 8-8.) Table 8-7 describes the fields in this window and provides the corresponding CLI commands.

Figure 8-8 Creating New HTTP User Agent Window



- Step 5** In the SubString field, enter a string to be used to match while searching for an alternative user-agent. The substring configuration must be unique. The Content Distribution Manager GUI returns an error message for duplicate substring configurations.
- Step 6** In the CacheString field, enter a string that is to form the objects key.
The following characters are invalid: acute accent (‘), double quote (“), pipe (|), question mark (?), and white space.
- Step 7** To save the configuration, click **Submit**. The window refreshes, and the HTTP User Agents window appears with the new item listed. You can create up to 32 entries.

- Step 8** Enable the Content Engine to use these settings by checking the **Enable Cache Vary User Agent** check box in the HTTP Caching for Content Engine window. (See the [“Configuring HTTP Cache Settings” section on page 8-6.](#))

Table 8-7 HTTP Cache User Agent Settings

GUI Parameter	Function	CLI Command
SubString	Configures the string to be used to match while searching for an alternative user-agent. Maximum characters is 16.	http cache-vary-user-agent sub-string <i>string</i> cache-string <i>string</i>
CacheString	Configures the cache-string that forms the objects key. Maximum characters is 16.	

Configuring HTTP Destination Port Restrictions

Cisco ACNS software allows you to configure HTTP destination port access restrictions from the Content Distribution Manager GUI.

To configure HTTP destination port restrictions, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure the destination port restrictions. The Device Home window for the chosen Content Engine appears.
- Step 3** In the Contents pane, choose **Web > HTTP > HTTP Destination Port Restrictions**. The HTTP Destination Port Restriction Settings window appears. By default the window shows the current settings for the chosen device.
- Step 4** To allow access to all ports, check the **Allow All Ports** check box.



Note By default, the Allow All Ports check box is checked, and this allows traffic to port range 80–87 and to any port above 1024.

- Step 5** In the Allow Port Range List field, enter the port numbers or port ranges to which you want to allow traffic. You can enter a maximum of eight port numbers or port ranges separated by a space.
- Step 6** To deny traffic to any of the ports, check the **Deny All Ports** check box.



Note Port numbers below port 1024 are denied by default, except for port range 80–87.

- Step 7** To allow traffic to these ports, uncheck the **Deny Port Range 1-79** check box.
- Step 8** To allow traffic to these ports, uncheck the **Deny Port Range 80-1024** check box.
- Step 9** In the Other Deny Port Range List field, enter the port numbers or port ranges of ports to which you want to deny traffic. You can enter up to eight port numbers or port ranges separated by a space.

**Note**

If you have both the Deny Port Range 1-79 and the Deny Port Range 80-1024 checked, you can enter only six more port numbers or port ranges in the Other Deny Port Range List field. When either the Deny Port Range 1-79 or the Deny Port Range 80-1024 is checked, the Other Deny Port Range List field accepts a maximum of seven ports instead of eight.

- Step 10** To save the settings, click **Submit**. A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking **Reset**. The Reset button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

To configure HTTP destination port restrictions from the CLI, use the **http destination-port** global configuration command.

Configuring HTTPS Settings

You can configure, modify, and view HTTPS settings for Content Engines and device groups by completing the following tasks:

- [Configuring HTTPS Certificates, page 8-27](#)
- [Configuring HTTPS Certificate Groups, page 8-31](#)
- [Configuring HTTPS Keys, page 8-34](#)
- [Configuring HTTPS Proxy Settings, page 8-37](#)
- [Configuring HTTPS Servers, page 8-42](#)
- [Configuring HTTP and HTTPS Outgoing Proxy Exclusion Settings, page 8-48](#)

Configuring HTTPS Certificates

A digital certificate is a credential that allows the Content Engine to be presented to an HTTPS client as the origin HTTPS server. You can create certificates, import a certificate from external sources, or remove existing certificates using the Content Distribution Manager GUI.

You can assign a certificate and associate a key with the HTTPS server assuming that you have configured the Content Engine using the Creating New HTTPS Server for Content Engine window or using the **https server** global configuration command. The Content Engine presents the certificate to HTTPS clients that make requests to the HTTPS server.

The Content Engine accepts certificates in Privacy-Enhanced Mail (PEM) format, which is used by Apache servers, and Public-Key Cryptography Standards (PKCS) #12 format, which is used by Microsoft Internet Information Server (IIS).

The Content Engine uses PEM format internally, and automatically converts certificates in PKCS #12 format to PEM format. If you need to use a certificate in a different format, first convert it to one of these supported formats.

To configure HTTPS certificates for the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure HTTPS certificates. The Device Home for Content Engine window appears.
- Step 3** In the Contents pane, choose **Applications > Web > HTTPS > Certificates**. The HTTPS Certificates for Content Engine window appears. The certificates created for device groups can only be viewed; they cannot be modified.
- Step 4** To configure an HTTPS digital certificate, click the **Create New HTTPS Certificate** icon in the taskbar. [Table 8-8](#) describes the fields that appear in this window and provides the corresponding CLI global configuration commands.
- Step 5** In the Certificate Name field, specify the name of the certificate. Certificate names must consist solely of alphanumeric, underscore, and hyphen characters. Certificate names can begin with a numeric character. Certificate names can have a maximum length of 64 characters.

You need to create a certificate before importing its value from an external source associating it with an HTTPS server. Only certificate names are stored in the CMS database; the actual certificates are stored on the Content Engine.



Note Once you have created a certificate, you cannot modify the name. You need to delete the existing certificate and create a new one.

- Step 6** In the Certificate URL field, specify the external location from which the certificate is to be imported. The URL must be a valid URL using HTTP, FTP, or HTTPS. The certificate value is imported after you specify the URL and submit changes.

A certificate name without imported values can be listed in the Content Distribution Manager GUI; however, only certificates with valid values can be associated with an HTTPS server or added to a certificate group. Two different certificates can be imported from the URL and two different HTTPS servers can be associated with the same certificate.

After a certificate is imported from the Content Distribution Manager GUI, a new icon (**Re-import**) appears in the taskbar. The **Re-import** icon can be used to reimport the certificates if the import operation did not succeed the first time.
- Step 7** In the Username field, specify the username required to access the external source from which the certificate is being imported. An entry is required in this field only if the external source is password-protected.
- Step 8** In the Password field, specify the password used to authenticate users who want to gain access to the external source from which the certificate is being imported. An entry is required in this field only if the external source is password-protected. Reenter the password in the Confirm Password field.
- Step 9** To save the settings, click **Submit**.

Table 8-8 **HTTPS Certificate Settings**

GUI Parameter	Function	CLI Command
Certificate Name	Name of the certificate.	https server name cert <i>cert_name</i>
Certificate URL	External location from which the certificate is to be imported. The URL must be a valid URL using HTTP, FTP, or HTTPS.	https server name host <i>ipaddress</i> or <i>FQDN</i>
Username	Username required to access the external source from which the certificate is being imported.	—
Password	Password used to authenticate users who want to gain access to the external source.	—
Confirm Password	Verifies the password entry.	—



Note For other operations that you can perform using the icons in the taskbar, see [Table 3-3, “Content Distribution Manager GUI Icons.”](#)

Filtering HTTPS Certificates for the Content Engine

You can use the Content Distribution Manager GUI to filter and display a subset of HTTPS certificates by certificate name and URL.

To set the filter criteria to display selected HTTPS certificates, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to filter HTTPS certificates. The Device Home for Content Engine window appears with the Contents pane on the left.
- Step 3** In the Contents pane, choose **Applications > Web > HTTPS > Certificates**. The HTTPS Certificates for Content Engine window appears. The certificates created for device groups can only be viewed; they cannot be modified.
- Step 4** Click the **Filter Table** icon in the taskbar. The Filtering HTTPS Certificates for Content Engine window appears.
- Step 5** Choose a filter operator from the Certificate Name drop-down list and enter the name of the certificate that must be matched. [Table 8-9](#) describes the filter operators for filtering HTTPS certificates.
- Step 6** Choose a filter operator from the Certificate URL drop-down list and enter the certificate location that must be matched.
- Step 7** Click **Submit**. The HTTPS Certificates for Content Engine window reappears, displaying a list of all certificates that match the specified filter criteria.

Table 8-9 Filtering HTTPS Certificates

Filter Operators	Description
Like	Include HTTPS certificates with a setting that is similar to the string entered. Use the percent sign (%) wildcard characters to match zero or more characters, and use the underscore (_) wildcard character to match exactly one character. If the string contains special characters such as ^, (,), or ', use the escape character, a backslash (\), to ignore the meaning of such special characters. Any character except a % or _ can follow the escape character. Wildcard characters following an escape character are treated as literal characters. For example, to display HTTPS certificates with names that begin with httpscer and are followed by a single character, enter httpscer_ in the Certificate Group Name field. Any certificate with the name Httpscer is not displayed.
Not like	Includes HTTPS certificates with a setting that is not similar to the string entered. Use the percent sign (%) wildcard characters to match zero or more characters, and use the underscore (_) wildcard character to match exactly one character. If the string contains special characters such as ^, (,), or ', use the escape character, a backslash (\), to ignore the meaning of such special characters. Any character except a % or _ can follow the escape character. Wildcard characters following an escape character are treated as literal characters. For example, to display HTTPS certificates except the ones with URLs that start with ftp://, enter %ftp:// in the Certificate URL field. Any certificate with the URL https://a.com or http://a.com is displayed.
In	Includes certificates with settings that are present in the list of specified strings. You can specify either a single string or a list of strings separated by commas. Use single quotes (') when the specified string contains commas or parentheses. For example, to display certificates with names CER1, CER2, and CER(3), enter (CER1,CER2,'CER(3)') or CER1,CER2,'CER(3)' in the Certificate Name field. Certificates with names such as CER4 and CER5 are not displayed.
Not in	Includes certificates with settings that are not present in the list of specified strings. You can specify either a single string or a list of strings separated by commas. Use single quotes (') when the specified string contains commas or parentheses. For example, to display certificates except the ones with names CER1, CER2, and CER(3), enter (CER1,CER2,'CER(3)') or CER1,CER2,'CER(3)' in the Certificate Name field. Certificates with names such as CER4 and CER5 are displayed

Table 8-9 Filtering HTTPS Certificates (continued)

Filter Operators	Description
Equal to	Includes certificates with a setting that is exactly the same as the string entered. For example, to display all certificates imported from the URL http://Location1, enter http://Location1 in the Certificate URL field. Certificates imported from http://Location^1 are not displayed. You cannot use keywords, special characters, or wildcard characters.
Not equal to	Includes certificates with a setting that is not exactly the same as the string entered. For example, to display all certificates not imported from the URL http://Location1, enter http://Location1 in the Certificate URL field. Certificates imported from http://Location^1 are displayed. You cannot use keywords, special characters, or wildcard characters.

Configuring HTTPS Certificate Groups

Certificate groups represent a trust relationship chain from root Certificate Authority to end entity. Each certificate in a certificate group (except the end entity's certificate) signs and trusts the next one in the chain. An end entity's certificate can be trusted only if all certificates in the certificate group leading to this certificate can be trusted. A certificate group can be used to represent an HTTPS server just like a single certificate, but with the added benefit that the client does not need to have all certificates locally. A certificate group can also be used to verify and authenticate an HTTPS server by comparing the server's certificates to the those of the certificate group.

To configure HTTPS certificate groups for the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure HTTPS certificate groups. The Device Home for Content Engine window appears.
- Step 3** In the Contents pane, choose **Applications > Web > HTTPS > Certificate Groups**. The HTTPS Certificate Groups for Content Engine window appears. The certificate groups created for device groups can only be viewed; they cannot be modified.
- Step 4** To configure an HTTPS digital certificate group, click the **Create New HTTPS Certificate Group** icon in the taskbar. The HTTPS Certificate Group for Content Engine window appears. You can create certificate groups with a given name, and add specified existing certificates to the current certificate group using this window.
- Step 5** In the Certificate Group Name field, specify the name of the certificate group to be used for certificate chains and server certificate authentication. Certificate group names must consist solely of alphanumeric, underscore, and hyphen characters. Certificate group names can begin with a numeric character. Certificate group names can have a maximum of 64 characters.

**Note**

Once you have created a certificate group name, you cannot it. You need to delete the existing certificate group and create a new one.

- Step 6** Choose the name of the certificate to be added to the certificate group from the Certificate drop-down lists. All certificates created using the HTTPS Certificates for Content Engine window are listed here. You can add a maximum of six certificates per certificate group.
- Two different certificate groups can have the same combination of HTTPS certificates. You can choose up to six certificates from the Certificate drop-down lists 1 through 6. These drop-down lists are set to an initial value of **Do not set**. You must choose these certificates in sequence when you add them to the certificate group and each associated certificate must be unique, or an error message is displayed when you click **Submit**.
- Step 7** To delete a certificate group, click the **Delete HTTPS Certificate Group** icon in the taskbar. You cannot delete a certificate group that is referenced by an HTTPS server. You must remove the association between the certificate group and HTTPS server before deleting.
- Step 8** To view a subset of the entire list of HTTPS certificate groups, click the **Filter Table** icon in the taskbar. See the next section “[Filtering HTTPS Certificate Groups for the Content Engine](#),” for more details.
- Step 9** To revert to the display of all configured HTTPS certificate groups, click the **View All HTTPS Certificate Groups** icon in the taskbar.
- Step 10** To save the settings, click **Submit**.

Table 8-10 *HTTPS Certificate Group Settings*

GUI Parameter	Function	CLI Command
Certificate Group Name	Name of the certificate group.	https server <i>name</i> certgroup chain <i>certgroup_name</i>
Certificate 1, 2, 3...	Certificate chain to use for the HTTPS server.	

Filtering HTTPS Certificate Groups for the Content Engine

You can use the Content Distribution Manager GUI to filter and display a subset of HTTPS certificate groups by certificate group name.

To set the filter criteria to display selected HTTPS certificate groups, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to filter HTTPS certificate groups. The Device Home for Content Engine window appears with the Contents pane on the left.
- Step 3** In the Contents pane, choose **Applications > Web > HTTPS > Certificate Groups**. The HTTPS Certificate Groups for Content Engine window appears. The certificate groups created for device groups can only be viewed; they cannot be modified.
- Step 4** Click the **Filter Table** icon in the taskbar. The Filtering HTTPS Certificate Groups for Content Engine window appears.

- Step 5** Choose a filter operator from the Certificate Group Name drop-down list, and enter the name of the certificate group that must be matched.
- Step 6** To revert to the display of all configured HTTPS certificate groups, click the **View All HTTPS Certificate Groups** icon in the taskbar. [Table 8-11](#) describes the filter operators for filtering HTTPS certificate groups.
- Step 7** Click **Submit**. The HTTPS Certificate Groups for Content Engine window reappears, displaying a list of all certificate groups that match the specified filter criteria.

Table 8-11 Filtering HTTPS Certificate Groups

Filter Operators	Description
Like	Includes HTTPS certificate groups with a setting that is similar to the string entered. Use the percent sign (%) wildcard characters to match zero or more characters, and use the underscore (_) wildcard character to match exactly one character. If the string contains special characters such as ^, (,), or ', use the escape character, a backslash (\), to ignore the meaning of such special characters. Any character except a % or _ can follow the escape character. Wildcard characters following an escape character are treated as literal characters. For example, to display HTTPS certificate groups with a name that begins with httpscer and are followed by a single character, enter httpscer_ in the Certificate Group Name field. Any certificate with the name Httpscer is not displayed.
Not like	Includes HTTPS certificate groups with a setting that is not similar to the string entered. Use the percent sign (%) wildcard characters to match zero or more characters, and use the underscore (_) wildcard character to match exactly one character. If the string contains special characters such as ^, (,), or ', use the escape character, a backslash (\), to ignore the meaning of such special characters. Any character except a % or _ can follow the escape character. Wildcard characters following an escape character are treated as literal characters. For example, to display HTTPS certificate groups except the ones with a name that begins with httpscer, enter httpscer in the Certificate Group Name field. Any certificate with the name Httpscer is displayed.
In	Includes certificate groups with settings that are present in the list of specified strings. You can specify either a single string or a list of strings separated by commas. Use single quotes (') when the specified string contains commas or parentheses. For example, to display certificate groups with names CER1, CER2, and CER(3), enter (CER1,CER2,'CER(3)') or CER1,CER2,'CER(3)' in the Certificate Name field. Certificate groups with names such as CER4 and CER5 are not displayed.

Table 8-11 Filtering HTTPS Certificate Groups (continued)

Filter Operators	Description
Not in	Includes certificate groups with settings that are not present in the list of specified strings. You can specify either a single string or a list of strings separated by commas. Use single quotes (') when the specified string contains commas or parentheses. For example, to display certificate groups except the ones with names CER1, CER2, and CER(3), enter (CER1,CER2,'CER(3)') or CER1,CER2,'CER(3)' in the Certificate Name field. Certificate groups with names such as CER4 and CER5 are displayed.
Equal to	Includes certificate groups with a name that is exactly the same as the string entered. For example, to display a certificate group with the name httpscer, enter httpscer in the Certificate Group Name field. A certificate group with the name Httpscer is not displayed. You cannot use keywords, special characters, or wildcard characters.
Not equal to	Includes certificate groups with a name that is not exactly the same as the string entered. For example, to display all certificate groups except the one with the name httpscer, enter httpscer in the Certificate Group Name field. A certificate group with the name Httpscer is displayed. You cannot use keywords, special characters, or wildcard characters.

Configuring HTTPS Keys

The private key is the secret half of a key pair used in a public key algorithm. Private keys are typically used to encrypt a symmetric session key, digitally sign a message, or decrypt a message that has been encrypted with the corresponding public key. PKCS #12 specifies a portable format for storing or transporting a user's private keys and certificate information. The private key that the Content Engine uses to act as an origin HTTPS server must match the selected certificate.

You can create keys with a given name, import keys from external sources, or remove existing keys using the Content Distribution Manager GUI. To configure HTTPS keys for the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure HTTPS keys. The Device Home for Content Engine window appears.
- Step 3** In the Contents pane, choose **Applications > Web > HTTPS > Keys**. The HTTPS Keys for Content Engine window appears. (The keys created for device groups can only be viewed; they cannot be modified.)
- Step 4** Click the **Create New HTTPS Key** icon in the taskbar to configure an HTTPS key. The Creating new HTTPS Key for Content Engine window appears.

- Step 5** In the Key Name field, specify the name of the private key to be imported from an internal or external certificate authority (CA). You need to create a private key before importing its value from an external source and associating it with an HTTPS server. Key names must consist solely of alphanumeric, underscore, and hyphen characters. Key names can begin with a numeric character. Key names can have a maximum length of 64 characters.



Note Once you have created the name of an HTTPS key, you cannot modify it. You need to delete the existing key and create a new one. Only HTTPS key names are stored in the CMS database. The actual certificates are stored on the Content Engine.

- Step 6** In the Key URL field, specify the external location from which the private key is to be imported. The URL must be a valid URL using HTTP, FTP, or HTTPS. The key value is imported when you specify the URL and submit changes.

A key name without imported values can be listed in the Content Distribution Manager GUI; however, only keys with valid values can be associated with an HTTPS server. Two different keys can be imported from the URL and two different HTTPS servers can be associated with the same key.

After a key is imported from the Content Distribution Manager GUI, a new icon (**Re-import**) appears in the taskbar. The **Re-import** icon can be used to reimport the key if the import did not succeed the first time.

- Step 7** In the Username field, specify the username required to access the external source from which the key is being imported. An entry is required in this field only if the external source is password-protected.
- Step 8** In the Password field, specify the password used to authenticate users who want to gain access to the external source from which the key is being imported. An entry is required in this field only if the external source is password-protected. Reenter the password in the Confirm Password field.
- Step 9** To view a subset of the entire list of HTTPS keys, click the **Filter Table** icon in the taskbar. See the next section “[Filtering HTTPS Keys for the Content Engine](#)” for more details.
- Step 10** To revert to the display of all configured HTTPS keys, click the **View All HTTPS Keys** icon in the taskbar.
- Step 11** To save the settings, click **Submit**.

Table 8-12 *HTTPS Key Settings*

GUI Parameter	Function	CLI Command
Key Name	Name of the key.	<code>https server <i>name</i> key <i>name</i> password <i>password</i></code>
Key URL	External location from which the private key is to be imported.	
Password	Password to decrypt the private key file.	

Filtering HTTPS Keys for the Content Engine

You can use the Content Distribution Manager GUI to filter and display a subset of HTTPS keys by the private key name and URL.

To set the filter criteria to display selected HTTPS keys, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to filter HTTPS keys. The Device Home for Content Engine window appears with the Contents pane on the left.
- Step 3** In the Contents pane, choose **Applications > Web > HTTPS > Keys**. The HTTPS Keys for Content Engine window appears. The keys created for device groups can only be viewed; they cannot be modified.
- Step 4** Click the **Filter Table** icon in the taskbar. The Filtering HTTPS Keys for Content Engine window appears.
- Step 5** Choose a filter operator from the Key Name drop-down list and enter the name of the key that must be matched.
- Step 6** Choose a filter operator from the Key URL drop-down list and enter the key location that must be matched.
- Step 7** To revert to the display of all configured HTTPS keys, click the **View All HTTPS Keys** icon in the taskbar. [Table 8-13](#) describes the filter operators for filtering HTTPS keys.
- Step 8** Click **Submit**. The HTTPS Keys for Content Engine window reappears, displaying a list of all keys that match the specified filter criteria.

Table 8-13 Filtering HTTPS Keys

Filter Operators	Description
Like	Includes HTTPS keys with a setting that is similar to the string entered. Use the percent sign (%) wildcard characters to match zero or more characters, and use the underscore (_) wildcard character to match exactly one character. If the string contains special characters such as ^, (,), or ', use the escape character, a backslash (\), to ignore the meaning of such special characters. Any character except a % or _ can follow the escape character. Wildcard characters following an escape character are treated as literal characters. For example, to display HTTPS keys with a name that begins with httpskey and is followed by a single character, enter httpskey_ in the Station Name field. Any key with the name Httpskey is not displayed.
Not like	Includes HTTPS keys with a setting that is not similar to the string entered. Use the percent sign (%) wildcard characters to match zero or more characters, and use the underscore (_) wildcard character to match exactly one character. If the string contains special characters such as ^, (,), or ', use the escape character, a backslash (\), to ignore the meaning of such special characters. Any character except a % or _ can follow the escape character. Wildcard characters following an escape character are treated as literal characters. For example, if you want to display HTTPS keys except the ones with a URL that starts with ftp://, you should enter %ftp:// in the Media field. Any key with the URL https://a.com or http://a.com is displayed.

Table 8-13 Filtering HTTPS Keys (continued)

Filter Operators	Description
In	Includes keys with settings that are present in the list of specified strings. You can specify either a single string or a list of strings separated by commas. Use single quotes (') when the specified string contains commas or parentheses. For example, to display keys with names KEY1, KEY2, and KEY(3), enter (KEY1,KEY2,'KEY(3)') or KEY1,KEY2,'KEY(3)' in the Key Name field. Keys with names such as KEY4 and KEY5 are not displayed.
Not in	Includes keys with settings that are not present in the list of specified strings. You can specify either a single string or a list of strings separated by commas. Use single quotes (') when the specified string contains commas or parentheses. For example, to display keys except the ones with names KEY1, KEY2, and KEY(3), enter (KEY1,KEY2,'KEY(3)') or KEY1,KEY2,'KEY(3)' in the Key Name field. Keys with the names such as KEY4 and KEY5 are displayed.
Equal to	Includes keys with a setting that is exactly the same as the string entered. For example, to display all keys imported from the URL http://Location1, enter http://Location1 in the Key URL field. Keys imported from http://Location^1 are not displayed. You cannot use keywords, special characters, or wildcard characters.
Not equal to	Includes keys with a setting that is not exactly the same as the string entered. For example, to display all keys not imported from the URL http://Location1, enter http://Location1 in the Key URL field. Keys imported from http://Location^1 are displayed. You cannot use keywords, special characters, or wildcard characters.

Configuring HTTPS Proxy Settings

Cisco ACNS 5.x software supports HTTPS in the following two scenarios:

- The Content Engine receives an HTTPS request sent by a web client configured to use the Content Engine as an HTTPS proxy server.
- The Content Engine in transparent mode intercepts a request sent by a web client to another HTTPS proxy server.

In both cases the Content Engine creates a connection to the origin server (directly or through another proxy server) and allows the web client and origin server to set up a Secure Socket Layer (SSL) tunnel through the Content Engine.



Note

A Domain Name System (DNS) server must be configured to support HTTPS proxy requests. See the [“Configuring the DNS Server for HTTP Proxy Caching”](#) section on page 8-78 for more information.

To configure HTTPS proxy parameters, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**. If you have created device groups, the Device Group window appears.



Note HTTPS proxy parameters can be configured for individual Content Engines or for device groups. This procedure describes the steps for configuring an HTTPS proxy for a device group.

- Step 2** Click the **Edit** icon next to the name of the device group that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **Applications > Web > HTTPS > HTTPS Proxy**. The HTTPS Proxy Settings for Content Engine window appears. (See [Figure 8-9](#) and [Figure 8-10](#).) [Table 8-14](#) describes the fields in these windows.

Figure 8-9 HTTPS Proxy Settings Window—Top

The screenshot shows the 'HTTPS Proxy Settings for Content Engine, CONTENTENGINE' window. The 'Contents' pane on the left shows the navigation tree with 'HTTPS Proxy' selected. The main settings area includes the following fields:

- Enable Incoming Proxy:** ☒
- List of Incoming Ports:** 1 2 3 4 (Space separated list of up to 8 ports)
- Enable TCP:** ☒
- TCP Read/Write Timeout:** 300 (seconds) (1-3600)
- Allow All ports:** ☒
- Allow port 443:** ☒
- Allow port 563:** ☒
- Other Allow Port List:** (Space separated list of up to 8 ports)
- Deny All ports:** ☐
- Deny Port List:** (Space separated list of up to 8 ports)
- Enable Outgoing Proxy:** ☒
- Use Origin Server:** ☒
- Outgoing Connection Timeout:** 3000 (microseconds) (1/200-5000000)

Buttons for 'Submit' and 'Cancel' are at the bottom right.

- Step 4** To allow the Content Engine to accept incoming requests on ports in addition to port 80, check the **Enable Incoming Proxy** check box.
- Step 5** In the List of Incoming Ports field, enter the port numbers used by the proxy server. Separate numbers with a space.
- Step 6** To configure the TCP server connection's read/write timeout, check the **Enable TCP** check box.
- Step 7** In the TCP Read/Write Timeout field, enter a value in seconds for the read/write timeout. The range is from 1 to 3600, and the default is 120 seconds.

- Step 8** To enable the use of an outgoing proxy, if needed, check the **Enable Outgoing Proxy** check box.
- Step 9** If HTTPS traffic is to be allowed on all ports, check the **Allow All Ports** check box.
- Step 10** If HTTPS traffic is to be allowed on port 443, check the **Allow port 443** check box.
- Step 11** If HTTPS traffic is to be allowed on port 563, check the **Allow port 563** check box.
- Step 12** In the Other Allow Port List field, enter a list of port numbers on which HTTPS traffic is to be allowed, and separate each number with a space. This field is required if an outgoing proxy is enabled. The range is from 1 to 65535. Up to eight ports can be configured. Ports 443 and 563 are allowed by default.
- Step 13** If HTTPS traffic is to be blocked on all ports, check the **Deny All Ports** check box.
- Step 14** In the Deny Port List field, enter a list of port numbers on which HTTPS requests must be rejected. Separate each number with a space. This field is required if an outgoing proxy is enabled. The range is from 1 to 65535. Up to 8 ports can be configured. Ports under 1024 are denied for HTTPS requests by default. This ensures that unwanted access to any HTTPS port is prevented when a request goes through the Content Engine.
- Step 15** To enable the outgoing proxy configuration fields, check the **Enable Outgoing Proxy** check box. (See Figure 8-10.)

Figure 8-10 HTTPS Proxy Settings Window—Bottom

The screenshot displays the 'HTTPS Proxy Settings for Content Engine, CONTENTENGINE' window. The left sidebar shows a tree view with 'HTTPS' selected under 'Web'. The main content area contains the following settings:

- Deny Port List:** (Empty text field)
- Enable Outgoing Proxy:** ☒
- Use Origin Server:** ☒
- Outgoing Connection Timeout:** 3000 (microseconds) (200-5000000)
- Outgoing Monitor Period:** 20 (seconds) (10-300)
- Outgoing Proxy Hostname 1:** 10.77.155.221
- Outgoing Proxy Port 1:** 8080
- Outgoing Proxy Hostname 2:** 1.2.3.4
- Outgoing Proxy Port 2:** 2
- Outgoing Proxy Hostname 3:** 1.2.3.41
- Outgoing Proxy Port 3:** 4
- Outgoing Proxy Hostname 4:** 1.2.3.42
- Outgoing Proxy Port 4:** 5
- Outgoing Proxy Hostname 5:** 1.2.3.43
- Outgoing Proxy Port 5:** 3
- Outgoing Proxy Hostname 6:** 1.2.3.45
- Outgoing Proxy Port 6:** 6
- Outgoing Proxy Hostname 7:** 1.2.3.47
- Outgoing Proxy Port 7:** 7
- Outgoing Proxy Hostname 8:** 1.2.3.48
- Outgoing Proxy Port 8:** 8

A note at the bottom states: "Note: * - Required Field". The 'Submit' and 'Cancel' buttons are at the bottom right.

- Step 16** To use the origin server if all outgoing proxies fail, check the **Use Origin Server** check box. If all the configured proxy servers fail, the Content Engine can optionally redirect HTTPS requests to the origin server specified in the HTTPS header if you have checked this check box. If this option is not enabled, the client receives an error response. Response errors and read errors are returned to the client, because it is not possible to detect whether these errors are generated at the origin server or at the proxy.

- Step 17** Enter a value in microseconds in the Outgoing Connection Timeout field. This value specifies the number of microseconds before the outgoing proxy servers are probed to see whether they are active.
- Step 18** Enter a value in the Outgoing Monitor Period field to specify how frequently the Content Engine polls the specified outgoing HTTPS proxy servers.
- A background process on the Content Engine monitors the state of the configured outgoing proxy servers. You can configure the Content Engine to poll the specified outgoing proxy servers at a specific time in order to monitor their availability. The outgoing monitor period is the time during which the proxy servers are polled. The monitoring period is specified in seconds, and can be from 10 to 300 seconds. The default monitoring period is 60 seconds. If one of the outgoing proxy servers is unavailable, the polling mechanism waits for the connect timeout (300000 microseconds) before polling the next outgoing proxy server.
- Step 19** In the Outgoing Proxy Hostname fields 1 through 8, specify an IP address or host name if an outgoing proxy is enabled. The server configured in the first field becomes the primary outgoing HTTPS proxy server for the Content Engine. This field is required if an outgoing proxy is enabled.
- You can configure up to eight HTTPS outgoing proxy servers for each Content Engine. At any one time, the Content Engine uses only one of the configured outgoing HTTPS proxy servers. They cannot be used simultaneously.
- Step 20** In the Outgoing Proxy Port fields 1 through 8, specify a port number that will be used by the outgoing proxy server to accept proxy HTTPS requests if an outgoing proxy is enabled. (A port is required for the primary server entered in the Hostname field.)
- Step 21** To save the settings, click **Submit**.
- A “Click to Submit” message appears in red next to the current settings when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The **Reset** button is visible only when you apply default or device group settings to change the current device settings but the settings have not yet been submitted.

Table 8-14 shows the GUI parameters and CLI commands associated with HTTPS proxy features. The order in which the CLI commands are entered is not important.

Table 8-14 HTTPS Proxy Settings

GUI Parameter	Function	CLI Command
Enable incoming proxy	Configures the Content Engine to accept incoming HTTPS traffic request on ports in addition to port 80.	https proxy incoming
List of incoming ports	Supports up to 8 incoming proxy ports.	https proxy incoming <i>port 1–65535, port, . . .</i>
Enable TCP	Enables the TCP Read/Write Timeout field.	—
TCP Read/Write Timeout	Configures the TCP server connection’s read/write timeout. The range is from 1 to 3600, and the default is 120 seconds.	https tcp-rw-timeout <i>seconds</i>
Enable outgoing proxy	Enables an outgoing HTTPS proxy server.	https proxy outgoing <i>host {hostname ipaddress} port 1–65535</i>
Allow All ports	Allows HTTPS traffic on all ports.	https destination-port allow all

Table 8-14 *HTTPS Proxy Settings (continued)*

GUI Parameter	Function	CLI Command
Allow port 443	Allows HTTPS traffic on port 443.	https destination-port allow 443
Allow port 563	Allows HTTPS traffic on port 563.	https destination-port allow 563
Other Allow Port List	List of port numbers on which HTTPS traffic is to be allowed (required, if an outgoing proxy is enabled). The range is 1–65535. Up to 8 ports can be configured. Ports 443 and 563 are allowed by default.	https destination-port allow <i>port_num port_num port_num . . .</i>
Deny All ports	Denies HTTPS traffic on all ports.	https destination-port deny all
Deny Port List	List of port numbers on which HTTPS requests must be rejected (required, if an outgoing proxy is enabled).	https destination-port deny ports
Enable Outgoing Proxy	Enables the fields in the GUI window for configuring the outgoing proxy settings.	—
Use Origin Server	Redirects HTTPS requests to the origin server specified in the HTTPS header.	https proxy outgoing origin-server
Outgoing Connection Timeout	Timeout period (in microseconds) to use when probing for outgoing proxy servers. The range is from 200 to 5000000. The default value is 300 microseconds.	https proxy outgoing connection-timeout <i>microsecs</i>
Outgoing Monitor Period	Specifies the interval in seconds at which to monitor one outgoing proxy server. If multiple outgoing proxy servers are configured, they are monitored sequentially. The default value is 200 seconds.	https proxy outgoing monitor <i>seconds</i>
Outgoing Proxy Hostname 1–8	IP address or host name of the HTTPS outgoing proxy server.	https proxy outgoing host { <i>hostname</i> <i>ipaddress</i> }
Outgoing Proxy Port 1–8	Port of the HTTPS outgoing proxy server.	https proxy outgoing host { <i>hostname</i> <i>ipaddress</i> } port <i>1–65535</i>

**Note**

HTTPS traffic is encrypted and cannot be interpreted by the Content Engine or any other device between the web client and the origin server. HTTPS objects are not cached.

**Note**

For other operations that you can perform using the icons in the taskbar, see [Table 3-3, “Content Distribution Manager GUI Icons.”](#)

Configuring HTTPS Servers

You can configure an HTTPS server and configure a caching solution to allow a Content Engine act as an origin HTTPS server. This configuration can reduce WAN traffic and increase data security because authorized clients from remote branch offices can access their own Content Engines, which are configured as HTTPS servers, by using HTTPS.

The Content Engine decodes HTTPS traffic from a client and perform normal HTTP operations on it, such as caching and request processing. The Content Engine will initiate HTTPS connections to an origin server and fetch content from origin servers upon a cache miss or cache validation.

Some restrictions exist concerning the usage of certificates, keys, and certificate groups.

- A certificate without valid imported values cannot be associated with a certificate group.
- A certificate group without added certificates cannot be associated with an HTTPS server.
- A certificate used by an HTTPS server cannot be modified or deleted.
- A certificate group used by an HTTPS server cannot be modified or deleted.
- A certificate added to a certificate group cannot be modified or deleted.
- A key used by an HTTPS server cannot be modified or deleted.

To configure HTTPS servers for the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure HTTPS servers. The Device Home for Content Engine window appears.
- Step 3** In the Contents pane, choose **Applications > Web > HTTPS > HTTPS Servers**. The HTTPS Servers for Content Engine window appears.
- Step 4** The Aggregate Settings **Yes** radio button is selected by default. This means that HTTPS servers are displayed for the Content Engine as well as for any associated device groups. The settings for the HTTPS servers for device groups cannot be modified or deleted; they are read-only. However, you can modify the HTTPS servers for the Content Engine. If you click the Aggregate Settings **No** radio button, you can view and modify the HTTPS servers for the Content Engine only. The HTTPS servers for any associated device groups are not even displayed.
- Step 5** To configure an HTTPS server, click the **Create New HTTPS Server** icon in the taskbar. The Creating new HTTPS Server for Content Engine window appears.

To configure HTTPS servers, you need to configure three sets of settings:

- HTTPS Servers
- Certificate Key/Associations
- General HTTPS Settings

Step 6 Under the HTTPS Server section, follow these steps:

- a. In the Name field, specify the name of the HTTPS server. Server names can consist of Arabic numerals and upper- and lowercase alphabetic, underscore, hyphen, and period characters. Server names must begin with an alphabetic character or underscore and have a limit of 15 characters.
- b. To enable caching and SSL termination for the HTTPS server, check the **Enable** check box. You must have entered the host name, certificate or certificate group, and private key before an HTTPS server can be enabled.
- c. In the Host Name field, specify the IP address or fully qualified domain name for the origin HTTPS server. Two different HTTPS servers cannot have the same host name.

Each active server has one or more IP addresses associated with it, specified either as an IP address or a domain name that is resolved to one or more IP addresses. An HTTPS request arriving at a Content Engine (either by WCCP intercept or by direct proxy) is matched with a server configuration that will be used to handle the request by comparing the target IP address of the request with the IP addresses associated with the active servers.

Step 7 Under the Certificate Key/Associations section, follow these steps:

- a. Choose the certificate group for the HTTPS server from the Server Certificate drop-down list. All certificate groups created using the HTTPS Certificate Groups for Content Engine window will be displayed here.

The certificate group is included in the SSL handshake that is sent to SSL clients. Only one certificate group can be selected for an HTTPS server at a time. SSL Version 2 does not support certificate groups.

- b. Choose the certificate chains and authentication needed to access the HTTPS server from the Authenticating Certificate Group drop-down list. All certificate groups created using the HTTPS Certificate Groups for Content Engine window are displayed here. Only one certificate group can be selected for an HTTPS server at a time.

SSL Version 2 does not support certificate groups.

- c. From the Certificate drop-down list, choose a specific SSL certificate so that the Content Engine can act as an origin HTTPS server while authenticating clients. All certificates created using the HTTPS Certificates for Content Engine window are displayed here. Only one certificate can be selected for an HTTPS server at a time.
- d. From the Private Key drop-down list, choose a private key so that the Content Engine can act as an origin HTTPS server while authenticating clients. This private key must match the selected certificate. All keys created using the HTTPS Keys for Content Engine window are displayed here. Only one key can be selected for an HTTPS server at a time.
- e. If the private key file is encrypted, in the Private Key Password field, specify the password required to decrypt it. Reenter the password in the Confirm Password field.

Step 8 Under the General HTTPS Settings section, follow these steps:

- a. In the SSL cache size field, specify the number of SSL session caching entries. This indicates the maximum number of entries that can be stored in the cache before the origin HTTPS server must be contacted.
- b. In the SSL cache Timeout field, specify the SSL session cache timeout in seconds. The Content Engine contacts the origin server to authenticate clients once this value is exceeded.
- c. Choose the protocol versions to be used for communication between clients and HTTPS servers from the Protocol Version drop-down list. [Table 8-15](#) shows the protocol versions available.
- d. To enable server certificate authentication for the HTTPS server, check the **Enable Authentication** check box. This option is enabled by default.

- e. To ignore errors caused by using the certificate before it is valid, check the **Ignore CertNotYetValid Error** check box.
- f. To ignore errors due to domain name mismatch, check the **Ignore DomainName Error**.
- g. To ignore certificate expiration errors, check the **Ignore ExpiredDate Error** check box.
- h. To ignore errors caused by an unrecognized Certificate Authority (CA), check the **Ignore InvalidCA Error** check box.

Table 8-15 Protocol Versions

Protocol Version	Description
SSL v2	Sets the use of the SSL Version 2 protocol only.
SSL v23 and TLS v1	Sets the use and understanding of the SSL Version 2 or Version 3 and Transport Layer Security (TLS) Version 1 protocol.
SSL v3	Sets the use and understanding of the SSL Version 3 protocol only.
TLS v1	Sets the use and understanding of the TLS Version 1 protocol only.

Step 9 To save the settings, click **Submit**.

Filtering HTTPS Servers for the Content Engine

You can use the Content Distribution Manager GUI to filter and display a subset of HTTPS servers by server name, status (enabled or disabled), host name, SSL cache size, SSL cache timeout, protocol version, status of server certificate authentication (enabled or disabled), and status of server certificate authentication error tracking (enabled or disabled).

To set the filter criteria to display selected HTTPS servers, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to filter HTTPS servers. The Device Home for Content Engine window appears with the Contents pane on the left.
- Step 3** In the Contents pane, choose **Applications > Web > HTTPS > HTTPS Servers**. The HTTPS Servers for Content Engine window appears.
- Step 4** Click the **Filter Table** icon in the taskbar. The Filtering HTTPS Server Entries for Content Engine window appears.
- Step 5** From the **Name** drop-down list, choose a filter operator, and enter the name of the server that must be matched.
- Step 6** From the **Enable** drop-down list, choose a filter operator, and from the second drop-down list, choose whether or not you want to view servers on which HTTPS caching is enabled.
- Step 7** From the **Host** drop-down list, choose a filter operator, and enter the IP address or fully qualified domain name of the host that must be matched.
- Step 8** From the **SSL cache Size** drop-down list, choose a filter operator, and enter the size of the SSL cache that must be matched.

- Step 9** From the **SSL cache Timeout** drop-down list, choose a filter operator, and enter the timeout (in seconds) that must be matched.
- Step 10** From the **Protocol Version** drop-down list, choose a filter operator, and enter the version of the communication protocol that must be matched.
- Step 11** From the **Enable Authentication** drop-down list, choose a filter operator, and choose whether or not you want to view HTTPS servers for which server certificate authentication is enabled.
- Step 12** From the **CertNotYetValid Error** drop-down list, choose a filter operator, and choose whether or not you want to view HTTPS servers for which errors caused by using the certificate before it is valid are ignored.
- Step 13** From the **Ignore DomainName Error** drop-down list, choose a filter operator, and choose whether or not you want to view HTTPS servers for which errors due to domain name mismatch are ignored.
- Step 14** From the **Ignore ExpiredDate Error** drop-down list, choose a filter operator, and choose whether or not you want to view HTTPS servers for which certificate expiration errors are ignored.
- Step 15** From the **Ignore InvalidCA Error** drop-down list, choose a filter operator, and choose whether or not you want to view HTTPS servers for which errors caused by an unrecognized Certificate Authority (CA) are ignored.
- Step 16** Click **Submit**. The HTTPS Servers for Content Engine window reappears, displaying a list of all servers that match the specified filter criteria.
- Step 17** To revert to the display of all configured HTTPS servers, click the **View All HTTPS Servers** icon in the taskbar.

Table 8-16 Filtering HTTPS Servers

Filter Operators	Description
Like	Includes HTTPS servers with a setting that is similar to the string entered. Use the percent sign (%) wildcard characters to match zero or more characters, and use the underscore (_) wildcard character to match exactly one character. If the string contains special characters such as ^, (,), or ', use the escape character, a backslash (\), to ignore the meaning of such special characters. Any character except a % or _ can follow the escape character. Wildcard characters following an escape character are treated as literal characters. For example, to display HTTPS servers with a name that begins with httpscer and are followed by a single character, enter httpscer_ in the Name field. Any server with the name Httpscer is not displayed.
Not like	Includes HTTPS servers with a setting that is not similar to the string entered. Use the percent sign (%) wildcard characters to match zero or more characters, and use the underscore (_) wildcard character to match exactly one character. If the string contains special characters such as ^, (,), or ', use the escape character, a backslash (\), to ignore the meaning of such special characters. Any character except a % or _ can follow the escape character. Wildcard characters following an escape character are treated as literal characters. For example, to display HTTPS servers except the ones with the URL that starts with httpscer, enter httpscer in the Name field. Any certificate with the name Httpscer is displayed.

Table 8-16 **Filtering HTTPS Servers (continued)**

Filter Operators	Description
In	Includes servers with names that are present in the list of specified strings. You can specify either a single string or a list of strings separated by commas. Use single quotes (') when the specified string contains commas or parentheses. For example, to display servers with the names CER1, CER2, and CER(3), enter (CER1,CER2,'CER(3)') or CER1,CER2,'CER(3)' in the Server Name field. Servers with names such as CER4 and CER5 are not displayed.
Not in	Includes servers with settings that are not present in the list of specified strings. You can specify either a single string or a list of strings separated by commas. Use single quotes (') when the specified string contains commas or parentheses. For example, to display servers except the ones with the names CER1, CER2, and CER(3), enter (CER1,CER2,'CER(3)') or CER1,CER2,'CER(3)' in the Server Name field. Servers with names such as CER4 and CER5 are displayed.
Equal to	Includes servers with a setting that is exactly the same as the string or the value entered, or with a setting that matches the configuration chosen from the drop-down list. For example, to display an HTTPS server with the host name host1, enter host1 in the Host field. An HTTPS server with the host name host2 is not displayed. You cannot use keywords, special characters, or wildcard characters.
Not equal to	Includes servers with a setting that is not exactly the same as the string or the value entered, or with a setting that does not match the configuration chosen from the drop-down list. For example, to display all HTTPS servers except the one with protocol version sslv2-only, enter SSL v2 in the Protocol Version field. Any HTTPS server with the protocol version sslv3-only is displayed. You cannot use keywords, special characters, or wildcard characters.
Greater than	Includes servers with SSL cache size, SSL cache timeout, and protocol version higher than the value entered. For example, to display all HTTPS servers for which the number of SSL session caching entries exceeds 1000, enter 1000 in the SSL cache size field. Any HTTPS server with a cache size lesser than or equal to 1000 is not displayed.
Less than	Includes servers with SSL cache size, SSL cache timeout, and protocol version lower than the value entered. For example, to display all HTTPS servers for which the SSL session cache timeout is less than 1000 seconds, enter 1000 in the SSL cache timeout field. Any HTTPS server with an SSL session cache timeout greater than or equal to 1000 seconds is not displayed.

Table 8-16 *Filtering HTTPS Servers (continued)*

Filter Operators	Description
Greater than or equal to	Includes servers with SSL cache size, SSL cache timeout, and protocol version higher than or equal to the value entered. For example, to display all HTTPS servers for which the number of SSL session caching entries exceeds 1000, enter 1000 in the SSL cache size field. Any HTTPS server with a cache size less than 1000 is not displayed.
Less than or equal to	Includes servers with SSL cache size, SSL cache timeout, and protocol version lower than or equal to the value entered. For example, to display all HTTPS servers for which the SSL session cache timeout in seconds is lesser than or equal 1000, enter 1000 in the SSL cache timeout field. Any HTTPS server with an SSL session cache timeout greater than 1000 seconds is not displayed.

Transparent HTTPS Caching Using SSL

Caching an HTTPS client request using SSL termination means that the Content Engine decrypts the SSL-encrypted data. By decrypting the SSL-encrypted data, the Content Engine can see the HTTPS client request in plain text, which allows the Content Engine to perform numerous HTTP processing tasks (for example, caching, rule processing, and filtering) on such requests.

Transparent HTTPS caching using SSL works as follows:

1. The Content Engine, configured as an HTTPS server, receives an HTTPS request redirected through a WCCP-enabled router.
2. The Content Engine sends back an SSL certificate (obtained from the origin server) to the requesting client to negotiate an SSL connection.
3. The client sends HTTPS requests inside the negotiated SSL connection.
4. The Content Engine examines the request, looks in its cache, and performs normal HTTP request processing.
5. If the Content Engine can fulfill the request from its local storage (cache hit), it sends the requested content back using the SSL connection.
6. If the Content Engine cannot fulfill the request from its local storage (cache miss), it initiates a connection to the origin server to retrieve the requested content through the SSL connection.
7. The Content Engine caches the requested content (if possible) and also sends a copy to the requesting client through the negotiated SSL connection.

For the SSL-termination feature to work properly, you must install the SSL certificate and private key of the origin HTTPS servers on the Content Engine. The SSL-termination feature works in both transparent mode and proxy mode if the Content Engine has the correct certificates and private keys installed. For specific requested content to be cached, you must import the proper certificates and keys for these origin HTTPS servers into the Content Engine and configure the Content Engine to cache content from these origin HTTPS servers. See the [“Configuring HTTPS Certificates”](#) section on page 8-27.

Configuring HTTP and HTTPS Outgoing Proxy Exclusion Settings

When in transparent caching mode, the Content Engine can intercept requests sent to another proxy and send these requests to one of the following two destinations:


- **Default server**—This is the default option. The Content Engine retrieves the objects from the web server itself, or if it is configured to use an outgoing proxy for this protocol, it forwards the request to its outgoing proxy. In this scenario, the client browser configuration is ignored, and the Content Engine configuration is used to retrieve the object from the server.
- **Original proxy**—The Content Engine forwards the request to the proxy to which the client had originally addressed the request. This proxy may be different from the Content Engine's own outgoing proxy for the specified protocol.

ACNS 5.x software also has an option that allows the administrator to specify a single domain name, host name, or IP address to be globally excluded from proxy forwarding. Domains are entered as an ASCII string, separated by spaces. The wildcard character * (asterisk) can be used for IP addresses (for instance, 172.16.*.*). Requests with a destination specified with wildcard characters bypass the Content Engine proxy as well as the failover proxies.

To configure proxy protocol parameters, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure.
- Step 3** From the Contents pane, choose **Applications > Web > Outgoing Proxy Exclusions**. The Outgoing Proxy Exclusions window appears. (See [Figure 8-11](#).) [Table 8-17](#) describes the fields in this window.

Figure 8-11 Outgoing Proxy Exclusions Settings Window



- Step 4** To intercept requests sent to another proxy and send these requests to either the origin server or a specified proxy server, check the **Enable Outgoing Proxy Exclusion** check box.
- Step 5** Enter domain names, IP addresses, or host names, each separated by a space, in the Outgoing Proxy Exclude List field. The item entered here is globally excluded from proxy forwarding. You can also specify the asterisk (*) wildcard character to match any IP address to be excluded.
- Step 6** To set the transparent mode behavior for proxy-style requests, check the **Enable Transparent Mode** check box.
- Step 7** Choose **default-server**, **original-proxy**, or **reset** from the Transparent Proxy Mode drop-down list. [Table 8-17](#) explains these options.

Table 8-17 Proxy Protocols Key Parameters

Key Parameter	Function	CLI Command
Default-server	The Content Engine retrieves objects from the web server itself. With this option, a proxy-style request can be sent to an outgoing proxy server if such a server is configured. If no outgoing proxy server is configured, then the request is served by the origin server and the Content Engine.	proxy-protocols transparent default-server

Table 8-17 Proxy Protocols Key Parameters (continued)

Key Parameter	Function	CLI Command
Original-proxy	The Content Engine forwards the request to the original proxy that the client addressed the request to.	proxy-protocols transparent original-proxy
Reset	This option resets the transparent proxy mode to use the default server. The request will be returned to the client during a cache miss. The client does not obtain the requested object in this mode.	proxy-protocols transparent reset

Step 8 To save the settings, click **Submit**.

A “Click to Submit” message appears in red next to the current settings when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The **Reset** button is visible only when you apply default or device group settings to change the current device settings but the settings have not yet been submitted.



Note For other operations that you can perform using the icons in the taskbar, see [Table 3-3, “Content Distribution Manager GUI Icons.”](#)

Configuring FTP Settings

You can configure, modify, and view FTP settings for Content Engines and device groups by completing the following tasks:

- [Configuring FTP-Over-HTTP Connection Settings, page 8-52](#)
- [Configuring FTP-Over-HTTP Cache Freshness, page 8-55](#)
- [Configuring Native FTP Connections, page 8-58](#)
- [Configuring the Client Side of Nontransparent Native FTP Caching, page 8-60](#)
- [Configuring Native FTP Cache Freshness, page 8-61](#)
- [Enabling Inetd FTP Service, page 8-61](#)
- [Configuring Native FTP Custom Messages, page 8-62](#)

You can configure the Content Engine for FTP caching in either of the following two usage modes:

- **FTP-over-HTTP mode.** The Content Engine (nontransparent forward proxy server) caches the contents of the specified FTP URLs that are sent to it directly by clients that are using the HTTP protocol. This mode allows users to use their browsers (HTTP protocol) to access files (to send and receive files) on remote FTP servers.
- **Native FTP mode.** The Content Engine caches the content of the FTP requests that are sent from clients in the native FTP protocol. In ACNS 5.5, native FTP caching is supported in transparent and nontransparent proxy mode. (Native FTP caching was only supported in transparent proxy mode in the ACNS 5.1 and 5.2 software releases.)

In both of these usage modes, the Content Engine uses the FTP protocol to retrieve and locally cache the content of the FTP requests. These two usage modes differ in the protocol used by the client to issue the FTP request. In FTP-over-HTTP mode, clients use their browsers (the HTTP protocol) to issue FTP requests. In native FTP mode, clients use the native FTP protocol to issue FTP requests.

**Note**

Transparent redirection of FTP requests is supported only by WCCP Version 2; transparent redirection through a Layer 4 switch is not supported.

**Note**

In ACNS 5.3.1 software, the following changes were made to FTP-related commands:

- **show ftp proxy** EXEC command has been replaced by **show ftp-native** and **show ftp-over-http**.
- **show statistics ftp** EXEC command has been replaced by **show statistics ftp-over-http** and **show statistics ftp-native**.
- **clear statistics ftp** EXEC command has been replaced by **clear statistics ftp-over-http** and **clear statistics ftp-native**.

FTP-over-HTTP Caching

ACNS 5.5 software supports proxying and caching of FTP URL client requests using proxy-mode HTTP requests when URLs specify the FTP protocol (for example, `ftp://ftp.mycompany.com/ftkdir/ftp_file`). For instance, the following is an example of an FTP-over-HTTP request that allows the end user to use a browser to access public files from an FTP server:

```
ftp://ftp.funet.fi/pub/cbm/crossplatform/converters/unix/
```

For these requests, the client uses HTTP as the transport protocol with the Content Engine, whereas the Content Engine uses FTP with the FTP server. When the Content Engine receives an FTP request from the web client, it first looks in its cache. If the object is not in its cache, it fetches the object from an upstream FTP proxy server (if one is configured), or directly from the origin FTP server.

The FTP proxy supports anonymous as well as authenticated FTP requests. Only base64 encoding is supported for authentication. The FTP proxy accepts all FTP URL schemes defined in RFC 1738. In the case of a URL in the form `ftp://user@site/dir/file`, the proxy sends back an authentication failure reply and the browser supplies a popup window for the user to enter login information.

The FTP proxy supports commonly used MIME types, attaches the corresponding header to the client, chooses the appropriate transfer type (binary or ASCII), and enables the browser to open the FTP file with the configured application. For unknown file types, the proxy uses binary transfer as the default and instructs the browser to save the downloaded file instead of opening it. The FTP proxy returns a formatted directory listing to the client if the FTP server replies with a known format directory listing. The formatted directory listing has full information about the file or directory and provides the ability for users to choose the download transfer type.

The Content Engine can apply the Rules Template to FTP requests based on server name, domain name, server IP address and port, client IP address, and URL.

The Content Engine logs FTP transactions in the transaction log, in accordance with Squid syntax.

Configuring FTP-Over-HTTP Connection Settings

The Content Engine can handle FTP-style requests over HTTP when configured in proxy mode. When the Content Engine receives an FTP request from a client, it processes the request by searching its cache. If the object is not in its cache, it fetches the object from an upstream FTP proxy server if this proxy server has been configured, or it fetches the object directly from the origin FTP server.

To configure FTP-over-HTTP connection settings by using the Content Distribution Manager GUI, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**. If you have created device groups, the Device Group window appears. (Alternatively, choose **Devices > Devices**.)
- Step 2** Click the **Edit** icon next to the name of the device group (or Content Engine) that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **Applications > Web > FTP > FTP Over HTTP Connections**. The FTP Connection Settings window appears. (See [Figure 8-12](#).) [Table 8-18](#) describes the fields in this window.

Figure 8-12 FTP-Over-HTTP Connection Settings for Device Group Window

FTP Over HTTP Connection Settings for Device Group, DG1

There are currently no settings for this Group.

FTP Over HTTP Connection Settings will be enabled when using Factory Defaults. The values shown in this page are in effect.

FTP Over HTTP Connection Settings	
Enable Proxy Caching:	☐
Enable Incoming Active Mode:	☐
Enable Incoming Proxy:	☐
List of Incoming Ports:	Space separated list of up to 8 ports
Anonymous FTP Password:	Confirm Anonymous FTP Password:
Enable Outgoing Proxy:	☐
Use Origin Server:	☐
Outgoing Connection Timeout:	800000 (microseconds) (200-5000000)
Outgoing Monitor Period:	30 (seconds) (10-300)
Outgoing Proxy Hostname 1:	Outgoing Proxy Port 1:
Outgoing Proxy Hostname 2:	Outgoing Proxy Port 2:
Outgoing Proxy Hostname 3:	Outgoing Proxy Port 3:
Outgoing Proxy Hostname 4:	Outgoing Proxy Port 4:

Submit Cancel

Tuesday, August 02, 2005 13:47:08

- Step 4** To activate the Anonymous FTP Password field, check the **Enable Proxy Caching** check box.
- Step 5** To enable the Content Engine to act in FTP active mode, check the **Enable Incoming Active Mode** check box. If active mode is enabled, the Content Engine attempts to fetch the object in active mode. If active mode fails, the Content Engine attempts to fetch the object again in passive mode.
- Step 6** To accept incoming FTP-over HTTP requests on ports in addition to port 80, check the **Enable Incoming Proxy** check box.

- Step 7** In the List of Incoming Ports field, enter port numbers used by the FTP proxy server or Content Engine to receive requests. The range is from 1 to 65535. Up to eight ports are allowed.
- Step 8** In the Anonymous FTP Password field, enter the anonymous password (for example, `wwwuser@cisco.com`) needed to access the FTP proxy server. The default is `anonymous@hostname`.
- Step 9** In the Confirm Anonymous FTP Password field, reenter the anonymous password needed to access the FTP proxy server.
- Step 10** To allow an outgoing FTP proxy server or another Content Engine to receive FTP cache miss request traffic, check the **Enable Outgoing Proxy** check box.
- Step 11** To use the origin server if all outgoing proxies fail, check the **Use Origin Server** check box. If all the configured proxy servers fail, the Content Engine can optionally redirect FTP requests to the origin server specified in the FTP header if you have checked this check box. If this option is not enabled, the client receives an error response. Response errors and read errors are returned to the client, because it is not possible to detect whether these errors are generated at the origin server or at the proxy.
- Step 12** Enter a value in microseconds in the Outgoing Connection Timeout field. This value specifies the number of microseconds before the outgoing proxy servers are probed to see whether they are active.
- Step 13** To specify how frequently the Content Engine polls the specified outgoing FTP proxy servers, enter a value in the Outgoing Monitor Period field.
- A background process on the Content Engine monitors the state of the configured outgoing proxy servers. You can configure the Content Engine to poll the specified outgoing proxy servers at a specific time in order to monitor their availability. The outgoing monitor period is the time during which the proxy servers are polled. The monitoring period is specified in seconds, and can be from 10 to 300 seconds. The default monitoring period is 60 seconds. If one of the outgoing proxy servers is unavailable, the polling mechanism waits for the connect timeout (300000 microseconds) before polling the next outgoing proxy server.
- Step 14** In the Outgoing Proxy Host Name field, enter the host name or IP address for the outgoing FTP proxy.
- Step 15** In the Outgoing Proxy Port field, enter a port number corresponding to the outgoing proxy host name from [Step 14](#). The range is from 1 to 65535, and up to eight ports are allowed. In proxy mode, the Content Engine accepts and serves FTP requests only on the ports specified. All FTP requests on other proxy-mode ports are rejected.
- The Content Engine supports both anonymous and authenticated FTP requests. In the case of the URL `ftp://user@site/dir/file`, the proxy sends back an authentication failure reply and the browser displays a popup window for you to enter login information.
- Step 16** To save the settings, click **Submit**.
- A “Click to Submit” message appears in red next to the current settings when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The Reset button is visible only when you apply default or device group settings to change the current device settings but the settings have not yet been submitted.

Table 8-18 *FTP-Over-HTTP Connection Settings*

GUI Parameter	Function	CLI Command
Enable Proxy Caching	Enables the Anonymous FTP Password field.	—
Enable Incoming Active Mode	Configures the Content Engine to support FTP active mode. Passive mode is the default.	ftp-over-http proxy active-mode enable

Table 8-18 FTP-Over-HTTP Connection Settings (continued)

GUI Parameter	Function	CLI Command
Enable Incoming Proxy	Configures the Content Engine to accept incoming FTP requests on ports in addition to port 80.	ftp-over-http proxy incoming <i>portnum</i>
List of Incoming Ports	Port numbers used by the proxy server or Content Engine to receive FTP requests. This number ranges from 1 to 65535. You can specify up to 8 ports.	
Anonymous FTP Password	Password required to access the outgoing FTP proxy server.	ftp-over-http proxy anonymous-pswd <i>password</i>
Confirm Anonymous FTP Password	Repeat the password required to access the outgoing FTP proxy server.	—
Enable Outgoing Proxy	Configures an outgoing proxy server or another Content Engine, without using ICP or WCCP, to receive FTP cache miss request traffic.	ftp-over-http proxy outgoing
Use Origin Server	When enabled, handles requests using the origin server if all outgoing proxy servers fail. If disabled, the user receives an error message if all outgoing proxy servers fail.	ftp-over-http proxy outgoing origin-server
Outgoing Connection Timeout	Timeout period (in microseconds) to use when probing for outgoing proxy servers. The range is from 200 to 5000000. The default value is 300 microseconds.	ftp-over-http proxy outgoing connection-timeout <i>microsecs</i>
Outgoing Monitor Period	Specifies the interval in seconds at which to monitor one outgoing proxy server. If multiple outgoing proxy servers are configured, they are monitored sequentially. The default value is 200 seconds.	ftp-over-http proxy outgoing monitor <i>seconds</i>
Outgoing Proxy Hostname 1–8	Configures multiple outgoing proxy servers. Enter the host name or IP address for the outgoing FTP proxy server.	ftp-over-http proxy outgoing host <i>{hostname ipaddress}</i>
Outgoing Proxy Port 1–8	Port number corresponding to the outgoing FTP proxy host names.	ftp-over-http proxy outgoing host <i>{hostname ipaddress}</i> <i>portnum</i>



Note For other operations that you can perform using the icons in the taskbar, see [Table 3-3, “Content Distribution Manager GUI Icons.”](#)

Configuring FTP-Over-HTTP Cache Freshness

You can configure FTP cache object freshness parameters for Content Engines on the ACNS network using the FTP Cache Freshness Settings window from the Content Distribution Manager GUI. These parameters can be configured for either directory listings or particular objects in the cache.

To configure FTP freshness parameters, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**. If you have created device groups, the Device Group window appears. (Alternatively, choose **Devices > Devices**.)
- Step 2** Click the **Edit** icon next to the name of the device group (or Content Engine) that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **Applications > Web > FTP > FTP Over HTTP Cache Freshness**. (See [Figure 8-13](#).) [Table 8-20](#) describes the fields in this window.

Figure 8-13 FTP-Over-HTTP Cache Freshness Settings Window

The screenshot displays the Cisco Application and Content Networking System GUI. The top navigation bar includes 'Devices', 'Device Groups', 'Locations', and 'Statistics'. The left sidebar shows a tree view with 'Contents' expanded, leading to 'Applications > Web > FTP > FTP Over HTTP Cache Freshness'. The main content area is titled 'FTP Over HTTP Cache Freshness Settings for Device Group, DG1' and contains the following settings:

- Enable:** A checkbox that is currently unchecked.
- Directory Listing Age Multiplier:** A text input field with the value '30'.
- File Object Age Multiplier:** A text input field with the value '30'.
- Max TTL Scale:** A dropdown menu set to 'days'.
- Max Directory Listing TTL:** A text input field with the value '3'.
- Max File Object TTL:** A text input field with the value '7'.
- Minimum TTL:** A text input field with the value '30' and a unit selector set to '(minutes)'.
- Enforce Max Object Size:** A checkbox that is currently unchecked.
- Max Object Size:** A text input field with the value '204800' and a unit selector set to '(KB)'.
- Re-eval Requests:** A dropdown menu set to 'none'.

A note at the bottom states: 'Note: * - Required Field'. The bottom right corner features 'Submit' and 'Cancel' buttons.

- Step 4** To enable FTP cache freshness settings on the Content Engine, check the **Enable** check box.
- Step 5** Specify a value for the Directory Listing Age Multiplier field.
The age multiplier value enables the Content Engine to guess the life of a directory by multiplying the time since the object was last modified by a percentage to obtain an approximate expiration date. After this date, the object is considered stale, and subsequent requests for the object cause a fresh retrieval by the Content Engine. Valid values range from 0 to 100 percent. The default value is 30 percent.
- Step 6** Specify a value for the File Object Age Multiplier field.

The age multiplier value enables the Content Engine to guess the life of an object by multiplying the time since the object was last modified by a percentage to obtain an approximate expiration date. After this date, the object is considered stale, and subsequent results cause a fresh retrieval by the Content Engine. Valid values range from 0 to 100 percent. The default value is 60 percent.

- Step 7** Choose **days**, **hours**, **minutes**, or **seconds** from the Max Time to Live (TTL) Scale drop-down list.
- The TTL sets a ceiling on estimated expiration dates. If an object has an explicit expiration date, this takes precedence over the configured TTL.
- Step 8** Specify a value for the Max Directory Listing TTL field. [Table 8-19](#) provides the valid range of values.

Table 8-19 Time To Live Range of Values for FTP Freshness

Scale	Range
Days	1–1825
Hours	1–43800
Minutes	1–2628000
Seconds	1–157680000

- Step 9** Specify a value for the Max File Object TTL field.
- See [Table 8-19](#) for a list of maximum values depending on the scale used.
- Step 10** Specify a value (in minutes) for the Minimum TTL field.
- A minimum TTL is the minimum Time To Live for objects in the cache. The range of values is 0 to 86400 minutes. The default value is 30 minutes.
- Step 11** To enforce the maximum object size specified in the next step, check the **Enforce Max Object Size** check box.
- Step 12** Specify a value in the Max Object Size field.
- This value represents the maximum object size (in kilobytes) that can be stored. The range of values is 1 to 2096128 kilobytes. The default value is set to the maximum value of 2096128 kilobytes.
- Step 13** Choose a method for handling reevaluation requests from the Re-eval Request drop-down list.
- You can apply these FTP cache freshness parameters to the directory listing (**directory-listing**) or to objects as well as the directory listing (**all**). You can also choose not to apply a TTL scale to these parameters (**none**). In this case, a TTL is not applied, and objects in the cache have no expiration dates.
- Step 14** To save the settings, click **Submit**.
- A “Click to Submit” message appears in red next to the current settings when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The **Reset** button is visible only when you apply default or device group settings to change the current device settings but the settings have not yet been submitted.

Table 8-20 *FTP-over-HTTP Cache Freshness Settings*

GUI Parameter	Function	CLI Command
Enable	Enables cache freshness settings.	—
Directory Listing Age Multiplier	Enables the Content Engine to guess the life of a directory by multiplying the time since the object was last modified by a percentage to obtain an approximate expiration date.	ftp-over-http age-multiplier directory-listing <i>dl_time file fo_time</i>
File Object Age Multiplier	Enables the Content Engine to guess the life of an object by multiplying the time since the object was last modified by a percentage to obtain an approximate expiration date.	
Max TTL Scale	Sets the unit of time for the maximum directory listing or file object TTL.	ftp-over-http max-ttl { days directory-listing <i>dlmax_days</i> file <i>fmax_days</i> hours directory-listing <i>dlmax_hours</i> file <i>fmax_hours</i> minutes directory-listing <i>dlmax_min</i> file <i>fmax_min</i> seconds directory-listing <i>dlmax_sec</i> file <i>fmax_sec</i> }
Max Directory Listing TTL	Sets a ceiling on estimated expiration dates for directory listings.	
Max File Object TTL	Sets a ceiling on estimated expiration dates for file objects.	
Minimum TTL	Minimum Time To Live for objects in the cache.	ftp-over-http min-ttl <i>min_minutes</i>
Enforce Max Object Size	Stops the Content Engine from caching objects that exceed the specified maximum size.	ftp-over-http object max-size <i>size</i>
Max Object Size	Maximum object size (in kilobytes) that can be stored.	
Re-eval Requests	Configures how reevaluation requests are to be handled.	ftp-over-http reval-each-request { all none directory-listing }



Note For other operations that you can perform using the icons in the taskbar, see [Table 3-3, “Content Distribution Manager GUI Icons.”](#)

Native FTP Caching

ACNS 5.5 software supports nontransparent native FTP caching, as well as transparent native FTP caching. With this capability, a Content Engine has the ability to handle native FTP requests from FTP clients in nontransparent mode.

When the Content Engine receives a native FTP request from an FTP client (for example, a native FTP request from a Reflection X or WS-FTP client or a UNIX or DOS command line FTP program), the Content Engine processes the request. If the requested content is already stored in its local cache, the Content Engine serves the content to the FTP client. If the content is not already stored in its local cache, the Content Engine has the ability to perform an FTP request to the origin FTP server to retrieve the content, and then store the content in its local cache. Native FTP requests are logged in the HTTP transaction log on the Content Engine.

Restrictions Regarding Native FTP Caching in ACNS 5.5 Software

Restrictions regarding native FTP caching support are as follows:

- Maximum FTP object size of 2096128 kilobytes (2 GB)
- No support for bandwidth control for FTP client requests and FTP server pulls
- No support for the Type of Service (ToS) bit for FTP client requests
- No support for pre-positioned files in cdnfs
- No support for the Internet Content Adaptation Protocol (ICAP)
- No support for proxy authentication
- No support for the Internet Cache Protocol (ICP)
- No support for healing mode
- No support for Layer 4 switch FTP redirection
- No support for FTP request proxy rules
- No support for MIN-TTL and AGING-HEURISTIC-TTL cache control knob configurations
- No support for any URL filtering schemes (good list, bad list, N2H2, Websense, and SmartFilter)
- No support for caching files from a Macintosh FTP server
- No support for “offline” operation for the FTP proxy server

Configuring the Content Engine for Transparent Native FTP Caching

To configure the Content Engine for native FTP caching in transparent mode, you must enable the WCCP ftp-native service (service 60) on the Content Engine. The ftp-native service is the WCCP caching service that permits WCCP Version 2 routers to redirect native FTP requests transparently to a single port on the Content Engine. The Content Engine retrieves the requested FTP content, stores a copy locally (performs native FTP caching), and serves the requested content to the client.

(For information about how to configure the WCCP service for transparent native FTP caching, see the [“Configuring WCCP Service Settings for the Content Engine” section on page 4-10.](#))

Configuring Native FTP Connections

To configure native FTP connections, follow these steps:

-
Step 1	From the Content Distribution Manager GUI, choose Devices > Devices (or Devices > Device Groups).
Step 2	Click the Edit icon next to the name of the Content Engine or device group that you want to configure.
Step 3	In the Contents pane, choose Applications > Web > FTP > Native FTP Connections . The Native FTP Connections Settings window appears. Table 8-21 describes the fields in this window.

- Step 4** To enable the Content Engine to act in native FTP active mode, check the **Enable Active Mode** check box. If active mode is enabled, the Content Engine attempts to fetch the object in active mode, if the FTP client issues an active-mode data transfer request. (See the [“About Native FTP Active and Passive Modes”](#) section on page 8-59.)
- Step 5** To enable native FTP proxy authentication, check the **Enable Authentication** check box. (See the [“About Authenticating Native FTP Requests”](#) section on page 15-15.)
- Step 6** To accept incoming native FTP requests on ports in addition to port 21 (the standard FTP port), check the **Enable Incoming Proxy** check box.
- Step 7** In the List of Incoming Ports field, list up to eight incoming port numbers separated by spaces (for example, 8021, 9021, and so forth).
- Step 8** To save the settings, click **Submit**.

Table 8-21 Native FTP Connection Settings

GUI Parameter	Function	CLI Command
Enable Active Mode	Enables the Content Engine to act in native FTP active mode.	ftp-native proxy active-mode enable
Enable Authentication	Enables native FTP proxy authentication. The default is disabled. (See the “Considerations When Configuring Native FTP Proxied Authentication” section on page 15-16.)	ftp-native proxy authentication enable
Enable Incoming Proxy	Enables incoming native FTP requests on ports in addition to port 21.	ftp-native proxy incoming <i>portnum</i>
List of Incoming Ports	Port numbers on which the Content Engine will accept incoming requests (native FTP requests) from FTP clients. Valid port numbers are 1 to 65535. You can specify up to eight incoming ports.	<i>portnum ...</i>

About Native FTP Active and Passive Modes

A Content Engine that is operating as an FTP proxy supports passive and active modes for fetching files and directories. In native FTP caching mode, if active mode is enabled, then the Content Engine uses the same mode with the FTP server for the data connection as the client used to access the Content Engine, which can be either active or passive.

If active mode is not enabled, the Content Engine uses passive mode with the origin FTP server for the data connection.

If the Content Engine adheres to the client’s mode (active or passive) for native FTP, the following occurs:

- The Content Engine (the native FTP proxy server) performs an active-mode data transfer to or from the origin FTP server if the FTP client issues an active-mode data transfer request.
- The Content Engine performs a passive-mode data transfer to or from the FTP server if the FTP client issues a passive-mode data transfer request.

Note that the format of the URL that the Content Engine (a nontransparent proxy server that is functioning as a native FTP proxy server) creates for a native FTP request depends on the FTP login name and the transfer mode (binary or ACSII file transfer mode).

- If the FTP login name is an actual username instead of “anonymous,” then the string “*user*:password*@" is included in the URL before the host.
- If the mode used to transfer the file is binary mode, then the string “;type=i” is included at the end of the URL. The following is an example of the URL format that the Content Engine creates for a specific user when binary mode is being used.

```
ftp://*user*:password*@10.100.200.5/home/myhome/mybinfile.obj;type=i
```

The URL for an “anonymous” user login and an ACSII file transfer mode will not have any fields embedded in the URL, as shown in the following example:

```
ftp://10.100.200.5/home/myhome/mytextfile.txt
```



Note

Passive mode transfer between the FTP client and the native FTP proxy cannot be performed if the Content Engine is behind a NAT because the NAT firewall blocks the Content Engine from discovering its conduit IP address.

Configuring the Client Side of Nontransparent Native FTP Caching

Content Engines that are acting as nontransparent FTP proxy servers can accept native FTP requests from such FTP clients as Reflection X clients, WS-FTP clients, and Unix or DOS command line FTP programs. This section provides some examples of how to configure these different types of FTP clients to send native FTP requests directly to the Content Engine (that is acting as a nontransparent FTP proxy server for these FTP clients). Refer to the *Cisco ACNS Software Configuration Guide for Locally Managed Deployments, Release 5.5* for information on how to configure the client side proxy FTP request to the Content Engine.

The following example shows how to use a UNIX command line FTP program to configure the client side proxy FTP request to the Content Engine that is acting as the nontransparent FTP proxy server. In this example, 10.9.192.10 is the IP address of the FTP client, and 10.9.192.11 is the IP address of the Content Engine. The IP address of the server, abchost.company.com, does not get displayed at the FTP client.

In this example, the FTP client is opening a local port (on the FTP client machine) to which the Content Engine connects for the data transfer. When the client opens a local port, the data transfer is said to be an “active-mode” transfer. (Not shown in this example is “passive-mode” data transfer, where the Content Engine opens a port on the Content Engine and requests the client to connect to that port.)

```
shell# ftp -d 10.9.192.11 8501
Connected to 10.9.192.11.
220-Welcome to FTP-proxy.
220-Login to origin server using the 'USER username@server-hostname' command,
or 220 Login to origin server using the 'SITE server-hostname' and
'USER username' commands.
Name (10.9.192.11:admin): smartuser@abchost.company.com
331 Password required for smartuser.
Password:
230 User smartuser logged in.
Remote system type is UNIX.
Using binary mode to transfer files.
ftp> dir
---> PORT 2,9,192,10,81,212
200 PORT command successful.
```

```

---> LIST
150 Opening ASCII mode data connection for /bin/ls.
total 8
-rw-r--r-- 1 nobody abc 5496 Jan 31 2002 boot
-rw-r--r-- 1 nobody abc 143 Oct 24 2001 cop
226 Transfer complete.
ftp> quit
---> QUIT
shell#

```

Configuring Native FTP Cache Freshness

To configure native FTP cache freshness settings, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices** (or **Devices > Device Groups**).
 - Step 2** Click the **Edit** icon next to the name of the Content Engine or device group that you want to configure.
 - Step 3** In the Contents pane, choose **Applications > Web > FTP > Native FTP Cache Freshness**. The Native FTP Cache Freshness Settings window appears.
 - Step 4** Specify a value in the Max Object Size field.
This value represents the maximum object size (in kilobytes) that can be stored. The range of values is 1 to 2096128 kilobytes. The default value is 2096128 kilobytes.
 - Step 5** To save the settings, click **Submit**.
-

To configure native FTP cache freshness settings from the CLI, use the **ftp-native object max-size** command.

Enabling Inetd FTP Service

Inetd (an Internet daemon pronounced eye net dee) is a program that listens for connection requests or messages for certain ports and starts server programs to perform the services associated with those ports.

To enable Inetd FTP service, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**. If you have created device groups, the Device Group window appears. (Alternatively, choose **Devices > Devices**.)
 - Step 2** Click the **Edit** icon next to the name of the device group (or Content Engine) that you want to configure. The Contents pane appears on the left.
 - Step 3** From the Contents pane, choose **Applications > Web > FTP > Inetd FTP**. The Inetd FTP Settings window appears.
 - Step 4** To enable Inetd FTP service on the Content Engine, check the **Inetd Enable FTP Service** check box.
 - Step 5** To save this setting, click **Submit**.
-

To enable Inetd FTP service from the CLI, use the **inetd enable ftp** global configuration command.

Configuring Native FTP Custom Messages

ACNS 5.5 software supports custom welcome messages and ACL-denied messages for FTP proxy. You can download a custom message file to the Content Engine and upload the message to a remote host server. The welcome message is the FTP response to the client when the client connects to the Content Engine in proxy mode. The ACL-denied message is the FTP response to the client when the client is denied access through a configured access control list. (For more information about configuring access control lists, see [Chapter 17, “Creating and Managing IP Access Control Lists.”](#))

About Writing FTP Custom Welcome Messages

If you intend to enable proxy authentication, then the welcome message should inform the user to authenticate with the proxy before logging in to the origin server, as shown in the following example:

```
Welcome to ce-Boxman. BigCorp's Content Engine for Native FTP Proxy. Please login to the proxy with your username and password.
```

Alternatively, if you do not intend to enable proxy authentication, then the welcome message should inform the user to log in to the origin server using either the USER or SITE method, as shown in the following examples:

```
Welcome to ce-Boxman. BigCorp's Content Engine for Native FTP Proxy. Please login to the origin server using the 'username@server-hostname' format.
```

```
Welcome to ce-Boxman. BigCorp's Content Engine for Native FTP Proxy. Please login to the origin server using the 'SITE server-hostname' command followed by the 'USER username' command.
```

Configuring FTP Message Download Settings

You can configure welcome messages and ACL-denied messages to an entire device group, or you can configure messages for an individual Content Engine. Only one of each message type can be configured for a Content Engine; therefore, if messages are configured for your Content Engines through the device group, you cannot configure a new FTP message for an individual Content Engine in the group. The Content Distribution Manager GUI issues a warning message when you attempt to create an FTP message where one already exists for the Content Engine.

To download a message to the Content Engine, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**.
 - Step 2** Click the **Edit** icon next to the device group that you want to configure.
 - Step 3** In the Contents pane, choose **Applications > Web > FTP > FTP Messages > Download**. The FTP Message Download for Device Group window appears.
 - Step 4** From the FTP Message drop-down list, choose the message type: **welcome** or **acl-denied**.
 - Step 5** In the Download URL field, enter a URL for the message file that you want the Content Engine to fetch. Only HTTP, HTTPS, and FTP protocols are supported.
 - Step 6** To save the settings, click **Submit**. A message appears stating that the size of the file should not be greater than 16 kilobytes. To continue, click **OK**. The Content Engine retrieves the message file from the URL.
-

To download native FTP custom messages from the CLI, use the **ftp-native custom-message download EXEC** command.

Viewing the Content Engine FTP Message Configurations

To view the FTP message configurations for an individual Content Engine, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon for the Content Engine that you want to view.
 - Step 3** In the Contents pane, choose **Applications > Web > FTP > FTP Messages > Download**. The FTP Message Download for Content Engine window appears.
 - Step 4** To view all of the FTP message configurations, including any that are configured through the Content Engine's device group and any configured on the Content Engine only, click the **Aggregate Settings: Yes** radio button.
 - Messages configured on the Content Engine only are identified with an Edit icon next to the listing.
 - Messages configured through the Content Engine's device group are identified with a View only icon next to the listing.
 - Step 5** To view only FTP messages configured on the individual Content Engine, click the **Aggregate Settings: No** radio button.
-

Resetting a Native FTP Custom Message

Use the reset (or delete) option when you want to change a native FTP custom welcome message or an ACL-denied message. To reset a message, follow these steps:

-
- Step 1** To reset a message configuration from an individual device, choose **Devices > Devices**, and click the **Edit** icon for the Content Engine.
 - Step 2** In the Contents pane, choose **Applications > Web > FTP > FTP Messages > Download**. The FTP Message Download for Content Engine window appears.
 - Step 3** To reset all messages configured for the individual Content Engine, click the **Trash** icon in the taskbar. To confirm your action, click **OK**.
 - Step 4** To reset a welcome message only, click the **Edit** icon for the welcome message. In the Modifying FTP Message for Content Engine window, click the **Trash** icon in the taskbar. To confirm your action, click **OK**.
 - Step 5** To reset an ACL-denied message only, click the **Edit** icon for the ACL-denied message. In the Modifying FTP Message for Content Engine window, click the **Trash** icon in the taskbar. To confirm your action, click **OK**.
 - Step 6** To reset message configuration for a device group, choose **Devices > Device Groups**.
 - Step 7** Click the **Edit** icon next to the device group that you want to configure.
 - Step 8** In the Contents pane, choose **Applications > Web > FTP > FTP Messages > Download**. The FTP Message Download for Device Group window appears.
 - Step 9** To reset all messages configured for the device group, click the **Trash** icon in the taskbar. To confirm your action, click **OK**.

- Step 10** To reset a welcome message only, click the **Edit** icon for the welcome message. In the Modifying FTP Message for Device Group window, click the **Trash** icon in the taskbar. To confirm your action, click **OK**.
- Step 11** To reset a ACL-denied message only, click the **Edit** icon for the ACL-denied message. In the Modifying FTP Message for Device Group window, click the **Trash** icon in the taskbar. To confirm your action, click **OK**.

To reset a native FTP custom message from the CLI, use the **ftp-native custom-message reset** {acl-denied | all | welcome} EXEC command.

Configuring Native FTP Custom Message Upload Settings



Note

Upload settings can only be configured after you have configured the download settings for the message file. Upload settings can only be configured for individual devices; the upload functionality is not supported for device groups.

To upload a custom message file from the Content Engine to a remote host server, configure the native FTP custom message upload settings by following these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the Content Engine that you want to configure.
- Step 3** In the Contents pane, choose **Applications > Web > FTP > FTP Messages > Upload**. The FTP Message Upload Settings for Content Engine window appears.
- Step 4** From the FTP Message drop-down list, choose a message type: **welcome** or **acl-denied**.
- Step 5** To upload the custom message file to the specified host, directory, and filename using the FTP protocol, complete the following settings:
- In the FTP Server Address field, enter the hostname or IP address of the remote host to which you want the message file uploaded.
 - In the FTP Server Directory field, enter the name of the directory on the host server.
 - In the FTP Server Filename field, enter the filename.
- Step 6** In the FTP Server Username field, enter the username for accessing the remote FTP server.
- Step 7** In the FTP Server Password field, enter the password for accessing the remote FTP server.
- Step 8** To confirm the password, re-enter the password in the Confirm Password field.
- Step 9** To save the settings, click **Submit**.
-

To upload native FTP custom messages from the CLI, use the **ftp-native custom-message upload** EXEC command.

Configuring TFTP Settings

The Trivial File Transfer Protocol (TFTP) gateway feature provides a way for Content Engines to serve content files requested by networking devices that use the native TFTP protocol. Content Engines that use ACNS software perform TFTP to HTTP or FTP translation and eliminate the need for the system administrator to configure and manage a dedicated TFTP server to serve TFTP requests. This feature is called TFTP gateway because it allows the Content Engine to accept native TFTP requests from the client at the front end and serve the request using HTTP or FTP protocol at the back end.

Content files include router software images, router configurations, set top box images, IP phone configuration files, and so forth. If the requested file is not available on the Content Engine, the Content Engine, acting on behalf of the requesting device, retrieves the file from the origin server and caches it. The Content Engine then forwards the file to the device that originated the request. Future requests by any devices for the same file are satisfied by forwarding the file from the Content Engine cache.

You can configure, modify, and view TFTP settings for Content Engines and device groups by completing the following tasks:

- [Configuring TFTP General Settings, page 8-65](#)
- [Configuring TFTP Proxy Server Settings, page 8-66](#)
- [Configuring TFTP Directory Settings, page 8-69](#)

Configuring TFTP General Settings

Content Engines serve TFTP requests from clients by accepting native TFTP requests and retrieving content using FTP or HTTP from locally configured directories or origin servers. Content devices such as set top boxes and routers make TFTP requests for downloading router software images, set top box images, and router configurations. Content Engines allow TFTP-to-HTTP or TFTP-to-FTP translation without the need for configuring and managing a dedicated TFTP server to serve TFTP requests.

To enable the TFTP service on the Content Engine to serve TFTP requests from networking devices, follow these steps:

-
Step 1	Choose Devices > Devices . The Devices window appears.
Step 2	Click the Edit icon next to the desired Content Engine. The Modifying Content Engine window appears with the Contents pane on the left.
Step 3	In the Contents pane, choose Applications > Web > TFTP > TFTP General Settings . The TFTP Settings for Content Engine window appears. Table 8-22 describes the fields in this window and provides the corresponding CLI global configuration commands.
Step 4	To enable the TFTP service application to run on the specified Content Engine, check the Inetd Enable TFTP Service check box. Checking this check box enables the Content Engine to serve TFTP requests by retrieving the requested content. If you uncheck this check box, TFTP service is disabled on the Content Engine.
Step 5	To configure the maximum number of TFTP requests that the Content Engine will service simultaneously, enter a value in the Maximum Number of TFTP Connections field. The default is 200 requests. The range is 0–300.
Step 6	To save the settings, click Submit .

A “Click to Submit” message appears in red next to the current settings when there are pending changes to be saved. You can also revert to the previously configured window settings by clicking **Reset**. The Reset button is visible only when you apply default or device group settings to change the current device settings, but the settings have not yet been submitted.



Note For other operations that you can perform using the icons in the taskbar, see [Table 3-3, “Content Distribution Manager GUI Icons.”](#)

Table 8-22 *TFTP General Settings*

GUI Parameter	Function	CLI Command
Inetd Enable TFTP Service	Enables TFTP service on the Content Engine.	inetd enable tftp
Maximum Number of TFTP Connections	Maximum number of TFTP requests that the Content Engine will service simultaneously.	tftp max-connections <i>number</i>

Configuring TFTP Proxy Server Settings

In the ACNS software, the Content Engine can act as a TFTP server and process TFTP requests sent to the Content Engine. In addition, the Content Engine can act as a TFTP gateway, allowing users to obtain files from HTTP or FTP servers in response to a TFTP request.



Note The Content Engine does not support transparent TFTP requests. It only accepts TFTP requests that explicitly contain the Content Engine hostname or IP address.

When a requested file does not exist on the Content Engine, the Content Engine makes an HTTP or FTP request to the origin server to retrieve the content. The requested content is cached and sent to the TFTP client. Future requests for the same content from TFTP clients are served by forwarding the file from the cfs or cdnfs directory on the Content Engine.



Note Authenticated objects are never cached by the Content Engine for HTTP or FTP. If you want to cache these objects, leave the username and password fields in the **tftp-server gw** command blank. When used with FTP, this configuration is equivalent to allowing anonymous access.

You can configure a maximum of two origin servers: a primary and a secondary backup server. The TFTP server application running on the Content Engine attempts to retrieve content from the secondary backup server, if configured, when the primary server fails. You can associate priorities with the two origin servers to categorize them as primary or secondary servers.

To configure the TFTP proxy server settings, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the desired Content Engine. The Modifying Content Engine window appears with the Contents pane on the left.
- Step 3** In the Contents pane, choose **Applications > Web > TFTP > TFTP Proxy**. The TFTP Proxy for Content Engine window appears.
- Step 4** Click the **Create New TFTP Proxy** icon in the taskbar. The Creating New TFTP Proxy for Content Engine window appears. (See [Figure 8-14](#).) [Table 8-23](#) describes the fields in this window and provides the corresponding CLI global configuration commands.

Figure 8-14 Creating New TFTP Proxy for Content Engine Window



Note

Only two TFTP proxy servers can be defined. When you attempt to create a third TFTP proxy server, the system displays a popup window stating that only two TFTP proxy servers can be defined because each server needs to be configured with a unique priority value.

- Step 5** From the Protocol drop-down list, choose the protocol (**HTTP** or **FTP**) that the proxy server will use to serve TFTP requests.
- Step 6** In the TFTP Gateway Server field, enter the host name or IP address of the TFTP proxy server.
- Step 7** To change the default port setting for the HTTP protocol, enter a port number in the Port field.

- Step 8** From the Priority drop-down list, choose the priority (**LOW** or **HIGH**) with which the origin servers are queried for content.
- Priority values categorize a server as a primary or a secondary backup server. The proxy server with the higher priority is the primary origin server and is queried first for content. You cannot configure two origin servers with the same priority, whether they are of the same or different protocols.
- Step 9** In the Directory Path field, enter the path to the content that exists on the proxy server. You can enter the complete path or a relative path that specifies the directory in which the requested content resides.
- Step 10** To specify user authentication details for logging in to the proxy server, check the **Configure User** check box. The user details cannot be entered if this check box is not checked.
- Step 11** In the User Name field, enter the login ID of the user to be used on the remote server.



Note Authenticated objects are never cached by the Content Engine for HTTP or FTP. If you want to cache these objects, leave the username and password fields in the **tftp-server gw** command blank. When used with FTP, this configuration is equivalent to allowing anonymous access.

- Step 12** In the Password field, enter the user password to be used to authenticate users who access the proxy server. Reenter the same password in the Confirm Password field. The password details are encrypted in the display.
- Step 13** To save the settings, click **Submit**.

Table 8-23 TFTP Proxy Settings

GUI Parameter	Function	CLI Command
Protocol	Protocol (HTTP or FTP) that the proxy server will use to serve TFTP requests.	tftp-server gw proto {http ftp} server {hostname ipaddress} pri priority [name name passwd password] [path directory]
TFTP Gateway Server	Host name or IP address of the TFTP proxy server.	tftp-server gw proto http server {hostname ipaddress} [port port_num]
Port	Port that the proxy server will use for the HTTP protocol. The default port for HTTP is port 80.	
Priority	priority (LOW [2] or HIGH [1]) with which the origin servers are queried for content.	
Directory Path	Path to the content that exists on the proxy server. You can enter the complete path or a relative path that specifies the directory in which the requested content resides.	
Configure User	Enables fields to specify user authentication details for logging in to the proxy server.	
User Name	Login ID of the user to be used on the remote server.	

Table 8-23 TFTP Proxy Settings (continued)

GUI Parameter	Function	CLI Command
Password	Password to authenticate users who access the proxy server.	tftp-server gw proto {http ftp} server {hostname ipaddress} pri priority [name name passwd password] [path directory]
Confirm Password	—	


Configuring TFTP Directory Settings

When a TFTP client makes a request for content, the TFTP server application running on the Content Engine searches the specified directory to serve the content. If the request does not contain any directory path, the default directory is used. You can configure a maximum of eight TFTP directory configurations that TFTP clients can access.

To configure the directories to be used by TFTP, follow these steps:

-
- Step 1** Choose **Devices > Devices**. The Devices window appears.
 - Step 2** Click the **Edit** icon next to the desired Content Engine. The Modifying Content Engine window appears with the Contents pane on the left.
 - Step 3** In the Contents pane, choose **Applications > Web > TFTP > TFTP Directory**. The TFTP Directory Settings for Content Engine window appears. (See [Figure 8-15](#).)

Figure 8-15 TFTP Directory Settings Window



- Step 4** In the TFTP Directory Settings section, check the **Configure TFTP Directories** check box to enable configuration of TFTP directories.

If this check box is unchecked, the remaining fields in the window are disabled. By default, this check box is unchecked. When the Configure TFTP Directories check box is unchecked, all configured TFTP directories are removed.

- Step 5** To configure the default directory, enter the complete absolute path to the directory to be used by the TFTP server application in the Directory 1 field. The first directory configured in the list of TFTP directories becomes the default directory.

You can configure seven more TFTP directories. You cannot duplicate directory paths.

- Step 6** To save the settings, click **Submit**.

A “Click to Submit” message appears in red next to the current settings when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The Reset button is visible only when you apply default or device group settings to change the current device settings, but the settings have not yet been submitted.



Note For other operations that you can perform using the icons in the taskbar, see [Table 3-3, “Content Distribution Manager GUI Icons.”](#)

To configure TFTP directories from the CLI, use the **tftp-server dir** *directory* global configuration command.

Configuring DNS Caching Name Service for WCCP Transparent Caching

For transparent interception of DNS requests using WCCP, administrators must configure the DNS caching name service (WCCP service group number 53) on the Content Engine and on a router that supports WCCP Version 2.

The DNS process interacts with the WCCP process in these ways:

- Maintaining the bypass lists.
- Monitoring the active state of the DNS process to make sure that the DNS can accept requests. If the DNS cache has no servers that are responsive, the cache will deregister the service until it has acceptable servers.
- Configuring and managing the WCCP service (the DNS caching name service).

In a centrally managed ACNS network that employs one or more Content Routers, configuring the DNS caching name service with WCCP interception on registered Content Engine causes a conflict with the Content Router because the Content Router and the Content Engine will be listening for DNS requests on the same port (port 53). Consequently, you should not configure DNS cache support with WCCP interception in centrally managed environments that employ Content Routers. You can, however, enable the standard DNS caching name service (without WCCP interception support) in such environments.

The advantages of having the Content Engines intercept DNS requests include the following:

- You can control the servers used through WCCP redirect lists, through bypass entries, or by simply having the Content Engine intercept requests to servers not in use and then reissuing them to others.
- DNS servers can be easily updated and changed as needed.
- Performance is improved (faster response time and lower bandwidth consumption) because of the DNS cache hits.

You can partition your network with a series of forwarders, each with its own DNS cache, which results in the following:

- You can effectively control the DNS traffic across a WAN.
- You can be better equipped to compensate for the loss of one or more servers because the Content Engine can issue the request to a different upstream server.

Typically, this partitioning also results in a performance improvement if a cluster of requests occurs, especially on high-latency interconnections.

Configuring DNS Caching Server Settings for the Content Engine

To configure the DNS caching server settings for the Content Engine, follow these steps:

-
Step 1	Choose Devices > Devices . The Devices window appears, listing all the devices registered in the ACNS network.
Step 2	Click the Edit icon of the Content Engine for which you want to configure DNS caching server settings. The Device Home for Content Engine window appears.

In the Contents pane, choose **Applications > DNS > DNS Caching Server**. The DNS Caching Server Settings for Content Engine window appears. [Table 8-25](#) describes the fields in this window and provides the corresponding CLI global configuration commands.

- Step 3** Check the **Enable DNS Caching** check box to start the DNS server for resolution of domain names to IP addresses after the listener port is configured. (See the “[Configuring DNS Server Bindings for the Content Engine](#)” section on page 8-75.) Enabling the DNS server creates an entry of 127.0.0.1 as the name server for the system and starts the memory-based DNS cache on port 53.



Note To enable transparent interception of DNS requests using WCCP, you must configure the WCCP DNS caching name service (the WCCP service with service ID 53) on the Content Engine and the WCCP Version 2 router.

- Step 4** Check the **Lookup Name Servers Recursively** check box to allow each configured name server to be queried in turn if the request to the primary name server fails. This option is relevant only if you have configured the DNS cache to use the DNS servers configured on the Content Engine as the original request contains only a single DNS server. (See [Step 6](#) for more details.)

- Step 5** Check the **Use Expired Cache Entries** check box to use a resource record in the DNS cache, even if it has expired.

When checked, the DNS cache uses an expired resource record, if the record resides in its cache, to serve the user request. At the same time, the DNS cache initiates a new lookup to revalidate the cached object. This serving of expired resource records is particularly useful when the upstream DNS is slow or unresponsive. The expired record will continue to be used until the entry is updated by a new record, or is purged from the cache using the Least Recently Used (LRU) algorithm.

- Step 6** Choose the method that the Content Engine (DNS caching name server) uses to handle transparent DNS requests from the Use Origin Server drop-down list. [Table 8-24](#) describes the options in this list.

You can configure the method that the DNS caching name server uses to handle transparent DNS requests redirected from a WCCP-enabled router. These transparent DNS requests contain the original destination server in the packet headers. However, because the Content Engine can have more than one DNS name server configured, the DNS caching name server can be configured to use either the original DNS server specified in the original request packet headers or the DNS name servers configured on the Content Engine, or a combination of both.

Table 8-24 Use Origin Server Options

Option	Description
Do Not Set	Uses only the DNS servers configured on the Content Engine. This option is selected by default.
Only	Uses only the DNS server identified in the original request.
After-configured	Uses the DNS name servers configured on the Content Engine. If these fail, uses the DNS server specified in the original request packet.
Before-configured	Uses the DNS server specified in the original request. If this fails, uses the DNS name servers configured on the Content Engine.

- Step 7** In the Maximum Size of Cache field, specify the maximum size of the DNS cache memory in megabytes. The default value for the allocated memory used by the cache is 5 MB. It is important that you impose a strict maximum memory limit within which the DNS server operates so as not to unduly tax the overall system resources.

- Step 8** In the Minimum TTL field, specify the minimum time to store a resource record in the cache. The default is 1 second. After the time-to-live elapses, the cached response must be revalidated. If a response is served from the cache, for that response to be valid, all appropriate resource records must be valid in the cache. Minimum time-to-live must be lower than maximum time-to-live. Otherwise, an error message is displayed on submit.
- Step 9** In the Maximum TTL field, specify the maximum time to store a resource record in the cache. The default is 86400 seconds. The cached resource record expires after this period. Maximum time-to-live must be greater than minimum time-to-live. Otherwise, an error message is displayed on submit.
- Because the DNS protocol is using UDP, packets can be lost or dropped. The burden of retransmitting DNS requests is on the requester. Typically, a retransmit is initiated every 3 seconds until a response is received, or if a response is not received, the request times out after 90 seconds. If a DNS server times out, then a new upstream server is selected to query. If there are no more servers to query upstream, then the server returns a DNS failed response to the requesting client.
- Step 10** In the Retry Period field, specify the time period before an unanswered request is discarded. The default is 90 seconds.
- Step 11** In the Retry Timeout field, specify the time between retransmission of UDP DNS requests sent to an upstream DNS server. The default is 3 seconds.
- Step 12** To save the settings, click **Submit**. A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking **Reset**. The **Reset** button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

Table 8-25 DNS Caching Server Settings

GUI Parameter	Function	CLI Command
Enable DNS Caching	Enables the DNS server for resolution of domain names to IP addresses after the listener port is configured.	dns enable
Lookup Name Servers Recursively	Allows each configured name server to be queried in turn if the request to the primary name server fails.	dns serial-lookup http dns-cache serial-lookup
Use Expired Cache Entries	DNS uses a resource record in the DNS cache, even if it has expired.	dns user-expired enable
Use Origin Server	Configures the method that the DNS caching name server uses to handle transparent DNS requests redirected from a WCCP-enabled router.	dns use-original-server {only after-configured before-configured}
Maximum Size of Cache	Maximum size (in megabytes) of the DNS cache memory.	dns max-cache-memory <i>mbytes</i>
Minimum TTL	Minimum time to store a resource record in the cache.	dns min-ttl <i>seconds</i>
Maximum TTL	Maximum time to store a resource record in the cache.	dns max-ttl <i>seconds</i>

Table 8-25 *DNS Caching Server Settings (continued)*

GUI Parameter	Function	CLI Command
Retry Period	Time period before an unanswered request is discarded.	dns retry-period <i>seconds</i>
Retry Timeout	Time between retransmission of UDP DNS requests sent to an upstream DNS server.	dns retry-timeout <i>seconds</i>

Cacheable Resource Records in the Content Engine DNS Cache

Content Engines with ACNS software can intercept DNS requests that have been issued to DNS servers (forwarders or recursive) that traverse the interception device. Upon receiving a request, the Content Engine transmits that response to the client if it has a current cached response. For a cached response to be valid, it must have all the appropriate resource records for the required sections (Answer, Authority, Additional). If any of these resource records have been purged or have expired the request is queued for transmission.

[Table 8-26](#) describes the resource records that can be cached in the Content Engine DNS cache.

Table 8-26 *Cacheable Resource Records*

Record	Description
A	Address record
HINFO	Host information
NULL	Null record
TXT	Text string
WKS	Well known service description
NS	Name server record
CNAME	Canonical name
MB	Mail box domain name
MD	Mail destination
MF	Mail forwarder
MG	Mail group member
MR	Mail rename domain
PTR	Pointer record
MX	Mail exchange
SOA	Source of authority

Although other resource records are not cached, they are passed through to the client requesting them. The DNS caching name service does support compression of the DNS response as defined by RFC 1035, including those responses generated by the cache. If the cache does not understand the server response to the query, it does not cache that response but passes it through to the client. This action reduces the chance of a nonstandard client or server being disrupted.

Configuring DNS Server Bindings for the Content Engine

To configure DNS server bindings for the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the devices registered in the ACNS network.
- Step 2** Click the **Edit** icon of the Content Engine for which you want to configure IP addresses and port numbers that the DNS cache uses to listen for requests. The Device Home for Content Engine window appears.
- Step 3** In the Contents pane, choose **Applications > DNS > DNS Server Binding**. The DNS Server Bindings for Content Engine window appears.
- Step 4** Click the **Create New Server Binding** icon in the taskbar to configure IP address and port number pair for the DNS server to listen for client requests. The Creating New DNS Server Binding for Content Engine window appears.



Note You can create a new DNS server binding only if the DNS server is not listening on configured interface IP addresses for client requests and if the DNS server binding entries have not reached the maximum limit of 16 entries.

If the DNS server is already listening on all the interface IP addresses, then the following warning message appears: “DNS Service is already listening on all IP Addresses.”

If the DNS server binding entries have already reached their maximum limit of 16, then the following warning message appears: “DNS Server Binding Entries have reached their maximum limit of 16.”

- Step 5** Configure the DNS caching name server to listen either on a specific interface IP address or to listen on all the IP addresses configured on the Content Engine.
 - To configure the DNS caching name server to listen on a specific interface IP address, click the **Listen on IP Address** radio button and choose an IP address from the drop-down list. All IP addresses associated with the primary interface or other interfaces configured on the Content Engine are listed here. You can choose an IP address of the interface from the drop-down only after you click the radio button.
 - To configure the DNS caching name server to listen on all the IP addresses configured on the Content Engine, click the **Listen on all IP Addresses** radio button. Once the host name has been resolved to an IP address, it is stored in the memory-based DNS cache.
- Step 6** In the Port field, specify the port on which the DNS caching server listens for client requests. The range is 1–65535. Ensure that the configured port number is not already used by any other application running on the Content Engine.

**Note**

If **Listen on IP Address** is configured, then the DNS caching server listens on the specific IP address and on this port. If **Listen on All IP Addresses** is configured, then the DNS caching server listens on all IP addresses configured on the Content Engine and on this port.

- Step 7** In the Hostname field, specify the host name of the listener to be mapped to an IP address.
- Step 8** To save the settings, click **Submit**. A dialog box alerts you that if the configured port number is being used by any other application running on the Content Engine, the DNS cache listener will fail.
- If you have configured **Listen on All IP Addresses**, a dialog box alerts you that configuring the DNS cache server to listen on all IP addresses for client requests will remove all other DNS server binding settings. Click **OK** to confirm.
- The DNS Server Bindings for Content Engine window appears with the configured DNS cache listener settings.

To configure the DNS caching name server to listen either on a specific interface IP address or to listen on all the IP addresses configured on the Content Engine from the CLI, use the following global configuration command:

```
dns listen {all | ipaddress} port port_num hostname hostname
```

Configuring DNS Address Record Mappings for the Content Engine

The DNS server must know the DNS name of the host on which it is being enabled and map the name to an IP address within its own cache. If the DNS cache listener does not match a DNS name, use the mapping type to pin an IP address to name mapping.

To configure the DNS host name to IP address mapping for the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon of the Content Engine for which you want to statistically map the IP addresses and host names. The Device Home for Content Engine window appears.
- Step 3** In the Contents pane, choose **Request Routing > DNS > A-Record Mapping**. The DNS A-Record Mappings for Content Engine window appears.
- Step 4** Click the **Create New DNS A-Record Mappings** icon in the taskbar. The Creating New DNS A-Record Mapping for Content Engine window appears.
- When you try to configure either a host name or a canonical name (canonical names are configured from the DNS CNAME Record Mappings window) and the total number of pinning configurations has reached its maximum limit of 256, an error message is displayed.
- Step 5** Choose the method to allow you to look up an IP address against a name within the cache from the Mapping Type drop-down list. [Table 8-27](#) describes the options in the drop-down list.

Table 8-27 Mapping Type Options

Option	Description
Forward	Maps the host name to the IP address
Reverse	Maps the IP address to the host name
Both	Maps in both forward and reverse directions

- Step 6** In the Hostname field, specify the host name to be mapped to an IP address. Host names cannot be duplicated, must begin with an alphabet character, and cannot exceed 256 characters.
- Exceptions to host name duplication are as follows:
- If a host name has been configured with forward mapping type, you can configure the same host name with reverse mapping type.
 - Conversely, if a host name has been configured with reverse mapping type, you can configure the same host name with forward mapping type.
- Step 7** In the IP Address fields, specify the IP address of bidirectional, forward, or reverse mapping. When a client query arrives, the configured host name will be resolved to any one of the IP addresses specified here. The first IP address is a required entry.



Note You must specify IP addresses in sequence. Otherwise, an error message is displayed on submit.

- Step 8** To save the settings, click **Submit**.
- If you click the **Delete DNS A-Record Mapping** icon to delete an address record mapped to a canonical name, an error message appears. You need to remove the mapping before you attempt to delete the address record.

To configure the DNS host name to IP address mappings from the CLI, use the following global configuration command:

```
dns pin {forward hostname ipaddress | reverse hostname ipaddress | both hostname ipaddress}
```

Configuring DNS Canonical Name Record Mappings for the Content Engine

The DNS server must know the DNS name of the host on which it is being enabled and map the name to an IP address within its own cache. If the DNS cache listener does not match a DNS name, map the canonical name resource record to the address record after pinning an IP address to name mapping.

To configure the mapping of DNS canonical name to address records for the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the devices registered in the ACNS network.
- Step 2** Click the **Edit** icon of the Content Engine for which you want to insert canonical name (CNAME) mapping. The Device Home for Content Engine window appears.
- Step 3** In the Contents pane, choose **Request Routing > DNS > CNAME Record Mapping**. The DNS CNAME Record Mappings for Content Engine window appears.

- Step 4** Click the **Create New DNS CNAME Record Mappings** icon in the taskbar. The Creating New DNS CNAME Record Mapping for Content Engine window appears.



Note When you try to configure a canonical name and the total number of pinning configurations has reached its maximum limit of 256, an error message is displayed.

- Step 5** In the Canonical Name field, specify the canonical name to be mapped to the address resource record. You can map a canonical name with a maximum of eight address records. Canonical names cannot be duplicated, must begin with an alphabet character, and cannot exceed 256 characters.

- Step 6** Choose the address record to be looked up against the canonical name within the cache from the A-Record drop-down list. The first A-record is a required option.

The A-record drop-down list contains all configured forward and bidirectional mapping type address records. You can associate one or more unique A-records with the CNAME record.



Note You must specify A-records in sequence. Otherwise, an error message is displayed on submit.

- Step 7** To save the settings, click **Submit**.

To configure the mapping of DNS canonical name to address records from the CLI, use the **dns pin cname records** global configuration command.

Configuring the DNS Server for HTTP Proxy Caching

DNS allows the network to translate domain names entered in requests into their associated IP addresses. To configure DNS caching on a Content Engine, you must complete the following tasks:

- Specify the list of DNS servers, which are used by the network to translate requested domain names into IP addresses that the Content Engine should use for domain name resolution.
- Specify the name of the local domain.
- Specify the maximum number of records that the DNS cache on the Content Engine should store. This setting defines the size of the DNS cache on the Content Engine.
- Enable DNS caching on the Content Engine.

By default, the DNS caching name service on the Content Engine uses a DNS server from its list of configured DNS servers instead of the original DNS server.

- Original DNS server—DNS server from the original request (referred to as the original DNS server).
- List of Configured DNS Servers—DNS servers that are used in the network and have been added to the list of DNS servers that the Content Engine should use for domain name resolution.

To configure DNS server settings for a Content Engine, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the Content Engine that you want to configure. The Contents pane appears on the left.

- Step 3** From the Contents pane, choose **General Settings > Services > DNS**. The HTTP DNS Cache Settings window appears. (See [Figure 8-16](#).) [Table 8-28](#) defines the fields in this window and provides the corresponding CLI global configuration commands.

Figure 8-16 HTTP DNS Cache Settings Window

The screenshot shows the 'HTTP DNS Cache Settings' window. The 'Contents' pane on the left has 'DNS' selected under 'Services'. The main area shows the following settings:

- Current applied settings from Content Engine, **CONTENTENGINE**
- Enable Caching of DNS Entries for Http Proxy: ☐
- Maximum Cached DNS Entries: (4096-65536)
- Enable Serial Lookup: ☐
- Local Domain Name: Space separated list of up to 3 Domain Names
- List of DNS Servers: Space separated list of up to 8 DNS Servers
- Note: * - Required Field
- Buttons: Submit, Cancel

- Step 4** To enable DNS caching, check the **Enable DNS Caching of DNS Entries for Http Proxy** check box.
- Step 5** In the Maximum Cached DNS Entries field, enter a value in megabytes to configure the size of the DNS cache hash table to be temporarily used for storing records. The range is from 5 to 512.
- Step 6** To allow each configured name server to be queried in turn if the request to the primary name server fails, check the **Enable Serial Lookup** check box.
- Step 7** In the Local Domain Name field, enter the name of the local domain that is to be DNS cached.
- Step 8** In the List of DNS Servers field, enter a list of DNS servers used by the network to resolve host names to IP addresses. You can configure up to eight DNS servers. Separate items in the list with a space.
- Step 9** To save the settings, click **Submit**. A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default and device group settings. To revert to the previously configured window settings, click **Reset**. The Reset button appears only when you have applied default or group settings to change the current device settings but the settings have not yet been submitted.

Table 8-28 HTTP Proxy DNS Cache Settings

GUI Parameter	Function	CLI Command
Enable DNS Caching of DNS Entries for HTTP Proxy	Enables DNS caching in the system.	dns enable
Maximum Cached DNS Entries	Defines the maximum size of the DNS cache hash table used for temporarily saving DNS records.	dns max-cache-memory
Enable Serial Lookup	Allows each configured name server to be queried in turn if the request to the primary name server fails.	dns serial-lookup
Local Domain Name	Name of the local domain to be DNS cached.	ip domain-name <i>name1 name2 name3</i>
List of DNS Servers	DNS servers that are used by the network to translate requested domains into IP addresses.	ip name-server

Configuring ICP Settings

Internet Cache Protocol (ICP) is a lightweight message format used for communicating among Content Engines and for supporting interoperability with older proxy protocols. ICP is used to exchange hints about the existence of URLs in neighboring Content Engines in a Content Engine farm. Content Engines exchange ICP queries and replies to gather information for use in selecting the most appropriate location from which to retrieve an object.

Although ICP has traditionally been a way to scale the overall size of a cluster of caches beyond a single unit, history has shown ICP to be a poor way of scaling a cache clustering solution. In fact, because of the way that traffic is currently directed toward a transparent network cache cluster, the requirement for ICP is all but negated for the majority of cache deployments.

The ICPv2 protocol is documented in two standards documents:

- *RFC 2186: Internet Cache Protocol (ICP), Version 2*
- *RFC 2187: Application of Internet Cache Protocol (ICP), Version 2*

**Note**

The ability to act as both an ICP server (servicing requests from neighboring Content Engines) and an ICP client (sending requests to neighboring Content Engines) is supported.

ICP Client Settings


You can configure your Content Engine farm to generate ICP queries before retrieving requested objects from the Internet using the ICP client functionality.

With ICP functionality, you can configure parent and sibling Content Engines in a hierarchy. You can configure a Content Engine to be either a parent or a sibling. Parent Content Engines are able to retrieve data during a cache miss, whereas sibling Content Engines cannot retrieve data and instead forward the request to the parent Content Engines.

To configure ICP client functionality, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure.
- Step 3** From the Contents pane, choose **Applications > Web > HTTP > ICP > Client**. The ICP Client Settings window appears. (See [Figure 8-17](#).) [Table 8-29](#) describes the fields in this window.

Figure 8-17 ICP Client Settings Window



- Step 4** To enable ICP client queries, check the **Enable** check box.
- Step 5** In the Maximum Reply Wait Time field, specify the timeout period in seconds for ICP responses.
- Step 6** In the Number of Failures field, specify the number of failures that each Content Engine allows an unresponsive ICP server before it removes that ICP server from the list of available servers.
- Step 7** Enter a domain to be excluded from ICP queries in the Excluded Domains field.
- Step 8** To save the settings, click **Submit**.

[Table 8-29](#) describes the ICP client settings in the Content Distribution Manager GUI and provides the corresponding global configuration CLI commands.

Table 8-29 ICP Client Parameters

GUI Parameter	Function	CLI Command
Enable	Enables the ICP client.	icp client enable
Maximum Reply Wait Time	Configures how long the Content Engine waits before retrieving the requested data directly from the Internet. The range is from 1 to 30 seconds, and the default value is 2 seconds.	icp client max-wait <i>timeout</i>
Number of Failures	Configures the number of failures that each Content Engine allows an unresponsive ICP server before it removes that ICP server from the list. The range is from 0 to 100, and the default value is 20 failures.	icp client max-fail <i>retries</i>
Exclude Domains	Excludes local domains from ICP caching services.	icp client exclude <i>domainnames</i>

ICP Remote Client Settings

To add a remote ICP client to the ICP client list using the Content Distribution Manager GUI, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **Applications > Web > HTTP > ICP > Remote Clients**.
- Step 4** Click the **Create New ICP Remote Client** icon. The Creating New ICP Remote Client window appears.
- Step 5** Enter the host name or IP address for the ICP remote client in the Client Name field.
- Step 6** To configure the Content Engine to act as a parent server to the designated client, check the **Fetch Misses** check box. If the Content Engine cannot satisfy the client's request, it forwards the request to another server or the Internet.
- Step 7** To save the settings, click **Submit**.

To set the ICP server remote client from the CLI, use the following global configuration command:

```
icp server remote-client {hostname | ipaddress} {fetch | no-fetch}
```


ICP Server Settings

You can also configure a Content Engine to act as an ICP server. This allows the Content Engine to probe the hierarchy of Content Engines by multicasting an ICP message to ICP parent and sibling clients in the hierarchy.

To configure ICP server functionality, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **Applications > Web > HTTP > ICP > Server**. The ICP Server Settings for Content Engine window appears. (See Figure 8-18.) Table 8-30 describes the fields in this window and provides the corresponding CLI global configuration commands.

Figure 8-18 ICP Server Settings Window



- Step 4** To enable ICP server queries, check the **Enable** check box.
- Step 5** In the ICP Port field, designate a port that will listen for ICP requests on the server.
- Step 6** In the HTTP Port field, specify the HTTP port that will listen for ICP-generated requests on the server.
- Step 7** To save the settings, click **Submit**.


Table 8-30 ICP Server CLI Command Summary

GUI Parameter	Function	CLI Command
Enable	Enables the ICP sever.	icp server enable
ICP Port	Configures an HTTP proxy port to listen for ICP-generated requests. The range is from 0 to 65535. The default port number is 3128.	icp server http-port <i>port</i>
HTTP Port	Configures an ICP server port on a Content Engine to listen for ICP requests. The range is from 0 to 65535. The default port number is 3130.	icp server port <i>icpport</i>

ICP Remote Server Settings

To add a remote ICP server to the ICP server list using the Content Distribution Manager GUI, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the device group that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **Applications > Web > HTTP > ICP > Remote Servers**.
- Step 4** Click the **Create New ICP Remote Servers** icon. The Creating New ICP Remote Server window appears. (See [Figure 8-19](#).) [Table 8-31](#) describes the fields in this window and provides the corresponding CLI global configuration commands.

Figure 8-19 Creating New ICP Remote Server Settings Window

- Step 5** In the Server Name field, enter the host name or IP address for the ICP remote server.
- Step 6** To enable the configured ICP server to act as a parent server and retrieve cache miss data, check the **Parent** check box. If this box is not checked, the configured ICP server will not retrieve cache miss data.
- Step 7** In the ICP Port field, enter the ICP port to which ICP queries are directed on this ICP server. The range is from 1 to 65535, and the default port is 3130.
- Step 8** In the HTTP Port field, enter the HTTP port on the ICP server to which proxy-style requests are forwarded. The range is from 1 to 65535, and the default port is 3128.
- Step 9** In the Excluded Domains field, enter a list of excluded domains if you want to limit ICP requests directed to this ICP server to a specific set of domains. Otherwise, all ICP requests (aside from those specified as "local domains") are forwarded to this ICP server.
- Step 10** To save the settings, click **Submit**.

Table 8-31 ICP Remote Server Settings

GUI Parameter	Function	CLI Command
Server Name	Host name or IP address for the ICP remote server.	icp client add-remote-server { <i>hostname</i> <i>ipaddress</i> } { parent sibling } icp-port <i>icpport</i> http-port <i>httpport</i> [restrict <i>domainnames</i>]
Parent	Enables the configured ICP server to act as a parent server and retrieve cache miss data.	
ICP Port	Port to which ICP queries are directed.	
HTTP Port	Port on the ICP server to which proxy-style requests are forwarded.	
Excluded Domains	Excludes the domains listed.	



Configuring Streaming Media Services

Streaming media services enable the delivery of digital media directly to the end user from a point of origin, such as an origin server, encoder, or Content Engine cache. Streaming media can be delivered as live content, or it can be delivered as on-demand content, such as video on demand (VoD). Cisco ACNS software supports several types of streaming media services, including RealNetworks RealMedia, Microsoft Windows Media Technologies (WMT), Apple QuickTime, and Cisco Streaming Engine.

This chapter discusses these supported streaming media services and explains how to enable these services on Content Engines in your ACNS network. This chapter contains the following sections:

- [Content Engine Streaming Overview, page 9-1](#)
- [Licensing and Enabling Streaming Servers, page 9-6](#)
- [Using the Real-Time Streaming Protocol Gateway, page 9-7](#)
- [Using Windows Media Technologies, page 9-15](#)
- [Configuring Weighted Load Balancing for Live Stream Splitting, page 9-38](#)
- [Configuring Bandwidth Settings for Content Services, page 9-40](#)



Note

Content Engines that use a Content Service Switch (CSS) to load balance streaming traffic cannot stream UDP traffic (such as RTSPU) because the Content Service Switch does not support UDP traffic.

Content Engine Streaming Overview

Content Engines deployed in a centrally managed ACNS network provide both live and recorded streaming media services. Live video streaming and recorded video streaming services (such as video-on-demand streaming) address various challenges in the field of corporate communications. This section describes some of these communication challenges and explains which streaming solution is best suited to meet your corporate communication needs.

Corporate communications can include company meetings, state of the business addresses, quarterly earnings broadcasts, web casts, emergency broadcasts, company television, and so forth. Live streaming events often include supplemental media, such as slide presentations, polling, Q&A, and live chat.

One of the primary challenges in today's corporate environment is to deliver video and audio communications to employees, partners, and company stakeholders globally at the same time, and to safely deploy such media-rich content without adversely impacting existing employee work applications.

While live video streaming enables corporate executives to provide the same message to all potential viewers at the same time and is an engaging, inexpensive means to reach all employees, large international companies can encounter challenges with live streaming because of time zone differences and bandwidth constrictions.

Live streaming is limited by the amount of available WAN bandwidth so that remote sites on low bandwidth WANs are not be able to participate in a live video streaming event. However, it is possible to reach low bandwidth (64-Kbps) sites with an audio-only stream because audio consumes only about 12 to 20 Kbps per client. Another way to resolve time zone and bandwidth challenges is to use cached video streaming or video on demand (VoD).

Live Streaming for the Non-Multicast Enabled Network


A publicly announced live streaming event is the single most taxing event on a non-multicast enabled network and on the origin servers. In a live streaming event that does not take advantage of network multicast or application-level stream splitting, a separate stream is created from every single client to the originating server. The strain caused by such excessive traffic between clients and servers can cause the network to collapse.

For example, in a network with 10,000 users, if the stream encoding rate is reasonably high (around 300 kbps), the media server or video server would have to deliver the video stream at a rate of 3 Gbps. In a data center deployment, the data center would require a large number of video servers to support the event, while in all probability, the rest of the network would collapse before the video servers could even receive all the requests. If a WAN exists between the users and the video servers, the strain is compounded. If collateral content (such as images, which can often be as large as 100 KB per object) is included for the event, and 10,000 users simultaneously request this content, the web servers that serve these types of objects would likely fail. (A typical web server can send 1000 objects per second.)

In such a network scenario, strategically placed Cisco ACNS Content Engines can minimize the impact of a live streaming event on the network by providing local delivery of those objects. Content Engines also support unicast live stream splitting.

Unicast Live Stream Splitting

Unicast live stream splitting is a solution for corporate networks that do not support multicasting or for administrators who must provide live streaming for unplanned events. A Content Engine placed with each WAN router provides unicast stream splitting for live streams that originate on the non-multicast enabled Internet or the within the corporate intranet. (See [Figure 9-1.](#))

Figure 9-1 Unicast Live Stream Splitting

Content Engine unicast stream splitting provides the following benefits:

- Companies do not have to upgrade network infrastructure to deploy IP multicast, if that is a challenge.
- Administrators do not have to create scheduled multicast events because Content Engines automatically split the stream.
- The Content Engines cause very little delay compared to delays incurred by the encoder, encoder buffer, and media player buffer.
- Users can view the broadcast as it occurs.


In all circumstances, however, live video quality is affected by the available WAN bandwidth. Also, when live streams are being delivered to remote sites, you must have a full duplex-switched Ethernet LAN infrastructure at the remote site because every client pulls a separate unicast stream.

Hybrid Unicast-to-Multicast Live Streaming

Unicast-to-multicast streaming is a solution similar to stream splitting, except that in final delivery segment the stream is converted to multicast to minimize the bandwidth usage impact on the network and to minimize the load on the Content Engine. With unicast-to-multicast streaming, if you have a Content Engine in the same location as your clients, the Content Engine can acquire the stream as a single unicast stream over the WAN and convert it to a multicast in the LAN.

If you have a multicast-enabled WAN, you can use a Content Engine to publish the multicast in the data center and stream the multicast to all your multicast-enabled clients over the WAN. Multicast-enabled network routers and switches replicate (or copy) the stream (packets) where necessary so that all users can view the video event simultaneously. (See [Figure 9-2](#).)

Figure 9-2 Hybrid Unicast to Multicast Live Streaming



Video On Demand Streaming

Video on demand (VoD) streaming can be used to deliver communications across large international company networks where WAN bandwidth and time zone differences present a challenge. With VoD you can deliver the same message to everyone in your organization, everywhere, anytime.

In a network with no Content Engines deployed, each client VoD request generates a separate stream from the video server across the WAN. The aggregate bandwidth usage of all clients must be less than the available WAN bandwidth or serious congestion will occur. When you place a Content Engine in the same location as the requesting client, you remove the WAN bottleneck, and the client can have an experience that is more media rich than the original live broadcast.

VoD streaming uses a recorded file that is usually generated from a live event. This file takes the form of a Windows Media file (wma, wmv, asf, and so forth), a RealMedia file (rm, ra, and so forth), or a standard video file (such as, mpg, mp4, mov). The video file can be administratively pre-positioned to the local Content Engine file system using the intelligent and controlled file distribution capabilities of


ACNS software, or the file can be retrieved and stored in the Content Engine file system as a result of a client request. Regardless, once the file is in the Content Engine file system, the file can be delivered locally across the LAN.

Video storage requirements are generally calculated in hours. If the stream encoding rate is 300 Kbps, then an hour of video consumes $300,000 \text{ bits/second} \times 3,600 \text{ seconds/hour} \times 8 \text{ bits/Byte}$ for a total of 135 MB. Based on these estimates, a Content Engine with a 40-GB capacity can store about 296 hours of video.

Demand-Cached Video

Demand caching (also called pull caching) allows the Content Engine to dynamically populate its cache in response to user requests. Caching is useful when the video encoding rate (or streaming rate) is less than the available WAN bandwidth, but not enough bandwidth exists for multiple clients to view the video simultaneously. Caching is the only solution if the organization that controls the Content Engine does not have access to the video files before a client requests the files. This situation can occur when the content originates outside the organization, from the Internet, or from a partner. A primary advantage to demand caching is that it does not require any administration. The disadvantage is that the first user requesting the video can experience poor quality if the network is congested. (See [Figure 9-3](#).)


Figure 9-3 Demand-Cached Video



Pre-Positioned Video

Pre-positioned video involves prepopulating Content Engines with the recorded video program. When a recorded video is company-owned or can be acquired, the company administrator can program a Content Engine to fetch the specifically named content or files from an origin server, set bandwidth distribution controls, and configure the Content Engine to automatically distribute the recorded files to Content Engines at the edge of the company LAN. Users can join a scheduled broadcast of the video, or the video can be made available on demand from a link published on the company website. (See [Figure 9-4](#).)

Figure 9-4 Pre-Positioned Video



Licensing and Enabling Streaming Servers

For your ACNS network devices to be able to serve streaming media, you must tell the ACNS network that a particular streaming media service is present, and you must accept a licence agreement for each type of media service that you intend to use. Media services available in ACNS 5.x software include Windows Media Server, RealProxy, RealSubscriber, and multicast client.



Note

Cisco Streaming Engine does not require a separate license agreement.

You can accept license agreements at the device level or at the device group level. As soon as any license agreement has been accepted for any device or device group, the acceptance is applied to all devices that are currently registered to the Content Distribution Manager and to all devices that register in the future.

When a Content Distribution Manager is upgraded to ACNS 5.5 software, the Content Distribution Manager will apply any accepted license agreements from the previous ACNS software release to all registered devices in the network. If in a previous ACNS 5.x Release, you accepted a particular licence agreement on some, but not all, devices, any device for which the license was not accepted prior to the Content Distribution Manager upgrade will have its license agreement accepted automatically.

**Note**

You cannot remove a license agreement once it has been accepted. If there is a particular media type that you do not want to serve through your ACNS network, you need to disable the corresponding media player so that your ACNS network devices do not attempt to serve that media type.

Using the media player mappings in the manifest file, you can customize your content type mappings from MIME content types or file extensions to configured media players. (See the [“Working with Manifest Files”](#) section on page A-6 for more information on configuring manifest files.)

**Note**

Modifying the configuration of any media players causes any ACNS network devices running that media player software to restart and register the configuration change. Depending on the number and location of devices that restart following a media player configuration change, you may experience a temporary interruption in service for the time it takes your devices to come back online—usually a few minutes.

Before enabling licenses for streaming media services, make sure that the clock and calendar settings on the ACNS network device are correct; otherwise, you will see an error message, and the services will fail to install. Use the **show clock EXEC** command to display the system clock. To set the system clock, use the **clock set EXEC** command.

Using the Real-Time Streaming Protocol Gateway

The Real-Time Streaming Protocol (RTSP) is a standard Internet streaming control protocol (RFC 2326). It is a widely used application-level protocol that controls the delivery of data with real-time properties, such as video and audio. Apple QuickTime, RealNetworks, and the Cisco Streaming Engine use RTSP as a streaming control protocol.

The RTSP gateway is a process running on the Content Engine that accepts an RTSP request and performs the initial RTSP contact with RTSP-based clients (RealMedia, QuickTime, and so on) on behalf of the RTSP-based streaming servers available on the Content Engine. From information derived from the initial request operation, the RTSP gateway queries the URL filtering, rules, authentication, and unified name space (UNS) agents to determine whether or not the content is pre-positioned.

**Note**

Pre-positioned content requests are only accepted and served on the RTSP gateway (for both RealMedia and the Cisco Streaming Engine) default port number 554. If the default RTSP port number is changed to any other port number on the Content Engine while the Content Engine is waiting for RTSP pre-positioned content from the origin server, the Content Engine will be unable to serve the content.

The RTSP gateway is always enabled and cannot be disabled. Based on the properties of the incoming request, including such properties as user agent, final destination, and media file type, the RTSP gateway performs these tasks:

- Forwards the incoming request to the RealProxy that resides on the same Content Engine as the RTSP gateway
- Forwards the incoming request to the RealSubscriber that resides on the same Content Engine as the RTSP gateway
- Forwards the incoming request to the Cisco Streaming Engine
- Redirects the incoming request
- Rejects the incoming request

To fulfill a client RTSP request for non-pre-positioned content, a Content Engine (CE1) forwards the request to the origin server. If the request is transparently intercepted by another Content Engine (CE2) on its way to the origin server, and the content is not pre-positioned on CE2, then the RTSP gateway application on CE2 issues an RTSP redirect response and adds a bypass entry in the WCCP bypass list.

The Content Engine can be configured to accept transparently redirected RTSP requests as well as traditional proxy-style RTSP requests from RealProxy client software. The RealProxy software is configured with the RealAdministrator GUI, accessed from the RealProxy window of the Content Engine management GUI. Detailed configuration, statistics, and reporting of RealProxy status are obtained through the RealAdministrator GUI.

To access the Content Engine management GUI from the Content Distribution Manager, choose **Devices > Devices**. Click the **Edit** icon next to the name of the Content Engine that you want to configure, and click the **Run CE Cache GUI** icon in the task bar. See the “[Configuring the Content Engine GUI for Secure or Nonsecure Access](#)” section on page 16-59.

The **Admin** button is active on the Content Engine management GUI when the RealProxy software is installed and enabled. You will be provided with a default username and password to access this administration window from the Content Engine GUI.

**Note**

You must configure disk space to include mediafs storage with the **disk config EXEC** command before you can run streaming media using RealProxy.

Configuring the RTSP Gateway

To configure RTSP gateway parameters, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**. If you have created device groups, the Device Group window appears. (Alternatively, choose **Devices > Devices**.)

**Note**

RTSP gateway parameters can be configured for individual Content Engines or for device groups. This procedure describes the steps for configuring RTSP gateway parameters for a device group.

- Step 2** Click the **Edit** icon next to the name of the device group that you want to configure. The Contents pane appears on the left.

- Step 3** From the Contents pane, choose **Applications > Streaming > RTSP Gateway**. The RTSP Gateway Settings window appears. (See [Figure 9-5](#).) [Table 9-1](#) describes the fields in this window and provides the corresponding CLI commands.

Figure 9-5 RTSP Gateway Settings Window

The screenshot shows the 'RTSP Gateway Settings' window for the 'CONTENTENGINE' device group. The 'Enable' checkbox is checked. The 'Incoming Port' is 554, 'IP-Address' is 114.116.115.112, 'Layer 4 Switch Support' is checked, 'Maximum Initial Setup Delay' is 10 seconds, and 'Maximum Request Rate' is 40 requests per second. The 'Bypass Gateway' is set to 'cisco-streaming-engine'. A note at the bottom states: '* RTSP Gateway is always enabled. Checking this box allows you to configure settings for the gateway.'

- Step 4** Check the **Enable** check box so that you can configure the RTSP gateway settings. Unchecking this box does *not* disable RTSP.
- Step 5** In the Incoming Port field, enter the port number used by the RTSP proxy to receive requests.
- Step 6** In the IP-Address field, enter the IP address of the RTSP gateway device. The RTSP gateway device is usually the Content Engine that will be serving the requested content.
- Step 7** To enable redirection using a Layer 4 switch, check the **Enable L4 Switch Support** check box.
- Step 8** In the Maximum Initial Setup Delay field, specify the maximum delay allowed (in seconds) between TCP accept and the first RTSP message from the client. The default is 10 seconds.
- Step 9** In the Maximum Request Rate field, specify the maximum number of incoming requests per second that the RTSP gateway allows. The default is 40 requests per second.

Table 9-1 RTSP Gateway Settings

GUI Parameter	Function	CLI Command
Incoming Port	Specifies the ports on which to listen for RTSP requests.	rtsp ports
IP-Address	Specifies the IP address of the RTSP gateway.	rtsp ip-address

Table 9-1 RTSP Gateway Settings (continued)

GUI Parameter	Function	CLI Command
Layer 4 Switch Support	Enables redirection with a Layer 4 switch, such as a Cisco CSS Series switch.	rtsp l4-switch enable
Maximum Initial Setup Delay	Maximum delay allowed (in seconds) between TCP accept and the first RTSP message from the client.	rtsp advanced max-initial-setup-delay seconds
Maximum Request Rate	Maximum number of incoming requests per second that the RTSP gateway allows.	rtsp advanced max-request-rate seconds

Enabling RealProxy

RealProxy is a licensed software program that requires the purchase of a license from Cisco. To enable this licensed software program, you must perform the following tasks:

- Accept the license agreement.
- Enter your Cisco license key or accept an evaluation license.
- Enable the software program.

Before enabling licenses for streaming media services, make sure that the clock and calendar settings on the ACNS network device are correct; otherwise, you will see an error message and the services will fail to install. Use the **show clock EXEC** command to display the system clock. To set the system clock, use the **clock set EXEC** command.



Note

RealProxy software is configured using the RealSystem administrator GUI, which is accessed from the RealProxy page of the Content Engine management GUI. To access the Content Engine management GUI from the Content Distribution Manager, choose **Devices > Devices**. Click the **Edit** icon next to the name of the Content Engine that you want to configure and click the **Run CE Cache GUI** icon in the taskbar.

To enable RealProxy for the devices in your ACNS network, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**. The Device Groups window appears. (Alternatively, choose **Devices > Devices**.)
- Step 2** Click the **Edit** icon next to the name of the device group (or Content Engine) that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **Applications > Streaming > Real Networks > Real Proxy License**. The Real Proxy License Agreement window appears if you have not previously read and accepted the license agreement for any device or device group in the network. After you have read and accepted the terms of this agreement, check the **Accept License** check box at the bottom of this window, and click **Submit**. The page refreshes itself, and the Real Proxy License Settings window appears. (See [Figure 9-6](#).)

**Note**

Accepting the RealProxy license agreement for any Content Engine or device group enables the RealProxy license for all the devices in the network.

Figure 9-6 Real Proxy License Settings for Device Group Window

The screenshot shows the Cisco Application and Content Networking System interface in Microsoft Internet Explorer. The main window is titled "Real Proxy License Settings for Device Group, DG3". It includes a "Contents" sidebar on the left with options like "Device Group Home", "Software Update", "Assignments", "Prepositioning", "Request Routing", "Request Processing", "Applications", "Default and Maximum Bandwidth", "Bandwidth Schedules", "Streaming", "Windows Media", "Real Networks", "Real Proxy License", "Real Subscriber License", "Cisco Streaming Engine", "RTSP Gateway", "Web", "DNS", "Set Top Box", "File Sharing", and "General Settings". The main content area displays the "Real Proxy License Settings" for the device group "DG3". It includes a message: "Real Proxy License Settings will be enabled when using Factory Defaults. The values shown in this page are in effect." and a table of license keys. The table has columns for "License Key", "Device Group", and "Validate". The "Enable" checkbox is currently unchecked. The "Evaluate" checkbox is also unchecked. The table lists various license keys (CE-501, CE-507, CE-510, CE-511, CE-560, CE-565, CE-566, CE-590, CE-2636DW, CE-2636SW) and their corresponding device groups. At the bottom right, there are "Submit" and "Cancel" buttons.

License Key	Device Group	Validate
CE-501 License Key:	0 CE-501 in DeviceGroup	Validate: ☐
CE-507 License Key:	1 CE-507 in DeviceGroup	Validate: ☐
CE-510 License Key:	1 CE-510 in DeviceGroup	Validate: ☐
CE-511 License Key:	0 CE-511 in DeviceGroup	Validate: ☐
CE-560 License Key:	0 CE-560 in DeviceGroup	Validate: ☐
CE-565 License Key:	0 CE-565 in DeviceGroup	Validate: ☐
CE-566 License Key:	0 CE-566 in DeviceGroup	Validate: ☐
CE-590 License Key:	0 CE-590 in DeviceGroup	Validate: ☐
CE-2636DW License Key:	0 CE-2636DW in DeviceGroup	Validate: ☐
CE-2636SW License Key:	0 CE-2636SW in DeviceGroup	Validate: ☐
CE-7305 License Key:	0 CE-7305 in DeviceGroup	Validate: ☐

**Note**

The RealProxy License Settings for Content Engine window (**Devices > Devices**) does not list all the Content Engine models. It has one License Key field for the particular Content Engine that you chose.

- Step 4** To activate the RealProxy license, check the **Enable** check box.
- Step 5** To obtain a permanent license, enter a license key number in the field next to the Content Engine model which corresponds to the license key that you purchased.

**Note**

Because license keys are model-specific, the License Settings for Device Group window allows you to enter a license key for each model type in the group. This window is populated with input fields for each model type known to the Content Distribution Manager, even if there are no Content Engines of that type in the network. (The License Settings window for device groups reports the count for each Content Engine model in the device group and reports 0 [zero] if there are no devices of a particular type in the group.) This allows you to enter a license key for a device before it is registered. A field for an unknown device license key allows you to enter a license key for a Content Engine model type that is unknown to the Content Distribution Manager.

Alternatively, to use a temporary 60-day license, check the **Evaluate** check box.

Step 6 To validate the license key that you entered, check the **Validate** check box.

Step 7 To save the settings, click **Submit**.

Enabling RealSubscriber

ACNS software includes RealServer Version 9 as an optional component that is used as a streaming media engine. When RealServer software is configured for subscriber-only mode, it is called RealSubscriber. RealSubscriber is a licensed software program that requires the purchase of a license from Cisco. To enable this licensed software program, you must perform the following tasks:

- Accept the license agreement.
- Enter your Cisco license key or accept an evaluation license.
- Enable the software program.

To enable RealSubscriber for the devices in your ACNS network, follow these steps:

Step 1 From the Content Distribution Manager GUI, choose **Devices > Device Groups**. The Device Groups window appears. (Alternatively, choose **Devices > Devices**.)

Step 2 Click the **Edit** icon next to the device group (or Content Engine) that you want to configure. The Contents pane appears on the left.

Step 3 Choose **Applications > Streaming > Real Networks > Real Subscriber License**.

The Real Subscriber License Agreement window appears if you have not previously read and accepted this license agreement for any device or device group in the network. After you have read and accepted the terms of this agreement, you must check the **Accept License** check box at the bottom of this window and click **Submit**. The page refreshes itself, and the Real Subscriber License Settings window appears. (See [Figure 9-7](#).)

**Note**

Accepting the Real Subscriber License Agreement for any Content Engine or device group enables the Real Subscriber License for all the devices in the network.

Figure 9-7 RealSubscriber License Settings for Device Group Window

Real Subscriber License Settings for Device Group, DG3

There are currently no settings for this Group. Real Subscriber License Settings will be enabled when using Factory Defaults. The values shown in this page are in effect.

Real Subscriber License Settings			
Enable:	☐		
Evaluate:	☒		
CE-501 License Key:		0 CE-501 in DeviceGroup	Validate: ☐
CE-507 License Key:		1 CE-507 in DeviceGroup	Validate: ☐
CE-510 License Key:		1 CE-510 in DeviceGroup	Validate: ☐
CE-511 License Key:		0 CE-511 in DeviceGroup	Validate: ☐
CE-560 License Key:		0 CE-560 in DeviceGroup	Validate: ☐
CE-565 License Key:		0 CE-565 in DeviceGroup	Validate: ☐
CE-566 License Key:		0 CE-566 in DeviceGroup	Validate: ☐
CE-590 License Key:		0 CE-590 in DeviceGroup	Validate: ☐
CE-2636DW License Key:		0 CE-2636DW in DeviceGroup	Validate: ☐
CE-2636SW License Key:		0 CE-2636SW in DeviceGroup	Validate: ☐
CE-7305 License Key:		0 CE-7305 in DeviceGroup	Validate: ☐

Submit Cancel

**Note**

The RealSubscriber License Settings for Content Engine window (**Devices > Devices**) does not list all the Content Engine models. It has one License Key field for the particular Content Engine that you chose.

Step 4 To activate a RealSubscriber license, check the **Enable** check box.

Step 5 To obtain a permanent license, enter a license key number in the field next to the Content Engine model which corresponds to the license key that you purchased.

**Note**

Because license keys are model-specific, the License Settings for Device Group window allows you to enter a license key for each model type in the group. This window is populated with input fields for each model type known to the Content Distribution Manager, even if there are no Content Engines of that type in the network. (The License Settings window for device groups reports the count for each Content Engine model in the device group and reports 0 [zero] if there are no devices of a particular type in the group.) This allows you to enter a license key for a device before it is registered. A field for an unknown device license key allows you to enter a license key for a Content Engine model type that is unknown to the Content Distribution Manager.

Alternatively, to use a temporary 60-day license, check the **Evaluate** check box.

Step 6 To validate the license key that you entered, check the **Validate** check box.

Step 7 To save the settings, click **Submit**.

To enable RealSubscriber for a device by using the CLI, follow these steps:

Step 1 View the RealSubscriber license agreement by entering the following command in EXEC mode:

```
ContentEngine# show rtsp server real-subscriber license-agreement
```

Step 2 After reading the license agreement, enter global configuration mode and accept the license agreement by entering the following command:

```
ContentEngine# configure
ContentEngine(config)# rtsp server real-subscriber enable accept-license-agreement
```

Step 3 Enter your Cisco license key by entering the following command:

```
ContentEngine(config)# rtsp server real-subscriber license-key licensekey
```

Alternatively, accept a 60-day evaluation license by entering the following command:

```
ContentEngine(config)# rtsp server real-subscriber evaluate
```

You can accept a RealSubscriber evaluation license that allows you to use the licensed software for a limited period. The evaluation license can be invoked only once and does not require a license key. The evaluation license allows the licensed program to be enabled for 60 days.

When the preset evaluation time has elapsed, RealSubscriber is automatically disabled. You can purchase a valid license key from Cisco Systems and enter the license key into the configuration at any time during the evaluation period for uninterrupted use of RealSubscriber.

Step 4 Enable RealSubscriber by entering the following command:

```
ContentEngine(config)# rtsp server real-subscriber enable
```

Enabling the Cisco Streaming Engine

ACNS 5.x software includes support for a standard RTSP server, also known as the Cisco Streaming Engine solution for RTSP streaming media protocol.



Note

You do not need a license to use the Cisco Streaming Engine to deliver RTSP streaming media.



Note

The Cisco Streaming Engine has been upgraded with a new streaming component that supports Third Generation Partnership Project (3GPP) streaming files (files with the .3gp file extension). This streaming service upgrade in the ACNS 5.5 release provides for the uniform delivery of rich multimedia content over broadband mobile networks to multimedia-enabled cellular phones.

To configure Cisco Streaming Engine parameters, follow these steps:

-
Step 1	From the Content Distribution Manager GUI, choose Devices > Devices . The Devices window appears.
Step 2	Click the Edit icon next to the Content Engine that you want to configure. The Contents pane appears on the left.
Step 3	From the Contents pane, choose Applications > Streaming > Cisco Streaming Engine .
Step 4	To allow use of the Cisco Streaming Engine for RTSP streaming media delivery, check the Enable check box.
Step 5	To save the settings, click Submit .
Step 6	To override the device group settings applied to the device with the factory default settings, click the Override Group Settings with Defaults icon in the taskbar. This icon appears only if you have applied the device group settings to the Content Engine.
-

Using Windows Media Technologies

Microsoft Windows Media Technologies (WMT) is a set of streaming solutions for creating, distributing, and playing back digital media files across the Internet. WMT includes the following applications:

- Windows Media player—End user application
- Windows Media server—Server and distribution application
- Windows Media encoder—Encodes media files for distribution
- Windows Media codec—Compression algorithm applied to live or on-demand content
- Window Media Rights Manager (WMRM)—Encrypts content and manages user privileges

ACNS 5.5 software supports the following:

- Windows Media Version 9 live streaming from a Windows Media encoder over a Microsoft Media server (MMS-over-HTTP and RTSP)
- Serving files encoded with the Version 9 codec
- Streaming live or on-demand content to Windows Media Player Version 9
- Serving pre-positioned Windows Media files over RTP or RTSP
- Configuring Fast Cache and Fast Start

Enabling Windows Media Services



Note

Licenses purchased for Windows Media Services (WMS) in the ACNS 5.5 software and later releases are not supported in ACNS 5.4.x software or earlier releases. However, a Content Distribution Manager running ACNS 5.4 software can configure licenses for both ACNS 5.4 and ACNS 5.5 devices. One caveat is that an ACNS 5.4 Content Distribution Manager cannot validate ACNS 5.5 license keys. If you have a mixture of ACNS 5.4 and ACNS 5.5 devices in your network and you want to configure licenses by device groups, we recommend that you group devices in the Content Distribution Manager separately by software release to avoid license incompatibility.

To disseminate live and pre-positioned Windows Media content on an ACNS network, Windows Media caching proxy and server capabilities must be enabled on the Content Engine. Windows Media services is a licensed software program that requires the purchase of a license from Cisco. To enable this licensed software program, you must perform the following tasks:

- Accept the license agreement.
- Enter your Cisco license key or accept an evaluation license.
- Enable the software program.

**Note**

You must configure disk space to include mediafs storage with the **disk config EXEC** command before you can run streaming media using Windows Media.

To enable Windows Media services for the devices in your ACNS network, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**. The Device Groups window appears. (Alternatively, choose **Devices > Devices**.)
- Step 2** Click the **Edit** icon next to the device group (or Content Engine) that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **Applications > Streaming > Windows Media > License**.

The Windows Media License Agreement for Content Engine window appears if you have not previously read and accepted this license agreement for any device or device group in the network. After you have read and accepted the terms of this agreement, check the **Accept License** check box at the bottom of this window and click **Submit**. The page refreshes itself, and the Windows Media License Settings window appears. (See [Figure 9-8](#).)

**Note**

Accepting the Windows Media License Agreement for any Content Engine or device group enables the Windows Media License for all the devices in the network.

Figure 9-8 Windows Media License Settings Window

Windows Media License Settings for Device Group, DG3

Windows Media License Settings will be enabled when using Factory Defaults. The values shown in this page are in effect.

Windows Media License Settings			
Enable Windows Media Services:	☐		
Enable Windows Media Proxy:	☒		
Evaluate:	☐		
CE-501 License Key:		0 CE-501 in DeviceGroup	Validate: ☐
CE-507 License Key:		1 CE-507 in DeviceGroup	Validate: ☐
CE-510 License Key:		1 CE-510 in DeviceGroup	Validate: ☐
CE-511 License Key:		0 CE-511 in DeviceGroup	Validate: ☐
CE-560 License Key:		0 CE-560 in DeviceGroup	Validate: ☐
CE-565 License Key:		0 CE-565 in DeviceGroup	Validate: ☐
CE-566 License Key:		0 CE-566 in DeviceGroup	Validate: ☐
CE-590 License Key:		0 CE-590 in DeviceGroup	Validate: ☐
CE-2636DW License Key:		0 CE-2636DW in DeviceGroup	Validate: ☐
CE-2636SW License Key:		0 CE-2636SW in DeviceGroup	Validate: ☐

Submit Cancel

**Note**

The Windows Media License Settings for Content Engine window (**Devices > Devices**) does not list all the Content Engine models. It has one License Key field for the particular Content Engine that you chose.

- Step 4** To unlock the configuration window, check the **Enable Windows Media Services** check box.
- Step 5** To enable the Windows Media caching proxy functionality, check the **Enable Windows Media Proxy** check box.
- Step 6** To obtain a permanent license, enter a license key number in the field next to the Content Engine model which corresponds to the license key that you purchased.

**Note**

Because license keys are model-specific, the License Settings for Device Group window allows you to enter a license key for each model type in the group. This window is populated with input fields for each model type known to the Content Distribution Manager, even if there are no Content Engines of that type in the network. (The License Settings window for device groups reports the count for each Content Engine model in the device group and reports 0 [zero] if there are no devices of a particular type in the group.) This allows you to enter a license key for a device before it is registered. A field for an unknown device license key allows you to enter a license key for a Content Engine model type that is unknown to the Content Distribution Manager.

Alternatively, to use a temporary 60-day license, check the **Evaluate** check box.

- Step 7** To validate the license key that you entered, check the **Validate** check box.
- Step 8** To apply default settings, click the **Apply Defaults** icon in the taskbar. The Reset button returns settings to the most recently saved settings.
- Step 9** To save the settings, click **Submit**.

Configuring Windows Media Proxy Settings for the Content Engine


The Content Engine acting as a Windows Media proxy supports a basic proxy feature—it accepts incoming Windows Media streaming requests from clients and acts on behalf of the clients communicating with an origin server. The Windows Media proxy accepts and serves streaming requests over the RTSP protocol, as well as the HTTP protocol (including MMS-over-HTTP). (See [Figure 9-9](#).)



Note

MMS is the protocol that Windows Media (prior to version 9) uses for communication between media players and servers. ACNS 5.5 software does not support the MMS protocol. ACNS 5.5 software supports RTSP and MMS-over-HTTP only.

Figure 9-9 Content Engine as Conventional Windows Media Caching Proxy



When possible, the Windows Media proxy also caches streaming content and serves the request from the Content Engine cache instead of from the origin server. It accepts transparently intercepted requests (through WCCP or Layer 4 [L4] redirect) as well as manual proxy requests (clients configured to use an upstream proxy).

The Windows Media proxy also supports multicasting and splitting of live streams. (That is, it splits a single inbound feed to multiple clients.)

Before you can enable Windows Media proxy service, your system must meet the following requirements:

- The Content Engine has ACNS 4.1 software or a later release.
- Windows Media proxy software is installed, and disk space is configured to include mediafs storage.
- You have a Microsoft Windows Media license key.
- You have the IP address of the Content Engine.

To configure Windows Media proxy settings, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices** or **Devices > Device Groups**.
- Step 2** Click the **Edit** icon next to the Content Engine (or device group) that you want to configure. The Contents pane appears on the left.

- Step 3** From the Contents pane, choose **Applications > Streaming > Windows Media > General Settings**. The Windows Media Settings for Content Engine window appears. (See [Figure 9-10](#) and [Figure 9-11](#).) [Table 9-2](#) describes the fields in this window and provides the corresponding CLI global configuration commands.

Figure 9-10 Windows Media Settings Window—Windows Media Proxy Settings

The screenshot shows the Cisco Application and Content Networking System interface in Microsoft Internet Explorer. The main window is titled "Windows Media Settings for Content Engine, CONTENTENGINE". The left sidebar contains a "Contents" pane with a tree view showing the navigation structure. The main content area is divided into two sections: "Windows Media Proxy Settings" and "Windows Media General Settings".

Windows Media Proxy Settings:

- Enable L4 Switch Support: ☐
- Enforce Max Object Size: ☐
- Max Object Size: (MB) (1-1000000)
- Enable Outgoing HTTP Proxy: ☐ Host Name: Port:
- Enable Outgoing RTSP Proxy: ☐ Host Name: Port:
- Enable Accelerate Proxy Cache Performance: ☒

Windows Media General Settings:

- Disable HTTP Windows Media Traffic: ☒ Disable RTSP VMT Traffic: ☐ Disable RTSP VMT Traffic: ☐
- Maximum Concurrent Connections: (1-8000)
- Enforce Maximum Outgoing Bitrate: ☐
- Maximum Outgoing Bitrate: (Kbps) (0-2147483647)
- Enforce Maximum Incoming Bitrate: ☐
- Maximum Incoming Bitrate: (Kbps) (0-2147483647)

At the bottom right, there are "Submit" and "Cancel" buttons.

- Step 4** Under the Windows Media Proxy Settings heading, check the **Enable L4 Switch Support** check box to enable redirection using a Layer 4 switch.
- Step 5** To allow the maximum size for an object to be cached, check the **Enforce Max Object Size** check box.
- Step 6** Enter a value (in megabytes) in the Max Object Size field. The range is from 1 to 1000000. The default value is 1024 megabytes.
- Step 7** To configure the outgoing MMS-over-HTTP proxy, check the **Enable Outgoing HTTP Proxy** check box, and then enter the host name or IP address and the port number of the outgoing proxy. The port number range is 1 to 65535.
- Step 8** To configure the outgoing RTSP proxy, check the **Enable Outgoing RTSP Proxy** check box. Then enter the host name or IP address and the port number of the outgoing proxy. The port number range is 1 to 65535.
- Step 9** To enable caching performance improvement of the Windows Media proxy, check the **Enable Accelerate Proxy Cache Performance** check box.
- Step 10** Under the Windows Media General Settings heading, disable Windows Media traffic as desired. (See [Figure 9-10](#).)

- To disallow streaming over HTTP protocol (http://), check the **Disable HTTP Windows Media Traffic** check box.
 - To disallow streaming over RTSP protocol (rtspt://), check the **Disable RTSP WMT Traffic** check box.
 - To disallow streaming over RTSPU protocol (rtspu://), check the **Disable RTSPU WMT Traffic** check box.
- Step 11** Specify the maximum number of concurrent sessions in the Maximum Concurrent Connections field. The default value is 2500 sessions.
- Step 12** To allow use of the maximum bit rate allowed for serving content, check the **Enforce Maximum Outgoing Bitrate** check box.
- Step 13** Specify the maximum bit rate allowed (in kilobytes per second) in the Maximum Outgoing Bitrate field.
- Step 14** To allow use of the maximum incoming bit rate allowed for receiving content, check the **Enforce Maximum Incoming Bitrate** check box.
- Step 15** Specify the maximum bit rate allowed (in kilobytes per second) in the Maximum Incoming Bitrate field.
- Step 16** To remove personalization information from the URL before using it for live splitting, check the **Enable Live URL Stripping** check box. (See [Figure 9-11](#).)

Figure 9-11 Windows Media Settings Window—Windows Media General and Advanced Settings

The screenshot displays the Cisco Application and Content Networking System configuration page for Windows Media Settings. The left sidebar shows a tree view with categories like Contents, Applications, Streaming, and Windows Media. The main content area is titled 'Windows Media Settings for Content Engine, CONTENTENGINE' and contains several sections:

- General Settings:** Includes checkboxes for 'Enable Live URL Stripping', 'Enable Accelerate Live-Split Performance', 'Enable Accelerate VOD Performance', and 'Restrict HTTP Allowed Extensions'. It also has a text field for 'HTTP Allowed Extensions' with a hint 'Space separated list of up to 20 extensions'.
- Fast Start Settings:** Includes checkboxes for 'Enable Fast Start Feature' and 'Enable Fast Cache', and a text field for 'Fast Cache Max Delivery Rate'.
- Windows Media Multicast Settings:** Includes a text field for 'Number of hops to live'.
- Windows Media Advanced Client Settings:** Includes text fields for 'Idle Timeout' and 'Maximum data packet size'.
- Windows Media Advanced Server Settings:** Includes a checkbox for 'Enable Log Forwarding'.

The bottom of the page features 'Submit' and 'Cancel' buttons. The status bar at the bottom right shows the IP address '151725'.

- Step 17** To enable performance improvement for Windows Media proxy live splitting, check the **Enable Accelerate Live Split Performance** check box. This feature is an extension to the performance improvement for the Windows Media proxy.

- Step 18** To enable performance improvement for Windows Media proxy video-on-demand, check the **Enable Accelerate VoD Performance** check box.
- Step 19** To make it possible to add to or delete from the list of permitted filename extensions, check the **Restrict HTTP Allowed Extensions** check box.
- Step 20** Specify the list of allowable extensions for HTTP in the HTTP Allowed Extensions field. The default extensions are automatically entered in this list.
- The list of default filename extensions includes asf, none, nsc, wma, and wmv. The “none” filename extension has been included to enable the Content Engine to serve files without filename extensions, such as live encoder URLs. The nsc filename extension is used for multicast media files.
- Step 21** To enable the Fast Start feature, check the **Enable Fast Start Feature** check box. (See the [“About Fast Start”](#) section on page 9-24.)
- Step 22** In the Fast Start Max Bandwidth field, enter a value in kbps to limit the speed at which the client can receive data at the fast start stage. The range is 1–65535.
- Step 23** To enable the Fast Cache feature, check the **Enable Fast Cache** check box. (See the [“About Fast Cache”](#) section on page 9-25.)
- Step 24** In the Fast Cache Max Delivery Rate field, enter a value to limit the total bandwidth that a client can consume as fast cached data. The range is 1–65535.
- The value is the number multiplied by the normal bit rate of the file. For example, if the file is encoded as 100 kbps and the maximum delivery rate for fast caching is configured as 5, the maximum playback speed would be 500 kilobits per second.
- Step 25** Under the Windows Media Multicast Settings heading, enter the number of hops to live for multicast Windows Media packets. The default is 5 hops. The range is 0 to 255 hops. For clients that are many router hops away from Content Engines functioning as multicast stations, the requested content is sometimes not delivered even though the delivering device has sent the content to the multicast address configured on the Content Engine. In such cases, configure a larger number of hops to allow a longer Time To Live for multicast packets being delivered to the Content Engine.
- Step 26** Under the Windows Media Advanced Client Settings heading, enter a value (in seconds) in the Idle Timeout field. The default is 120 seconds.
- Step 27** In the Maximum data packet size field, enter the maximum data packet size in bytes. The range is 512–2048.
- Step 28** Under the Windows Media Advanced Server Settings heading, to enable log forwarding to an upstream Content Engine or Windows Media server, check the **Enable Log Forwarding** check box.
- Step 29** To save the settings, click **Submit**.

Table 9-2 Windows Media Settings Window

GUI Parameter	Description	CLI Command
Windows Media Proxy Settings		
Enable L4 Switch Support	Enables redirection with a Layer 4 switch, such as a Cisco CSS 11000 series switch.	wmt l4-switch enable
Enforce Max Object Size	Enforces use of the maximum object size.	wmt cache max-obj-size
Max Object Size (MB)	Specifies the maximum allowable size of streaming media objects in megabytes (MB).	wmt cache max-obj-size <i>max_obj_size</i>

Table 9-2 Windows Media Settings Window (continued)

GUI Parameter	Description	CLI Command
Enable Outgoing HTTP Proxy	Enables the use of outgoing MMS-over-HTTP.	wmt proxy outgoing http host
Host Name	Host name or IP address of the outgoing HTTP proxy.	wmt proxy outgoing http host <i>hostname</i>
Port	Port of the outgoing proxy.	wmt proxy outgoing http host <i>hostname port</i>
Enable Outgoing RTSP Proxy	Enables use of outgoing RTSP.	wmt proxy outgoing rtsp host
Host Name	Host name or IP address of the outgoing RTSP proxy server.	wmt proxy outgoing rtsp host <i>hostname</i>
Port	Port of the outgoing proxy server.	wmt proxy outgoing rtsp host <i>hostname port</i>
Enable Accelerate Proxy Cache Performance	Enables caching performance improvement of the Windows Media proxy.	wmt accelerate proxy-cache enable
Windows Media General Settings		
Disable HTTP Windows Media Traffic	Disallows streaming over HTTP protocol (http://).	wmt disallowed-client-protocols http
Disable RTSPT WMT Traffic	Disallows streaming over RTSPT protocol (rtspt://).	wmt disallowed-client-protocols rtspt
Disable RTSPU WMT Traffic	Disallows streaming over RTSPU protocol (rtspu://).	wmt disallowed-client-protocols rtspu
Maximum Concurrent Connections	Maximum number of unicast clients that can be served. The range depends on the physical limitations of the platform. The default value is 2500.	wmt max-concurrent-sessions <i>num</i>
Enforce Maximum Outgoing Bitrate	Enforces the maximum stream bit rate that can be served.	bitrate wmt outgoing
Maximum Outgoing Bitrate (kbps)	Maximum streaming bit rate that can be served. The range is 0 to 2147483647 kbps. There is no default value.	bitrate wmt outgoing <i>max_bitrate</i>
Enforce Maximum Incoming Bitrate	Enforces the maximum incoming bit rate for receiving content.	bitrate wmt incoming
Maximum Incoming Bitrate (kbps)	Maximum streaming bit rate that can be received.	bitrate wmt incoming <i>max_bitrate</i>
Enable Live URL Stripping	Removes personalization information from the URL before using it for live splitting.	wmt live-url-stripping enable

Table 9-2 Windows Media Settings Window (continued)

GUI Parameter	Description	CLI Command
Enable Accelerate Live Split Performance	Enables performance improvements in live splitting for the Windows Media proxy.	wmt accelerate live-split enable
Enable Accelerate VoD Performance	Enables performance improvements in video-on-demand for the Windows Media proxy.	wmt accelerate vod enable
Restrict HTTP Allowed Extensions	Lets you add or remove permitted extensions.	—
HTTP Allowed Extensions	List of allowable extensions for HTTP. You can add or delete filename extensions from this list with the following restrictions: - Each extension must be alphanumeric, with the first character in the extension being an alphabetic character. - You cannot have more than 10 characters in a filename extension. - You cannot add more than 20 filename extensions to the allowed list.	wmt http allow extension <i>ext1 ext2 ...</i>
Enable Fast Start Feature	Enables Fast Start for MMS-over-HTTP or RTSP.	wmt fast-start enable
Fast Start Max Bandwidth	Maximum bandwidth allowed per player when Fast Start is used to serve packets to this player.	wmt fast-start max-bandwidth <i>kbits</i>
Enable Fast Cache	Enables Fast Cache for MMS-over-HTTP or RTSP.	wmt fast-cache enable
Fast Cache Max Delivery Rate	Maximum delivery rate allowed per player when Fast Cache is used to deliver packets to this player.	wmt fast-cache max-delivery-rate <i>num</i>
Windows Media Multicast Settings		
Number of hops to live	Number of hops to live for multicast WMT packets. The default is 5 hops. The range is 0 to 255 hops.	wmt multicast time-to-live <i>ttl</i>
WMT Advanced Client Settings		
Idle Timeout	Number of seconds to timeout when the client connection is idle. The default is 120. The range is 30 to 300 seconds.	wmt advanced client idle-timeout <i>seconds</i>
Maximum data packet size	Maximum packet size (in bytes) allowed. The default is 1500 bytes. The range is 512–2048.	wmt advanced client max-packet-size <i>bytes</i>

Table 9-2 *Windows Media Settings Window (continued)*

GUI Parameter	Description	CLI Command
WMT Advanced Server Settings		
Enable Log Forwarding	Enables log forwarding to an upstream Content Engine or Windows Media server.	wmt advanced server log-forwarding enable

About Fast Start

Windows Media Player must buffer a certain amount of data before it can start rendering content. If clients are using Windows Media Player for Windows XP or a later version of the player, you can use Fast Start to provide data directly to the buffer at speeds higher than the bit rate of the content requested. This feature enables users to start receiving content more quickly. After the initial buffer requirement is fulfilled, on-demand and broadcast content streams at the bit rate defined by the content stream format.



Note

In ACNS 5.2.x software, Fast Start is available only for MMS-over-HTTP requests with Windows Media Services Version 9.0.

In ACNS 5.3 software and later releases, support for Fast Start is also available for RTSP requests from a Windows Media 9 player.

The Fast Start feature is used only by clients that connect to a unicast stream.

Using Fast Start enables users to have a better experience when playing back the content. Users can fast-forward and rewind content without additional delay and rebuffering. A Windows Media Player that connects through broadband networks starts playing the content more quickly, making the experience much more like viewing a television program or listening to a radio broadcast. Users notice that server-side playlists streaming from your publishing point switch smoothly and seamlessly between content items. Additionally, the prebuffering of data makes the Windows Media Player resistant to playback errors due to lost packets or other network issues. The benefit of Fast Start is not available to the first client connecting to a live stream.

The increased bandwidth that the Fast Start feature initially uses to send data to the Windows Media Player can overburden a network if many players connect to the stream at the same time. To reduce the risk of network congestion caused by the Fast Start feature, you can limit the amount of bandwidth the Fast Start feature uses to stream to each player.

To configure Fast Start on the Content Engine using the CLI, follow these steps:

-
- Step 1** To enable Fast Start on the Content Engine, enter the following command:
- ```
ce-565(config)# wmt fast-start enable
```
- Step 2** To set the maximum burst bandwidth allowed per media player when Fast Start is used to serve packets to the media player, enter the following command:
- ```
ce-565(config)# wmt fast-start max-bandwidth number
```

Setting a maximum bandwidth reduces the risk of network congestion by limiting the amount of bandwidth that Fast Start can use to stream content to each media player.

Step 3 To verify that Fast Start is enabled, enter the following command:

```
ce-565# show wmt
```

About Fast Cache

When used with the Windows Media Player 9 Series, Fast Cache provides a way to stream content to clients faster than the data rate specified by the stream format. For example, using Fast Cache, the server can transmit a 128-kilobit per second (kbps) stream at 700 kbps. The stream is still rendered in Windows Media Player at the specified data rate, but the client is able to buffer a much larger portion of the content before rendering it. This allows the client to handle variable network conditions without a perceptible impact on the playback quality of either on-demand or broadcast content.

This ability is useful in the following situations:

- When the available network bandwidth of the client exceeds the required bandwidth of the content; for example, clients that use a cable modem, Digital Subscriber Line (DSL) connection, or corporate intranets
- When the network connectivity is intermittent or has high latency, for example, in wireless networks
- When the quality of the content received is of paramount importance, for example, in businesses that provide pay-per-view movies



Note

In ACNS 5.2.x software, Fast Cache is available only for MMS-over-HTTP requests with Windows Media Services Version 9.0. In ACNS 5.3 software and later releases, support for Fast Cache is also available for WMT RTSP requests from Windows Media 9 players.

Fast Cache is valid only for the RTSP and HTTP protocols; it is not valid for the MMS protocol.

Fast Cache is not supported for live streams and requests that result in a cache miss.

The following scenario illustrates bandwidth control and Fast Cache speed adjustment. Client A is a client PC that is running the Windows Media 9 player, which has an IP address of 10.77.140.133, and is configured as follows:

```
<?xml version="1.0"?>
<BandwidthSpec>
  <BandwidthRule>
    <ClientNetwork>10.77.140.133/32</ClientNetwork>
    <description>(Apply to my PC)</description>
    <Allow limit="1100" service="wmt"/>
  </BandwidthRule>
  ...
</BandwidthSpec>
```

If Client A requests media content through RTSP (RTSP TCP mode) and the bit rate of the requested file is 500 kbps, the following occurs. Because the available bandwidth for Client A is 1100 kbps, the Fast-Cache speed for Client A is restricted to 2 (because a speed of 2 would consume 1000 kbps [500 x 2 = 1000 kbps]). The bandwidth consumed by Client A is less than the total available bandwidth limit of 1100 kbps.

To configure Fast Cache on the Content Engine using the CLI, follow these steps:

Step 1 To enable Fast Cache on the Content Engine, enter the following command:

```
ce-565(config)# wmt fast-cache enable
```

Step 2 To set the maximum delivery rate allowed per media player when Fast Cache is used to serve packets to the player, enter the following command:

```
ce-565(config)# wmt fast-cache max-delivery-rate number
```



Note When Fast Cache is configured, the Content Engine serves content using either the bit rate specified in the client request or the maximum delivery rate configured for Fast Cache in the Content Engine, whichever is smaller.

Step 3 To verify that Fast Cache is enabled, enter the following command:

```
ContentEngine# show wmt
```

Configuring the Windows Media Player to Use the Content Engine as a Proxy

To configure the Windows Media Player to use the Content Engine as the Windows Media proxy, follow these steps:

Step 1 Open the Windows Media Player.

Step 2 Choose **Tools > Options**.

Step 3 Click the **Network** tab.

Step 4 Click the **Multicast**, **UDP**, **TCP**, and **HTTP** radio buttons if they are not selected.

Step 5 Click **RTSP** under Proxy Settings, and click **Configure**. The Configure Protocol Page in WMT Media Player page appears.

Step 6 Click **Use the following proxy server**.

Step 7 Enter the IP address of the Content Engine in the Address field.

Step 8 In the Port field, enter the port number that you entered in [Step 2](#).


Step 9 Click **OK**.

Configuring Windows Media Live Splitting for Unicast and Multicast Transmissions

Splitting is a method of sending live data transmissions across the Internet to a proxy edge device. The live stream is then fanned out to multiple downstream clients using unicast or multicast. (See [Figure 9-12](#).)

The Windows Media-enabled Content Engine acts as a proxy for the streaming server or broadcast server and replicates the stream to local users. Splitting the stream at the Content Engine proxy is better than splitting it at the server because the proxy is closer to the clients and potentially saves considerable network bandwidth between the client and the server.

Figure 9-12 Live Splitting



When a client requests a publishing point on a server (without specifying an ASF file), then the Content Engine dynamically creates an alias file that references the remote server. All further requests to that station are served by splitting the stream. Note that when the first client (Client 1) that requested the original stream disconnects from the network, the proxy continues to serve the other clients (Client 2 and Client 3), until all clients disconnect from the network.



Note

Live splitting is supported for different data packet transport protocols (HTTP, RTSP TCP [RTSPT], RTSP UDP [RTSPU], and IP multicast).

Live Stream Splitting Scenarios

Based on the capabilities and limitations of the network, a Content Engine can receive and deliver Windows Media streaming content through IP multicast or unicast transmissions in the following four combinations:

- Unicast-In Multicast-Out
- Multicast-In Multicast-Out
- Multicast-In Unicast-Out
- Unicast-In Unicast-Out

When you combine the two Windows Media multicast-out features, you can receive and deliver Windows Media streaming media content through IP multicasting and do conversions from multicast to unicast and from unicast to multicast.



Note

You must accept a Windows Media license and enable Windows Media on the Content Engine and device groups before you can enable Windows Media multicasting in your ACNS 5.x network. See the [“Enabling Windows Media Services”](#) section on page 9-15.

Each station needs a multicast IP address. You must enter a valid Class D IP address multicast address in the range 224.0.0.0 to 239.255.255.255, except for the reserved IP ranges based on RFC 1700 and related documents, as follows:

- 224.0.0.0 to 224.0.6.255
- 224.0.13.0 to 224.0.13.255
- 224.1.0.0 to 224.2.255.255
- 232.0.0.0 to 232.255.255.255

**Note**

You must choose a multicast IP address that does not conflict internally within the same multicast-enabled network configuration. This multicast IP address is not related to the IP address of the Content Engine. (See the “[Unusable IP Multicast Addresses](#)” section on page B-1.)

The allowed multicast port range is 1 through 65535. However, the multicast-enabled network may impose certain restrictions on your choice of port. Normally, port numbers below 1024 should be avoided, but the Content Engine does not enforce any restrictions.

**Note**

If a live stream is interrupted on the server side, you must stop the multicast station and then restart the same station to resume live multicasting.

Configuring Unicast-In Multicast-Out

The unicast-in multicast-out multicast delivery feature enables you to distribute streaming media efficiently by allowing different devices on the IP multicast to receive a single stream of media content from the Content Engine simultaneously. This can save significant network bandwidth because a single stream is sent to many devices, rather than sending a single stream to a single device every time that this stream is requested. This multicast delivery feature is enabled by setting up a multicast address on the Content Engine to which different devices, configured to receive content from the same channel, can subscribe. The delivering device sends content to the multicast address set up at the Content Engine, from which it becomes available to all subscribed receiving devices.

To enable Windows Media services for unicast-in multicast-out, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **Applications > Streaming > Windows Media > Multicast Stations**. The Windows Media Multicast Station for Content Engine window appears.
- Step 4** In the Aggregate Settings section, the **Yes** radio button is chosen by default, which means that the broadcast aliases created for the Content Engine as well as associated device groups are displayed and applied. Alternatively, to apply the settings of the Content Engine only, click the **No** radio button.
- Step 5** Click the **Create New Windows Media Multicast Stations** icon. The Creating New Windows Media Station for Content Engine window appears. (See [Figure 9-13](#).)

Figure 9-13 Windows Media Multicast Station Configuration Example

The screenshot shows the Cisco Application and Content Networking System configuration interface. The main heading is "Creating new Windows Media Station for Content Engine, CONTENTENGINE". The interface is divided into several sections:

- Windows Media Station:**
 - Station Name:
 - Address:
 - Port:
 - Media:
 - Repeat Forever: ☐
 - Unicast URL:
 - Schedule Start Now: ☐
- Failover Settings:**
 - Retry Count:
 - Retry Interval: minutes
 - Alternate Source: Please specify up to 8 URLs, separated by space
- Windows Media Multicast Logging Settings:**
 - Enable Windows Media Multicast Logging: ☐
 - Enable Local Multicast Logging: ☐
 - Multicast Logging URL:

At the bottom right, there are "Submit" and "Cancel" buttons. The interface is viewed through a Microsoft Internet Explorer browser window.

- Step 6** Under the Windows Media Station heading, configure the following:
- In the Station Name field, enter a name for the multicast station.
 - In the Address field, enter a Class D IP address to be used as the multicast station IP address.
 - In the Port field, enter the port number to be used by this station.
 - In the Media field, enter the source of the Windows Media streaming media.
 - If you want this source to play the media files without interruption, check the **Repeat Forever** check box. Otherwise, the multicast stream stops playing when the end of the *source.nsc* file is reached.

In the example in [Figure 9-13](#), a multicast station named *Multicast Station 1* is configured and used by the Content Engine as the multicast source. Its Class D IP address is 233.33.33.333, and the multicast port is 3333.



Note This source file *source.asf* can be located on any Windows Media server, including a Windows server, or on the Content Engine. In the case of the Content Engine, pre-positioned media files should be stored in the */local1/wmt_vod* directory. In this scenario, the media source is represented by *rtsp://CEIPaddress/wmt_vod/source.asf*, where *CEIPaddress* is the IP address of the Content Engine.

- In the Unicast URL field, enter the URL to allow unicast live splitting for clients who cannot be reached by multicast.

For viewers who are unable to receive Windows Media multicast streams or for virtual private network (VPN) users who are not accessible to multicasting, the unicast published URL can be used for unicast live splitting. The Windows Media Player will fall back to unicast in the event of multicast failure. The unicast published URL is made available inside the multicast description metafile (.nsc) to viewers for automatic fallback to unicast live streaming. The unicast URL serves as a backup URL if the request is not intercepted by the ACNS network, when the ACNS network is unable to route multicast streams, or when the client is unable to receive multicast streams.

- g. To enable the automatic restart of the active multicast stations after a device reloads, check the **Schedule Start Now** check box.

Step 7 Under the Failover Settings heading, configure the following:

- a. To set the number of times for the multicast station to retry to multicast the stream from the configured source, enter a number in the Retry Count field.
- b. To configure the period of time between retries, enter the number of minutes in the Retry Interval field.
- c. If the primary source of the multicast fails, you can configure up to eight alternative sources for the multicast by entering the URLs in the Alternate Source field.

Step 8 Under the Windows Media Multicast Logging Settings heading, enable multicast logging.

- To enable the Content Engine to collect statistics regarding the multicast stream using the .nsc file, check the **Enable Windows Media Multicast Logging** check box.

A Windows Media server supports a large number of media players for a particular stream during multicast, providing scalability without overloading the server. When Windows Media multicast logging is enabled, you can also receive feedback on certain statistical data, such as the number of times that buffering occurred while the stream was played, the number of packets lost during transmission, the browser type used when the Windows Media Player is embedded in a browser, and the protocol used to access the stream. These details are helpful for business analysts and program creators to learn about the effectiveness of the program multicast statistics. The media player obtains the IP address and port to monitor for content. When the stream stops playing, the Windows Media Player automatically collects and sends the statistics to the multicast logging URL using the HTTP POST request method. You can choose to specify the URL to which Windows media transaction log files must be sent.

- Alternatively, to allow Windows Media Players to send multicast statistics to the Windows Media server configured on the Content Engine, check the **Enable Local Multicast Logging** check box, and enter a URL in the Multicast Logging URL field.

Step 9 To save the settings, click **Submit**.

Step 10 Open your Windows Media Player and choose **File > Open URL**. Enter the following URL:

http://CEIPAddress/MulticastStation1.nsc



Note

The source file for streaming can be located on any Windows Media server, including a Windows server, or on the Content Engine. In the case of the Content Engine, pre-positioned media files must be stored in the /local1/wmt_vod directory. In this case, the source of the media file is represented by the URL `rtsp://CEIPAddress/wmt_vod/asf_file`.

Step 11 Click **OK**. The Windows Media Player should retrieve the multicast description .nsc file and join the multicast station that was specified in [Step 6](#).

**Note**

The use of port 80 is implied in the URL for Windows Media multicasting. An equivalent URL is `http://CEIPaddress:80/test1.nsc`.

To configure the multicast station settings from the CLI, use the **wmt multicast station-configuration** global configuration command.

```
wmt multicast {station-configuration name [{dest_addr dest_port media_source] [log {local | webserver webserver_url}] [play-forever] [unicast-url url] | failover {alternate-source alt_src_url | retry-count number | retry-interval number} | schedule-start {minute hour day month | now}] | time-to-live ttl}
```

Configuring Multicast-In Multicast-Out

The multicast-in multicast-out multicast receive feature enables you to receive multicast Windows Media streams delivered through IP multicasting and then relay them to end users through another delivery channel (unicast or multicast).

In this multicasting scenario, a description file, **.nsc*, is created that is accessible through multicast-out to clients. This is similar to the unicast-in multicast-out scenario except that the input source is multicast. The clients use this description file to subscribe to the multicast.

To enable Windows Media services for multicast-in multicast-out, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **Applications > Streaming > Windows Media > Multicast Stations**. The Windows Media Multicast Station for Content Engine window appears.
- Step 4** In the Aggregate Settings section, the **Yes** radio button is chosen by default, which means that the broadcast aliases created for the Content Engine as well as associated device groups are displayed and applied. Alternatively, to apply the settings of the Content Engine only, click the **No** radio button.
- Step 5** Click the **Create New Windows Media Multicast Stations** icon. The Creating New Windows Media Station for Content Engine window appears. (See [Figure 9-14](#).)

Figure 9-14 Windows Media Multicast Station Configuration Example

The screenshot shows the Cisco ACNS configuration interface. The main configuration area is titled "Creating new Windows Media Station for Content Engine, CONTENTENGINE". It contains the following sections:

- Windows Media Station:**
 - Station Name: Multicast Station 2
 - Address: 233.33.33.334
 - Port: 3334
 - Media: source.nsc
 - Repeat Forever: ☐
 - Unicast URL:
 - Schedule Start Now: ☐
- Failover Settings:**
 - Retry Count:
 - Retry Interval: minutes
 - Alternate Source: Please specify up to 8 URLs, separated by space
- Windows Media Multicast Logging Settings:**
 - Enable Windows Media Multicast Logging: ☐
 - Enable Local Multicast Logging: ☐
 - Multicast Logging URL:

Buttons for "Submit" and "Cancel" are located at the bottom right of the configuration area.

- Step 6** Under the Windows Media Station heading, configure the following:
- In the Station Name field, enter a name for the multicast station.
 - In the Address field, enter a Class D IP address to be used as the multicast station IP address.
 - In the Port field, enter the port number to be used by this station.
 - In the Media field, enter the filename of the multicast *source*.nsc* file.
 - To play the media files without interruption, check the **Repeat Forever** check box if you want this source. Otherwise, the multicast stream stops playing when the end of the *source*.nsc* file is reached.

In this example, a multicast station named *Multicast Station 2* is configured and used by the Content Engine as the multicast source. Its Class D IP address is 233.33.33.334, and the multicast port is 3334.



Note This source file *source.nsc* can be located on any Windows Media server, including a Windows server, or on the Content Engine. In the case of the Content Engine, pre-positioned media files should be stored in the */local1/wmt_vod* directory. In this scenario, the media source is represented by *rtsp://CEIPaddress/wmt_vod/source.asf*, where *CEIPaddress* is the IP address of the Content Engine.

- In the Unicast URL field, enter the URL to allow unicast live splitting for clients who cannot be reached by multicast.

For viewers who are unable to receive Windows Media multicast streams or for VPN users who are not accessible to multicasting, the unicast published URL can be used for unicast live splitting. The Windows Media Player will fall back to unicast in the event of multicast failure. The unicast published URL is made available inside the multicast description metafile (.nsc) to viewers for automatic fallback to unicast live streaming. When it is unable to route multicast or when the client is unable to receive multicast streams, the unicast URL serves as a backup URL if the request is not intercepted by the ACNS network.

- g. To enable the automatic restart of the active multicast stations after a device reloads, check the **Schedule Start Now** check box.

Step 7 Under the Failover Settings heading, configure the following:

- a. To set the number of times for the multicast station to retry to multicast the stream from the configured source, enter a number in the Retry Count field.
- b. To configure the period of time between retries, enter the number of minutes in the Retry Interval field.
- c. If the primary source of the multicast fails, you can configure up to eight alternative sources for the multicast by entering the URLs in the Alternate Source field.

Step 8 Under the Windows Media Multicast Logging Settings heading, enable multicast logging.

- To enable the Content Engine to collect statistics regarding the multicast stream using the .nsc file, check the **Enable Windows Media Multicast Logging** check box.

A Windows Media server supports a large number of media players for a particular stream during multicast, providing scalability without overloading the server. When Windows Media multicast logging is enabled, you can also receive feedback on certain statistical data, such as the number of times that buffering occurred while the stream was played, the number of packets lost during transmission, the browser type used when the Windows Media media player is embedded in a browser, and the protocol used to access the stream. These details are helpful for business analysts and program creators to learn about the effectiveness of the program multicast statistics. The media player obtains the IP address and port to monitor for content. When the stream stops playing, the Windows Media Player automatically collects and sends the statistics to the multicast logging URL using the HTTP POST request method. You can choose to specify the URL to which Windows media transaction log files must be sent.

- Alternatively, to allow Windows Media Players to send multicast statistics to the Windows Media server configured on the Content Engine, check the **Enable Local Multicast Logging** check box.

Step 9 In the Multicast Logging URL field, enter the URL of the Windows Media server to which the multicast statistics need to be sent, using the HTTP POST request method. In case the Windows Media server is configured on any Windows Media server other than the Content Engine, you need to explicitly specify the URL.

Step 10 To save the settings, click **Submit**.

Step 11 Open your Windows Media player and choose **File > Open URL**. Enter the following URL:

http://CEIPAddress/MulticastStation2.nsc



Note

The source file for streaming can be located on any Windows Media server, including a Windows server, or on the Content Engine. In the case of the Content engine, pre-positioned media files must be stored in the /local1/wmt_vod directory. In this case, the source of the media file is represented by the URL, rtsp://CEIPAddress/wmt_vod/nsc_file.

- Step 12** Click **OK**. The Windows Media player should retrieve the multicast description *.nsc* file and join the multicast station that was specified in Step 6a.



Note The use of port 80 is implied in the URL for Windows Media multicasting. An equivalent URL is `http://CEIPaddress:80/test1.nsc`.

To configure the multicast station settings from the CLI, use the **wmt multicast station-configuration** global configuration command.


```
wmt multicast { station-configuration name [{ dest_addr dest_port media_source } [ log { local | webserver webserver_url } ] [ nsc-reference { ref_URL } ] [ play-forever ] [ unicast-url url ] | failover { alternate-source alt_src_url | retry-count number | retry-interval number } | schedule-start { minute hour day month | now } ] | time-to-live ttl }
```

Configuring Multicast-In Unicast-Out

The multicast-in unicast-out configuration enables you to deliver an incoming stream live to requesting clients using multicast as the source of the streaming media. In this scenario, a unicast-out publishing point is created to deliver the incoming multicast stream live to requesting clients by converting the multicast stream to unicast.

To enable Windows Media services for multicast-in unicast-out, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **Applications > Streaming > Windows Media > Broadcast Alias**. The Windows Media Broadcast Alias for Content Engine window appears.
- Step 4** In the Aggregate Settings section, the **Yes** radio button is chosen by default, which means that the broadcast aliases created for the Content Engine and its associated device groups are displayed and applied. Alternatively, to apply the settings of the Content Engine only, click the **No** radio button.
- Step 5** Click the **Create New Windows Media Broadcast Aliases** icon. The Creating New Windows Media Broadcast Alias for Content Engine window appears. (See Figure 9-15.)

Figure 9-15 Windows Media Broadcast Alias Window—Multicast-In Unicast-Out Configuration

- Step 6** In the Broadcast Alias field, enter an alias for the source URL or station name to be used for the broadcast.

When a broadcast alias is configured, a client makes a request to the Content Engine as if it were the Windows Media Server, and the Content Engine checks to see whether the incoming stream is present. If it is, then the Content Engine joins the stream and splits it to the new client. If the request is the first client request for this stream, the Content Engine sends the request out to the server and then serves it to the new client.

For example, in [Figure 9-15](#) a unicast publishing point with the alias name *broadcast1* is configured.

- Step 7** In the Source URL field, enter the URL of the station to be used for the broadcast. For example:
http://172.16.30.21/station.nsc

In this example, the multicast source URL contains the file *station.nsc*. A source URL in the format `http://server/file.nsc` tells the Content Engine that the input source is sending out Windows Media multicast streams.

- Step 8** To save the settings, click **Submit**.

- Step 9** Open your Windows Media player, and choose **File > Open URL**. Enter the following URL:
rtsp://CEIPaddress/broadcast1

where CEIPaddress is the IP address or domain name of the Content Engine and broadcast1 is the alias name used in [Step 6](#).

In this scenario an RTSP URL is used to access the streaming media and that only the *alias name* is specified instead of the **.nsc* files as in the multicast-out scenarios. This configuration converts the multicast stream to unicast and sends it to the requesting client.

- Step 10** Click **OK**. The Windows Media player should receive content from the multicast source specified in [Step 7](#).
-

To configure the broadcast alias settings from the CLI, use the following global configuration command:

```
wmt broadcast {alias-name name source url}
```

Configuring Unicast-In Unicast-Out

In this scenario, unicast-in unicast-out provides a point-to-point connection between the client and the Content Engine. The Content Engine in turn makes a single connection to the media server. Multiple requests for the same stream can be split by the Content Engine so that each client receives a distinct data stream directly from the Content Engine while the Content Engine maintains its single stream connection to the media server. The client can request a stream that contains either stored or live content.


The advantage of unicasting when streaming media over a network is that only a single stream needs to be distributed over the network between the origin server and Content Engine, but that stream can be delivered to multiple clients in a nonmulticast environment. A server running Windows Media Services can provide a unicast video stream to multiple clients through a single stream delivered to the Content Engine. Those clients can take advantage of the VCR-like controls in the Windows Media player to pause the stream or to skip backward or forward (in the case of stored content [video on demand]).

There are two ways to configure unicast-in unicast-out:

- By live splitting without any configuration. In this case, the Content Engine acts as a proxy. When clients request the same unicast URL, the Content Engine proxy automatically splits the stream from the source to the clients.
- By configuring the Content Engine with a broadcast alias. In this case, a client makes the request to the Content Engine as if it were the Windows Media Server, and the Content Engine checks to see whether the incoming stream is present. If it is, then the Content Engine joins the stream and splits it to the new client. If the request is the first client request for this stream, the Content Engine sends the request out to the server and then serves it to the new client.

To enable Windows Media services for unicast-in unicast-out by configuring the Content Engine with a broadcast alias, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **Applications > Streaming > Windows Media > Broadcast Alias**. The Windows Media Broadcast Alias for Content Engine window appears.
- Step 4** In the Aggregate Settings section, the **Yes** radio button is chosen by default, which means that the broadcast aliases created for the Content Engine and associated device groups are displayed and applied. Alternatively, to apply the settings of the Content Engine only, click the **No** radio button.
- Step 5** Click the **Create New Windows Media Broadcast Alias** icon. The Creating New Windows Media Broadcast Alias for Content Engine window appears. (See [Figure 9-16](#).)

Figure 9-16 Windows Media Broadcast Alias Window—Unicast-In Unicast-Out Configuration

- Step 6** In the Broadcast Alias field, enter an alias name for the source URL.
For example, in [Figure 9-16](#) a unicast publishing point with the alias name *cotv* is configured.
- Step 7** In the Source URL field, enter the source URL of the content to be used for the transmission. For example:
rtsp://wms.company.com/cotv
where *wms.company.com* is the name of the Windows Media Server, and *cotv* is the name used when the broadcast alias is created. The RTSP protocol is used to retrieve the stream.
- Step 8** To save the settings, click **Submit**.
- Step 9** Open your Windows Media player, and choose **File > Open URL**. Enter the following URL:
rtsp://CEIPaddress/cotv
where *CEIPaddress* is the IP address or domain name of the Content Engine, and *cotv* is the alias name used in [Step 6](#).
- Step 10** Click **OK**. The Windows Media player should receive content from the unicast source specified in [Step 7](#).

To configure the broadcast alias settings from the CLI, use the following global configuration command:

```
wmt broadcast {alias-name name source url}
```


Configuring Windows Media Multicast Station Schedules

To configure Windows Media multicast station schedules for the Content Engine, follow these steps:

-
Step 1	Choose Devices > Devices . The Content Engine window appears.
Step 2	Click the Edit icon next to the Content Engine whose station schedule you want to set. The Contents pane appears on the left.
Step 3	From the Contents pane, choose Applications > Streaming > Windows Media > Multicast Station Schedules . The Windows Media Multicast Station Schedule for Content Engine window appears.
Step 4	Click the Create New Windows Media Station Schedule icon. The Creating New Windows Media Station Schedule for Content Engine window appears.
Step 5	Choose a multicast station from the Station Name drop-down list.
Step 6	In the Month field, enter the month (1–12) that you want the schedule to start.
Step 7	In the Day field, enter the day (1–31) that you want the schedule to start.
Step 8	In the Hour field, enter hour (0–23) that you want the schedule to start.
Step 9	In the Minute field, enter the minute (0–59) that you want the schedule to start.
Step 10	To save the settings, click Submit .
-

To configure multicast station schedules from the CLI, use the following global configuration command:

wmt multicast schedule-start *name minute hour day month*

Configuring Weighted Load Balancing for Live Stream Splitting

The live stream routing (LSR) application that operates on each Content Engine establishes the splitting tree route for live streaming URLs from the home Content Engine to the root Content Engine. Each live streaming URL belongs to a channel in the ACNS network.

The LSR application gathers a list of all Content Engines that are on the location path from the home Content Engine to the root Content Engine that are subscribed to the channel. Then, it builds a Content Engine path by selecting one Content Engine from each location on the location path. This Content Engine path is called live splitting path for this particular URL from this Content Engine.

This path is encoded in a special URL format and is sent to the next hop Content Engine on the path. The next hop Content Engine decodes the special URL and looks for its next hop. It then sends a request to its next hop and so on, until the root Content Engine is finally reached.

The LSR application obtains a matrix of Content Engines from the channel routing library for each channel that contains the URL. The matrix contains an ordered list for each location on the path from the home Content Engine to the root Content Engine. Each ordered list is composed of a number of Content Engines within the location that are assigned to the specified channel.

On receiving the matrix, the LSR application selects one Content Engine from each ordered list, forming a path from the home Content Engine to the root Content Engine. This path is used for live streaming URLs. The LSR application selects Content Engines from the list in the same order that they are listed. A Content Engine in the middle of the list is selected only after all preceding Content Engines have failed.

Because the LSR application always selects one Content Engine from the ordered list, when there are many different live streaming URLs, the load within one location is randomly distributed among all assigned Content Engines. Content Engines of different capacities (such as the CE-7325 and the CE-565) have an equal probability of being selected as the splitting Content Engine on the location path.

With a weighted load balancing scheme, users can influence the order in which Content Engines are listed. By adjusting the live split load weight, you can influence the probability of one Content Engine receiving live stream splitting traffic over another Content Engine in much the same way that you can influence location leader and forwarder selections for channel routing by assigning weight and priority values.

When the channel routing application generates the ordered list of Content Engines for each location, the load weight assigned to one Content Engine is compared against the weights of other Content Engines in the same location. The Content Engines are then ordered so that the Content Engine with the higher load weight has a greater probability of being ordered at the beginning of the list.

The weight is a probabilistic value; therefore, it is possible for a Content Engine with a lower weight to be at the beginning of the list, while a Content Engine with a higher weight could be at the end of the list. Weighted load balancing is mainly useful when there are multiple live stream URLs.

See also, the [“Configuring Location Leader Preference and Forwarder Selection Probability” section on page 5-21](#) and the [“Configuring the Forwarder Lookup Level” section on page 5-22](#).

To configure the live split load weight for a Content Engine, follow these steps:

-
- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
 - Step 2** To specify the relative probability that a Content Engine is likely to receive live stream splitting traffic as compared to other Content Engines, click the **Edit** icon next to the Content Engine that you want to configure. The Modifying Content Engines window appears.
 - Step 3** In the Contents pane, choose **Applications > Streaming > Live Splitting**. The Live Splitting Settings for Content Engine window appears.
 - Step 4** Under the Live Splitting Settings section, to specify the relative probability of a Content Engine being selected as a live stream splitter, perform one of the following:
 - To associate a weight with the Content Engine being selected for live stream splitting for any URL, click the **Weighted** radio button.
 - To adjust the probability of live stream splitter selection, click and drag the Weighted slider control across the calibrated ruler. The range is 1 to 20. The default weight value is 1. The value corresponding to the position of the slider is displayed in the field next to it. The higher the weight, the higher the probability of the Content Engine being chosen as a live stream splitter.
 - To specify that the Content Engine should not be selected as a live stream splitter unless there is no other Content Engine in the location subscribed to the channel, or all other Content Engines assigned to the channel in the same location have failed, click the **Last Resort** radio button. This setting is associated with a weight of 0.
 - To specify that the Content Engine should never be selected as a live stream splitter in any case, click the **Never Use** radio button. This setting is associated with a weight of -1.
 - Step 5** In the Content Engines in Location section, view information about other Content Engines that are assigned to the same location. This information is displayed in a standard table format with columns for Content Engine, Status, Model, Software Version, and Live Splitting Probability.

Step 6 To save the settings, click **Submit**. A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved. You can also revert to the previously configured settings by clicking **Reset**. The Reset button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

If you try to leave this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box only appears if you are using the Internet Explorer browser.

Related CLI Commands

The live splitting probability configured for the Content Engine can be viewed using the **show distribution channels** command. However, this parameter can be configured only using the Content Distribution Manager GUI for each Content Engine. Default values are assumed if you do not manually configure this parameter.

The **show distribution location live-load-weight** command displays the live-load-weight value of Content Engines that are assigned to the channel within the location.

Configuring Bandwidth Settings for Content Services

Bandwidth settings for content services control the amount of bandwidth that each content service uses to stream content to users.

Content services bandwidth includes bandwidth allocation for Windows Media, RealProxy, RealServer, and Cisco Streaming Engine services. Windows Media bandwidth settings apply to Windows Media streaming of live, cached, and pre-positioned content. RealServer bandwidth settings apply to RealServer streaming of pre-positioned and live content that has been specified in the manifest file for a channel. RealProxy bandwidth settings apply to RealProxy streaming of cached and live content that has not been specified in the manifest file for a channel. Cisco Streaming Engine bandwidth settings apply to the standard RTSP server streaming of pre-positioned content only.

For each type of bandwidth, you can specify the amount of bandwidth to be used for a particular time interval. This is called scheduled bandwidth. Default bandwidth is the amount of bandwidth associated with each content services type when there is no scheduled bandwidth. When a Content Engine is a member of a device group, and if the default bandwidth is specified for that Content Engine, that bandwidth setting overrides the settings at the device group level. However, if there are no settings defined for a Content Engine, the device group settings become applicable for the Content Engine. When a Content Engine is a member of multiple device groups, the settings of the device group that have been most recently modified are used.

Maximum bandwidth specifies the upper limit beyond which any specified bandwidth is not allowed. The total bandwidth configured for all content services must not exceed the bandwidth limits specified for a particular Content Engine model. In addition, the license keys configured for Windows Media further restrict the maximum bandwidth, which must be less than the allowable maximum bandwidth for each Content Engine model.

Configuring Content Services Default and Maximum Bandwidth Settings

To configure content services default and maximum bandwidth settings, follow these steps:

-
- Step 1** Choose **Devices > Devices**. The Devices window appears.
 - Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to view. The Modifying Content Engine window appears.
 - Step 3** In the Contents pane, choose **Applications > Default and Maximum Bandwidth**. The Default and Maximum Bandwidth for Content Engine window appears.
 - Step 4** For Windows Media Incoming, in the Default Bandwidth field, specify the amount of default bandwidth, in kbps, allowed for incoming Windows Media traffic from end users. In the Max Bandwidth field, specify the maximum bandwidth, in kbps, permitted by the system license.
 - Step 5** For Windows Media Outgoing, in the Default Bandwidth field, specify the amount of default bandwidth, in kbps, allowed for outgoing Windows Media traffic from Content Engines. In the Max Bandwidth field, specify the maximum bandwidth, in kbps, permitted by the system license.
 - Step 6** For Real Proxy Incoming, in the Default Bandwidth field, specify the amount of default bandwidth, in kbps, allowed for incoming RealProxy traffic from end users. In the Max Bandwidth field, specify the maximum bandwidth, in kbps, permitted by the system license. RealProxy traffic refers to RealMedia traffic that has been cached in response to requests from end users.
 - Step 7** For Real Proxy Outgoing, in the Default Bandwidth field, specify the amount of default bandwidth, in kbps, allowed for outgoing RealProxy traffic from Content Engines. In the Max Bandwidth field, specify the maximum bandwidth, in kbps, permitted by the system license.
 - Step 8** For Real Server, in the Default Bandwidth field, specify the amount of default bandwidth, in kbps, for streaming RealServer traffic to end users. In the Max Bandwidth field, specify the maximum bandwidth, in kbps, permitted by the system license. RealServer traffic refers to RealMedia content that has been pre-positioned.
 - Step 9** For Cisco Streaming Engine, in the Default Bandwidth field, specify the amount of default bandwidth, in kbps, allowed for the Cisco Streaming Engine to stream content to end users. In the Max Bandwidth field, specify the maximum bandwidth, in kbps, permitted by the system license.
 - Step 10** For HTTP, in the Default Bandwidth field, specify the amount of default bandwidth, in kbps, associated with HTTP traffic. In the Max Bandwidth field, specify the maximum allowed bandwidth, in kbps, to handle HTTP requests.
 - Step 11** Click **Submit** to save your settings. A “Click Submit to Save” message appears in red next to the current settings when there are pending changes to be saved after you have applied default and device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The **Reset** button is visible only when you apply default or group settings to change the current device settings but the settings have not yet been submitted.



Note For other operations that you can perform using the icons in the taskbar, see [Table 3-3, “Content Distribution Manager GUI Icons.”](#)

To configure content services default and maximum bandwidth settings from the CLI, use the following global configuration command:

```
bandwidth {real-proxy {incoming kbits | outgoing kbits} | real-server kbits | wmt {incoming kbits | outgoing kbits} | cisco-streaming-engine kbits | http kbits} {default | max-bandwidth}
```

Configuring Scheduled Bandwidth Settings

To configure scheduled bandwidth settings for content services, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **Applications > Bandwidth Schedules**. The Bandwidth Schedule for Content Engine window appears.
- Step 4** In Aggregate Settings, the **Yes** radio button is chosen by default, which means that the bandwidth configurations of the Content Engine as well as the device groups to which it is assigned are displayed and applied. To apply and view the bandwidth settings of the Content Engine only, click the **No** radio button.
- Step 5** Click the **Create New Bandwidth Schedule** icon. The Create New Content Service Bandwidth Settings window appears. (See Figure 9-17.) Table 9-3 describes each field displayed in this window.

Figure 9-17 Create New Content Service Bandwidth Settings Window

The screenshot shows the Cisco Application and Content Networking System GUI in Microsoft Internet Explorer. The main window is titled "Create New Content Service Bandwidth Settings for Content Engine, CONTENTENGINE". The left sidebar shows a tree view with "Contents" expanded, and "Bandwidth Schedules" selected. The main content area is divided into two sections: "Bandwidth Configuration" and "Bandwidth Time Settings".

Bandwidth Configuration:

- Bandwidth Type:** A dropdown menu with the text "Please make a choice".
- Bandwidth Rate:** A text input field followed by "Kbps".

Bandwidth Time Settings:

- Time Selection:** A section with two radio buttons: "Use specific days" (selected) and "Specify day range".
- Start Time:** A text input field.
- End Time:** A text input field.
- Full Week:** A checkbox.
- Days:** A row of checkboxes for Sun, Mon, Tue, Wed, Thu, Fri, and Sat.
- Start Day:** A dropdown menu with "Sunday" selected.
- End Day:** A dropdown menu with "Sunday" selected.

Note: * - Required Field

At the bottom right of the window are "Submit" and "Cancel" buttons.

- Step 6** In the Bandwidth Configuration section, choose a bandwidth type from the Bandwidth Type drop-down list.
- Step 7** In the Bandwidth Rate field, specify the bandwidth rate allowed in kbps.
- Step 8** Under the Bandwidth Time Settings headings, the current time is displayed in UTC. For Time Selection, enter the start time and end time for bandwidth usage (in hh:mm) in the appropriate fields.



Note For a schedule from 8:00 p.m. to 8:00 a.m., the administrator must configure two schedules in order to span the two days; one from 8:00 p.m. to 11:59 p.m. (2000 to 2359) and another from 12:00 a.m. to 8:00 a.m. (0000 to 0800).

- Step 9** To specify whether the bandwidth settings apply for an entire week or for specific days of the week, click the **Use Specific days** radio button. Check the appropriate check boxes to denote the days on which the configured settings apply.
- Step 10** Alternatively, to choose a starting day of the week and an ending day of the week for the bandwidth configurations, click the **Specific day range** radio button.
- Step 11** To save your settings, click **Submit**.

Table 9-3 Streaming Bandwidth Configuration Settings

GUI Parameter	Function	CLI Command
Bandwidth Type	Windows Media Incoming is for incoming Windows Media streaming content requests from end users. Windows Media Outgoing is for outgoing Windows Media media from Content Engines. Real Proxy Incoming is for incoming RealProxy traffic from end users. Real Proxy Outgoing is for outgoing RealProxy traffic from Content Engines. Real Server is for streaming RealMedia pre-positioned content to end users. Cisco Streaming Engine is for streaming content in response to RTSP requests from end users HTTP is for sending content in response to HTTP requests from end users	<pre>bandwidth {real-proxy {incoming kbits outgoing kbits} real-server kbits wmt {incoming kbits outgoing kbits} cisco-streaming-engine kbits http kbits} {start-time weekday time end-time weekday time}</pre>
Bandwidth Rate	Maximum amount of bandwidth you want to allow in kilobits per second.	
Start Time	Time of day for the bandwidth rate setting to start, using a 24-hour clock in local time (hh:mm).	
End Time	Time of day for the bandwidth rate setting to end (hh:mm).	

Table 9-3 Streaming Bandwidth Configuration Settings (continued)

GUI Parameter	Function	CLI Command
Use specific days	Days of the week on which configured bandwidth settings apply. - Full Week—Specifies that the bandwidth settings are applicable for an entire week - Sun, Mon, Tue, Wed, Thu, Fri, and Sat—Specifies the days of the week on which the bandwidth settings apply	bandwidth { real-proxy incoming <i>kbits</i> outgoing <i>kbits</i> } real-server <i>kbits</i> wmt incoming <i>kbits</i> outgoing <i>kbits</i> } cisco-streaming-engine <i>kbits</i> http <i>kbits</i> } { start-time <i>weekday time</i> end-time <i>weekday time</i> }
Specific day range	Range of days of the week on which configured bandwidth settings apply. - Start day—Day of the week to start for allowable bandwidth - End day—Day of the week to end for allowable bandwidth	

You can configure scheduled bandwidth settings with overlapping timings. When you use the aggregate method of device and device group configuration, it is possible that the list of settings configured for devices and device groups might have overlapping schedules. However, you cannot configure bandwidths of the same type with overlapping bandwidth settings. This configuration is permissible only for bandwidths of different types.

For example, assume that a Content Engine is associated with a device group and two bandwidth configurations have been configured—one, a Windows Media incoming bandwidth with a schedule from 1300 to 1500 on Mondays and at a rate of 20 Mbps for the Content Engine, and another, a RealProxy outgoing bandwidth with a schedule from 1200 to 1700 on Mondays and at a rate of 20 Mbps for the device group. In this case, the configurations are successful because the bandwidth types are different. Assuming the same scenario, if you attempt to configure a Windows Media incoming bandwidth for the device with time settings that overlap those of the Content Engine, the configuration will fail.

The actual bandwidth settings might result in a large number of overlapping settings. As a result of such complex configurations, you might want to view a pictorial display of the allowable bandwidth settings against the days of the week. For this purpose, you can view a graph that represents the content services bandwidth settings. (See the next section, “[Displaying a Graphical Representation of the Content Service Bandwidth Settings](#).”)

When you choose the Aggregate Settings **Yes** radio button, the bandwidth settings that have been previously configured for device groups to which the Content Engine belongs cannot be modified or deleted in the Content Service Bandwidth Settings for Content Engine window. In other words, you can only view the bandwidth configurations created for the device groups.

Displaying a Graphical Representation of the Content Service Bandwidth Settings

You can view a graphical representation of the bandwidth settings for content services configured on a Content Engine. The vertical axis of the graph represents the amount of bandwidth in kbps, and the horizontal axis represents the days of the week. The units shown on the vertical axis are determined dynamically based on the bandwidth rate for a particular bandwidth type. The units shown on the horizontal axis represent 24 hours per each day of the week. Each type of bandwidth is represented by a unique color. A legend at the bottom of the graph maps colors to the corresponding bandwidth type.

To view the graph that displays the content services bandwidth, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the desired Content Engine. The Modifying Content Engine window appears.
- Step 3** In the Contents pane, choose **Applications > Default and Maximum Bandwidth**. The Default and Maximum Bandwidth Settings for Content Engine window appears.
- Step 4** Click the **Display Graph** icon in the taskbar. The Content Service Bandwidth for Content Engine popup window appears, displaying the bandwidth configuration graph. (See [Figure 9-18](#).)

To view the graph by bandwidth type, detailed or composite view, or days of the week, click a view option in the text at the top of the window. [Table 9-4](#) describes the view options.

Figure 9-18 Content Services Bandwidth Graph




Table 9-4 Viewing Options for Content Services Bandwidth Graph

Option	Description
View specific servers	Displays the bandwidth settings for the corresponding bandwidth type (server) selected.
Windows Media In	Displays the bandwidth settings for incoming Windows Media traffic.
Windows Media Out	Displays the bandwidth settings for outgoing Windows Media traffic.
Real Proxy In	Displays the bandwidth settings for incoming RealProxy traffic.
Real Proxy Out	Displays the bandwidth settings for outgoing RealProxy traffic.
Real Server	Displays the bandwidth settings for RealServer traffic.
Cisco Streaming Engine	Displays the bandwidth settings for Cisco Streaming Engine pre-positioned content.
HTTP	Displays the bandwidth settings for HTTP requests.
All Servers	Displays a consolidated view of all configured bandwidth types. This is the default view and is combined with the Full Week view.
View mode	Displays detailed and composite bandwidth settings.
Show Detailed Bandwidth/Show Effective Bandwidth	Toggles between the two options: Show Detailed Bandwidth—Displays detailed bandwidth settings for the Content Engine and its associated device groups. The bandwidth settings of the device and device groups are shown in different colors for easy identification. Show Effective Bandwidth—Displays the composite (aggregate) bandwidth settings for the Content Engine and its associated device groups.
Show Aggregate View/Show Non-Aggregate View	Toggles between the two options: Show Aggregate View—Displays the bandwidth settings configured for the corresponding device groups. Non-Aggregate View—Displays the bandwidth settings configured for the Content Engine.
View by day	Displays the bandwidth settings for a particular day or all days of the week.
Sun, Mon, Tues, Wed, Thurs, Fri, Sat	Displays the bandwidth settings for the corresponding day of the week.
Full Week	Displays the bandwidth settings for the entire week. This is the default view and is combined with the All Servers view.

- Step 5** To display the latest bandwidth settings, click **Refresh**.
- Step 6** After you have finished viewing the settings, click **Close**.
-

Configuring the Windows Media Incoming Bandwidth Bypass List

Incoming bandwidth refers to the bandwidth between a local Content Engine and the origin server. When the Content Engine is configured for Windows Media proxy services, incoming bandwidth usage for video-on-demand (VoD) content is unpredictable. This unpredictability is because the consumption of incoming bandwidth for VoD content can be triggered arbitrarily by an end user requesting the content. If the VoD content is not found in the Content Engine cache, a cache miss occurs, and the Windows Media proxy has to fetch the content from the origin server. The Content Engine administrator cannot predict the incoming bandwidth usage for such events, so a large number of cache-miss VoD requests can consume all of the incoming bandwidth.

The Windows Media incoming bandwidth bypass configuration allows the administrator to configure a list of hosts that will bypass the incoming bandwidth limitation. Content from a source that is listed as a host in this configuration is allowed to bypass the incoming bandwidth check for broadcast alias or multicast station content. This feature is particularly useful when the administrator wants to configure a broadcast alias or multicast station for a mission-critical live event, such as a message from the company CEO.

To configure the list of hosts for bypassing incoming bandwidth limits, follow these steps:

-
- Step 1** In the Content Distribution Manager GUI, Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the desired Content Engine. The Modifying Content Engine window appears with the Contents pane on the left.
- Step 3** In the Contents pane, choose **Applications > Streaming > Windows Media > Bypass List**. The Local Content Services Settings for Content Engine window appears.
- Step 4** In the Windows Media BW Incoming Bypass List field, enter the IP addresses or the host names of the hosts for which you wish to bypass the incoming bandwidth check. Separate each entry with a space.
- Step 5** To save the settings, click **Submit**. A “Click Submit to Save” message appears in red next to the Current Settings when there are pending changes to be saved after you have applied default of device group settings. To revert to the previously configured window settings, click **Reset**. The **Reset** button is visible only after you change the current device settings and before you click **Submit**.



Note You can configure a maximum of four hosts in the Windows Media BW incoming bypass list. An error message appears if you try to configure more.



Managing Caching Performance and Differentiated Services

This chapter explains how to configure TCP stack parameters for increased performance and throughput and how to configure Type of Service (ToS) and differentiated services code point (DSCP) settings to mark content for different classes of service. It contains the following sections:

- [Configuring TCP Settings for Increased Cache Performance, page 10-1](#)
- [Configuring IP Differentiated Services, page 10-9](#)
- [Enabling the MTU Discovery Utility, page 10-10](#)

Configuring TCP Settings for Increased Cache Performance

Although Content Engines are typically used to increase the throughput of HTTP streams over TCP end to end, adjusting TCP parameters for better performance is often overlooked when configuring content caching. For data transactions and queries between client and servers, the size of windows and buffers is important, and fine-tuning the TCP stack parameters therefore becomes the key to maximizing cache performance. The relevant TCP parameters to maximize cache performance and throughput include the ability to tune timeout periods, client and server receive and send buffer sizes, and TCP window scaling behavior.



Note

Because of the complexities involved in TCP parameters, care is advised in tuning these parameters. In nearly all environments, the default TCP settings are adequate. Fine tuning of TCP settings is for network administrators with adequate experience and full understanding of TCP operation details.

To configure TCP settings for the Content Engine, follow these steps:

-
- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
 - Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure TCP settings. The Device Home for Content Engine window appears.
 - Step 3** In the Contents pane, choose **General Settings > Network > TCP**. The TCP Settings for Content Engine window appears.

To configure TCP parameters, you need to configure three sets of settings in this window:

- TCP General Settings
- TCP Client Settings
- TCP Server Settings

See [Figure 10-1](#) and [Figure 10-2](#). [Table 10-1](#) describes the fields in this window and provides the corresponding CLI global configuration commands.

Figure 10-1 TCP Settings Window—General Settings

The screenshot shows the Cisco Application and Content Networking System (ACNS) interface in Microsoft Internet Explorer. The main title is "Cisco Application and Content Networking System". The left sidebar contains a "Contents" menu with options like Device Home, Device Activation, Assignments, Prepositioning, Request Routing, Request Processing, Applications, General Settings, Login Access Control, Authentication, Services, Storage, Network, Network Interfaces, External IP, Port Channel Settings, TCP, IP General Settings, IP Routes, IP ACL, IP ACL Feature Usage, Notification and Tracking, Troubleshooting, and Device Monitoring. The main content area is titled "TCP Settings for Content Engine, CONTENTENGINE". It displays "Current applied settings from Content Engine, CONTENTENGINE". Under "TCP General Settings", the following settings are shown: Enable Explicit Congestion Notification (checked), Enable Type Of Service (checked), Congestion Window Size (2 segments), Retransmit Time Multiplier (1), Initial Slow Start Threshold (2), Keepalive Probe Count (4), Keepalive Probe Interval (75 seconds), and Keepalive Timeout (90 seconds). Under "TCP Server Settings", the following settings are shown: Enable Server Satellite (unchecked), Server Max Segment Size (1460 bytes), and Server Receive Buffer Size (32 KB). The bottom right corner has "Submit" and "Cancel" buttons.

TCP General Settings	
Enable Explicit Congestion Notification:	☒
Enable Type Of Service:	☒
Congestion Window Size:	2 (segments) (1-10)
Retransmit Time Multiplier:	1 (1-3)
Initial Slow Start Threshold:	2 (2-10)
Keepalive Probe Count:	4 (1-10)
Keepalive Probe Interval:	75 (seconds) (1-120)
Keepalive Timeout:	90 (seconds) (1-120)

TCP Server Settings	
Enable Server Satellite:	☐
Server Max Segment Size:	1460 (bytes) (512-1460)
Server Receive Buffer Size:	32 (KB) (1-512)

Figure 10-2 TCP Settings Window—Server and Client Settings

TCP Server Settings		
Enable Server Satellite:	☐	
Server Max Segment Size:	1460	(bytes) (512-1460)
Server Receive Buffer Size:	32	(KB) (1-512)
Server Read/Write Timeout:	120	(seconds) (1-3600)
Server Send Buffer Size:	8	(KB) (8-512)

TCP Client Settings		
Enable Client Satellite:	☐	
Client Max Segment Size:	1432	(bytes) (512-1460)
Client Receive Buffer Size:	28	(KB) (1-512)
Client Read/Write Timeout:	360	(seconds) (1-3600)
Client Send Buffer Size:	32	(KB) (8-512)

TCP Memory Limit Settings		
TCP Memory Limit Low Water Mark:	25	(MB) (4-600)
TCP Memory Limit High Water Mark - Pressure:	28	(MB) (5-610)

Step 4 Under the TCP General Settings section, follow these steps:

- a. To enable reduction of delay and packet loss in data transmissions, check the **Enable Explicit Congestion Notification** check box. It provides TCP support for RFC 2581. (See the [“About Explicit Congestion Notification”](#) section on page 10-7.)
- b. To enable the TCP Type of Service, check the **Enable Type Of Service** check box. It is disabled by default.
- c. In the Congestion Window Size field, specify the initial congestion window value in segments. The default is 2 segments. (See the [“About Congestion Windows”](#) section on page 10-7.)
- d. In the Retransmit Time Multiplier field, specify the factor used to modify the length of the retransmit timer by 1 to 3 times the base value determined by the TCP algorithm. The default is 1, which leaves the times unchanged. (See the [“About the Retransmit Time Multiplier”](#) section on page 10-7.)



Note

Modify this factor with caution. It can improve throughput when TCP is used over slow reliable connections but should never be changed in an unreliable packet delivery environment.

- e. In the Initial Slow Start Threshold field, specify the threshold for slow start in segments. The default is 2 segments. (See the [“About TCP Slow Start”](#) section on page 10-8.)
- f. In the Keepalive Probe Count field, specify the number of times the Content Engine can retry a connection before the connection is declared unsuccessful. The default is 4 attempts.

- g. In the Keepalive Probe Interval field, specify the length of time that the Content Engine keeps an idle connection open. The default is 300 seconds.
- h. In the Keepalive Timeout field, specify the length of time that the Content Engine keeps a connection open before disconnecting. The default is 90 seconds.

Step 5 Under the TCP Client Settings section, follow these steps:

- a. To set client TCP compliance to the RFC 1323 standard, check the **Enable Client Satellite** check box. (See the [“About TCP-Over-Satellite Extensions”](#) section on page 10-8.)



Note If you enable a client satellite, the server satellite is automatically enabled, and if you enable the server satellite, the client satellite is automatically enabled.

- b. In the Client Max Segment Size field, specify the maximum packet size sent to clients. The default is 1432 bytes.
- c. In the Client Receive Buffer Size field, specify the TCP receiving buffer size (in kilobytes) for incoming TCP packets. The default is 32 KB.
- d. In the Client Read/Write Timeout field, specify the period after which the Content Engine times out trying to read or write to the network. The default is 120 seconds.
- e. In the Client Send Buffer Size field, specify the TCP sending buffer size (in kilobytes) for outgoing TCP packets. The default is 32 KB.

Step 6 Under the TCP Server Settings section, follow these steps:

- a. To set server TCP server compliance to the RFC 1323 standard, check the **Enable Server Satellite** check box. (See the [“About TCP-Over-Satellite Extensions”](#) section on page 10-8.)



Note If you enable a server satellite, the client satellite is automatically enabled, and the reverse.

- b. In the Server Max Segment Size field, specify the maximum packet size sent to the server. The default is 1460 bytes.
- c. In the Server Receive Buffer Size field, specify the TCP receiving buffer size (in kilobytes) for incoming TCP packets. The default is 32 KB.
- d. In the Server Read/Write Timeout field, specify the period after which the Content Engine times out trying to read or write to the network. The default is 120 seconds.
- e. In the Server Send Buffer Size field, specify the TCP sending buffers size (in kilobytes) for outgoing TCP packets. The default is 32 KB.

Step 7 To save the settings, click **Submit**. A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking **Reset**. The **Reset** button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

If you try to leave this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box only appears if you are using the Internet Explorer browser.

Table 10-1 TCP Settings

GUI Parameter	Function	CLI Command
TCP General Settings		
Enable Explicit Congestion Notification	Enables reduction of delay and packet loss.	tcp ecn enable
Enable Type Of Service	Enables Type of Service.	tcp type-of-service enable
Congestion Window Size	Congestion window size in segments.	tcp cwnd-base <i>segments</i>
ReTransmit Time Multiplier	Factor used to modify the length of the retransmit timer.	tcp increase-xmit-timer-value <i>value</i>
Initial Slow Start Threshold	Threshold for slow start in segments.	tcp init-ss-threshold <i>value</i>
Keepalive Probe Count	Number of times the Content Engine can retry a connection before it is considered unsuccessful.	tcp keepalive-probe-cnt <i>count</i>
Keepalive Probe Interval	Length of time that the Content Engine keeps an idle connection open.	tcp keepalive-probe-interval <i>seconds</i>
Keepalive Timeout	Length of time that the Content Engine keeps a connection open before disconnecting.	tcp keepalive-timeout <i>seconds</i>
TCP Client Settings		
Enable Client Satellite	Sets client TCP compliance to the RFC 1323 standard.	tcp client-satellite
Client Max Segment Size	Maximum packet size sent to clients.	tcp client-mss <i>max_seg_size</i>
Client Receive Buffer Size	TCP receiving buffer size in kilobytes (1-512) for incoming TCP packets.	tcp client-receive-buffer <i>kbytes</i>
Client Read/Write Timeout	Period after which the Content Engine times out when attempting to read or write to the network.	tcp client-rw-timeout <i>seconds</i>
Client Send Buffer Size	TCP sending buffers size in kilobytes (1-512) for outgoing TCP packets.	tcp client-send-buffer <i>kbytes</i>
TCP Server Settings		
Enable Server Satellite	Sets TCP server compliance to the RFC 1323 standard.	tcp server-satellite
Server Max Segment Size	Maximum packet size sent to the server.	tcp server-mss <i>max_seg_size</i>
Server Receive Buffer Size	TCP receiving buffer size in kilobytes (1-512) for incoming TCP packets.	tcp server-receive-buffer <i>kbytes</i>

Table 10-1 TCP Settings (continued)

GUI Parameter	Function	CLI Command
Server Read/Write Timeout	Period after which the Content Engine times out when attempting to read or write to the network.	tcp server-rw-timeout <i>seconds</i>
Server Send Buffer	TCP sending buffer size in kilobytes (1-512) for outgoing TCP packets.	tcp server-send-buffer <i>kbytes</i>

Configuring TCP Memory Limits

The TCP memory limit settings allow you to control the amount of memory that can be used by the TCP subsystem send and receive buffers.



Caution

Do not modify the default values unless you know what you are doing. The default values are device dependent and have been chosen after extensive testing. They should not be changed under normal conditions. Increasing these values can result in the TCP subsystem using more memory, which might render the system unresponsive. Decreasing these values can result in increased response times and lower performance.

To configure the TCP memory limit settings from the Content Distribution Manager GUI, choose **Devices > Devices (or Device Groups) > General Settings > Network > TCP**. The TCP Settings window appears. (See [Figure 10-2](#).)

[Table 10-2](#) lists the TCP memory limit settings in the Content Distribution Manager GUI and the corresponding CLI commands.

Table 10-2 TCP Memory Limit Settings

Content Distribution Manager GUI Parameter	Function	CLI Command
TCP Limit Low Water Mark	The lower limit (in MBytes) of memory pressure mode, below which TCP enters into normal memory allocation mode. The range is 4–600.	tcp memory-limit low-water-mark <i>megabytes</i> high-water-mark-pressure <i>megabytes</i>
TCP Memory Limit High Water Mark–Pressure	The upper limit (in megabytes) of normal memory allocation mode, beyond which TCP enters into memory pressure mode. The range is 5–610.	high-water-mark-absolute <i>megabytes</i>
TCP Memory Limit High Water Mark–Absolute	The absolute limit (in MBytes) on TCP memory usage. The range is 6–620.	

Table 10-3 describes the default values for each command parameter, which are based on the total amount of memory for the device.

Table 10-3 *Default TCP Memory Limit Settings*

Total System Memory	Low	Pressure	Absolute
1 GB, 2 GB, or 4 GB	360 MB	380 MB	400 MB
512 MB	180 MB	190 MB	200 MB
256 MB	25 MB	28 MB	30 MB

The following conditions must be satisfied whenever these default values are changed:

- The low water mark must be a number that is less than the high water mark pressure setting.
- The high water mark pressure must be a number that is less than the high water mark absolute setting:

```
low-water-mark < high-water-mark-pressure < high-water-mark-absolute
```

About Explicit Congestion Notification

The TCP Explicit Congestion Notification (ECN) feature allows an intermediate router to notify the end hosts of impending network congestion. It also provides enhanced support for TCP sessions associated with applications that are sensitive to delay or packet loss, including Telnet, web browsing, and transfer of audio and video data. The major issue with ECN is the need to change the operation of both the routers and the TCP software stacks to accommodate the operation of ECN.

About Congestion Windows

The congestion window (*cwnd*) is a TCP state variable that limits the amount of data that a TCP sender can transmit onto the network before receiving an acknowledgment (ACK) from the receiving side of the TCP transmission. The TCP *cwnd* variable is implemented by the TCP congestion avoidance algorithm. The goal of the congestion avoidance algorithm is to continually modify the sending rate so that the sender automatically senses any increase or decrease in available network capacity during the entire data flow. When congestion occurs (manifested as packet loss), the sending rate is first lowered and then gradually increased as the sender continues to probe the network for additional capacity.

About the Retransmit Time Multiplier

The TCP sender uses a timer to measure the time that has elapsed between sending a data segment and receiving the corresponding ACK from the receiving side of the TCP transmission. When this retransmit timer expires, the sender (according to the RFC standards for TCP congestion control) must reduce its sending rate. However, because the sender is not reducing its sending rate in response to network congestion, the sender is not able to make any valid assumptions about the current state of the network. Therefore, in order to avoid congesting the network with an inappropriately large burst of data, the sender implements the slow start algorithm, which reduces the sending rate to one segment per transmission. (See the next section, “[About TCP Slow Start](#).”)

You can modify the sender's retransmit timer by using the Retransmit Time Multiplier field in the Content Distribution Manager GUI or the **tcp increase-xmit-timer-value** global configuration command in the CLI. The retransmit time multiplier modifies the length of the retransmit timer by one to three times the base value, as determined by the TCP algorithm that is being used for congestion control.

When making adjustments to the retransmit timer, be aware that they affect performance and efficiency. If the retransmit timer is triggered too early, the sender pushes duplicate data onto the network unnecessarily; if the timer is triggered too slowly, the sender remains idle for too long, unnecessarily slowing data flow.

About TCP Slow Start

Slow start is one of four congestion control algorithms used by TCP. The slow start algorithm controls the amount of data being inserted into the network at the beginning of a TCP session, when the capacity of the network is not known.

For example, if a TCP session began by inserting a large amount of data into the network, much of the initial burst of data would probably be lost. Instead, TCP initially transmits a modest amount of data that has a high probability of successful transmission. TCP then probes the network by sending increasing amounts of data as long as the network does not show signs of congestion.

The slow start algorithm begins by sending packets at a rate that is determined by the congestion window, or *cwnd* variable. (See [“About Congestion Windows.”](#)) The algorithm continues to increase the sending rate until it reaches the limit set by the slow start threshold (*ssthresh*) variable. (Initially, the value of the *ssthresh* variable is adjusted to the receiver's maximum window size [RMSS]. However, when congestion occurs, the *ssthresh* variable is set to half the current value of the *cwnd* variable, marking the point of the onset of network congestion for future reference.)

The starting value of the *cwnd* variable is set to that of the sender maximum segment size (SMSS), which is the size of the largest segment that the sender can transmit. The sender sends a single data segment, and because the congestion window is equal to the size of one segment, the congestion window is now full. The sender then waits for the corresponding ACK from the receiving side of the transmission. When the ACK is received, the sender increases its congestion window size by increasing the value of the *cwnd* variable by the value of one SMSS. Now the sender can transmit two segments before the congestion window is again full and the sender is once more required to wait for the corresponding ACKs for these segments. The slow start algorithm continues to increase the value of the *cwnd* variable and therefore increase the size of the congestion window by one SMSS for every ACK received. If the value of the *cwnd* variable increases beyond the value of the *ssthresh* variable, then the TCP flow control algorithm changes from the slow start algorithm to the congestion avoidance algorithm.

About TCP-Over-Satellite Extensions

The Content Engine has the ability to turn on TCP-over-satellite extensions (as documented in RFC 1323) to maximize performance and end-to-end throughput over satellite-type connections.

The large number of satellites available to network infrastructures has increased the amount of bandwidth available in the air. Taking advantage of these connections through satellite-type connections has created new challenges in the use of TCP transactions and acknowledgments:

- Latency—Round trip times to satellites orbiting 24,000 miles above the earth are 550 milliseconds for a single satellite hop. Buffer size must be set to prevent low-throughput connections.
- Bit errors—Packet loss can occur in a land-based device-to-satellite connection in addition to the losses caused by regular network congestion.

- Asymmetric bandwidth—Return bandwidth from satellites can be narrower than receiving bandwidth, thereby affecting performance.

Use the fields provided under the TCP Server Settings and TCP Client Settings sections to set the TCP connection so that it complies with RFC 1323.

Configuring IP Differentiated Services

In the ACNS network, you can configure Content Engines, Content Routers, and Content Distribution Managers for Type of Service (ToS) or differentiated services code point (DSCP) either through the CLI or through the Content Distribution Manager GUI.

The differentiated services architecture is based on a simple model where traffic entering a network is classified and possibly conditioned at the boundaries of the network. The class of traffic is then identified with a differentiated services (DS) code point or bit marking in the IP header. Within the core of the network, packets are forwarded according to the per-hop behavior associated with the DS code point.

To set the global ToS or DSCP values for the IP header from the CLI, use the **ip dscp** global configuration command:

```
ip dscp {client {cache-hit {match-server | set-dscp dscp-packets | set-tos tos-packets} | cache-miss {match-server | set-dscp dscp-packets | set-tos tos-packets}} | server {match-client | set-dscp dscp-packets | set-tos tos-packets}}
```

To configure TOS and DSCP values using the Content Distribution Manager GUI, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices** (or **Devices > Device Groups**).
 - Step 2** Click the **Edit** icon next to the name of the device or device group that you want to configure.
 - Step 3** In the Contents pane, choose **General Settings > Network > IP General Settings**. The IP General Settings window appears.
 - Step 4** Under the IP General Settings heading, you can choose to enable the MTU discovery utility by checking the **Enable Path MTU Discovery** check box. (For more information, see the next section, “[Enabling the MTU Discovery Utility](#).”)
 - Step 5** To configure DSCP or TOS responses to the client for a cache hit, check the **Set IP DSCP Client Cache-Hit** check box to enable the options in the GUI, and choose one of the following methods for marking the IP packets:
 - **Match Server**
 - **DSCP**
 - **TOS**
 - Step 6** To configure DSCP or TOS responses to the client for a cache miss, check the **Set IP DSCP Client Cache-Miss** check box, and choose a method for marking the IP packets.
 - Step 7** To configure DSCP or TOS for outgoing requests to the server, check the **Set IP DSCP Server** check box, and choose a method for marking the IP packets.
 - Step 8** To save the settings, click **Submit**.
-

Enabling the MTU Discovery Utility

Cisco ACNS software supports the IP Path Maximum Transmission Unit (MTU) Discovery mechanism, as defined in RFC 1191. When enabled, the Path MTU Discovery utility discovers the largest IP packet size allowable between the various links along the forwarding path and automatically sets the correct value for the packet size. By using the largest MTU that the links will bear, the sending device can minimize the number of packets it must send.

IP Path MTU Discovery is useful when a link in a network goes down, forcing the use of another, different MTU-sized link. IP Path MTU Discovery is also useful when a connection is first being established and the sender has no information at all about the intervening links.



Note

IP Path MTU Discovery is a process initiated by the sending device. If a server does not support IP Path MTU Discovery, the receiving device will have no mechanism available to avoid fragmenting datagrams generated by the server.

To enable this autodiscovery utility for the Content Engine or Content Router from the CLI, use the **ip path-mtu-discovery enable** global configuration command. To enable it using the Content Distribution Manager GUI, check the **Enable Path MTU Discovery** check box in the IP General Settings window for the device (**Devices > Devices > General Settings > Network > IP General Settings**). By default, this feature is disabled. With the feature disabled, the sending device uses a packet size that is the lesser of 576 bytes and the next hop MTU. Existing connections are not affected when this feature is turned on or off.



Viewing Content Replication Status

The replication status feature allows you to view the status of content replication using the Content Distribution Manager GUI, output from CLI commands, or an API file. This chapter explains how to monitor content replication status from the Content Distribution Manager GUI.

In the Content Distribution Manager GUI, replication status is provided under the Content tab and the Devices tab. You can view the replication status by channel or by device. The Content Distribution Manager GUI also allows you to obtain a detailed view for a specific Content Engine or channel. You can also view detailed replication status for a particular content item in a channel.

This chapter contains the following sections describing views that you can obtain under the Content tab:

- [Viewing the System-Wide Replication Status by Channel, page 11-1](#)
- [Viewing the Replication Status for a Specific Channel, page 11-5](#)
- [Viewing the Detailed Replication Status of Content Items for a Specific Channel, page 11-9](#)
- [Viewing the Detailed Replication Status of Content Items for a Channel by Device, page 11-10](#)
- [Viewing the Detailed Replication Status of Content Items Across All Content Engines in a Channel, page 11-12](#)

This chapter contains the following sections describing views that you can obtain under the Devices tab:

- [Viewing the System-Wide Replication Status by Content Engine, page 11-13](#)
- [Viewing the Replication Status for a Specific Content Engine, page 11-14](#)
- [Viewing the Detailed Replication Status of Content Items for a Selected Channel by Content Engine, page 11-15](#)



Note

Individual Content Engines report their known status, but the Content Distribution Manager has a global view of the channel and may modify the state that the Content Engine reports. If the root Content Engine is processing an updated manifest file or has found new content of which a receiver Content Engine is unaware, the receiver Content Engine reports the replication status as “Completed,” yet the Content Distribution Manager reports the status as “Update Pending.” If the root Content Engine fails to update files, the status of all the receiver Content Engines is reported as “In Process.”

Viewing the System-Wide Replication Status by Channel

The Channels window lists all channels on the system and displays the replication status information for each channel. This display summarizes the replication status of all Content Engines associated with a specific channel in a given state.

To view system-wide replication status for each channel, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels** to display the Channels window. (See [Figure 11-1](#).)

Figure 11-1 Channels Window—System-Wide Replication Status by Channel



- Step 2** View the replication status information for each channel. [Table 11-1](#) describes the status information that is displayed in this window.

Table 11-1 System-Wide Replication Status by Channel—Channels Window

Column Heading	Description
Channel	Name of the channel.
Type	Type of channel. The channel types are Live and Content.
Website	Name of the website assigned to the channel.

Table 11-1 *System-Wide Replication Status by Channel—Channels Window (continued)*

Column Heading	Description
Status	Graphical display indicating acquisition, replication, and device errors. Status lights represent the highest level of errors encountered: Green—No errors encountered. Yellow—Only minor errors encountered. Red—At least one critical error encountered, such as an acquisition failure, a content replication failure, or a failed or nonresponsive Content Engine. For details of the errors, click the status light for a particular channel, which takes you to the Replication Status for Channel window. (See Table 11-2 for a description of status errors and their corresponding status lights.)
State	State of the channel. States are reported for the root Content Engine and for receiver Content Engines. (See Table 11-3 for a definition of the different channel states.) The state is also a link to the Replication Status for Channel window that provides a more detailed view of the replication status for the channel. (See Figure 11-3.)
Manifest State	State of the manifest file. States reported are as follows: Fetching—The manifest file is being fetched. Fail Fetching—The manifest file has failed to be fetched. Parsing—The manifest file is being parsed. Fail Parsing—The manifest file has failed to be parsed. Completed—The manifest file was successfully fetched and parsed. No Status Reported—Root Content Engine is in a Pending or Disabled state.

Table 11-2 *Channel Status Errors*

Status Light	Error	Description
Yellow	Manifest retrieval error	The root Content Engine cannot retrieve the manifest file for 1 or 2 consecutive attempts.
Red	Manifest retrieval error	The root Content Engine cannot retrieve the manifest file for 3 consecutive attempts.
Red	Manifest syntax error	The root Content Engine fails to parse the manifest file.
Red	Crawl job processing error	The root Content Engine encounters problems while crawling for content.
Red	Acquisition or content replication error	The Content Engine fails to obtain the content.
Red	Disk quota exceeded error	The Content Engine cannot store or process the content because there is no more disk space available.

Table 11-2 *Channel Status Errors (continued)*

Status Light	Error	Description
Yellow	Replication status update error	Content replication failed for 1 or 2 consecutive attempts.
Red	Replication status update error	Content replication failed for 3 or more consecutive attempts.
Red	Content Engine unreachable error	The Content Engine is offline or the Content Engine has not responded to replication status requests for 3 consecutive polling periods.
Red	Root Content Engine failover	The root Content Engine has failed over to a temporary root Content Engine. Receiver Content Engines have not identified a valid root Content Engine.
Red	Receiver Content Engine device or channel error	Receiver Content Engine is not reporting replication status or any other content replication problem.

Table 11-3 *Channel States in Replication Status for Channel Window*


State	Description
Completed	All receiver Content Engines are in the Completed state, and the root Content Engine is either in the Completed, Re-checking, Retrieving Manifest, or Processing Manifest state. (See Table 11-6 for a description of Content Engine states.) When the root Content Engine in the Re-checking state determines that new content needs to be acquired, the channel state changes to In Process.
In Process	In Process can mean: - The root Content Engine is in the Retrieving Manifest, Processing Manifest, Acquiring Content, or Re-checking Content state. - Any receiver Content Engine is in the Pending Update, Replicating, or Recovering from Failure state. - The root Content Engine has failed and receiver Content Engines are still reporting status.
Failed	Failed can mean: - An acquisition or content replication error has occurred. (See Table 11-2.) - A Content Engine has gone offline or has not reported status in 3 consecutive polling periods. - The channel has more than 1 root Content Engine - The channel has no root Content Engine, but has receiver Content Engines reporting replication status.

Viewing the Replication Status for a Specific Channel

To view the replication status for a particular channel, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Services > Web > Channels**. When you roll your mouse over an alarm light in the Status column for a particular channel, the content replication status for that channel is displayed. (See [Figure 11-2](#).)

Figure 11-2 Channels Window



- Step 2** Click an alarm light in the Status column or a link in the State column. The Replication Status for Channel window appears. This window displays the replication status for the particular channel chosen. (See [Figure 11-3](#).)


[Table 11-4](#) describes the acquisition status information that is displayed in this window, and [Table 11-5](#) describes the replication status information for individual devices that is displayed in this window.



Note

Alternatively, you can choose **Services > Web > Channels**, click the **Edit** icon next to the channel for which you want to view the replication status, and choose **Replication Status** from the Contents pane.

Figure 11-3 Replication Status for Channel Window



This window also allows you to do the following:

- See a detailed view of replication status using search criteria. (See the “[Viewing the Detailed Replication Status of Content Items for a Specific Channel](#)” procedure on page 11-9.)
- Query the replication status of content items (by pattern) for a selected Content Engine in the channel. (See the “[Viewing the Detailed Replication Status of Content Items for a Channel by Device](#)” procedure on page 11-10.)

Table 11-4 Acquisition Status for a Channel

Field	Description
User Selected Root CE	Name of the user-selected root Content Engine.
Current Root CE	Name of the current root Content Engine. The current root Content Engine will be the same as the user-selected root Content Engine as long as the user-selected root is active; if it fails for any reason, the temporary root Content Engine becomes the current root Content Engine.
Disk Quota Used	Amount of available disk space used for the channel.
Status	State of the root Content Engine. (For a description of root Content Engine states, see Table 11-6 .)

Table 11-4 Acquisition Status for a Channel (continued)

Field	Description
Manifest Last Modified Time	Time when the manifest file was last saved, as recorded on the Content Engine.
Manifest Last Checked Time	Time when the root Content Engine last checked the manifest file for changes.

Table 11-5 Replication Status for Devices Assigned to a Channel

Field	Description
Device	Name of the Content Engine assigned to the channel.
Type	Type of Content Engine: root, receiver, or temporary root.
Status	Graphical display indicating acquisition, replication, and device errors. Status lights represent the highest level of errors encountered: Green—No errors encountered. Yellow—Only minor errors encountered. Red—At least one critical error encountered, such as an acquisition failure, a content replication failure, or a failed or nonresponsive Content Engine. For details of the errors, click the status light. A dialog box pops up, listing the errors. (See Table 11-2 for a description of status errors and their corresponding status lights.)
State	State of either the root Content Engine or receiver Content Engines. (See Table 11-6 for a description of Content Engine states.)
Last Report Time	Time when the last report from the Content Engine was received by the Content Distribution Manager. This time stamp uses the Content Distribution Manager clock.
File Count	
Completed	Number of files that the Content Engine has successfully acquired or received.
In Process	Number of new files to be acquired or replicated. Includes only files for which no acquisition or replication attempts have previously been made.
Failed	For the root Content Engine: Number of files that failed to be acquired in at least 1 attempt. For receiver Content Engines: Number of files that failed to be replicated in at least 1 attempt. Note The failure count for the receiver Content Engine has no relationship to the failure count for the root Content Engine. If the root Content Engine fails to replicate an item, the receiver counts this item as “In Process.”
Total	Total number of Completed, In Process, and Failed files.

Table 11-6 **Content Engine States**

State	Description
Root Content Engine	
Retrieving Manifest	The root Content Engine is retrieving the manifest file from the origin server or rechecking the manifest file for changes.
Processing Manifest	The root Content Engine has retrieved the manifest file and is parsing it.
Acquiring Content	The root Content Engine has processed the manifest file and is crawling or fetching content.
Re-checking Content	The root Content Engine is checking the content or crawl job freshness.
No Status Reported	No Status Reported can mean: • The root Content Engine is unreachable for 3 consecutive polling periods. • The root Content Engine is offline. • The Content Distribution Manager has recently restarted and has not yet received a report from the root Content Engine.
Completed	The root Content Engine is not in the Retrieving, Processing, Acquiring Re-checking, or Unreachable state.
Receiver Content Engine	
Pending Update from Root	The receiver Content Engine is not synchronized with the root Content Engine.
Replicating	The receiver Content Engine is synchronized with the root Content Engine and is replicating content.
Completed	The receiver Content Engine has finished replicating all the content with no errors.
Recovering from Failure	The receiver Content Engine has not identified the root Content Engine. This state occurs during a failover from the root Content Engine to a temporary root Content Engine.
No Status Reported	No Status Reported can mean: • The receiver Content Engine is unreachable for 3 consecutive polling periods. • The receiver Content Engine is offline. • The Content Distribution Manager has recently restarted and has not yet received a report from the receiver Content Engine.

Viewing the Detailed Replication Status of Content Items for a Specific Channel

To view the detailed replication status of content items for a channel, follow these steps:

- Step 1** Choose **Services > Web > Channels**. The Channels window appears.
- Step 2** Click the **Edit** icon next to the channel for which you want to view the replication status.
- Step 3** In the Contents pane, choose **Replication Status**. The Replication Status for Channel window appears.
- Step 4** In the View Detailed Replication Status section (see [Figure 11-3](#)), enter a string in the Get Detailed Status using Search Criteria field. Use an asterisk (*) to match one or more characters, or a question mark (?) to match only a single character.

The criteria are matched against the relative *cdn-url* attribute specified in the <item> tag in the manifest file. It is recommended that you start the search criteria by specifying wildcards such as *.htm or *.clip.mpeg.
- Step 5** To retrieve the specified items, click the **Go** button. The Replication Items window appears, displaying replication status for content items for a specific channel. (See [Figure 11-4](#).) [Table 11-7](#) describes the information that is displayed in the window.

Figure 11-4 Replication Status for Content Items for a Specific Channel




Table 11-7 Replication Status of Items for a Channel

Column Heading	Description
Url	URL of the origin web server that stores the website content.
Size	Size of the file to be acquired or crawled.
Status	Status of replication of content in the channel. The status will be shown as Complete if the replication is complete on all Content Engines assigned to the channel.
Replied CEs	Number of Content Engines that have replicated this item.

Table 11-7 *Replication Status of Items for a Channel*

Column Heading	Description
Playtime	Duration of playback of the file.
Modification Time	Time stamp of the earliest update for that channel from an active Content Engine.

Viewing the Detailed Replication Status of Content Items for a Channel by Device

Queries to determine the detailed replication status of a content item trigger extensive CPU cycles and high consumption of memory, because all the Content Engines assigned to a channel need to be polled, and the retrieved replication status is cached in the memory of the Content Distribution Manager. This results in performance degradation. To optimize the use of memory resources without compromising the need to obtain detailed replication status of a particular content item, you can choose a Content Engine assigned to a channel and generate a query.

To view the detailed replication status for a channel by device, follow these steps:

- Step 1** Choose **Services > Web > Channels**. The Channels window appears.
- Step 2** Click the **Edit** icon next to the channel that you want to view.
- Step 3** In the Contents pane, choose **Replication Status**. The Replication Status for Channel window appears.
- Step 4** In the Devices Assigned to Channel section (see [Figure 11-3](#)), click the radio button next to the name of the device that you want to view.
- Step 5** In the View Detailed Replication Status for Channel by Device section, do the following:
 - a. Choose content items (**all**, **replicated**, or **nonreplicated**) from the Get drop-down list.
 - b. In the Search Criteria for Selected Device field, enter a string that specifies the type of content items that you want displayed.
 - c. To retrieve the specified items, click the **Go** button.

For example, enter Get **replicated** content items using Search Criteria ***.wmv** for selected Device.



Note Use an asterisk (*) to match one or more characters, or a question mark (?) to match only a single character.

The Replication Items window for the selected device appears. (See [Figure 11-5](#).) [Table 11-8](#) describes the fields displayed in this window.

Figure 11-5 Replication Items for a Selected Device




Table 11-8 Replication Status of Items for a Channel by Device

Column Heading	Description
Url	URL of the origin web server that stores the website content.
Size	Size of the file to be acquired or crawled.
Status	Status of replication of content for the selected Content Engine.
Playtime	Duration of playback of the file.
Modification Time	Time stamp in Greenwich mean time (GMT) of the latest update to the content item as recorded on the origin server.

**Note**

When you click the **Force replication information refresh** icon in the taskbar of the Replication Items window for the selected device, the system displays a dialog box asking you to confirm whether you want to refetch the information from Content Engines assigned to this channel. To continue with the refresh process, click **OK**. You are notified that the request has been queued and asked to check back later.

Step 6 To refine your search from this window, do the following:

- Make a choice from the Get drop-down list.
- Enter a search string in the Search Criteria field.
- To retrieve the specified items, click the **Go** button.

Step 7 To return to the Replication Status for Channel window, click the straight blue **Back** arrow in the Content Distribution Manager GUI taskbar (not the **Back** button for your browser).

Viewing the Detailed Replication Status of Content Items Across All Content Engines in a Channel

You can choose the URL of the origin server and request a detailed replication status of content items across all Content Engines to trigger an aggregated replication status request. The detailed item replication status window shows a time stamp indicating the time in Greenwich mean time (GMT) at which the replication status was last cached to provide information about the freshness of the detailed content replication status.


To view the detailed replication status of an item across all Content Engines in a channel, follow these steps:

- Step 1** Choose **Services > Web > Channels**. The Channels window appears.
- Step 2** Click the **Edit** icon next to the channel for which you want to view the replication status.
- Step 3** In the Contents pane, choose **Replication Status**. The Replication Status for Channel window appears.
- Step 4** In the View Detailed Replication Status section (see [Figure 11-3](#)), enter a string in the Get Detailed Status using Search Criteria field. Use an asterisk (*) to match one or more characters, or a question mark (?) to match only a single character.

The criteria are matched against the relative *cdn-url* attribute specified in the <item> tag in the manifest file. We recommend that you start the search criteria by specifying wildcards such as *.htm or *.clip.mpeg.

- Step 5** To retrieve the specified items, click the **Go** button. The Replication Items window appears, displaying replication status for content items for a specific channel. (See [Figure 11-4](#).)
- Step 6** Click the **View** icon next to the URL that you want to view. The Replication Item window appears, displaying the replication status of the content item for each Content Engine in the channel. (See [Figure 11-6](#).) [Table 11-9](#) describes the fields in this window.

Figure 11-6 Replication Status of the Content Item for All Content Engines in the Channel



The screenshot shows the Cisco Application and Content Networking System interface. The 'Replication Item' window is open, displaying a table of replication items for the channel 'int54-camiant'. The table has columns for CE, Size, Status, Playtime, and Modification Time. The data shows three items, all with a status of 'Complete' and a size of 816,947 KB.

CE	Size	Status	Playtime	Modification Time
ce-13	816,947 KB	Complete	00:00:21	16:34:51 12-19-2002
ce-40	816,947 KB	Complete	00:00:21	16:34:51 12-19-2002
ce-39	816,947 KB	Complete	00:00:21	16:34:51 12-19-2002



Note

The Replication Item window is specifically designed to limit listings to 5000 objects for scalability reasons. These are system limits and not specifically enforced for replication status reporting.

Table 11-9 Replication Status of an Item for All Content Engines in a Channel


Column Heading	Description
CE	Name of the Content Engine to which the item has been replicated.
Size	Size of the file to be acquired or crawled.
Status	Status of replication of content in the channel. Status is shown as Complete if replication is complete on all Content Engines assigned to the channel.
Playtime	Duration of playback of the file.
Modification Time	Time stamp of the latest update for the content item as recorded on the origin web server.

- Step 7** To return to the Replication Status for Channel window, click the straight blue **Back** arrow in the Content Distribution Manager GUI taskbar (not the **Back** button for your browser).

Viewing the System-Wide Replication Status by Content Engine

To view the system-wide replication status by device, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Statistics > Replication Status**. The Device Replication Status window appears.
- This window summarizes the replication status of all channels associated with a specific Content Engine in a given state. (See [Figure 11-7](#).)

Figure 11-7 Device Replication Status Window

- Step 2** View the replication status information for each device. [Table 11-10](#) describes the status information that is displayed in this window.

Table 11-10 **Device Replication Status Window**

Column Heading	Description
Device	Name of the Content Engine.
Status	Graphical display indicating acquisition, replication, and device errors. Status lights represent the highest level of errors encountered: Green—No errors encountered. Yellow—Only minor errors encountered. Red—At least one critical error encountered, such as an acquisition failure, a content replication failure, or a failed or nonresponsive Content Engine. (See Table 11-2 for a description of status errors and their corresponding status lights.)
Channel Count	Number of channels reporting Content Engines in a particular state. (See Table 11-6 for a description of Content Engine states.)
Completed	Number of channels reporting this Content Engine in a Completed state.
In Process	In Process can mean: - Number of channels reporting this Content Engine (as a root Content Engine) in the Retrieving Manifest, Processing Manifest, Acquiring Content, or Re-checking Content state. - Number of channels reporting this Content Engine (as a receiver Content Engine) in the Pending Update, Replicating, or Recovering from Failure state.
Failed	Number of channels reporting this Content Engine in the Failed or Failed Update state.
Unknown	Number of channels reporting this Content Engine in the No Status Reported state.

Step 3 To view the replication status of individual channels assigned to a particular Content Engine, proceed to the next section, “[Viewing the Replication Status for a Specific Content Engine](#).”


Viewing the Replication Status for a Specific Content Engine

To view the replication status for a particular Content Engine, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Statistics > Replication Status**. The Device Replication Status window appears.
- Step 2** Click the **View** icon next to the Content Engine that you want to view. The Replication Status for Device window appears.
- The Replication Status for Device window displays the replication status of individual channels assigned to a particular Content Engine. You can filter the information by channels or by any of the other categories displayed in this window. (See [Figure 11-8](#).)

Table 11-5 describes the status information that is displayed in this window with the exception that the first column in this window displays the names of the channels to which the Content Engine is assigned rather than the names of the devices in the channel, as described in Table 11-5.

Figure 11-8 Replication Status for Device Window



- Step 3** To view the Content Engine forwarder path for a selected channel, click the **View** icon next to the name of the channel. To return to the Replication Status for Device window, click **Replication Status** in the Contents pane.
- Step 4** To view replication details for the selected channel, see the next section, “[Viewing the Detailed Replication Status of Content Items for a Selected Channel by Content Engine](#).”

Viewing the Detailed Replication Status of Content Items for a Selected Channel by Content Engine

A request for a detailed replication status triggers an aggregated replication status update request. This ensures that the cache status is refreshed on both the Content Engine and the Content Distribution Manager simultaneously.

To view the status of replication of content items for a selected channel by Content Engine, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Statistics > Replication Status**. The Device Replication Status window appears.
- Step 2** Click the **View** icon next to the Content Engine that you want to view. The Replication Status for Device window appears. (See [Figure 11-8](#).)
- Step 3** To view replication details for the selected channel, click the radio button next to a channel name.
- Step 4** From the **Get** drop-down list, choose the type of items to display (**all**, **replicated**, or **non replicated**).
- Step 5** In the Search Criteria for Selected Device field, enter a regular expression (such as ***.html**, ***.mpg**, ***.jpg**, or ***.***). Use an asterisk (*) to match one or more characters, and a question mark (?) to match exactly one character.
- Step 6** To retrieve the specified items, click the **Go** button. The Replication Items - Channel: Content Engine window appears. (See [Figure 11-9](#).) [Table 11-11](#) describes the fields displayed in this window.



Note

The Replication Items - Channel: Content Engine window is specifically designed to limit listings to 5000 objects for scalability reasons. These are system limits and not specifically enforced for replication status reporting.

Figure 11-9 Replication Items - Channel: Content Engine Window

Uri	Size	Status	Playtime	Modification Time
http://2.10.4.15/industrial.vmv	1.209 MB	Complete	00:00:34	15:34:47 12-19-2002
http://2.10.4.15/welcome3.asf	328.291 KB	Complete	00:00:17	19:56:47 09-29-1999
http://2.10.4.15/pinball.vmv	816.947 KB	Complete	00:00:21	16:34:51 12-19-2002
http://2.10.4.15/welcome2.asf	2.042 MB	Complete	00:01:07	22:03:47 07-09-1999
http://2.10.4.15/welcome1.asf	353.671 KB	Complete	00:00:17	19:56:32 09-29-1999

Table 11-11 Replication Status of Items for Content Engines in a Selected Channel

Column Heading	Description
URL	URL of the origin web server that stores the website content.
Size	Size of the file to be acquired or crawled.
Status	Status of replication of content from the root Content Engine.
Playtime	Duration of playback of the file.
Modification Time	Time stamp of the earliest update for that channel from an active Content Engine.

- Step 7** To further qualify your search, change the item type from the drop down list, if you wish, or specify another file type (such as ***.html**, ***.mpg**, or ***.jpg**) in the Search Criteria field. To retrieve the specified items, click the **Go** button.
- Step 8** To forcibly refetch the latest content replication information, click the **Force replication information refresh** icon in the taskbar. You are asked to confirm whether or not you wish to refetch the information from the Content Engine assigned to the particular channel.
- Step 9** To continue with the refresh process, click **OK**. You are notified that your request has been sent and prompted to check back after a few minutes.
- Step 10** To return to the Replication Status for Device window, click the straight blue **Back** button in the taskbar.

Using CLI Commands for Replication Status

You can use the Content Engine CLI to view the content replication status for a channel by using the following commands:

- Use the **show distribution channel channel-id** command to see if there is an unfinished job listed for the channel. The last line in this example shows that there are no unfinished jobs.

```
ContentEngine# show distribution channel channel-id 158
Channel ID: 158
Channel Name: sz300
Website Name: website3
Channel Priority: 500
ID of Configured Root CE: 140
Name of Configured Root CE: sz590b
IP of Configured Root CE: 10.1.1.170
This CE's Role: Not a Root CE
In Full Reload: No
Mcast Receiving: No
Mcast Sending: No
Metadata-Forwarder ID: 140
Metadata-Forwarder Name: sz590b
Metadata-Forwarder IP: 10.1.1.170
Ucast-Forwarder ID: 140
Ucast-Forwarder Name: sz590b
Ucast-Forwarder IP: 10.1.1.170
Last gen-id Switch: Never
ID of Effective Root CE: 140
Current root-ce-uid: 1042665163
Current low-water-marker: 1
Current max-gen-id: 300
Current max-del-gen-id: 0
Has Unfinished Job: No
```



Note The Has Unfinished Job line is only available if the Content Engine is *not* a root Content Engine. It is only available on a receiver Content Engine.

- Use the **show statistics replication channels** command to view the replication status of the channel. The **show statistics replication** commands display the progressive file count status during acquisition and replication.

The following sample output is from the Content Distribution Manager show statistics replication command.

```

CDM# show statistics replication channels selected-channel ws1 chl1c
Channel: chl1c
State: Completed
Status: Green
User Selected Root CE: ce-s4
Current Root CE: ce-s4
Receiver CEs Completed: 1
Receiver CEs In Progress: 0
Receiver CEs Failed: 0
Receiver CEs Not Responding: 0

```

```

Displaying Device Acquisition Replication Status
Device: ce-s4 (CE ID: 92)
Channel: chl1c (Channel ID: 246)
Website: ws1
Type: Root
State: Completed
Status: Green
Completed: 1
To Do: 0
Failed: 0
Total: 1
Last Report Time: Mon Apr 19 20:37:43 GMT 2004
Disk Quota Used: 281.171 KB
Manifest Last Modified: Mon Sep 29 14:59:01 GMT 2003
Manifest Last Check: Wed Apr 14 18:19:12 GMT 2004
Manifest State: Completed

```

```

Displaying Device Distribution Replication Status
Device: CE2 (CE ID: 97)
Channel: chl1c (Channel ID: 246)
Website: ws1
Type: Receiver
State: Completed
Status: Green
Completed: 1
To Do: 0
Failed: 0
Total: 1
Last Report Time: Mon Apr 19 20:32:16 GMT 2004

```

Detailed Error for CE2

The following sample output shows the replication status for the Content Engines of a selected channel.

```

CDM# show statistics replication content-engines selected-channel ws1 chl1c
Displaying Device Acquisition Replication Status
Device: ce-s4 (CE ID: 92)
Channel: chl1c (Channel ID: 246)
Website: ws1
Type: Root
State: Completed
Status: Green
Completed: 1
To Do: 0
Failed: 0
Total: 1
Last Report Time: Mon Apr 19 20:42:38 GMT 2004
Disk Quota Used: 281.171 KB
Manifest Last Modified: Mon Sep 29 14:59:01 GMT 2003
Manifest Last Check: Wed Apr 14 18:19:12 GMT 2004
Manifest State: Completed

```

```

Displaying Device Distribution Replication Status
Device: CE2 (CE ID: 97)
Channel: chlC (Channel ID: 246)
Website: ws1
Type: Receiver
State: Completed
Status: Green
Completed: 1
To Do: 0
Failed: 0
Total: 1
Last Report Time: Mon Apr 19 20:42:12 GMT 2004

```

Detailed Error for CE2

The following sample output is from the Content Engine show statistics replication command.

```

ContentEngine# show statistics replication channels selected-channel ws1 chlC
Displaying Device Acquisition Replication Status
Device: ce-s4 (CE ID: 92)
Channel: 5001Item (Channel ID: 412)
Website: ws1
Type: Root
State: Re-checking Content
Status: Green
Completed: 5001
To Do: 0
Failed: 0
Total: 5001
Last Report Time: Mon May 03 14:09:45 GMT 2004
Disk Quota Used: 1.340 GB
Manifest Last Modified: Wed Dec 18 15:15:49 GMT 2002
Manifest Last Check: Thu Apr 29 12:02:45 GMT 2004
Manifest State: Completed

```

```

Displaying Device Acquisition Replication Status
Device: ce-s4 (CE ID: 92)
Channel: syntaxError (Channel ID: 280)
Website: ws1
Type: Root
State: Processing Manifest
Status: Red
Completed: 0
To Do: 0
Failed: 0
Total: 0
Last Report Time: Mon May 03 14:09:45 GMT 2004
Disk Quota Used: 0 Bytes
Manifest Last Modified: Thu Oct 03 11:00:19 GMT 2002
Manifest Last Check: Mon May 03 14:08:13 GMT 2004
Manifest Error: Syntax Error: Parser Error at (line 5, char 35): Datatype error:
Type:InvalidDatatyp
eValueException, Message:The unary operation node had a binary node type.
Manifest State: Fail Parsing Manifest
7.3.4.2Get selected channel
CE#sh statistics replication channels selected-channel ws1 chlC
Displaying Device Acquisition Replication Status
Device: ce-s4 (CE ID: 92)
Channel: chlC (Channel ID: 246)
Website: ws1
Type: Root
State: Completed
Status: Green
Completed: 1

```



```
To Do: 0
Failed: 0
Total: 1
Last Report Time: Mon May 03 14:08:50 GMT 2004
Disk Quota Used: 281.171 KB
Manifest Last Modified: Mon Sep 29 14:59:01 GMT 2003
Manifest Last Check: Wed Apr 14 18:19:12 GMT 2004
Manifest State: Completed
```



PART 4

Managing ACNS Network Devices





Configuring Login Authentication, Configuration Authorization, and Accounting

This chapter explains how to configure login (user) authentication and configuration authorization support for Content Distribution Managers, Content Engines, and Content Routers deployed in a centrally managed environment. It describes how to configure devices to use local authentication, RADIUS, or TACACS+ login authentication methods.

This chapter contains the following sections:

- [Understanding Authentication, Authorization, and Accounting, page 12-1](#)
- [Configuring Devices for Login Authentication and Authorization, page 12-7](#)
- [Configuring User Accounts for Centrally Managed Devices, page 12-14](#)
- [Managing User Accounts and Privilege Profiles for the Content Distribution Manager, page 12-15](#)
- [Configuring AAA Accounting, page 12-27](#)
- [Viewing Audit Trail Logs, page 12-30](#)



Note

Login (user) authentication and authorization is independent of request authentication and authorization. For information about request authentication and authorization, see [Chapter 15, “Configuring Request Authentication and Authorization.”](#)

Understanding Authentication, Authorization, and Accounting

Authentication determines who the user is and whether that user should be allowed access to the network or a particular device. It allows network administrators to bar intruders from their networks. It may use a simple database of users and passwords. It can also use one-time passwords.

Authorization determines what the user is allowed to do. It allows network managers to limit which network services are available to different users.

Accounting tracks what users did and when they did it. It can be used for an audit trail or for billing for connection time or resources used (bytes transferred).

Collectively, authentication, authorization, and accounting are sometimes referred to as AAA. Central management of AAA means the information is in a single, centralized, secure database, which is much easier to administer than information distributed across numerous devices.

Understanding Login Authentication and Authorization

In the ACNS network, login authentication and authorization are used to control user access and configuration rights to Content Distribution Managers, Content Engines, and Content Routers. Login authentication is the process by which ACNS devices verify whether the person who is attempting to login to the device has a valid username and password. The person logging in must have a user account registered with the device. User account information serves to authorize the user for login and configuration privileges. The user account information is stored in an authentication, authorization, and accounting (AAA) database, and ACNS devices must be configured to access the particular authentication server (or servers) where the AAA database is lodged. When the user attempts to login to a device, the device compares the person's username, password, and privilege level to the user account information that is stored in the database.

Understanding User Accounts and Privilege Profiles

In a centrally managed ACNS network, user accounts can be created for access to the Content Distribution Manager and, independently, for access to the Content Engines and Content Routers that are registered to the Content Distribution Manager.

A privilege profile must be assigned to each new user account. ACNS software uses privilege profiles to determine which tasks a user can perform and the level of access provided.

There are two types of predefined privilege profiles:

- Normal user—User has read access and can see some of the Content Engine, Content Router, or Content Distribution Manager settings.
- Superuser—User has administrative privileges such as creating new users and modifying the Content Engine, Content Router, or Content Distribution Manager settings.

Users with administrative privileges can add, delete, or modify user accounts through the Content Distribution Manager GUI or the device CLI. For example, if an administrator logs in to a Content Engine with the predefined ACNS software superuser account (root administrator), the Content Engine grants that user the highest privilege level (level 15), which allows that user to perform any Content Engine administrative task during that login session.

For instance, that user could perform any of the following administrative tasks:

- Configure the Content Engine.
- Obtain statistical information that the Content Engine has collected.
- Reload the device.

Accessing a Device for the First Time

A Content Distribution Manager, Content Engine, or Content Router that is running ACNS software comes with a predefined superuser account that can be used initially to access the device and add other users.

When the system administrator logs in to a device before authentication and authorization have been configured, the administrator can access the device by using the default user account and password. Devices that are using ACNS software come with a single, predefined superuser user account (root administrator) that can be used for initial access to the device to initially configure the device and then add other user accounts. The username for this predefined superuser user account is *admin* and the default password is *default*. (If these defaults have been changed by another ACNS system administrator, you need to obtain the new username and password.)

Understanding AAA Database Management

User authentication data and user configuration privilege data can be maintained in any combination of these three databases in a centrally-managed ACNS network:

- Content Distribution Manager (local database)
- RADIUS server (external database)
- TACACS+ server (external database)

During Content Distribution Manager setup, you can choose to use an external access server or the internal (local) Content Distribution Manager-based authentication, authorization, and accounting (AAA) system for user access management. You can choose one method or any combination of the three methods.

If all three databases are enabled, then all databases are queried; if the user data cannot be found in the first database queried, then the second and third databases are queried. By default, Content Distribution Managers, Content Engines, and Content Routers use the local database to process login requests, with TACACS+ and RADIUS both disabled for login and configuration.

When the administrator logs in to a Content Distribution Manager, Content Engine, or Content Router, the device checks the specified authentication database to verify the user's username and password and to determine the access rights that this particular user should be granted during this login session.

**Note**

In TACACS+ there is an “enable password” feature that allows an administrator to define a different enable password for each user. If an ACNS user logs in to the Content Engine with a normal user account (privilege level of 0) instead of an admin or admin-equivalent user account (privilege level of 15), the user must enter the admin password to access privileged-level EXEC mode. This caveat applies even if these ACNS users are using TACACS+ for login authentication.

Default Login Authentication and Authorization Configuration

Table 12-1 lists the default configuration for login authentication and authorization.

Table 12-1 Default Configuration for Login Authentication and Authorization

Feature	Default Value
Login authentication	Enabled
Configuration authorization	Enabled
Authentication server failover because the authentication server is unreachable	Disabled
TACACS+ login authentication (console and Telnet)	Disabled
TACACS+ authorization (console and Telnet)	Disabled
TACACS+ key	None specified
TACACS+ server timeout	5 seconds
TACACS+ retransmit attempts	2 times
RADIUS login authentication (console and Telnet)	Disabled
RADIUS authorization (console and Telnet)	Disabled
RADIUS server IP address	None specified

Table 12-1 *Default Configuration for Login Authentication and Authorization (continued)*

Feature	Default Value
RADIUS server UDP authorization port	Port 1645
RADIUS key	None specified
RADIUS server timeout	5 seconds
RADIUS retransmit attempts	2 times

You can change these defaults through the Content Distribution Manager GUI or through the device CLI, as described in the [“Configuring Devices for Login Authentication and Authorization”](#) section on page 12-7.

Understanding Login Authentication Failover

By default, ACNS devices fail over to the secondary method of login authentication whenever the primary login authentication method fails. In ACNS software Release 5.0.5 and later, you can change this default login authentication failover method. For centrally managed deployments, you can use the Content Distribution Manager GUI (for example, choose **Devices > Devices > General Settings > Authentication > Login Authentication** and check the **Enable Failover Server Unreachable** box) or the CLI (use the **authentication fail-over server unreachable** global configuration command) to enable the failover due to unreachable server option.

Failover occurs only if the specified authentication server is unreachable. The following example sets login authentication failover to occur if the authentication server is unreachable. In this case, the Content Engine will query the next authentication method only if the login authentication server is unreachable.

```
ContentEngine(config)#authentication fail-over server-unreachable
ContentEngine(config)#
```

To use the login authentication failover feature, you must set TACACS+, or RADIUS as the primary login authentication method, and local as the secondary login authentication method.

If the failover due to unreachable server option is enabled, then note the following:

- Only two login authentication schemes (a primary and secondary scheme) are allowed on the device.
- The device will fail over from the primary authentication scheme to the secondary authentication scheme only if all specified authentication servers of the primary authentication scheme are unreachable.

For example, if the failover due to server unreachable option is enabled and RADIUS is set as the primary login authentication scheme and local is set as secondary login authentication scheme, the following occurs:

1. When the device receives a login request, it queries all configured RADIUS authentication servers.
2. If one of the configured RADIUS servers is reachable, the device uses this RADIUS database to authenticate the user.
3. If all configured RADIUS servers are not reachable, the device tries the secondary authentication scheme (that is, it queries its local authentication database) to authenticate the user.

**Note**

The local database will be contacted for authentication only if this RADIUS server is not reachable. In any other case (for example, if the authentication fails in the RADIUS server), then the local database is not contacted for authentication.

Conversely, if the failover due to unreachable server option is disabled, then the device contacts the secondary authentication database, regardless of the reason that the authentication failed with the primary authentication database.

If all the authentication databases are enabled for use, then all the databases are queried in the order of priority selected and based on the failover reason. If no failover reason is specified, then all databases are queried in the order of their priority. For example, first the primary authentication database is queried, then the secondary authentication database is queried, and finally the tertiary database is queried.

The local and the remote databases (TACACS+ and RADIUS) can be enabled or disabled through the Content Distribution Manager GUI or through the device CLI. The device verifies whether or not all are disabled and if so, sets the system to the default state (that is, local database is queried for authentication). (For information about this default state, see the [“Default Login Authentication and Authorization Configuration”](#) section on page 12-3.)

For more information on the various types of login authentication and authorization schemes, see the following sections:

- [Understanding Local Login Authentication and Authorization, page 12-5](#)
- [Understanding RADIUS Authentication and Authorization, page 12-6](#)
- [Understanding TACACS+ Authentication and Authorization, page 12-6](#)

For information about how to configure login authentication and authorization on Content Distribution Managers, Content Engines, and Content Routers, see the [“Configuring Devices for Login Authentication and Authorization”](#) section on page 12-7 and the

Understanding Local Login Authentication and Authorization

ACNS 5.x software provides authentication, authorization, and accounting (AAA) support for users who need a local access database with AAA features. Local login authentication can be configured on individual Content Engines using the Content Engine GUI or CLI, or when devices are managed centrally, local login can be configured using the Content Distribution Manager GUI.

When the primary login server and primary configuration server are set to local, usernames and passwords are local to each device and are not mapped to individual usernames. Local authentication and authorization uses locally configured login and passwords to authenticate login attempts.

By default, local login authentication is enabled first. You can disable local login authentication only after enabling one or more of the other login authentication methods. However, when local login authentication is disabled, if you disable all other login authentication methods, local login authentication is reenabled automatically.

Understanding RADIUS Authentication and Authorization

RADIUS is a client/server authentication and authorization access protocol used by a network access server (NAS) to authenticate users attempting to connect to a network device. The NAS functions as a client, passing user information to one or more RADIUS servers. The NAS permits or denies network access to a user based on the response it receives from one or more RADIUS servers. RADIUS uses the User Datagram Protocol (UDP) for transport between the RADIUS client and server.

You can configure a RADIUS key on the client and server. If you configure a key on the client, it must be the same as the one configured on the RADIUS servers. The RADIUS clients and servers use the key to encrypt all RADIUS packets transmitted. If you do not configure a RADIUS key, packets are not encrypted. The key itself is never transmitted over the network.



Note

For more information about how the RADIUS protocol operates, refer to RFC 2138, *Remote Authentication Dial In User Service (RADIUS)*.

RADIUS authentication is disabled by default. You can enable RADIUS authentication and other authentication methods at the same time. You can also specify which method to use first. For more information about configuring RADIUS authentication, see the [“Configuring Devices for Login Authentication and Authorization” section on page 12-7](#).

Understanding TACACS+ Authentication and Authorization

TACACS+ controls access to network devices by exchanging Network Access Server (NAS) information between a network device and a centralized database to determine the identity of a user or an entity. TACACS+ is an enhanced version of TACACS, a User Datagram Protocol (UDP)-based access-control protocol specified by RFC 1492. TACACS+ uses TCP to ensure reliable delivery and encrypt all traffic between the TACACS+ server and the TACACS+ daemon on a network device.

TACACS+ works with many authentication types, including fixed password, one-time password, and challenge-response authentication. TACACS+ login authentication occurs when a user first logs on to the locally deployed Content Engine.

When a user requests restricted services, TACACS+ encrypts the user password information using the MD5 encryption algorithm and adds a TACACS+ packet header. This header information identifies the packet type being sent (for example, an authentication packet), the packet sequence number, the encryption type used, and the total packet length. The TACACS+ protocol then forwards the packet to the TACACS+ server.

A TACACS+ server can provide authentication, authorization, and accounting functions. These services, while all part of TACACS+, are independent of one another, so a given TACACS+ configuration can use any or all of the three services.

When the TACACS+ server receives a packet, it does the following:

- Authenticates the user information and notifies the client that login authentication has either succeeded or failed.
- Notifies the client that authentication will continue and that the client must provide additional information. This challenge-response process can continue through multiple iterations until login authentication either succeeds or fails.

You can configure a TACACS+ key on the client and server. If you configure a key on the Content Engine, it must be the same as the one configured on the TACACS+ servers. The TACACS+ clients and servers use the key to encrypt all TACACS+ packets transmitted. If you do not configure a TACACS+ key, packets are not encrypted.

TACACS+ authentication is disabled by default. You can enable TACACS+ authentication and local authentication at the same time. For more information about configuring TACACS+ as a login authentication scheme, see the next section.

Configuring Devices for Login Authentication and Authorization

To configure login authentication and authorization for a Content Distribution Manager, Content Engine or Content Router, complete the following tasks:

1. Determine the login authentication scheme that you want to configure the device to use when authenticating login requests. (For example, use the local database as the primary login database and your RADIUS server as the secondary authentication database.)
2. Configure the remote authentication server settings on the device (if a remote authentication database is to be used). (See the [“Configuring Remote Authentication Server Settings for the Device”](#) section on page 12-7.)

For example, specify the IP address of the remote RADIUS servers or TACACS+ servers that the device should use to authenticate log-in requests.

3. Specify which authentication databases the device should check to process a login request. (See the [“Configuring Login Authentication and Configuration Authorization Schemes for the Device”](#) section on page 12-12.)
 - Specify the login authentication server failover scheme.
 - Specify the login server scheme.
 - Specify the config server scheme.



Caution

Make sure that RADIUS or TACACS+ authentication is configured and operating correctly before disabling local authentication and authorization. If you disable local authentication and RADIUS or TACACS+ is not configured correctly, or if the RADIUS or TACACS+ server is not online, you may be unable to log in to the device.

When local authentication is disabled, if you disable all other authentication methods, local authentication is reenabled automatically.

Configuring Remote Authentication Server Settings for the Device

If you have determined that your login authentication scheme is to include one or both external authentication servers, you must configure these server settings before you can configure the authentication scheme in the Content Distribution Manager GUI. The following sections provide steps for configuring RADIUS and TACACS+ server settings using the Content Distribution Manager GUI.

Configuring RADIUS Server Settings

RADIUS authentication clients reside on devices running ACNS 5.x software. When enabled, these clients send authentication requests to a central RADIUS server, which contains user authentication and network service access information.

**Tip**

The Content Distribution Manager does not cache user authentication information. Therefore, the user is reauthenticated against the RADIUS server for every request. To prevent performance degradation caused by many authentication requests, install the Content Distribution Manager in the same location as the RADIUS server, or as close as possible to it, to ensure that authentication requests can occur as quickly as possible.

To configure RADIUS server settings using the Content Distribution Manager GUI, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the device that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **General Settings > Authentication > RADIUS Server**. The RADIUS Server Settings window appears. (See [Figure 12-1](#).) [Table 12-2](#) describes the fields in this window.

Figure 12-1 RADIUS Server Settings Window

The screenshot shows the 'RADIUS Server Settings' window for the 'CONTENTENGINE' device. The 'Contents' pane on the left is expanded to 'Authentication > RADIUS Server'. The main area contains the following fields:

- Enable RADIUS Servers*:** ☒
- Time to wait*:** (seconds)
- Number of retransmits*:**
- Enable Redirect:** ☐
- Redirect Message 1*:**
- Location 1*:**
- Redirect Message 2*:**
- Location 2*:**
- Redirect Message 3*:**
- Location 3*:**
- Shared Encryption Key*:**
- Server 1 Name*:**
- Server 1 Port*:**
- Server 2 Name:**
- Server 2 Port:**
- Server 3 Name:**
- Server 3 Port:**
- Server 4 Name:**
- Server 4 Port:**
- Server 5 Name:**
- Server 5 Port:**

At the bottom, there is a note: '* To use RADIUS for Request Authentication, please go to the Authentication Scheme Settings page.' and buttons for 'Submit' and 'Cancel'.

- Step 4** To enable RADIUS authentication, check the **Enable RADIUS Servers** check box.
- Step 5** In the Time to wait field, specify how long the Content Engine should wait before timing out. The default value is 5 seconds.
- Step 6** In the Number of Retransmits field, specify the number of attempts allowed to connect to a RADIUS server.
- Step 7** To enable RADIUS redirection, check the **Enable Redirect** check box.

- Step 8** Enter a redirect message for the user in the Redirect Message field. Three redirect messages are allowed.
- Step 9** In the Location field, enter a location where the redirect message should be sent. Three separate locations are allowed.
- Step 10** In the Shared Encryption Key field, enter the secret key that is used to communicate with the RADIUS server.
- Step 11** Enter an IP address or host name information in the Server Name field. Five different hosts are allowed.
- Step 12** In the Server Port field, enter a UDP port number on which the RADIUS server is listening. You must specify at least one port. Five different ports are allowed.
- Step 13** To save the settings, click **Submit**.

Table 12-2 RADIUS Server Settings

Key Parameter	Description	CLI Command
Enable RADIUS Servers	Enables HTTP authentication using RADIUS servers.	radius-server enable
Time to wait	Number of seconds to wait for a response before timing out on a connection to a particular RADIUS server. The range is from 1 to 20 seconds. The default value is 5 seconds.	radius-server timeout
Number of retransmits	Number of attempts allowed to connect to a RADIUS server. The default value is 2 times.	radius-server retransmit
Enable redirect	Redirects an authentication response to a different authentication server if an authentication request using the RADIUS server fails.	radius-server redirect enable
Redirect Message	Message sent to the user if redirection occurs.	radius-server redirect message
Location	Sets an HTML page location. This is the URL destination of the redirect message that is sent when authentication fails.	radius-server redirect message <i>reply</i> location <i>url</i>
Shared Encryption Key	Encryption key shared with the RADIUS server.	radius-server key <i>keyword</i>
Server Name	IP address or host name of the RADIUS server.	radius-server host <i>hostname</i> or <i>ip-address</i>
Server Port	Port number on which the RADIUS server is listening.	radius-server host <i>auth-port</i> <i>port</i>

Configuring TACACS+ Server Settings

The TACACS+ database validates users before they gain access to a Content Engine. TACACS+ is derived from the United States Department of Defense (RFC 1492) and is used by Cisco Systems as an additional control of nonprivileged and privileged mode access. ACNS 5.x software supports TACACS+ only and not TACACS or Extended TACACS.

**Tip**

The Content Distribution Manager does not cache user authentication information, so the user is reauthenticated against the TACACS+ server for every request. To prevent performance degradation caused by many authentication requests, install the Content Distribution Manager in the same location as the TACACS+ server, or as close as possible to it, to ensure that authentication requests can occur as quickly as possible.

To configure TACACS+ server settings using the Content Distribution Manager GUI, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the device that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **General Settings > Authentication > TACACS+ Server**. The TACACS+ Server Settings window appears. (See [Figure 12-2](#).) [Table 12-3](#) describes the fields in this window.

Figure 12-2 TACACS+ Server Settings Window

The screenshot shows the 'TACACS+ Server Settings' window for the 'CONTENTENGINE' device. The 'Contents' pane on the left shows the navigation path: **General Settings > Authentication > TACACS+ Server**. The main configuration area includes the following fields:

- Enable TACACS+ Servers:** ☒
- Use ASCII Password Authentication:** ☐
- Time to wait:** 0
- Number of retransmits:** 0
- Security Word:** (masked with dots)
- Primary Server:** 172.9.26.12
- Secondary Server:** 10.83.133.65
- Tertiary Server:** (empty)

Below the fields, there are two informational notes:

- * To use TACACS+ for Request Authentication, please go to the Authentication Scheme Settings page.
- * To use TACACS+ for Login or Configuration Authentication, please go to the Login Authentication page.

A note at the bottom states: **Note: * - Required Field**. The window has 'Submit' and 'Cancel' buttons at the bottom right.

- Step 4** To enable TACACS+ authentication, check the **Enable TACACS+ Servers** check box.
- Step 5** To use the ASCII password type for authentication, check the **Use ASCII Password Authentication** check box. The default password type is PAP (Password Authentication Protocol). However, you can change the password type to ASCII when the authentication packets are to be sent in ASCII clear text format.

- Step 6** In the Time to wait field, specify how long the Content Engine should wait before timing out. The default value is 5 seconds.
- Step 7** In the Number of Retransmits field, specify the number of attempts allowed to connect to a TACACS+ server. The default value is 2.
- Step 8** In the Security Word field, enter the secret key that is used to communicate with the TACACS+ server.
- Step 9** In the Primary Server field, enter an IP address or host name information for the primary TACACS+ server.
- Step 10** In the Secondary Server field, enter an IP address or host name information for a secondary TACACS+ server.
- Step 11** In the Tertiary Server field, enter an IP address or host name information for a tertiary TACACS+ server.
- Step 12** To save the settings, click **Submit**.

Table 12-3 TACACS+ Server Key Parameter Settings

Key Parameter	Description	CLI Command
Enable TACACS+ Servers	Enables TACACS+ authentication.	tacacs enable
Use ASCII Password Authentication	Changes the default password type from PAP (Password Authentication Protocol) to ASCII clear text format.	tacacs password ascii
Time to wait	Number of seconds to wait for a response before timing out on a connection to a particular TACACS+ server. The range is from 1 to 20 seconds. The default value is 5 seconds.	tacacs timeout
Number of retransmits	Number of times that a connection to a TACACS+ server is allowed to be attempted before a connection is made. The default value is 2 times. The range is 1 to 3 times.	tacacs retransmit
Security Word	Encryption key shared with the TACACS+ server.	tacacs key
Primary Server	IP address or host name of the primary TACACS+ server.	tacacs host ip-address or hostname [primary]
Secondary Server Tertiary Server	IP address or host name of the backup TACACS+ server. 2 backup servers are allowed.	tacacs host ip-address or hostname

Configuring Login Authentication and Configuration Authorization Schemes for the Device

To configure the login authentication and configuration authorization schemes for the device, follow these steps. In the following example, local is configured as the primary database, TACACS+ as the secondary database, and RADIUS as the tertiary database.



Note

You must configure TACACS+ and RADIUS server settings for the device before you configure and submit these settings. See the [“Configuring TACACS+ Server Settings”](#) section on page 12-9 and the [“Configuring RADIUS Server Settings”](#) section on page 12-7 for information on how to configure these servers.

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the device that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **General Settings > Login Access Control > Login Authentication**. The Login Authentication Settings window appears. (See [Figure 12-3](#).)

Figure 12-3 Login Authentication Settings Window

The screenshot shows the Cisco Application and Content Networking System interface in Microsoft Internet Explorer. The main window is titled "Login Authentication Settings for Content Engine, CONTENTENGINE". On the left is a "Contents" pane with a tree view showing the navigation path: General Settings > Login Access Control > Login Authentication. The main area contains the following settings:

- Current applied settings from Content Engine, CONTENTENGINE**
- Enable Failover Server Unreachable:** ☒
- Authentication Login Servers:**
 - ☐ (checkbox)
 - Primary Login Server:** local
 - Secondary Login Server:** Do Not Set
 - Tertiary Login Server:** Do Not Set
- Authentication Config Servers:** ☒ (checkbox)
 - Primary Config Server:** TACACS+
 - Secondary Config Server:** local
 - Tertiary Config Server:** Do Not Set

At the bottom, there is a note: "Note: * - Required Field". The "Submit" and "Cancel" buttons are at the bottom right.

- Step 4** To query the secondary authentication database if the primary authentication server is unreachable, check the **Enable Failover Server Unreachable** check box.

To use this feature, you must set TACACS+ or RADIUS as the primary authentication method and local as the secondary authentication method.

- Step 5** To enable authentication privileges using the local, TACACS+, or RADIUS databases, check the **Authentication Login Servers** check box.
 - Step 6** From the Primary Login Server drop-down list, choose **local**.
 - Step 7** From the Secondary Login Server drop-down list, choose **TACACS+**.
 - Step 8** From the Tertiary Login Server drop-down list, choose **RADIUS**.
 - Step 9** To enable authorization privileges using the local, TACACS+, or RADIUS databases, check the **Authentication Config Servers** check box.
 - Step 10** From the Primary Config Server drop-down list, choose **local**.
 - Step 11** From the Secondary Config Server drop-down list, choose **TACACS+**.
 - Step 12** From the Tertiary Config Server drop-down list, choose **RADIUS**.
 - Step 13** To add authentication and authorization settings for the Content Engine, click **Submit**.
-

Changing the Admin Account Password for a Device

The default user account, *admin*, comes with the ACNS software and is preassigned the password *default*. This account allows access to all services and entities in the system. Cisco ACNS software allows any user with a role that gives them access to the Admin Password Setting window in the Content Distribution Manager GUI to configure a new password for the *admin* user account on individual Content Engines and Content Routers.



Note

Each Content Engine or Content Router must have only one user account with the username *admin*. Therefore, the admin account password configuration is done per device, rather than at the system level.

To configure the admin password for a Content Engine or Content Router, follow these steps:

- Step 1** From the Content Manager GUI, choose **Devices > Devices** (or **Devices > Device Groups**). The Devices (or Device Group) window appears.
 - Step 2** Click the **Edit** icon next to the device or device group for which you want to configure the admin account password.
 - Step 3** In the Contents pane, choose **General Settings > Login Access Control > Users > Admin Password**. The Admin Password Setting window appears.
 - Step 4** In the Password field, enter a password.
 - Step 5** In the Confirm Password field, reenter the password.
 - Step 6** Click **Submit** to save the new password.
-

Configuring User Accounts for Centrally Managed Devices

Cisco ACNS software allows you to create, modify, and delete user accounts for login access to individual devices (Content Engines and Content Routers) or device groups that are registered with a Content Distribution Manager.

Creating User Accounts for Centrally Managed Devices

To create a new user account from the Content Distribution Manager GUI, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices** (or **Devices > Device Groups**).
 - Step 2** Click the **Edit** icon next to the device (or device group) for which you want to create new user account. The Device Home window for the chosen device appears.
 - Step 3** In the Contents pane, choose **General Settings > Login Access Control > Users > Usernames**. The Usernames window for the chosen device (or device group) appears.

The Aggregate Settings **Yes** radio button is selected by default, and usernames are displayed for the Content Engine, as well as for any associated device groups. The settings for the Usernames for device groups cannot be modified or deleted; they are read-only. However, you can modify the Usernames for the Content Engine. If you click the Aggregate Settings **No** radio button, you can view and modify the Usernames for the Content Engine only. The Usernames for any associated device groups are not displayed.
 - Step 4** Click the **Create New Username** icon from the taskbar. The Creating New Local User for the device window appears.
 - Step 5** In the Username field, enter the username.
 - Step 6** In the Password field, enter the password for the user account.
 - Step 7** In the Confirm Password field, retype the password.
 - Step 8** In the Cifs Password field, enter the Windows user password.
 - Step 9** In the Confirm Cifs Password field, retype the Windows user password.
 - Step 10** From the Privilege drop-down list, choose **Super User** to create a user account with admin privileges. Alternatively, to create a user account without admin privileges, choose **Normal User**.
 - Step 11** To save the settings, click **Submit**.
-

This GUI procedure corresponds to the following global configuration CLI command:

```
username name { cifs-password | password | privilege }
```

Modifying and Deleting User Accounts for Centrally Managed Devices

To modify an existing user account, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices** (or **Devices > Device Groups**).
 - Step 2** Click the **Edit** icon next to the device or device group for which you want to modify the user account information.

- Step 3** In the Contents pane, choose **General Settings > Login Access Control > Users > Usernames**. The Usernames window for the chosen device (or device group) appears.
- Step 4** In the Username field, edit the name of the user, if needed.
- Step 5** Edit other information or settings as needed.
- Step 6** To save your settings, click **Submit**.

**Note**

Deleting a user account from the CLI does *not* disable the corresponding user account in the CMS database. Consequently, the user account remains active in the CMS database. User accounts created in the Content Distribution Manager GUI should always be deleted from the Content Distribution Manager GUI.

To delete a user account from the Content Distribution Manager GUI, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices** (or **Devices > Device Groups**).
- Step 2** Click the **Edit** icon next to the device or device group for which you want to delete a user account.
- Step 3** In the Contents pane, choose **General Settings > Login Access Control > Users > Usernames**. The Usernames window for the chosen device (or device group) appears.
- Step 4** In the taskbar, click the **Trash** icon.
- Step 5** To confirm the action, click **OK**.

Managing User Accounts and Privilege Profiles for the Content Distribution Manager

In a service provider environment, it is necessary to manage user authorization to services and access to domains (sets of Content Engines, device groups, and content providers). Typically, a service provider creates a user account for itself with administrator-level privileges, and then it creates user accounts for its customers. Service provider customers then have the ability to set up and manage user accounts on their own, which inherit the same level of rights as their main customer account unless the customers choose to limit user account privileges further.

ACNS 5.x software provides services for Content Distribution Manager user administration, role management, domain management, and accounting. These services can be configured in the Content Distribution Manager GUI.

Tasks include the following:

- [Creating and Managing User Accounts for the Content Distribution Manager, page 12-16](#)
- [Creating and Managing Roles, page 12-20](#)
- [Configuring and Managing Domains, page 12-23](#)
- [Changing Your CLI User Password, page 12-20](#)

ACNS 5.5 software supports a *user manager* account for creating and managing Content Distribution Manager user accounts. The user manager is responsible for creating a group of Content Distribution Manager user accounts, assigning roles and domains, as well as modifying any other user accounts in the Content Distribution Manager. User managers must be assigned the admin role, which authorizes all services and access to all domains. Only Content Distribution Manager users with administrator-level authorization are allowed to create and manage Content Distribution Manager user accounts.

After an upgrade to ACNS 5.5 software, any user accounts that do not meet the above requirement are prevented from performing user manager services. In other words, these users can log in to the Content Distribution Manager GUI but do not have access to the User Management pages (System > Users). These accounts can access the User Management pages by having their authorization increased to the administrator role. This can only be done by those Content Distribution Manager users with the administrator role.

Creating and Managing User Accounts for the Content Distribution Manager



Note

This section is addressed to users with administrator-level privileges (admin users) only.

When you create a user account, you enter information about the user. A user account contains a username, the name of the individual who owns the account, contact information, job title, and department. All user account information is stored in an internal database on the Content Distribution Manager.

Each user account can then be assigned to a role and a domain. A *role* defines which Content Distribution Manager GUI configuration pages the user can access and which services the user has authority to configure or modify. A *domain* defines which entities in the network the user can access and configure or modify. You can assign a user account to zero or more roles, and to zero or more domains.

Two default user accounts are preconfigured in the Content Distribution Manager and come with the ACNS software. The first account, called *admin*, is assigned the administrator role that allows access to all services and access to all entities in the system. This account cannot be deleted from the system, but it can be modified. Only the username and the role for this account are unchangeable.

The second preconfigured user account is called *default*. Any user account that is authenticated but has not been registered in the Content Distribution Manager obtains the access rights (role and domains) assigned to the default account. This account is configurable, but it cannot be deleted nor its username changed.

Using the Content Manager Distribution GUI, you can perform these tasks:

- Create a new user account.
- Modify and delete existing user accounts.
- View all user accounts in the system.

Creating New User Accounts Using the Content Distribution Manager GUI

When you create a new user account in the Content Distribution Manager GUI, you have the option to create the user account in the CLI at the same time. Using this GUI option to create the new account in the CLI provides the following benefits:

- The user account is created in the primary Content Distribution Manager management database and in the CLI from one central point.
- The user account is also created in the standby Content Distribution Manager management database and in the CLI from one central point.
- Users can change their passwords, and the password change will propagate to standby Content Distribution Managers.

If you choose to create the user account from the GUI *without* creating the user account in the CLI at the same time from the GUI, the following results apply:

- The user account is created in the primary and standby Content Distribution Manager management databases.
- No user account is created in the CLI, and the user *cannot* log in to the Content Distribution Manager GUI until an account is created from the CLI.
- Local users cannot change their passwords using the GUI.
- Local users can change their passwords using the CLI; however, the password change is not propagated from the CLI to the CMS when the CLI user option is enabled in the GUI.

If a user account has been created from the CLI only, when you log in to the Content Distribution Manager for the first time, the Centralized Management System (CMS) automatically creates a user account (with the same username as configured in the CLI) with default authorization and access control. However, to change the password in this scenario, the user account must be explicitly configured from the GUI with CLI user option enabled.

To create a new user account, follow these steps:

- Step 1** Choose **System > AAA > Users**. Click the **Create New User Accounts** icon. The Creating New User Account window appears. (See [Figure 12-4](#).)



Note This window can only be accessed by users with administrator-level privileges.

Figure 12-4 Creating New User Account Window

The screenshot shows the 'Creating New User Account' window in the Cisco Application and Content Networking System GUI. The window is titled 'Cisco Application and Content Networking System - Microsoft Internet Explorer'. The main content area is divided into three sections: 'Account Information', 'User Information', and 'Comments'. The 'Account Information' section contains fields for 'Username', 'Create CLI User' (checkbox), 'Password', 'Confirm Password', and 'Privilege Level' (dropdown menu). The 'User Information' section contains fields for 'First Name', 'Last Name', 'Phone Number', 'Email Address', 'Job Title', and 'Department'. The 'Comments' section is a large text area. At the bottom of the window, there are 'Submit' and 'Cancel' buttons. A note at the bottom left states: 'Note: * - Required Field'.

- Step 2** In the Username field, enter the user account name.
- Step 3** If you want to create a local user account with password and privilege level from the Content Distribution Manager GUI, click the **Create CLI User** check box. The user account is created automatically in the CLI. To deny creating a CLI user account from the GUI, leave the box unchecked.
- Step 4** In the Password field, enter a password for the CLI user account, and reenter the same password in the Confirm Password field.
- Step 5** From the drop-down list, choose a privilege level for the CLI user account. The choices are 0 (zero) (normal user) or 15 (superuser). The default value is 0.



Note A superuser can use privileged EXEC-level commands, whereas a normal user can use only user-level EXEC commands.

- Step 6** In the Username fields, enter the following information about the user: First Name, Last Name, Phone Number, Email Address, Job Title, and Department.
- Step 7** In the Comments field, enter any additional information about this account.
- Step 8** To save your settings, click **Submit**.

Modifying and Deleting User Accounts



Note Modifying a user account from the CLI does not update the Centralized Management System (CMS) database.

To modify an existing user account, follow these steps:

Step 1 In the Content Distribution Manager GUI, choose **System > AAA > Users**. The User Accounts window appears.

Step 2 Click the **Edit** icon next to the user account that you want to modify. The Modifying User Account window appears.



Note This window can only be accessed by users with administrator-level privileges.

Step 3 In the Username field, edit the name of the user, if needed.

Step 4 To make the password fields and privilege level drop-down list available for modification, check the **Change CLI User Password**.

Step 5 Edit other information or settings as needed.

Step 6 To save your settings, click **Submit**.



Note Deleting a user account from the CLI does *not* disable the corresponding user account in the CMS database. Consequently, the user account remains active in the CMS database. User accounts created in the Content Distribution Manager GUI should always be deleted from the Content Distribution Manager GUI.

To delete a user account from the Content Distribution Manager GUI, follow these steps:

Step 1 In the Content Distribution Manager GUI, choose **System > AAA > Users**. The User Accounts window appears.

Step 2 Click the **Edit** icon next to the user account that you want to delete. The Modifying User Account window appears.



Note This window can only be accessed by users with administrator-level privileges.

Step 3 In the taskbar, click the **Trash** icon.



Note If the CLI user account was created using the GUI, the corresponding CLI user account is removed from the CLI and is also deleted from all standby Content Distribution Managers.

Step 4 To confirm the action, click **OK**.

Changing Your CLI User Password

If you are a user *without* admin privileges and you are logged in to the Content Distribution Manager GUI, you can change your own CLI user password if you meet the following requirements:

- Your CLI user account and password were created in the Content Distribution Manager GUI and not in the CLI.
- You are authorized to access the password window.



Note

We do not recommend changing the CLI user password from the CLI. Any changes to CLI user passwords from the CLI are *not* updated in the management database and are not propagated to the standby Content Distribution Manager. Therefore, passwords in the management database will not match a new password configured in the CLI.



Note

The advantage of initially setting passwords from the Content Distribution Manager GUI is that both the primary and the standby Content Distribution Managers will be synchronized, and GUI users will not have to access the CLI to change their password.

To change your CLI user password, follow these steps:

- Step 1** In the Content Distribution Manager GUI, choose **System > Password**. The Changing Password for User Account window appears.
- Step 2** In the New Password field, enter the changed password.
- Step 3** In the Confirm New Password field, reenter the password for confirmation.
- Step 4** To save your settings, click **Submit**.

Viewing User Accounts

To view all user accounts, choose **System > AAA > Users** from the Content Distribution Manager GUI. The User Accounts window displays all the user accounts in the management database.

Creating and Managing Roles

The Content Distribution Manager provides many types of services. Not all users have access to all services. Users are assigned a role, which indicates the services to which they have access. A *role* is a set of enabled services.

Each user account can be assigned to zero or more roles. Roles are not inherited or embedded. Using the Content Distribution Manager GUI, you can perform these tasks:

- Create new roles.
- Modify and delete existing roles.
- Assign roles to user accounts.
- View all roles in the system.

The Content Distribution Manager provides one predefined role, known as the admin role. The admin role has access to all services and all ACNS network entities. To create a role, follow these steps:

- Step 1** Choose **System > AAA > Roles**. The Roles listing window appears.
- Step 2** In the taskbar, click the **Create New Role** icon. The Creating New Role window appears. (See [Figure 12-5](#).)

Figure 12-5 Creating New Role Window

The screenshot shows the 'Creating New Role' window in the Cisco Application and Content Networking System. The window is titled 'Creating New Role' and contains a 'Role Definition' section. In this section, there is a 'Name' field with an asterisk indicating it is required. Below the name field is a list of 'Enabled Services' with checkboxes next to each item: Devices, Services, System, All Domains, API, System-Wide Monitoring, and System Status. There is also a 'Comments' section with a text area. At the bottom of the window, there is a 'Note' that says '* - Required Field'. The window has 'Submit' and 'Cancel' buttons at the bottom right. The background shows the Cisco Systems logo and the 'Application and Content Networking System' header.

- Step 3** In the Name field, enter the name of the role.
- Step 4** To expand the listing of services under that category, click a folder, and then check the check box next to the service or services that you want to enable for this role. To choose all the services under one category simultaneously, check the check box next to the top-level folder for those services.
- Step 5** Enter any comments about this role in the Comments field.
- Step 6** To save your settings, click **Submit**.

Modifying and Deleting Roles



Note

The admin user account, by default, is allowed access to all services and cannot be modified.

To modify a role, follow these steps:

-
- Step 1** Choose **System > AAA > Roles**. The Roles window appears.
 - Step 2** To modify an existing role, click the **Edit** icon next to the name of the role you want to change. The Modifying Role window appears.
 - Step 3** In the Name field, enter a new name or change the existing name of the role.
 - Step 4** To enable a service for this role, check the check box next to the services that you want. To disable a previously selected service, uncheck the check box next to the service you want to disable. To choose all the services under one category simultaneously, check the check box next to the top-level service.
 - Step 5** In the Comments field, enter any comments about this role.
 - Step 6** To save your settings, click **Submit**.
-

Viewing Role Settings

You might want to view role settings before assigning a role to a particular user account.

To view role settings, follow these steps:

-
- Step 1** Choose **System > AAA > Users**. The User Accounts window appears with all configured user accounts listed.
 - Step 2** Click the **Edit** icon next to the user account for which you want to assign roles. The Modifying User Account window appears.
 - Step 3** In the Contents pane, choose **Role Management**. The Role Management for User Account User window appears.
 - Step 4** In the Role Management for User Account User window, click the **View** (eyeglass) icon next to the role name to display a new popup window called the Viewing Role window.
The role names, comments about this role, and services that are enabled for this role are displayed.
 - Step 5** After you have finished viewing the settings, click **Close**.
-

Assigning Roles to User Accounts



Note

The admin user account, by default, is assigned to the role that allows access to all domains and all entities in the system. It is not possible to change the role for this user account.

To assign roles to user accounts, follow these steps:

-
- Step 1** Choose **System > AAA > Users**. The User Accounts window appears with all configured user accounts listed.
 - Step 2** Click the **Edit** icon next to the user account for which you want to assign roles. The Modifying User Account window appears.

- Step 3** In the Contents pane, choose **Role Management**. The Role Management for User Account User window appears with all configured role names listed.
- Step 4** Click the **Assign** icon (blue cross mark) that appears next to the role name that you wish to assign to the selected user account.
- Step 5** To unassign a previously assigned user account role, click the **Unassign** (green tick mark) next to the role name.



Note Click the **Assign all Roles** icon in the taskbar to assign all roles in the current window to a user account. Alternatively, click the **Remove all Roles** icon to unassign all roles associated with a user account.

- Step 6** To save your settings, click **Submit**. A green tick mark appears next to the assigned roles and a blue cross mark appears next to the unassigned roles. The roles assigned to this user account will be listed in the Roles section in the Modifying User Account window.

Configuring and Managing Domains

A *domain* is a set of ACNS network entities or objects that make up the ACNS network. Whereas a role defines which services a user can perform in the ACNS network, a domain defines which entities the user has access to. The Content Distribution Manager GUI provides three predefined entities. An *entity* can be a Content Engine, content provider, or device group. These predefined entities are treated like services and can be enabled or disabled when you set up user roles.

When you configure a domain, you can choose to include Content Engines, content providers, or device groups in the domain. Using the Content Distribution Manager GUI, you can perform these tasks:

- Create new domains.
- Modify and delete existing domains.
- View all domains in the system.

To create a new domain, follow these steps:

- Step 1** Choose **System > AAA > Domains**. The Domains listing window appears.
- Step 2** In the taskbar, click the **Create New Domain** icon. The Creating New Domain window appears. (See [Figure 12-6](#).)

Figure 12-6 Creating New Domain Window

The screenshot shows the 'Creating New Domain' window in the Cisco Application and Content Networking System. The window is divided into two main sections: 'Domain Information' and 'Comments'. The 'Domain Information' section contains a 'Name' field (marked as required) and an 'Entity Type' dropdown menu. The dropdown menu is open, showing four options: 'Please make a choice', 'Content Engines', 'Content Providers', and 'Device Groups'. The 'Comments' section contains a large text area for entering comments. At the bottom of the window, there is a 'Note: * - Required Field' and two buttons: 'Submit' and 'Cancel'.


- Step 3** In the Name field, enter the name of the domain.
- Step 4** From the Entity Type drop-down list, choose the entity type that you want to assign to the domain. Entity choices include Content Engines, content providers, or device groups.
- Step 5** In the Comments field, enter any comments about this domain.
- Step 6** To save your settings, click **Submit**. If the entity type you chose has not already been assigned to the domain, then a message indicating that the entity type has not been assigned appears.

Adding an Entity to a Domain

To add an entity to a domain, follow these steps:

- Step 1** Choose **System > AAA > Domains** and, click the **Edit** icon next to the name of the domain that you want to modify.
- Step 2** In the Contents pane, choose **Entity Management**. The *Entity_name* Assignments for Domain window for the current domain appears. (See Figure 12-7.) In Figure 12-7, the entity is a content provider.

Figure 12-7 Entity Assignment to Domain Window—Content Provider Entity Shown



- Step 3** To add an entity to the current domain, click the **Assign** icon (blue cross mark) next to the name of the entity that you want to add. A green tick mark appears next to the selected entity when you submit the settings.
- Alternatively, to add all entities to the selected domain, click the **Assign all** icon in the taskbar.
- Step 4** To remove an entity from the current domain, click the **Unassign** icon (green tick mark) next to the name of the entity that you want to remove from the domain. A blue cross mark appears next to the unassigned entity after you submit the settings.
- Alternatively, to remove all entities from the domain, click the **Remove all** icon in the taskbar.
- Step 5** To save your settings, click **Submit**.

Modifying and Deleting Domains

To modify or delete an existing domain, follow these steps:

- Step 1** Choose **System > AAA > Domains**. The Domains window appears.
- Step 2** Click the **Edit** icon next to the domain that you want to modify. The Modifying Domain window appears.
- Step 3** Modify the settings as desired.
- Step 4** To save your settings, click **Submit**.
- Step 5** To delete the domain, click the **Trash** icon in the taskbar.

- Step 6** To confirm the action, click **OK**.

Viewing Domains

To view the domain configuration for a particular user account, follow these steps:

- Step 1** Choose **System > AAA > Users**. The User Accounts window appears with all configured user accounts listed.
- Step 2** Click the **Edit** icon next to the user account for which you want to view the domain configuration. The Modifying User Account window appears.
- Step 3** In the Contents pane, choose **Domain Management**. The Domain Management for User Account User window appears.
- Step 4** To display a new popup window called the Viewing Domain window, click the **View** (eyeglass) icon next to the domain name.
- The domain name, entity type, comments about this domain, and entities assigned to this domain are displayed.
- Step 5** After you have finished viewing the settings, click **Close**.

Assigning Domains to User Accounts

To assign domains to user accounts, follow these steps:

- Step 1** Choose **System > Users**. The User Accounts window appears with all configured user accounts listed.
- Step 2** Click the **Edit** icon next to the user account for which you want to assign domains. The Modifying User Account window appears.
- Step 3** In the Contents pane, choose **Domain Management**. The Domain Management for User Account User window appears with all configured domains and their entity types listed.
- Step 4** Click the **Assign** icon (blue cross mark) that appears next to the domain name that you wish to assign to the selected user account.
- Step 5** To dissociate an already associated domain from the user account, click the **Unassign** (green tick mark) next to the domain name.



Note To assign all domains in the current window to a user account, click the **Assign all Domains** icon in the taskbar. Alternatively, to unassign all domains associated with a user account, click the **Remove all Domains** icon.

- Step 6** To save your settings, click **Submit**. A green tick mark appears next to the assigned domains, and a blue cross mark appears next to the unassigned domains. The domains assigned to a user account are listed in the Domains section in the Modifying User Account window.

Configuring AAA Accounting

Accounting tracks all user actions and when the actions occurred. It can be used for an audit trail or for billing for connection time or resources used (bytes transferred).

The ACNS accounting feature uses TACACS+ server logging. Accounting information is sent to the TACACS+ server only, not to the console or any other device. The syslog file on the Content Engine logs accounting events locally. The format of events stored in the syslog is different from the format of accounting messages.

The TACACS+ protocol allows effective communication of AAA information between Content Engines and a central server. It uses TCP for reliable connections between clients and servers. Content Engines send authentication and authorization requests, as well as accounting information to the TACACS+ server.


**Note**

Before you can configure the AAA accounting settings for any device, you must first configure a TACACS+ server for the device. (See the [“Configuring TACACS+ Server Settings”](#) section on [page 12-9](#).)

To configure AAA accounting settings, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the device that you want to configure. The Device Home window appears.
- Step 3** In the Contents pane, choose **General Settings > Authentication > AAA Accounting**. The AAA Accounting Settings window appears. (See [Figure 12-8](#).) [Table 12-4](#) describes the fields in this window and provides the corresponding CLI global configuration commands that are available on the Content Distribution Manager.

Figure 12-8 AAA Accounting Settings Window



- Step 4** To activate accounting for system events, choose a keyword from the **System Events** drop-down list that indicates when accounting is to take place.
- Step 5** To activate accounting for EXEC mode processes, choose a keyword from the **Exec Shell and Login/Logout Events** drop-down list that indicates when accounting is to take place.
- Step 6** To activate accounting for all commands at the normal user level, choose a keyword from the **Normal User Commands** drop-down list that indicates when accounting is to take place.
- Step 7** To activate accounting for all commands at the administrative user level, choose a keyword from the **Administrative User Commands** drop-down list that indicates when accounting is to take place.

**Caution**

Before using the **wait-start** option, make sure that your Content Engine is configured with the TACACS+ server and is able to successfully contact the server. If the Content Engine cannot contact a configured TACACS+ server, it might become unresponsive.

Table 12-4 AAA Accounting Settings

GUI Parameter	Function	CLI Command
Event Type		
System Event	Accounting for all system-level events not associated with users, such as reloads.	aaa accounting system default {start-stop stop-only} tacacs

Table 12-4 AAA Accounting Settings (continued)

GUI Parameter	Function	CLI Command
Exec Shell and Login/Logout Events	Accounting for EXEC shell and user login and logout events on the Content Engine. Reports include username, date, start and stop times, and Content Engine IP address.	aaa accounting exec default {start-stop stop-only wait-start} tacacs
Normal User Commands	Accounting for all commands at the normal user privilege level.	aaa accounting commands 0 default {start-stop stop-only wait-start} tacacs
Administrative User Commands	Accounting for all commands at the administrative user privilege level.	aaa accounting commands 15 default {start-stop stop-only wait-start} tacacs
Options		
stop-only	The Content Engine sends a stop record accounting notice at the end of the specified activity or event to the TACACS+ accounting server.	aaa accounting {commands exec system} default stop-only tacacs
start-stop	The Content Engine sends a start record accounting notice at the beginning of an event and a stop record at the end of the event to the TACACS+ accounting server. The start accounting record is sent in the background. The requested user service begins regardless of whether or not the start accounting record was acknowledged by the TACACS+ accounting server.	aaa accounting {commands exec system} default start-stop tacacs
wait-start	The Content Engine sends both a start and a stop accounting record to the TACACS+ accounting server. However, the requested user service does not begin until the start accounting record is acknowledged. A stop accounting record is also sent.	aaa accounting {commands exec} default wait-start tacacs
Do Not Set	Accounting is disabled for the specified event.	no aaa accounting

Step 8 To save the settings, click **Submit**.

Viewing Audit Trail Logs

The Content Distribution Manager logs user activity in the system. The only activities that are logged are those that change the ACNS network. This feature provides accountability for users actions, for example, which user did what and when. Logged activities include the following:

- Creation of ACNS network entities
- Modification and deletion of ACNS network entities
- System configurations

To view audit trail logs, follow these steps:

Step 1 Choose **System > Logs > Audit Trail Logs**. The Audit Log window appears. (See [Figure 12-9](#).) All logged transactions in the Content Distribution Manager are listed by date and time, user, actual transaction that was logged, and the IP address of the machine that was used.

Figure 12-9 Audit Log Window



Step 2 To determine the number of rows that you want to display, choose a number from the Rows drop-down list.



Working with Device Configurations

An ACNS network is a coordinated system made up of three types of devices: Content Distribution Managers, Content Engines, and, if necessary, Content Routers. To deliver requested content using these devices, proper configuration is necessary. ACNS 5.x software also allows you to configure Content Engines into device groups using the Content Distribution Manager GUI so that content services can be configured for an entire group of Content Engines at one time. Device groups and Content Engines share the same configuration features and options.

This chapter provides information about configuring ACNS network device properties. It contains the following sections:

- [Working with Content Distribution Managers, page 13-1](#)
- [Working with Device Locations, page 13-5](#)
- [Working with Content Engines, page 13-8](#)
- [Working with Content Routers, page 13-22](#)
- [Working with Device Groups, page 13-26](#)
- [Working with Baseline Groups, page 13-31](#)
- [Local Device Configuration Versus Central Management, page 13-35](#)
- [Repurposing an ACNS Network Device, page 13-36](#)

Working with Content Distribution Managers

You can work with Content Distribution Managers by performing these tasks:

- [Viewing the Content Distribution Manager Configuration, page 13-2](#)
- [Configuring a Warm Standby Content Distribution Manager, page 13-2](#)
- [Switching a Content Distribution Manager from Warm Standby to Primary, page 13-3](#)
- [Deleting a Warm Standby Content Distribution Manager, page 13-5](#)

Viewing the Content Distribution Manager Configuration

To view the Content Distribution Manager configuration, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon next to the name of the Content Distribution Manager that you want to view. The Device Home window appears.
 - Step 3** To display the entire table of contents, click the **Show All** button above the Contents pane.
 - Step 4** In the Contents pane, choose **Device Activation**. The Modifying Content Distribution Manager window appears, which shows the basic network settings that were configured at initial setup through the CLI.
 - Step 5** To view configuration settings, choose **Monitoring > Show/Clear Commands > Show Commands** in the Contents pane.
 - Step 6** From the drop-down list, choose a **show** command, and click **Submit**.
A popup window appears, showing the CLI output for the command.
-

Configuring a Warm Standby Content Distribution Manager



Note

Primary and standby Content Distribution Managers must be running the same ACNS software release. If they are not, the standby Content Distribution Manager detects this problem and shuts down the Centralized Management System (CMS) until the problem is corrected. We recommend that you upgrade your standby Content Distribution Manager first and then upgrade your primary Content Distribution Manager.

To configure a device to be a warm standby Content Distribution Manager, follow these steps:

-
- Step 1** If the device is not already configured as a Content Distribution Manager, set the device mode by using the **device mode** command.


```
DeviceName# configure
DeviceName(config)# device mode content-distribution-manager
```

See the [“Repurposing an ACNS Network Device”](#) section on page 13-36.
 - Step 2** Obtain the primary Content Distribution Manager clock setting by using the **show clock** command, as shown in this example:


```
PrimaryCDM# show clock
Fri Jan 31 15:02:28 UTC 2003
```
 - Step 3** Set the standby Content Distribution Manager clock to match the primary Content Distribution Manager clock by using the **clock set EXEC** command.


```
DeviceName(config)# exit
DeviceName# clock set 15:02:28 31 january 2003
```
 - Step 4** Configure the Content Distribution Manager to be in standby mode by using the global configuration **cdm role standby** command.


```
DeviceName# configure
```

```
DeviceName(config)# cdm role standby
```

- Step 5** Specify the primary Content Distribution Manager IP address by using the **cdm ip ip_addr** global configuration command.

```
DeviceName(config)# cdm ip ip_addr
```

- Step 6** Register the device and enable management services by using the global configuration **cms enable** command.

```
DeviceName(config)# cms enable
```

Verify that the registration was successful. The output from this command should say whether or not the registration was successful.

- Step 7** Before saving the configuration to flash memory, use the **show cms info** command to verify the time of the last update.

```
CDM# show cms info
Registration information :
Device Mode               = cdm
Model                     = CDM4630
Node Id                   = 96
Time of last update       = Fri Jan 31 15:59:28 2003
```

To be current, the update time should be between 1 and 5 minutes old. You are verifying that the standby Content Distribution Manager has fully replicated the primary Content Distribution Manager configuration. If the update time is not current, determine whether or not there is a connectivity problem or if the primary Content Distribution Manager is down. Fix the problem if necessary, and wait until the configuration has replicated as indicated by the time of the last update.

Also, use the **show cms processes** command to verify that all management services are running, as shown in this example:

```
CDM# show cms processes
Service cms_httpd running
Service cms_cdm running
```

- Step 8** Save the configuration to flash memory by using the **copy running-config startup-config** EXEC command.

```
DeviceName(config)# exit
DeviceName# copy running-config startup-config
```

Switching a Content Distribution Manager from Warm Standby to Primary

If your primary Content Distribution Manager becomes inoperable for some reason, you can manually reconfigure one of your warm standby Content Distribution Managers to be the primary Content Distribution Manager.

To switch a warm standby Content Distribution Manager to be a primary Content Distribution Manager, follow these steps:

- Step 1** Check the status of recent updates from the primary Content Distribution Manager. Use the **show cms info** EXEC command, and check the time of the last update.

```
DeviceName# show cms info
```

To be current, the update time should be between 1 and 5 minutes old. You are verifying that the standby Content Distribution Manager has fully replicated the primary Content Distribution Manager configuration. If the update time is not current, check whether there is a connectivity problem or if the primary Content Distribution Manager is down. Fix the problem, if necessary, and wait until the configuration has replicated as indicated by the time of the last update.

- Step 2** Make sure that both Content Distribution Managers have the same Coordinated Universal Time (UTC) configured. Use the **show clock detail** command to check the time on each device, as shown in this example:

```
CDM1# show clock detail

Local time: Thu Mar  6 19:03:25 UTC 2003

      UTC time: Thu Mar  6 19:03:25 UTC 2003

Epoch: 1046977405 seconds
UTC offset: 0 seconds (0 hours 0 minutes)
```

If the UTC time does not match on the two devices, use the **clock EXEC** command to configure the time and then use the **clock timezone UTC** global configuration command to match the UTC offset., as shown in this example:

```
CDM2# clock set 19:03:25 mar 06 2003
CDM2# configure
CDM2(config)# clock timezone UTC 0 0
```

- Step 3** Configure the new role by using the global configuration **cdm role primary** command.

```
DeviceName# configure
DeviceName(config)# cdm role primary
```

This command changes the role from standby to primary and restarts the management service to recognize the change.

- Step 4** Save the configuration to flash memory by using the **copy running-config startup-config** EXEC command.

```
DeviceName(config)# exit
DeviceName# copy running-config startup-config
```

If you switch a warm standby Content Distribution Manager to primary while your primary Content Distribution Manager is still online and active, both Content Distribution Managers detect each other, automatically shut themselves down, and disable management services. The Content Distribution Managers are switched to halted role, which is automatically saved in flash memory.

To return halted Content Distribution Managers to online status, follow these steps:

-
- Step 1** Choose a role (active or standby) for each Content Distribution Manager.
- Step 2** Change the role using the **cdm role** command.
- Step 3** Enable management services using the **cms enable** command.
- Step 4** Save the configuration using the **copy running-config startup-config** command.
-

Deleting a Warm Standby Content Distribution Manager

You can delete a warm standby Content Distribution Manager from the ACNS network at any point after you have registered the device and before the device has come on line as the primary Content Distribution Manager. Once the device has been called into use as the primary Content Distribution Manager, however, you cannot delete it using the Content Distribution Manager GUI.

Delete a warm standby Content Distribution Manager when the device is experiencing unresolvable problems or when its network address or configuration has changed and you need to add the device back to the ACNS network using its new address and configuration information.

When deleting a warm standby Content Distribution Manager from the ACNS network, you are effectively removing that device and the content it contains from the routing scheme that the ACNS software uses to fill user requests. Should your primary Content Distribution Manager fail during the time that the warm standby is deleted, there will be no failover for the Content Distribution Manager.

To delete a warm standby Content Distribution Manager, follow these steps:

-
Step 1	Log in directly to the Content Distribution Manager CLI, and enter the cms deregister command. If for some reason the deregistration fails, you can force the deregistration by using the cms deregister force command.
Step 2	From the Content Distribution Manager GUI, choose Devices > Devices . The browser refreshes, listing the Content Distribution Managers on your ACNS network. The warm standby Content Distribution Manager is identified as <i>Standby</i> .
Step 3	Click the Edit icon next to the name of the warm standby Content Distribution Manager that you wish to delete. The Device Home window appears.
Step 4	In the Contents pane, choose Device Activation . The browser window refreshes, displaying the Modifying Content Distribution Manager window.
Step 5	In the taskbar, click the Trash icon. You are prompted to confirm your decision.
Step 6	To execute your request, click OK .
-

Working with Device Locations

You can work with locations by performing these tasks:

- [Creating and Modifying Locations, page 13-5](#)
- [Deleting Locations, page 13-8](#)
- [Viewing the Location Tree, page 13-8](#)

Creating and Modifying Locations

Locations are set up in the Content Distribution Manager GUI, and they organize and group Content Engines into virtual networks for distribution of content through channels.

Locations need to be configured before you can activate Content Engines and Content Routers and bring them online in the ACNS network.

To create a new location or modify an existing one, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Locations**. The Locations window appears. (See [Figure 13-1](#).)

Figure 13-1 Locations Window



- Step 2** In the taskbar, click the **Create New Location** icon. The Creating New Location window appears. (See [Figure 13-2](#).)

Figure 13-2 **Creating New Location Window**

The screenshot displays the 'Creating new Location' window within the Cisco Application and Content Networking System. The window is titled 'Creating new Location' and contains the following elements:

- Location Information:**
 - Name:** A text input field with an asterisk (*) indicating it is a required field.
 - Parent Location:** A drop-down menu currently set to 'None'.
 - Level:** A text input field containing the value '1'.
- Comments:** A large text area for entering comments.
- Note:** A small note at the bottom left states: 'Note: * - Required Field'.
- Buttons:** 'Submit' and 'Cancel' buttons are located at the bottom right of the window.

The background interface shows the 'Cisco Application and Content Networking System' header, a 'System Status' bar indicating '2 Devices, Critical' and '18 Channels, Minor', and a navigation menu with options like 'Devices', 'Device Groups', 'Locations', and 'Statistics'.

Step 3 In the Name field, enter a location name.



Note Fields marked with an asterisk are required.

Step 4 From the Parent Location drop-down list, choose a parent location (or choose **None**).

A location with no parent location assigned is a level 1 location. A location with a level 1 parent becomes a level 2 location, and so forth. The location level is displayed after you choose a parent location (or choose **None**) and click **Submit** to save the configuration.

Step 5 In the Comments field, enter comments about the location. This step is optional.

Step 6 To save your settings, click **Submit**.

Step 7 To modify a location, go to the Locations window, and click the **Edit** icon next to the name of the location you want to modify.

Step 8 Change the location name, parent location, or description as needed, and click **Submit** to save the changes.

Deleting Locations

You can delete locations as needed, as long as they are not the root locations of activated Content Engines or Content Routers.

**Note**

If a location has a Content Engine or Content Router assigned to it, you can first assign the Content Engine or Content Router to another location and then delete the original location.

To delete a location, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Locations**. The Locations window appears.
- Step 2** Click the **Edit** icon next to the location that you want to delete.
The Modifying Location window appears.
- Step 3** In the taskbar, click the **Trash** icon. You are asked to confirm your decision to delete the location.
- Step 4** To confirm the action, click **OK**. The location is deleted.
-

Viewing the Location Tree

The location tree represents the network topology you configured when you assigned a parent to each location. The Content Distribution Manager GUI graphically displays the relationships between the locations configured in your ACNS network.

To view the location tree, choose **Devices > Locations**. Then in the taskbar, click the **Location Trees** button. The location tree is displayed.

Working with Content Engines

You can work with Content Engines by performing these tasks:

- [Modifying Content Engine Properties, page 13-9](#)
- [Rebooting a Content Engine, page 13-12](#)
- [Deleting a Content Engine, page 13-12](#)
- [Viewing Content Engine Channel Assignments, page 13-14](#)

Modifying Content Engine Properties

Use the Content Distribution Manager GUI to make changes to Content Engine properties.

To modify the Content Engine properties, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to modify. The Device Home window appears.
- Step 3** In the Contents pane, chose **Device Activation**. The Device Activation window appears with fields for editing the properties of the selected Content Engine. (See [Figure 13-3](#).)

Figure 13-3 Device Activation Window for Modifying Content Engine Properties

The screenshot shows the Cisco Application and Content Networking System GUI in a Microsoft Internet Explorer browser window. The main navigation bar includes 'Devices', 'Services', and 'System'. The 'Devices' section is expanded, showing 'Device Groups', 'Locations', and 'Statistics'. The 'Contents' pane on the left lists various configuration options, with 'Device Activation' selected. The 'Device Activation, CONTENTENGINE' window is displayed, showing configuration fields for 'General Configuration' and 'Locality'. The 'General Configuration' section includes fields for Name (CONTENTENGINE), Status (Offline), Activate (checked), cdnfs (0 MB), and Set Default Coverage Zone File (checked). The 'Locality' section includes a Location dropdown (ContentEngine-location). Below these are sections for 'Request Routing' and 'NAT Configuration'. The 'Request Routing' section includes fields for Enable Request Routing (unchecked), Coverage Zone File (None), Enable PAC File Server (unchecked), and Enable Least Loaded (unchecked). The 'NAT Configuration' section includes a checkbox for 'Use CE's primary IP Address' (checked), a NAT Address field, and a Port field (443). A 'Comments' section at the bottom contains a text area with 'n/a'. A 'Note' at the bottom indicates that fields with an asterisk are required. The bottom right corner of the window has 'Submit' and 'Cancel' buttons.

- Step 4** Under the General Configuration heading, set or modify Content Engine properties.
 - To change the hostname of the Content Engine, enter a new name in the Name field.
 - To activate or deactivate the Content Engine check or uncheck the **Activate** check box.
 - To set the default coverage zone, check the **Set Default Coverage Zone File** check box.
- Step 5** Under the Locality heading, set or change the location information by choosing a new location from the **Location** drop-down list.

Step 6 Under the Request Routing Heading, set or modify routing properties.

- To enable routing on the Content Engine, check the **Enable Request Routing** check box.
- If you wish to have a local coverage zone file overwrite the ACNS network-wide coverage zone file on this routing-enabled Content Engine, you can choose a file from the **Coverage Zone File** drop-down list.

The Coverage Zone File drop-down list contains the destination paths of all valid coverage zone files registered to your ACNS network. When you select a coverage zone file from this list, the Content Engine uses this coverage zone file to redirect the user's request. (For information about registering and applying coverage zones, see the [“Registering and Applying a User-Defined Coverage Zone”](#) section on page 4-44.)

- To enable the Content Engine as a PAC file server, check the **Enable PAC File Server** check box.
- To enable load-based content routing on a routing-enabled Content Engine, check the **Enable Least Loaded** check box. (See the [“Configuring Load-Based Content Routing Using a Content Router”](#) section on page 4-55.)

Step 7 Under the NAT Configuration heading, configure the NAT settings.

- To use the IP address configured on the primary interface of the Content Engine as the management IP address, check the **Use CE's primary IP Address** check box.
- To allow the Content Distribution Manager to communicate with devices in the ACNS network that are behind the NAT firewall using an explicitly configured IP address, enter the NAT address of the Content Engine in the NAT Address field.
- In the Port field, enter the port number for the NAT address.



Note If the Content Distribution Manager cannot contact a device using the primary IP address, it will attempt to communicate using NAT IP address.

[Table 13-1](#) describes the Content Engine configuration fields that you can either view or modify.

Table 13-1 Content Engine Configuration Properties

Content Engine Configuration Property	Description
General Configuration	
Name	Identifier for the Content Engine. This name is used by content routing and must conform to the following rules: • The name must use only alphanumeric characters and hyphens (-). • The first and last character must be a letter or a digit. • Maximum length is 63 characters. • Names are case-insensitive. Note The following characters are considered illegal and cannot be used when naming a device: @, #, \$, %, ^, &, *, (), !, \, '"/>
Status	Read only. Content Engine status: online, pending, or inactive.

Table 13-1 Content Engine Configuration Properties (continued)

Content Engine Configuration Property	Description
Activate	When checked, activates the Content Engine for centralized management through the Content Distribution Manager GUI.
cdnfs	Read only. Current setting for cdnfs pre-positioned content disk space allocation. This value cannot be modified from this window. (See the “Updating Storage Capacity Through the Content Distribution Manager GUI” section in the <i>Cisco ACNS Software Update and Maintenance Guide</i> .)
Set Default Coverage Zone file	When checked, generates a default coverage zone with the Content Engine serving the local subnet it resides on. The coverage zone is an ACNS network-wide mapping of user client addresses to the Content Engine addresses that should respond to requests for that client.
Locality	
Location ¹	Location of the Content Engine. Network topology locations organize and group Content Engines into virtual networks for content distribution. Locations must be configured beforehand. (See the “Creating and Modifying Locations” section on page 5-3.)
Request Routing	
Enable Request Routing	Enables the Content Engine to act as a simplified hybrid content router as well as a Content Engine.
Coverage Zone File	Allows you to select an imported coverage zone file for simplified hybrid routing.
Enable PAC File Server	Enables the Content Engine to function as a proxy autoconfiguration file server.
NAT Configuration	
Use CE’s primary IP Address	Enables the Content Distribution Manager to use the IP address configured on the primary interface of the Content Engine to communicate with devices in the ACNS network that are behind a NAT firewall.
Use IP Address	When checked lets you manually configure a management IP address for the Content Distribution Manager to communicate with devices behind NAT firewall. If the Content Distribution Manager cannot contact a device using the current management IP address, it will not attempt to communicate using alternative IP addresses.
Port	Port number to enable communication between the Content Distribution Manager and devices behind a NAT firewall.

1. This field is required.

Step 8 To save your settings, click **Submit**.

Rebooting a Content Engine

Using the Content Distribution Manager GUI, you can reboot a Content Engine or group of Content Engines remotely. This option causes the Content Distribution Manager to perform a controlled shutdown of all Content Engines and then restarts the operating system on each device.

To reboot an individual Content Engine, follow these steps:

-
Step 1	From the Content Distribution Manager GUI, choose Devices > Devices .
Step 2	Click the Edit icon next to the name of the Content Engine that you wish to reboot. The Device Home window appears.
Step 3	In the taskbar, click the Reload Content Engine icon. You are prompted to confirm your decision.
Step 4	To begin rebooting the Content Engine, click OK .
-

To reboot an entire device group, follow these steps:

-
Step 1	From the Content Distribution Manager GUI, choose Devices > Device Groups .
Step 2	Click the Edit icon next to the name of the device group that you want to reboot. The Modifying Device Group window appears.
Step 3	In the taskbar, click the Reboot All Devices in Device Group icon. You are prompted to confirm your decision.
Step 4	To begin rebooting each Content Engine in the device group, click OK .
-

Deleting a Content Engine

Delete a Content Engine when the device is experiencing unresolvable problems or when its network address or configuration has changed and you need to add the device back to the ACNS network using its new address and configuration information.

When you delete a Content Engine from the ACNS network, you are removing that node and the content it contains from the routing scheme that the ACNS software uses to fill user requests. Although ACNS 5.x software routes requests around Content Engines that are busy, offline, or missing, removing a Content Engine may affect the speed at which the ACNS network can serve user requests.

**Note**

You cannot delete a Content Engine if it is the last node assigned to a location that is designated as the root location for a channel. If you receive an error referencing the root location for a channel, add more Content Engines to that location, or change the root location for the channel before attempting to delete the Content Engine again.

Before a node can be removed from the ACNS network, the following conditions must be met:

- The device must have been activated in the Content Distribution Manager. (See the [“Activating Devices in the Content Distribution Manager GUI”](#) section on page 2-22.)
- The Content Distribution Manager must be operating.

- The device must have the correct Content Distribution Manager IP address or host name configured.
- The Content Distribution Manager IP address or host name must be that of the primary Content Distribution Manager.

Deleting a Content Engine from the ACNS network involves using the CLI to remove the registration information from the Content Engine itself and removing the registration record from the Content Distribution Manager GUI.

**Note**

Do not use the Content Distribution Manager GUI to delete a Content Engine while the Content Engine is still active and registered. The GUI delete feature removes only the Content Engine's registration record from the Content Distribution Manager; it does not deregister the device. The Content Engine retains its registration information and continues to contact the Content Distribution Manager; however, the Content Distribution Manager no longer recognizes the Content Engine.

If for some reason the Content Distribution Manager loses the registration record of a device, use the **cms deregister force** command on the device to remove all of its registration information. Then use the **cms enable** command to reregister the device with the Content Distribution Manager as though it were a new node in the ACNS network.

If the Content Engine is being removed because of system disk failure, see the "Recovering from Missing Disk-Based Software" section in the *Cisco ACNS Software Update and Maintenance Guide*.

To delete a Content Engine, follow these steps:

Step 1 Open a Telnet session to the Content Engine CLI.

Step 2 Enter the **no cms enable** command in global configuration mode.

```
CE-507# configure
CE-507(config)# no cms enable
```

**Note**

Issuing the **no cms enable** command does not disable a primary sender Content Engine and activate the backup sender (or the reverse) if the backup sender is the active Content Engine being removed from the network; however, issuing the **cms deregister** command *does* disable the communication between the two multicast senders. The **cms deregister** command disables the CMS, all acquisition and distribution services, and all routing communications to and from this device.

Step 3 Enter the **cms deregister** command in EXEC mode.

```
CE-507(config)# exit
CE-507# cms deregister
```

**Note**

The **cms deregister** command cleans up the database automatically. You do not need to use the **cms database delete** command.

Step 4 If for some reason the deregistration fails, you can force the deregistration by using the **cms deregister force** command.

```
CE-507# cms deregister force
```



Note Ignore any messages stating that the deregistration failed. The **cms deregister force** command forces the deregistration to continue.

- Step 5** To add the device back in to the ACNS network, reregister the device with the Content Distribution Manager by using the **cms enable** command in global configuration mode.

```
CE-507# configure
CE-507(config)# cms enable
```

Viewing Content Engine Channel Assignments

You can use the Device Home page to view the number of channels to which a particular Content Engine is assigned. You can view channel assignments for an overview of how a given Content Engine is deployed across your ACNS network (to which channel it belongs, the channel quota, the website, and the content provider).

To view channel assignments for a Content Engine, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine for which you wish to view channel assignments. The Device Home window appears.
- Step 3** In the Device Home window, under Assignments: This *Devicename* belongs to *n* Channels, click the **Channels** link, or in the Contents pane, choose **Assignments**. The Channel assignments for Content Engine window appears, listing all the channels to which the Content Engine is assigned.

Commands Related to Content Acquisition and Distribution

Table 13-2 lists and describes the commands related to viewing content acquisition information. For complete descriptions of these commands, including usage guidelines and examples, refer to the *Cisco ACNS Software Command Reference, Release 5.5* publication.

Table 13-2 Acquisition-Related Commands

Command	Syntax	Description
acquirer check-time-for-old-content start-channel stop-channel test-url	acquirer { check-time-for-old-content [channel-id <i>channel-num</i> channel-name <i>channel-name</i>] [correct [channel-id <i>channel-num</i> channel-name <i>channel-name</i>]] start-channel { channel-id <i>channel-num</i> channel-name <i>channel-name</i> } stop-channel { channel-id <i>channel-num</i> channel-name <i>channel-name</i> } test-url <i>url</i> [use-http-proxy <i>url</i> use-ntlm-domain <i>domain-name</i> use-smb-options <i>smb-options</i>]}	Starts or stops content acquisition on a specified acquirer channel.
show acquirer channels progress proxy authentication	show acquirer [channels [channel-id <i>channel-num</i> channel-name <i>channel-name</i>] progress [channel-id <i>channel-num</i> channel-name <i>channel-name</i> <i>streams</i>] proxy authentication]	Displays the acquirer information for the channels specified for which the Content Engine is the root Content Engine.
show statistics acquirer channel-id channel-name contents errors job-list	show statistics acquirer [channel-id <i>channel-num</i> channel-name <i>channel-name</i> contents { channel-id <i>channel-num</i> channel-name <i>channel-name</i> } errors { channel-id <i>channel-num</i> channel-name <i>channel-name</i> } job-list { channel-id <i>channel-num</i> channel-name <i>channel-name</i> }]	Displays acquirer channel statistics for all the channels for which the Content Engine is the root Content Engine.

Table 13-3 lists and describes the commands related to viewing content distribution information. For complete descriptions of these commands, including usage guidelines and examples, refer to the *Cisco ACNS Software Command Reference, Release 5.5* publication.

Table 13-3 **Distribution-Related Commands**

Command	Syntax	Description
show distribution channels forwarder-list location mcast-data-receiver mcast-data-sender object-status processes remote	show distribution [channels [channel-id <i>channel-num</i> channel-name <i>name</i>]] show distribution { forwarder-list [channel-id <i>channel-num</i> [detail] channel-name <i>name</i> [detail] detail]} show distribution location { forwarder-load-weight live-load-weight location-loader-preference }[channel-id <i>channel-num</i> channel-name <i>channel-name</i>] show distribution mcast-data-receiver show distribution mcast-data-sender show distribution object-status object-url show distribution processes show distribution remote ip-address { metadata-center channel-id <i>channel-num</i> [start-generation-id <i>gen-id</i> end-generation-id <i>gen-id</i>] unicast-sender channel-id <i>channel-num</i> { cdn-url probe relative-cdn-url } <i>cdn-url</i> } show distribution remote traceroute { forwarder-next-hop channel-id <i>channel-num</i> { max-hop <i>maxhop_num</i> trace-till-good trace-till-root } unicast-sender channel-id <i>channel-num</i> { cdn-url probe relative-cdn-url } <i>cdn-url</i> { max-hop <i>maxhop_num</i> trace-till-good trace-till-root }}	Displays the distribution information for a specified channel.

Table 13-3 Distribution-Related Commands (continued)

Command	Syntax	Description
show statistics distribution all errors mcast-data-receiver mcast-data-sender metadata-receiver metadata-sender unicast-data-receiver unicast-data-sender	show statistics distribution { all errors mcast-data-receiver mcast-data-sender metadata-receiver metadata-sender unicast-data-receiver unicast-data-sender }	Displays the statistics of the content distribution components.
debug distribution all mcast-data-receiver mcast-data-sender metadata-receiver metadata-sender unicast-data-receiver unicast-data-sender	debug option	Monitors and records ACNS software distribution functions. Note We recommend that the debug command be used only at the direction of Cisco Systems technical support personnel. Performance is impacted when you run the debug command.

Configuring the Content Engine Common Internet File System Server

When the Content Engine Common Internet File System (CIFS) server is enabled, Windows clients can request pre-positioned content from the Content Engine by using the Windows file-sharing protocol Server Message Blocks (SMB) or CIFS. The Content Engine CIFS server uses the SMB protocol to serve Windows clients.

You can configure, modify, and view the CIFS server and control access of pre-positioned website content for Content Engines and device groups using the Content Distribution Manager GUI.



Note

All Content Engine settings in this section can also be configured for device groups by choosing **Devices > Device Groups** instead of **Devices > Content Engines** in the Content Distribution Manager GUI.


Configuring CIFS Server Settings

To configure the CIFS server settings for the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure CIFS settings. The Content Engine Device Home window appears with the Contents pane on the left.

- Step 3** In the Contents pane, choose **Applications > File Sharing > CIFS Server Configuration**. The CIFS Server Settings for Content Engine window appears. (See [Figure 13-4](#).)

Figure 13-4 CIFS Server Settings Window



Under the CIFS Server Settings panel, the currently applied settings details, such as whether or not default settings have been applied or settings have been derived from the device or device group, are displayed to enable you to identify the source information for the settings. The selected settings will be applied to the CIFS server configurations.

- Step 4** To enable the CIFS server on the Content Engine, check the **Enable CIFS Server** check box, and set the default mode of authentication for file sharing to ce-user.
- Step 5** From the Authentication Mode drop-down list, choose the Windows file sharing authentication mode. See [Table 13-4](#) for a description of the options displayed in the list.



Note The login authentication method must be set to local on the Content Engine to use the ce-user mode. By default, the CIFS authentication mode is set to ce-user.



Note CIFS authentication modes ce-user, password-server, and ldap-server do not work well with Windows operating systems. If you have a problem authenticating Windows-based users in these modes, we recommend that you use public or cifs-user authentication mode on the Content Engine.

The domain name and IP address of the NTLM servers that are configured in the NTLM Server Settings for Content Engine window become applicable when the password-server authentication mode is selected.

Table 13-4 Authentication Mode Settings

Authentication Mode	Description
ce-user	Enables the default CIFS authentication mode in ACNS software and allows a user to log in using the user account configured on the Content Engine to access file sharing. This mode is different from the CIFS user mode in earlier releases of ACNS software. You need not maintain two different passwords for the same user account.
public	Enables any user who is able to access the Content Engine using IP to browse all unprotected pre-positioned content on the cdnfs.
password-server	Enables NTLM authentication mode, in which the Content Engine relies on the password database stored on an NTLM password server to authenticate users. Password servers are Windows domain controllers. Choosing this option defines the domain or workgroup in which the Content Engine and Windows clients are configured and the server that stores the password database.
ldap-server	Enables the LDAP server for Windows file sharing authentication.

Step 6 In the Max Number of Connections field, enter the maximum number of simultaneous connections for pre-positioned content sharing. The default value is 16 to suit basic browsing requirements. The range is 1 to 65536.

Step 7 To save the settings, click **Submit**.

A “Click to Submit” message appears in red next to the Current Settings when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The **Reset** button is visible only when you apply default or device group settings to change the current device settings but the settings have not yet been submitted.

Configuring CIFS Server Website Access Control Settings

If no website is explicitly configured to be shared, all websites that have been previously defined are shared by default. In other words, all cdnfs content is shared through CIFS file sharing. The maximum number of websites that can be shared is 16.

If you choose the access level as **share** (see [Figure 13-5](#)) and set the authentication mode as **public** (see [Figure 13-4](#)), a warning message displays because all protected content will be accessible to any user who is able to access the Content Engine using IP.

To configure access control for certain pre-positioned websites on the cdnfs, follow these steps:

-
- Step 1** Choose **Devices > Devices**. The Devices window appears.
 - Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure access control. The Content Engine Device Home window appears with the Contents pane on the left.
 - Step 3** In the Contents pane, choose **General Settings > File Sharing > CIFS Server Access Control**. The CIFS Server Website Access Control for Content Engine window appears, listing all websites for which access control has been defined.

**Note**

When a device is associated with one or many device groups that have been configured with the settings displayed in this window, if you want to apply the group settings to the device, you can select the device group name from the Select a Device Group drop-down list that appears in the taskbar. A “Click Submit to Save” message appears in red next to the current settings when there are pending changes to be saved after you have applied default and device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The **Reset** button is visible only when you apply default or device group settings to change the current device settings when the settings have not yet been submitted.

The currently applied settings details, such as whether or not default settings have been applied or whether or not settings have been derived from the device or device group, display so that you can identify the source information for the settings.

- Step 4** If you want all websites that have been created to be shared, check the **Share All WebSites** check box. All individual websites for which access levels have been specified will be deleted. Click **Submit**. The system displays a dialog box prompting you to confirm your decision. Click **OK** to continue.
- Step 5** In the taskbar, click the **Create New CIFS Website** icon. The Creating New CIFS Website for Content Engine window appears. (See [Figure 13-5](#).)

Figure 13-5 Creating New CIFS Website Window

The screenshot displays the Cisco Application and Content Networking System web interface. The main window is titled 'Creating New CIFS Website for Content Engine, CONTENTENGINE'. The interface includes a left sidebar with a 'Contents' menu, a top navigation bar with 'Devices', 'Services', and 'System' tabs, and a main form area for 'CIFS Website Settings'. The form contains fields for 'Website' (a dropdown menu), 'Path' (a text input), 'Access Level' (a dropdown menu with 'share' selected), and 'Share Name' (a text input). A 'Note' indicates that fields with an asterisk are required. At the bottom right are 'Submit' and 'Cancel' buttons.

- Step 6** From the Website Name drop-down list, choose the websites for which you wish to limit access to pre-positioned cdnfs content. The websites must have been previously defined in the Creating Web Site window.
- Step 7** In the Path field, enter the first-level component of the selected websites. This component will be combined with the FQDN of the origin server specified in the Creating Web Site window. The base name of the path or the first-level component of a particular website must be unique. For example, if the path is entered as `www.domain.com/subdomain`, *subdomain* refers to the base name of the path and must not be the same for two websites.
- Step 8** From the Access Level drop-down list, choose the level of access to be allowed for the websites. Choose **share** if you wish to share protected cdnfs content. Otherwise, choose **protect** to deny access to protected cdnfs content.
- Step 9** To save the settings, click **Submit**.

Customizing the CIFS Share Name

ACNS 5.5 software allows you to change the CIFS share name by using the **share-name** option of the **network-file-system server cifs share-web-site** global configuration command.



Note

This option first became available in the ACNS 5.3 software release. You cannot use the **share-name** option in ACNS 5.0, 5.1, and 5.2 software.

The command syntax is as follows:

```
network-filesystem server cifs share-web-site website [{protect-auth-content |  
share-auth-content}] [share-name sharename]
```

The **share-name** option is shown in the following example:

```
CE(config)# network-filesystem server cifs share-web-site www.cisco.com share-name cco
```



Note

Share names must be limited to 64 characters or less and may contain only alphanumeric characters. The special characters permitted are “-”, “_”, and “.”.

Working with Content Routers

You can work with Content Routers by performing these tasks:

- [Modifying Content Router Properties, page 13-22](#)
- [Deleting a Content Router, page 13-24](#)

Modifying Content Router Properties

You use the Devices tab in the Content Distribution Manager GUI to make changes to a Content Router.

You can modify the following Content Router properties:

- Location
- Coverage zone file
- Description

To modify a Content Router, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears, listing all the Content Routers registered in your ACNS network.
 - Step 2** Click the **Edit** icon next to the name of the Content Router that you want to modify. The Content Router Device Home window appears.
 - Step 3** In the Contents pane, choose **Device Activation**. The Modifying Content Router window appears. (See [Figure 13-6](#).)

Figure 13-6 **Modifying Content Router Window**

The screenshot shows the Cisco ACNS web interface in Microsoft Internet Explorer. The main title is 'Cisco Application and Content Networking System'. The 'System Status' bar indicates '3 Devices, Critical' and '17 Channels, Minor'. The navigation menu includes 'Devices', 'Device Groups', 'Locations', and 'Statistics'. The 'Contents' sidebar lists 'Device Home', 'Device Activation', 'Offline Operation', 'General Settings', and 'Monitoring'. The main content area is titled 'Modifying Content Router, ContentRouter' and contains the following sections:

- General Configuration:**
 - Name: ContentRouter
 - Location: ContentEngine-location (dropdown)
 - Status: Online
 - Activate: ☒
 - Dynamic Content Routing: ☐
- Routing Configuration:**
 - Coverage Zone File: None (dropdown)
- Comments:**
 - n/a

A note at the bottom states: 'Note: * - Required Field'. At the bottom right are 'Submit' and 'Cancel' buttons.

- Step 4** If the Content Router is in Inactive status, check the **Activate** check box. If the Content Router is in Online status, this check box will be checked by default.
- Step 5** If you want a Content Router in Active status to be inactivated, click the **Deactivate Device** icon in the taskbar.
- Step 6** To change the location, choose a new location from the Locations drop-down list.
- Step 7** To have a local coverage zone file overwrite the ACNS network-wide coverage zone file on this router, you can choose a file from the Coverage Zone File drop-down list. Otherwise, choose **None**.
- The Coverage Zone File drop-down list contains the destination paths of all valid coverage zone files registered to your ACNS network. When you choose a coverage zone file from this list, the Content Router uses this coverage zone file to redirect the user's request. (For information about registering and applying coverage zones, see the [“Registering and Applying a User-Defined Coverage Zone”](#) section on page 4-44.)
- Step 8** To modify the description used to identify the Content Router, enter a new description in the Comments field.
- Step 9** To save your settings, click **Submit**. A message indicating that the transaction was successfully completed appears at the bottom of the window.

Deleting a Content Router

You can delete a Content Router if the device is experiencing unresolvable problems. If you delete the only Content Router in your ACNS network, you are removing the ability of your ACNS network to fill user requests unless you have WCCP or proxy routing configured on your ACNS network.

Removing the node from the ACNS network involves using the CLI to shut down ACNS network services and deregister the node. If you are removing the node because of hardware failure and it cannot be accessed through its CLI, you can remove the node by using the Content Distribution Manager GUI; however, the node continues to store its registration information until you deregister it by using the CLI.

To remove and deregister a Content Router, follow these steps:

Step 1 Open a Telnet session to the Content Router CLI.

Step 2 Enter the **no cms enable** command in global configuration mode.

```
CR# configure
CR(config)# no cms enable
```



Note Issuing the **no cms enable** command does not disable acquisition and distribution services on the device; however, issuing the **cms deregister** command does. The **cms deregister** command disables the CMS, all acquisition and distribution services, and all routing communications to and from this device.

Step 3 Enter the **cms deregister** command in EXEC mode.

```
CR(config)# exit
CR# cms deregister
```



Note The **cms deregister** command cleans up the database automatically. You do not need to use the **cms database delete** command.

Step 4 If for some reason the deregistration fails, you can force the deregistration by using the **cms deregister force** command.

```
CR# cms deregister force
```



Note Ignore any messages stating that the deregistration failed. The **cms deregister force** command forces the deregistration to continue.

Step 5 To add the device back in to the ACNS network, reregister the device with the Content Distribution Manager by using the **cms enable** command in global configuration mode.

```
CR# configure
CR(config)# cms enable
```

In case of a hardware failure, you might need to remove the node from the ACNS network routing scheme by using the Content Distribution Manager GUI.

Before a node can be removed from the ACNS network through the Content Distribution Manager GUI, the following conditions must be met:

- The device must have been activated in the Content Distribution Manager. (See the [“Activating Devices in the Content Distribution Manager GUI”](#) section on page 2-22.)
- The Content Distribution Manager must be running.
- The device must have the correct Content Distribution Manager IP address or host name configured.
- The Content Distribution Manager IP address or host name must point to the primary Content Distribution Manager.

To delete a Content Router using the Content Distribution Manager GUI, follow these steps:

-
- Step 1** From the Cisco Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears, listing the Content Routers in your ACNS network. The online status of the device is listed in the Status column.
- Step 2** Click the **Edit** icon next to the name of the Content Router that you wish to delete. The Content Router Device Home window appears.
- Step 3** In the taskbar, click the **Delete Device (Trash)** icon. You are prompted to confirm your decision.
- Step 4** To execute your request, click **OK**. The Content Router is removed from the Content Distribution Manager GUI.
- Step 5** If possible, access the Content Router CLI to deregister the Content Router.
- Step 6** In the CLI, enter the **cms deregister force** command.



Note You must use the **cms deregister force** command after deleting a device in the Content Distribution Manager GUI because once the device has been deleted, the Content Distribution Manager no longer has a record of the device.

- Step 7** To add the device back in to the ACNS network, reregister the device with the Content Distribution Manager by using the **cms enable** command in global configuration mode.
-

Configuring Services on Content Routers

The process of configuring services on Content Routers differs from that of configuring services on Content Engines and device groups in that only platform configuration and authentication services need to be enabled on Content Routers.

The configuration windows for these features are the same windows as those used for Content Engines and device groups. See the following sections and pages for configuration information.

For platform configuration, see the following sections:

- [Configuring NTP Settings, page 20-11](#)
- [Configuring CDP Settings, page 20-7](#)

For authentication and authorization, see the following sections:

- [Configuring Devices for Login Authentication and Authorization, page 12-7](#)
- [Configuring Request Authentication on Centrally Managed Content Engines, page 15-17](#)

For more information regarding the use of Content Routers in the ACNS network environment, see [Chapter 4, “Setting Up Content Request Routing in the ACNS Network.”](#)

Working with Device Groups

A device group is a set of similar devices (such as Content Engines) that share common qualities and capabilities. Some common qualities might include disk capacity, distribution minimum bandwidth, or routing properties.

Device groups are assigned to channels through the Content Distribution Manager GUI. Whenever a channel is created and additional device groups are added, or a channel assignment to the device group changes, devices in the group are notified of the changes and assignments.

Using the device group feature of the Content Distribution Manager, you can create a new device group definition, modify information about a device group, or remove a device group from the system.

You can work with device groups by performing the following tasks:


- [Creating a Device Group, page 13-26](#)
- [Deleting a Device Group, page 13-30](#)
- [Viewing Device Groups, page 13-30](#)

Creating a Device Group

To create or modify a device group, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**. The Device Groups window appears. (See [Figure 13-7](#).)

Figure 13-7 Device Groups Window



Step 2 In the taskbar, click the **Create New Device Group** icon. The Creating New Device Group window appears. (See [Figure 13-8](#).)

Figure 13-8 **Creating New Device Group Window**

The screenshot shows the Cisco Application and Content Networking System web interface. The main navigation bar includes 'Devices', 'Services', and 'System'. The 'Device Groups' tab is selected. The 'Creating New Device Group' window is displayed, featuring the following elements:

- Device Group Information:**
 - Name:** A required text field (indicated by an asterisk).
 - Device Type:** Set to 'Content Engine'.
 - Automatically assign all newly activated devices to this group:** An unchecked checkbox.
 - Group Role:**
 - Regular Group:** Selected with a radio button.
 - Baseline Group:** An unselected radio button.
 - Web, Video, Platform:** Three unchecked checkboxes.
 - Pages configured for this device group:** A link with a right-pointing arrow.
 - Select pages to hide from table of contents for this device group:** A link with a right-pointing arrow.
- Comments:** A text area for optional comments.
- Note:** A small note indicating that fields with an asterisk are required.
- Buttons:** 'Submit' and 'Cancel' buttons at the bottom right.

- Step 3** In the Name field, enter the name of the device group.
- The name must be unique and should be a name that is useful in distinguishing the device group from others on your system.
- Step 4** To set the device group as the default device group for all newly activated devices, check the **Automatically assign all newly activated devices to this group** check box.
- Step 5** To define a regular device group, click the **Regular Group** radio button.
- (To define a baseline group, see the [“Configuring a Baseline Group for a Service”](#) section on page 13-31.)
- Step 6** To view configuration information about an existing device group, click the **Pages configured for this device group** arrow. The window displays a list of GUI pages that are configured for that device group.
- Step 7** To customize the Contents pane for the device group, click the **Select pages to hide from table of contents for this device group** arrow, and check the windows that you want to hide.

**Note**

Use this feature to remove from view any configuration windows that you do not need for that particular device group. Hiding a configured window to keep it from showing in the Contents pane does not affect the settings that are configured already.

- Step 8** You may enter comments about the group in the Comments field. This field is optional.

**Note**

If you want to modify the list of Content Engines assigned to the device group, see the [“Assigning Content Engines to the Device Group” section on page 13-29](#).

- Step 9** To save the settings, click **Submit**.
- Step 10** To modify a device group, go to the Device Groups window, and click the **Edit** icon next to the name of the device group that you want to modify. The Modifying Device Group window appears.
- Step 11** Change the device group properties as needed, and click **Submit**.

Assigning Content Engines to the Device Group

**Note**

In ACNS 5.x software, device groups support only Content Engines. Content Routers cannot be assigned to device groups.

To add or remove Content Engines from a device group, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**.
- Step 2** Click the **Edit** icon next to the name of the device group that you want to edit. The Modifying Device Group window appears.
- Step 3** In the Contents pane, choose **Assignments > Devices**. The Content Engine assignment to DG window appears, displaying the Content Engines assigned to various locations.
- Step 4** Click the **Assign** icon (blue cross mark) next to the name of the Content Engine that you wish to add to the device group. Alternatively, to perform a bulk addition of all Content Engines in various locations to the device group, click the **Assign all Content Engines** icon in the taskbar.
- Step 5** To add the Content Engines to the device group, click **Submit**. Upon submission, a green tick mark appears next to the assigned Content Engine.
- Step 6** To remove Content Engines, click the **Unassign** icon (green tick mark) next to the name of the Content Engine that you want to remove from the device group. Alternatively, you can click the **Unassign all Content Engines** icon in the taskbar to perform a bulk removal of all Content Engines from the selected device group.
- Step 7** To remove the Content Engines from the device group that you are working on, click **Submit**. The list of Content Engines in the device group shows a blue cross mark that appears next to the name of the removed Content Engine.
- Step 8** To display a subset of Content Engines assigned to device groups, click the **Filter Table** icon in the taskbar. You can choose to filter Content Engines by name, location, and status. After setting the filtering criteria, click **Submit**.
- Step 9** To revert to the display of all Content Engines in the ACNS network, click the **View All Content Engines** icon in the taskbar.

Enabling Device Group Overlap

If you want the ability to assign a device to more than one device group, you must enable device group overlap. Content Engines can be assigned to multiple device groups when device group overlap is enabled in the Content Distribution Manager GUI.

To enable device group overlap, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **System > Configuration**. The Config Properties window appears.
 - Step 2** Click the **Edit** icon next to the property name DeviceGroup.overlap. The Modifying Config Property, *DeviceGroup.overlap* window appears.
 - Step 3** From the **Value** drop-down list, choose **true**. (The default is **false**.)

You can return to the default at any time by choosing **false** from the **Value** drop-down list. When you disable device group overlap, existing overlapping device groups are retained and continue to be handled as though overlap were enabled; however, any newly added groups do not allow overlapping, and new devices cannot be added to the existing overlapping groups.
 - Step 4** To save the settings, click **Submit**.
-

Deleting a Device Group

To delete a device group, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**. The Device Groups window appears. (See [Figure 13-7](#).)
 - Step 2** Click the **Edit** icon next to the name of the device group that you want to delete. The Modifying Device Group window appears.
 - Step 3** In the taskbar, click the **Trash** icon. You are prompted to confirm your decision to delete the device group.
 - Step 4** To confirm your decision, click **OK**. The device group is removed from the ACNS network.
-

Viewing Device Groups

To view all the device groups in your ACNS network, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**. The Device Groups window appears. (See [Figure 13-7](#).)
 - Step 2** To refresh the table, click the **View All Device Groups** icon in the taskbar.

The table also displays device type and comments. You can sort the columns in ascending or descending order by clicking the column heading.

Step 3 To print the table data, click the **Printer** icon in the taskbar.

Working with Baseline Groups

A baseline group is a special type of device group used for configuring a particular service, such as web, video, or platform. Cisco ACNS software has three predefined baseline groups for your convenience, one for each ACNS service area:

- WebBaselineGroup
- VideoBaselineGroup
- PlatformBaselineGroup

You can use these predefined baseline groups to configure the elements that are necessary for setting up the particular service. When a baseline group is configured for a service, all the devices in the group can be configured for that service and share the same baseline settings. You can associate a baseline group with one or more services. However, a particular service can have only one baseline group associated with it at any given time.

A baseline group is configured in the same way and behaves in the same manner as a device group. You configure the baseline group and then assign Content Engines to the group. Next, you can configure or modify service settings through the baseline service group and apply these settings to all the Content Engines that are associated with the group.

You can also switch the baseline group that is associated with a particular service, if you need to, so that, for example, you can run a particular service on a different subnet of Content Engines in your network and use service settings, such as bandwidth or request filtering, that meet the requirements of that segment of your network.

When you are working with baseline groups, you can perform the following tasks:

- [Configuring a Baseline Group for a Service, page 13-31](#)
- [Assigning Content Engines to the Baseline Group, page 13-33](#)
- [Configuring the Service Settings for the Baseline Group, page 13-34](#)
- [Forcing the Baseline Group Service Settings on All Content Engines in the Baseline Group, page 13-34](#)
- [Switching the Baseline Group for a Service, page 13-35](#)
- [Viewing Configured Service Settings for a Device Group, page 13-35](#)

Configuring a Baseline Group for a Service

Because ACNS software contains three preconfigured baseline groups, before you can configure a new one, you must first delete the existing baseline group. (See the [“Deleting a Device Group”](#) section on [page 13-30](#).)

The ACNS 5.5 Content Distribution Manager GUI provides several navigation paths to the baseline group configuration window. You can configure a new baseline group from a service Baseline Settings configuration window in the Services tab or from a Device Groups configuration window in the Devices tab.

To configure a baseline group from a service configuration window, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Web** or **Video** or **Platform > Baseline Settings**. The Selecting the (Web, Video, or Platform) Baseline Group window for the chosen service appears.
- If a baseline group is already configured for the service, the Modify Device Group Settings window appears. If you want to delete the baseline group and configure a new one, click the **Trash** icon in the taskbar.
 - If there is no baseline group configured for the service, the Content Distribution Manager displays a message that there is no baseline group configured for that service and provides a drop-down list from which you can choose an existing device group as the baseline group for the service, or you can create a new one.
- Step 2** Choose a device group, or choose the **Create New Device Group** option from the drop-down list, and click **Submit**.
- If you choose an existing group as baseline group for the service, the Modify Device Group Settings window for the chosen device group opens. From the Modify Device Group Settings window, you can choose to modify the setting options, or leave them alone. If you make changes to any of the settings, click **Submit** to save the changes.
 - If you choose to create a new device group, the Create New Device Group window opens.
- Step 3** In the Name field, enter the name of the baseline group.
- The name must be unique and should be a name that is useful in distinguishing the baseline group from others on your system. Names with spaces are invalid.
- Step 4** To set the baseline group as the default device group for all newly activated devices, check the **Automatically assign all newly activated devices to this group** check box.
- Step 5** To define a baseline group, click the **Baseline Group** radio button, and choose a service area by checking either the **Web**, **Video**, or **Platform** check box.
-
-  **Note** Once a baseline group has been chosen for a service, that selection is no longer available.
-
- Step 6** To customize the Contents pane for the baseline group, click the **Select pages to hide from table of contents for this device group** arrow, and check the check boxes for the windows that you want to hide.
-
-  **Note** Use this feature to remove from view any configuration windows that you do not need for that particular device group. Hiding a configured window to keep it from showing in the Contents pane does not affect the settings that are configured already.
-
- Step 7** To give the baseline group a description, enter comments about the group in the Comments field. This field is optional.
- Step 8** To save the settings, click **Submit**.
- After you click **Submit**, the Modify Device Group Settings window appears. The baseline group has been configured. Proceed to the next section, [“Assigning Content Engines to the Baseline Group.”](#)
-

**Note**

To configure a baseline group from a Device Groups configuration window, choose **Devices > Device Groups**. Then click the **Create New Device Group** icon in the taskbar. The Creating New Device Group window appears. (See [Figure 13-8](#).) Proceed as described in the preceding steps.

Assigning Content Engines to the Baseline Group

**Note**

In ACNS 5.x software, device groups support only Content Engines. Content Routers cannot be assigned to device groups.

To add or remove Content Engines from a baseline group, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**.
- Step 2** Click the **Edit** icon next to the name of the baseline group that you want to edit. The Modifying Device Group window appears.
- Step 3** In the Contents pane, choose **Assignments > Devices**. The Content Engine assignment to DG window appears, displaying the Content Engines assigned to various locations.
- Step 4** Click the **Assign** icon (blue cross mark) next to the name of the Content Engine that you wish to add to the baseline group. Alternatively, to perform a bulk addition of all Content Engines in various locations to the baseline group, click the **Assign all Content Engines** icon in the taskbar.
- Step 5** To add the Content Engines to the baseline group, click **Submit**. Upon submission, a green tick mark appears next to the assigned Content Engine.
- Step 6** To remove Content Engines, click the **Unassign** icon (green tick mark) next to the name of the Content Engine that you want to remove from the baseline group. Alternatively, you can click the **Unassign all Content Engines** icon in the taskbar to perform a bulk removal of all Content Engines from the selected baseline group.
- Step 7** To remove the Content Engines from the baseline group that you are working on, click **Submit**. The list of Content Engines in the baseline group shows a blue cross mark that appears next to the name of the removed Content Engine.
- Step 8** To display a subset of Content Engines assigned to baseline groups, click the **Filter Table** icon in the taskbar. You can choose to filter Content Engines by name, location, and status. After setting the filtering criteria, click **Submit**.
- Step 9** To revert to the display of all Content Engines in the ACNS network, click the **View All Content Engines** icon in the taskbar.

You can now configure the service settings for this baseline group. (See the next section, “[Configuring the Service Settings for the Baseline Group](#).”)

Configuring the Service Settings for the Baseline Group

To configure service settings, for a baseline group, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Services > Web** or **Video** or **Platform > Baseline Settings**. The Modifying Device Group window for the chosen service appears.
 - Step 2** Alternatively, choose **Devices > Device Groups**. Click the **Edit** icon next to the name of the baseline group that is associated with the service that you want to configure. The Modifying Device Group window appears.
 - Step 3** To view which GUI windows have already been configured for the baseline group, click the **Pages configured for this device group** arrow button. The GUI displays the list of pages that are configured for that device group. If you are creating a new baseline group or if there are no pages configured for that device group, the GUI displays, Null.
 - Step 4** If you have not already done so, you can hide any configuration windows in the Contents pane that you do not want to include as part of the baseline group service settings. To customize the Contents pane for the baseline group, click the **Select pages to hide from table of contents for this device group** arrow, and check the check boxes for the windows that you want to hide.
 - Step 5** Use the Contents pane to navigate to each configuration window that you want to include as part of the baseline group settings. After a particular setting is configured, the configuration window is listed under **Pages configured for this device group** in the Modifying Device Group window.
-

Forcing the Baseline Group Service Settings on All Content Engines in the Baseline Group

Cisco ACNS software allows you to force the complete set of configurations made for a device group to all devices that are associated with that group.

To force the device group configurations across all devices in a device group, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Device Groups**. The Device Groups listing window appears.
 - Step 2** Click the **Edit** icon of the device group that you want to configure.
 - Step 3** In the taskbar, click the **Force Group Settings** icon.
The Content Distribution Manager returns the following message:
The action will apply all settings configured for this device group to all the Content Engines assigned to it. Do you wish to continue?
 - Step 4** To force group settings across all devices in the device group, click **Yes**.
 - Step 5** To save the changes, click **Submit**.
-

Switching the Baseline Group for a Service

Cisco ACNS software allows you to switch the baseline group that is associated with a particular service, if you need to, so that, for example, you can run a particular service on a different subnet of Content Engines in your network and use service settings, such as bandwidth or request filtering, that meet the requirements of that segment of your network.

When you switch a baseline group, you must choose a regular device group to take its place. During the switch, the regular device group that you choose is converted to a baseline group, and the baseline group that you remove is converted to a regular device group.

To remove a baseline group from a service and associate another baseline group in its place, follow these steps:

-
- Step 1** Choose **Services > Web** or **Video** or **Platform > Baseline Settings**. The Modify Device Group window for the chosen service appears.
- Step 2** From the taskbar, click the **Switch Baseline Group** icon. The Content Distribution Manager returns the following message:
- This action will remove this device group as the Baseline Group for this service. You can then select another device group or create a new one to the Baseline Group for this service. Do you wish to Continue?
- Step 3** To remove the device group from the service, click **OK**. The Content Distribution Manager displays the Selecting the Baseline Group window for the chosen service.
- Step 4** From the **Select a Device Group to be the Baseline Group** drop-down list, choose a device group, or choose the **Create New Device Group** option.
- If you choose a device group to be the baseline group for that service, the Content Distribution Manager takes you to the Modify Device Group window for that device group.
 - If you choose the **Create New Device Group** option, the Content Distribution Manager takes you to the Create New Device Group window.
-

Viewing Configured Service Settings for a Device Group

To view configuration information about an existing device group, click the **Pages configured for this device group** arrow. The window displays a list of GUI pages that are configured for that device group.

Local Device Configuration Versus Central Management

Cisco ACNS 5.x software uses certain CLI commands that are particularly important for ACNS network functionality, which is likely to be managed by using device groups. If you configure any of these commands from the Content Engine or Content Router, either through the Content Engine GUI or CLI, then the settings are *not* stored as part of the Content Distribution Manager's ACNS network-wide configuration data and are likely to be overwritten by the Content Distribution Manager.

Furthermore, if you configured any of these settings on your Content Engine or Content Router before you registered it with your Content Distribution Manager, you must reconfigure these settings in the Content Distribution Manager GUI.

We recommend that you do *not* configure the following centrally managed settings from the local device GUI or CLI:

authentication	logging	tacacs
bandwidth	multicast	rule
bypass	ntlm	transaction-logs
cdp	ntp	multicast accept-license-agreement
dns-cache	proxy-auto-config	wmt
dns enable	proxy-protocols	excluding:
error-handling	ldap	wmt proxy
ftp	radius-server	wmt
http	rtsp	accept-license-agreement
https	excluding:	wmt live-url-stripping
icurl-filter	rtsp server real-subscriber accept-license-agreement	
ldap	rtsp proxy media-real accept-license-agreement rtsp-proxy	

Repurposing an ACNS Network Device

Some hardware models allow you to change the device mode and redeploy the device in your network for a different purpose. For example, you might repurpose a Content Engine as a Content Distribution Manager for failover and redundancy.

The following hardware models support device mode changes:

- CE-7305
- CE-565
- CE-566

To change the device mode of your eligible device, you must be prepared to configure the disk space allocations, as required by the different device modes, and to reboot the device for the new configuration to take effect.

Devices that can be reconfigured using the **device mode** command are shipped from the factory by default as Content Engines. When you change the device mode of a Content Engine to a Content Router or a Content Distribution Manager, you need to configure the sysfs. (See the [“Configuring Disk Space”](#) section on page 2-19.)

If, however, you are changing the device mode of a Content Router or a Content Distribution Manager back to a Content Engine, you must configure disk space allocations for the sysfs, cfs, mediafs, and cdnfs that are used on the Content Engine, as shown in this example:

```
DeviceName# disk config sysfs 10% cfs 20% mediafs 10% cdnfs 60%
```

(See the *Cisco ACNS Software Update and Maintenance Guide*.)

**Note**

Enabling content routing on a Content Engine is *not* considered a device mode change. You do not need to follow this procedure to enable the content routing functionality. (To enable content routing on a Content Engine, see the [“Modifying Content Engine Properties”](#) section on page 13-9.)

To change the device mode of a device that is already operating as part of your ACNS network, follow these steps:

Step 1 Access the device CLI.

Step 2 Enter the **show device-mode current** command to view the current device mode.

```
DeviceName# show device-mode current
Current device mode: content-engine
```

Step 3 Configure the new device mode by using the **device mode** command in global configuration mode, as shown in this example:

```
DeviceName# configure
DeviceName(config)# device mode content-router
```

Step 4 Deregister the device by using the **cms deregister force** command.

```
DeviceName(config)# cms deregister force
```

Step 5 Save the configuration by issuing the **write memory** command or the **copy running-config startup-config** command.

```
DeviceName# write memory
DeviceName# copy running-config startup-config
```

Step 6 Use the **restore factory-default preserve basic-config** command to reload the software and apply the device mode configuration.

```
DeviceName# restore factory-default preserve basic-config
```

**Note**

This command erases all existing content on the device; however, your network settings are preserved, and the device is accessible through Telnet and Secure Shell (SSH) after it reboots.

Step 7 Verify the new configuration. Check that the current and configured device modes are the same, as shown in this example:

```
DeviceName# show device-mode configured
Configured Device mode: content-router
DeviceName# show device-mode current
Current device mode: content router
```

Step 8 Configure the disk space allocation based on the new device mode and needs of your network by using the **disk config** command.

- Step 9** Follow the directions for configuring, registering, and activating new devices in [Chapter 2, “Getting Started.”](#)
-



Configuring Network Interfaces

This chapter describes the steps for configuring additional network interfaces to support network traffic. During initial setup you chose an initial interface and either configured it for DHCP or gave it a static IP address. (See the [“About Selecting Static IP Addresses or Using Interface-Level DHCP”](#) section on page 2-12.) This chapter describes how to configure additional interfaces using options for redundancy, load balancing, performance optimization, and so forth.

This chapter contains the following sections:

- [Configuring Multiple Network Interfaces and Standby Interfaces, page 14-1](#)
- [Configuring Multiple IP Addresses on a Single Interface, page 14-6](#)
- [Modifying Fast Ethernet or Gigabit Ethernet Interface Settings, page 14-7](#)
- [Configuring the Fibre Channel Interface, page 14-9](#)
- [Configuring EtherChannel, page 14-11](#)
- [Configuring Interfaces for DHCP, page 14-13](#)
- [Configuring IP Addresses for Network Address Translation, page 14-14](#)
- [Configuring Static IP Routes, page 14-15](#)

Configuring Multiple Network Interfaces and Standby Interfaces

You can configure multiple network interfaces as either active-active interfaces or as active-standby interfaces. You can configure multiple interfaces as active-active by using the **interface** global configuration command and by assigning an IP address to each interface. When multiple interfaces are configured, they are active simultaneously. This configuration is used to achieve better performance.

```
ContentEngine# configure
ContentEngine(config)# interface FastEthernet 0/0
ContentEngine(config-if)# ip address 10.10.10.10 255.0.0.0
```

In an active-standby configuration, the interfaces are configured by using the **standby** command, and they remain inactive unless an active interface fails. When an active network interface fails (because of cable trouble, Layer 2 switch failure, high error count, or other failure), and that interface is part of a standby group, a standby interface can become active and take the load off the failed interface. With active-standby interface configuration, only one interface is active at a given time. Active-standby is used mainly for fault tolerance purposes.

```
ContentEngine# configure
```



```

ContentEngine(config)# interface FastEthernet 0/1
ContentEngine(config-if)# standby ?
    <1-4> Standby group number
ContentEngine(config-if)# standby 3 ?
    errors    Set the maximum number of errors allowed on this interface
    ip        Set the IP address of a standby group
    priority  Set the priority of an interface for the standby group
ContentEngine(config-if)# standby 3 errors ?
    <0-4294967295> Max. no. of errors allowed on this interface for the standby
                    group
ContentEngine(config-if)# standby 3 ip ?
    A.B.C.D   IP address of the standby group
ContentEngine(config-if)# standby 3 priority ?
    <0-4294967295> Priority of this interface for the standby group

```

To configure standby interfaces, you must assign each physical interface to a standby group. The following rules define standby group relationships:

- A standby group comprises two or more physical interfaces.
- The maximum number of standby groups on a Content Engine is four.
- Each interface is assigned a unique IP address, and each standby group is assigned a unique standby IP address, shared by all members of the group.
- Configuring the duplex and speed settings of the standby group member interfaces provides better reliability.
- IP ACLs can be configured on physical interfaces that are members of a standby group.
- Each interface in a standby group is assigned a priority. The operational interface with the highest priority in a standby group is the active interface. Only the active interface uses the group IP address.
- If the active interface fails, the operational interface in its standby group that is assigned the next highest priority becomes active. However, when the interface with the higher priority recovers, it does not become active again without manual intervention.
- If all the members of a standby group fail and then one recovers, the ACNS software brings up the standby group on the operational interface.
- The priority of an interface in a standby group can be changed at runtime. The interface that has the highest priority after this change becomes the new active interface. (The default action is to preempt the currently active interface if an interface with higher priority exists.)
- The maximum number of errors allowed on the active interface before the interface is shut down and the standby is brought up is configured with the **errors** option, which is disabled by default.

**Note**

Interface IP addresses and standby group IP addresses must be on different subnets to ensure reliable operation. You can use dummy IP addresses in the private address space to serve as interface primary IP addresses and use the real Content Engine IP address to serve as the standby group IP address in a different subnet to satisfy this requirement.

**Note**

Make sure to configure the default gateway by using the **ip default-gateway** global configuration command instead of the **ip route** global configuration command.

Example

This example configures three interfaces to be part of the same standby group, with interface 3/0 as the active interface.

```
Console(config)# interface fastEthernet 3/0 standby 1 ip 172.16.10.10 255.255.254.0
```

```

Console(config)# interface fastEthernet 3/1 standby 1 ip 172.16.10.10 255.255.254.0
Console(config)# interface fastEthernet 3/2 standby 1 ip 172.16.10.10 255.255.254.0

Console(config)# interface fastEthernet 3/0 standby 1 priority 300

Console(config)# interface fastEthernet 3/1 standby 1 priority 200

Console(config)# interface fastEthernet 3/2 standby 1 priority 100

Console(config)# interface fastEthernet 3/0 standby 1 errors 10000

Console(config)# interface fastEthernet 3/1 standby 1 errors 10000

Console(config)# interface fastEthernet 3/2 standby 1 errors 10000

```

Use the **show standby EXEC** command to view your standby interface configuration.

```

Console# show standby
Standby Group:1
IP address: 172.16.10.10, netmask: 255.255.254.0
Maximum errors allowed on the active interface: 10000
  Member interfaces:
    FastEthernet 3/0      priority: 300
    FastEthernet 3/1      priority: 200
    FastEthernet 3/2      priority: 100
Active interface: FastEthernet 3/0

```

Configuring a Standby Interface Using the Content Distribution Manager GUI

In this procedure, you are configuring a logical interface called a standby interface. After you set up the parameters for this logical interface, you must associate physical interfaces with the standby interface to create the standby group. (A standby group consists of two or more physical interfaces.) In the Content Distribution Manager GUI, you create the standby group by assigning a standby group priority to the physical interface. (See the next section, “[Configuring the Interface Priority Setting](#).”)

To configure a standby interface using the Content Distribution Manager GUI, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the device for which you want to configure a standby interface. The Device Home window appears.
- Step 3** To display the entire table of contents, click the **Show All** button above the Contents pane.
- Step 4** In the Contents pane, choose **General Settings > Network > Network Interfaces**. The Network Interfaces window for the device appears.
- Step 5** In the taskbar, click the **Create New Interface** icon. The Creating New Network Interface window appears. [Table 14-1](#) describes the fields in this window and provides the corresponding CLI global configuration commands.
- Step 6** From the Port Type drop-down list, choose **Standby**. The window refreshes with fields for configuring the standby group settings.
- Step 7** From the Standby Group Number drop-down list, choose a group number (1–4).
- Step 8** In the Address field, specify the IP address of the standby group.
- Step 9** In the Netmask field, specify the netmask of the standby group.

- Step 10** In the Standby Group Number of Errors field, enter the maximum number of errors allowed on this interface. The range is 0 to 4294967295.
- Step 11** In the Gateway field enter the default gateway IP address. If an interface is configured for DHCP, then this field is read only.
- Step 12** To save the settings, click **Submit**.

Table 14-1 Standby Interface Settings

GUI Parameter	Function	CLI Command
Port Type	Makes settings available for configuring logical interfaces (PortChannel or Standby).	interface { FastEthernet GigbitEthernet } <i>slot/port</i> standby
Standby Group Number	Group number (1–4) for the standby interface.	interface { FastEthernet GigbitEthernet } <i>slot/port</i> standby <i>groupnum</i>
Address	IP address of the standby group.	interface { FastEthernet GigbitEthernet } <i>slot/port</i> standby <i>groupnum</i> ip <i>ipaddress</i>
Netmask	Netmask of the standby group.	interface { FastEthernet GigbitEthernet } <i>slot/port</i> standby <i>groupnum</i> ip <i>ipaddress</i> <i>netmask</i>
Standby Group Number of Errors	Maximum number of errors allowed on this interface.	interface { FastEthernet GigbitEthernet } <i>slot/port</i> standby <i>groupnum</i> errors <i>number</i>
Gateway	Default gateway IP address.	ip default-gateway <i>ipaddress</i>

Configuring the Interface Priority Setting

Once you have configured a logical standby interface by using the Content Distribution manager GUI, you then configure the standby group by setting a priority for each physical interface that you want to be associated with that standby group. The interface priority setting defines the active interface in a particular standby group and the order in which other interfaces in the standby group will become active if the active interface fails. The operational interface with the highest priority in a standby group is the active interface. Only the active interface uses the standby group IP address. You must have a standby interface configured before you can enter the priority settings in the Content Distribution Manager GUI. (See the previous section, “[Configuring a Standby Interface Using the Content Distribution Manager GUI](#).”)

To configure the priority of the interface and associate it with a particular standby group, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the device for which you want to configure a standby interface. The Device Home window appears.
- Step 3** In the Contents pane, choose **General Settings > Network > Network Interfaces**. The Network Interfaces window for the device appears.
- Step 4** Click the **Edit** icon next to the name of a physical interface (FastEthernet or GigabitEthernet) to which you want to assign a standby priority. The Modifying Network Interface window appears. (See Figure 14-1.)



Note Do not choose a logical interface (standby or portchannel) in this step. You cannot assign a standby priority to a logical interface.

Figure 14-1 Modifying Network Interface Window—Standby Group Priority Settings

The screenshot shows the 'Modifying Network Interface' window for FastEthernet 0/0/0. The 'Contents' pane on the left lists various configuration options, with 'Network Interfaces' selected. The main configuration area includes fields for AutoSense (checked), Speed (10), Mode (half-duplex), MTU (1500 bytes), Address (10.1.1.21), Netmask (255.255.255.0), and four secondary address/netmask pairs. The 'Use DHCP' checkbox is unchecked, and the 'Gateway' is set to 10.1.1.50. The 'Join Standby Group' section is expanded, showing four groups with priority fields. The 'Inbound ACL' and 'Outbound ACL' are both set to 'Do Not Set'. A 'Note' at the bottom indicates that fields with an asterisk are required.

- Step 5** In the Standby Group Priority fields, specify a priority level number (0–4294967295) to set the priority of the interface for the standby group (1–4) with which the interface is to be associated.

A Standby Group Priority field only becomes available when you have previously configured that standby group. You can configure up to four standby groups for each Content Engine.

**Tip**

If an interface belongs to more than one standby group, you can configure the interface with a different priority in each standby group for better load balancing. For example, interfaces FE 0/0 and FE 0/1 are both in standby group 1 and in standby group 2. If you configure FE 0/0 with the highest priority in standby group 1 and configure FE 0/1 with the highest priority in standby group 2, standby group 1 will use FE 0/0 as the active interface, while standby group 2 will use FE 0/1 as the active interface. This configuration allows each interface to back up the other one, if one of them fails.

Step 6 To save the settings, click **Submit**.

To configure the priority of the interface for a particular standby group from the CLI, use the following global configuration command:

```
interface {FastEthernet | GigabitEthernet} slot/port standby groupnum priority priority
```

Configuring Multiple IP Addresses on a Single Interface

You can configure more than one IP address on the same interface by using the **interface secondary** interface configuration command. You can configure up to four secondary IP addresses on a single interface. This configuration allows the device to be present in more than one subnet and can be used to optimize response time, because it allows the content to go directly from the Content Engine to the client that is requesting the information without being redirected through a router. The Content Engine becomes visible to the client because both are configured on the same subnet, as shown in this example:

```
ContentEngine# configure
ContentEngine(config)# interface FastEthernet 0/0
ContentEngine(config-if)# ip address 10.10.10.10 255.0.0.0 secondary
```

To configure multiple IP addresses on a single interface using the Content Distribution Manager GUI, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the device for which you want to configure interface settings. The Device Home window appears.
- Step 3** In the Contents pane, choose **General Settings > Network > Network Interfaces**. The Network Interfaces listing window appears.
- Step 4** Click the **Edit** icon for the physical interface (FastEthernet or GigabitEthernet) that you want to modify. The Modifying Network Interface window appears. (See [Figure 14-2](#).)

**Note**

Do not choose a logical interface (standby or portchannel) in this step. You cannot configure multiple interfaces on a logical interface.

Figure 14-2 Modifying Network Interface Window

The screenshot shows the Cisco ACNS web interface. The left sidebar contains a 'Contents' menu with options like Device Home, Device Activation, Assignments, Prepositioning, Request Routing, Request Processing, Applications, General Settings, Login Access Control, Authentication, Services, Storage, Network, Network Interfaces, External IP, Port Channel Settings, TCP, IP General Settings, IP Routes, IP ACL, IP ACL Feature Usage, Notification and Tracking, Troubleshooting, and Device Monitoring. The top navigation bar has tabs for 'Devices', 'Services', and 'System'. The main configuration area is titled 'Modifying Network Interface, FastEthernet 0/0 for Content Engine, CONTENTENGINE'. It contains the following fields:

- Slot: 0
- Port: 0
- Port Type: FastEthernet
- Port Channel Number:
- Description: This is the primary interface
- Use CDP: ☒
- Shutdown: ☐
- AutoSense: ☒
- Speed: 10
- Mode: half-duplex
- MTU: 1500 bytes
- Address: 10.1.1.21
- Netmask: 255.255.255.0
- Secondary Address 1:
- Secondary Netmask 1:
- Secondary Address 2:
- Secondary Netmask 2:
- Secondary Address 3:
- Secondary Netmask 3:
- Secondary Address 4:
- Secondary Netmask 4:
- Use DHCP: ☐
- Gateway: 10.1.1.50

At the bottom right, there are 'Submit' and 'Cancel' buttons.

- Step 5** In the Secondary Address and Secondary Netmask fields 1 through 4, you can enter up to four different IP addresses and secondary netmasks for the interface.
- Step 6** To save the settings, click **Submit**.

Modifying Fast Ethernet or Gigabit Ethernet Interface Settings



Note

Fast Ethernet, Gigabit Ethernet, and Fibre Channel interfaces can only be modified using the Content Distribution Manager GUI.

To modify the settings of an existing Fast Ethernet or Gigabit Ethernet interface for the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to modify the Fast Ethernet or Gigabit Ethernet interface settings. The Device Home for Content Engine window appears with the Contents pane on the left.

- Step 3** In the Contents pane, choose **General Settings > Network > Network Interfaces**. The Network Interfaces for Content Engine window appears, listing the network interfaces configured on particular slots and ports.
- Step 4** Click the **Edit Network Interface** icon next to the Fast Ethernet or Gigabit Ethernet interface that you want to modify. The Modifying Network Interface for Content Engine window appears, displaying the interface configurations on a particular slot and port. [Table 14-2](#) describes the configurable fields in this window and provides the corresponding CLI interface configuration commands.



Note Some of the fields in the window are not available. Interface configurations for slot, port, and port type are set for physical interfaces during initial startup or by using the **interface** configuration command on the Content Engine. The port channel number can be configured for a port channel interface when you create this type of logical interface in the Content Distribution Manager GUI; however, this field is not available when you are modifying a physical interface. (See [Configuring Port Channel Settings Using the Content Distribution Manager GUI](#), page 14-11.)

- Step 5** To enable Cisco Discovery Protocol (CDP) on an interface, check the **Use CDP** check box. When enabled, CDP obtains protocol addresses of neighboring devices and discovers the platform of those devices. It also shows information about the interfaces used by your router.
- Using the **cdp enable** command or configuring CDP from the CDP Settings window enables CDP globally on all the interfaces. If you want to control CDP behavior per interface, then use the interface-level **cdp enable** command or enable CDP from the Modifying Network Interface window. Note that the interface-level control overrides the global control.
- Step 6** To shut down the hardware interface, check the **Shutdown** check box.
- Step 7** To set the interface to autonegotiate the speed and mode, check the **AutoSense** check box. Checking this check box disables the manual Speed and Mode drop-down list settings.



Note Cisco router Ethernet interfaces do not negotiate duplex settings. If the Content Engine is connected to a router directly with a crossover cable, the Content Engine interface must be manually set to match the router interface settings. Disable autosense before configuring an Ethernet interface. When autosense is on, manual configurations are overridden. You must reboot the Content Engine to start autosensing.

- Step 8** To manually configure the interface transmission speed and mode settings, follow these steps:
- Uncheck the **AutoSense** check box.
 - From the Speed drop-down list, choose a transmission speed (**10**, **100**, or **1000** Mbps).
 - From the Mode drop-down list, choose a transmission mode (**full-duplex** or **half-duplex**).
- Full duplex transmission allows data to travel in both directions at the same time through an interface or a cable. A half-duplex setting ensures that data only travels in one direction at any given time. Although full duplex is faster, the interfaces sometimes cannot operate effectively in this mode. If you encounter excessive collisions or network errors, configure the interface for half-duplex rather than full duplex.
- Step 9** To set the interface Maximum Transmission Unit (MTU) size, specify a value (in bytes) in the MTU field. The range is 68–1500 bytes. The MTU is the largest size of IP datagram that can be transferred using a specific data link connection.
- Step 10** To change the interface IP address, enter a new IP address in the Address field.

Step 11 To change the interface netmask, enter a new netmask in the Netmask field.

Step 12 To save any network interface configuration changes, click **Submit**.

Table 14-2 *Network Interface Settings*

GUI Parameter	Function	CLI Command
Use CDP	Enables CDP on the individual interface.	interface { FastEthernet GigabitEthernet } <i>slot/port</i> cdp enable
Shutdown	Shuts down the hardware interface.	interface { FastEthernet GigabitEthernet } <i>slot/port</i> shutdown
Autosense	Enables autonegotiation of transmission speed and mode.	interface { FastEthernet GigabitEthernet } <i>slot/port</i> autosense
Speed	Manually configures the interface transmission speed when autosense is disabled.	interface { FastEthernet GigabitEthernet } <i>slot/port</i> bandwidth { 10 100 1000 }
Mode	Manually configures the interface transmission mode when autosense is disabled.	interface { FastEthernet GigabitEthernet } <i>slot/port</i> [full-duplex half-duplex]
MTU	Largest size of IP datagram that can be transferred using a specific data link connection.	interface { FastEthernet GigabitEthernet } <i>slot/port</i> mtu <i>size</i>
Address	IP address of the interface.	interface { FastEthernet GigabitEthernet } <i>slot/port</i> ip address <i>ipaddress netmask</i>
Netmask	Netmask of the interface.	
Gateway	Default gateway IP address.	ip default-gateway <i>ipaddress</i>

Configuring the Fibre Channel Interface

ACNS 5.x software supports Fibre Channel interfaces. Fibre Channel is the chosen technology for interconnecting storage devices and servers in a storage area network (SAN). In a SAN, the storage need not be directly attached to the server, and data transfer happens over a high-throughput, high-availability network. Fibre Channel is capable of operating at speeds of 1 gigabit per second (Gbps) and 2 Gbps.

To detect the presence of Fibre Channel storage, the Fibre Channel array must be configured to assign storage space for the Content Engine, and the Content Engine must be reloaded before it can detect the storage assignment. To confirm whether or not the Content Engine has detected the storage assignment, use the **show disks** and the **show disks details** commands.

To configure the Fibre Channel interface on the Content Engine, use the **interface FibreChannel slot/port** command in interface configuration mode, as shown in this example:

```
DeviceName# configure
DeviceName(config)# interface FibreChannel 0/0
DeviceName(config-if)#?
```



```

exit      Exit from this submode
mode      Change the fibre channel interface operating mode
no        Negate a command or set its defaults
speed     Change the fibre channel interface speed
DeviceName(config-if)# mode ?
  autosense      Use this mode to have the CE autosense
  direct-attached Use this mode when the CE is directly connected to storage array
  switched       Use this mode when the CE is connected to a switch
DeviceName(config-if)# speed ?
  1              1Gbps
  2              2Gbps
  autosense      autosense

```

For a complete description of the **interface FibreChannel** command syntax and usage, refer to the *Cisco ACNS Software Command Reference, Release 5.5* publication.

For information regarding which Fibre Channel storage arrays are supported by Cisco Systems, refer to the *Release Notes for Cisco ACNS Software, Release 5.5*.

Modifying Fibre Channel Interface Settings for the Content Engine



Note

Logical interfaces, such as standby and Port Channel interfaces, can be created, removed, and modified by using the Content Distribution Manager GUI. Physical interfaces, such as Fast Ethernet, Gigabit Ethernet, and Fibre Channel interfaces, can only be modified.

To modify the settings of an existing Fibre Channel interface for the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to modify the Fibre Channel interface settings. The Device Home for Content Engine window appears with the Contents pane on the left.
- Step 3** In the Contents pane, choose **General Settings > Network > Network Interfaces**. The Network Interfaces for Content Engine window appears, listing the network interfaces configured on particular ports and slots.
- Step 4** Click the **Edit Network Interface** icon next to the Fibre Channel interface that you want to modify. The Modifying Network Interface for Content Engine window appears, displaying the interface configurations on a particular port and slot.

The fields in the window that are *not* available are the interface configurations set by using the **interface** configuration command on the Content Engine. [Table 14-3](#) describes these fields.

Table 14-3 Unavailable Fields

GUI Parameter	Function
Slot	Slot number for the selected interface
Port	Port number for the selected interface
Port Type	Type of port <i>Fibre Channel</i> configured on the Content Engine

- Step 5** From the Speed drop-down list, choose the speed of the specified interface. [Table 14-4](#) describes the speed options.

Table 14-4 *Interface Speed Options*

Speed	Description
autosense	Sets the Fibre Channel interface to automatically sense the interface speed.
1	Sets the Fibre Channel interface speed to 1 Gbps.
2	Sets the Fibre Channel interface speed to 2 Gbps.

- Step 6** From the Mode drop-down list, choose the mode of operation of the interface. [Table 14-5](#) describes the mode options:

Table 14-5 *Interface Mode of Operation Options*

Mode	Description
autosense	Sets the operation mode of the Content Engine to autosense.
direct-attached	Sets the operation mode when the Content Engine is directly connected to a storage array.
switched	Sets the operation mode when the Content Engine is connected to a switch.

- Step 7** To save your settings, click **Submit**.

Configuring EtherChannel

EtherChannel for ACNS 5.x software supports the grouping of up to four same-speed network interfaces into one virtual interface. This grouping capability allows the setting or removing of a virtual interface that consists of two, three, or four Fast Ethernet interfaces or two Gigabit Ethernet interfaces. EtherChannel also provides interoperability with Cisco routers, switches, and other networking devices or hosts supporting EtherChannel; load balancing; and automatic failure detection and recovery based on each interface's current link status. To configure an EtherChannel, you use the **PortChannel** interface configuration command, so the EtherChannel is also referred to as a port channel.

Configuring Port Channel Settings Using the Content Distribution Manager GUI

To configure port channel settings using the Content Distribution Manager GUI, follow these steps:

- Step 1** From the Content Distribution Manager GUI choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the device for which you want to configure interfaces. The Device Home window appears.
- Step 3** In the Contents Pane, choose **General Settings > Network > Network Interfaces**. The Network Interfaces window appears, listing all the interfaces for the chosen device.
- Step 4** In the taskbar, click the **Create New Interface** icon. The Creating New Network Interface window appears.

- Step 5** From the Port Type drop-down list, choose **Port Channel**. The window refreshes and provides fields for configuring the network interface settings.
- Step 6** In the Port Channel Number field, enter either **1** or **2** for the port channel interface number.
- Step 7** To shut down this interface, check the **Shutdown** check box.
- Step 8** In the Gateway field, enter the default gateway IP address.
- Step 9** In the Address field, specify the IP address of the interface.
- Step 10** In the Netmask field, specify the netmask of the interface.
- Step 11** If you wish, from the Inbound ACL drop-down list, choose an IP ACL to apply to inbound packets. The drop-down list contains all the IP ACLs that you configured in the system.
- Step 12** If you wish, from the Outbound ACL drop-down list, choose an IP ACL to apply to outbound packets.
- Step 13** To apply the interface settings, click **Submit**.

To create an EtherChannel from the CLI, use the **interface PortChannel** *number* command in interface configuration mode, as shown in this example:

```
CE# configure
CE(config)# interface PortChannel 2
CE(config-if)# ip address 10.10.10.10 255.0.0.0
CE(config-if)# exit
```

To remove a PortChannel, use the **no** form of the command:

```
CE(config)# interface PortChannel 2
CE(config-if)# no ip address 10.10.10.10 255.0.0.0
CE(config-if)# exit
CE(config)# no interface portchannel 2
```

To add or remove ports from an EtherChannel, use the commands in the following examples. These commands add a physical Fast Ethernet port to a previously created Fast EtherChannel. The channel number is the same as the channel number specified in the **interface PortChannel** command. You can use either the Fast Ethernet or the Gigabit Ethernet ports to form an EtherChannel; however, an EtherChannel cannot contain both Fast Ethernet and Gigabit Ethernet interfaces. Note that a physical interface can be added to a EtherChannel subject to the device configuration.

Add an interface to an EtherChannel, as shown in this example:

```
CE# configure
CE(config)# interface FastEthernet 1/1
CE(config-if)# channel-group 2
CE(config-if)# exit
```

Remove an interface from an EtherChannel, as shown in this example:

```
CE(config)# interface FastEthernet 1/1
CE(config-if)# no channel-group 2
CE(config-if)# exit
```

Configuring a Load-Balancing Method for Interfaces within a Channel Group Using the Content Distribution Manager GUI

To configure load balancing using the Content Distribution Manager GUI, follow these steps:

-
- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the device with the port channel that you want to configure for load balancing.
- Step 3** In the Contents Pane, choose **General Settings > Network > Port Channel Settings**.
- Step 4** From the Load Balancing Method drop-down list, choose a load-balancing method:
- **dst-ip**—Destination IP address
 - **dst-mac**—Destination MAC address
 - **round robin**—Each interface in the channel group
- Round robin allows traffic to be distributed evenly among all interfaces in the channel group. The other balancing options give you the flexibility to choose specific interfaces (by IP address or MAC address) when sending an Ethernet frame.
- Step 5** To save the settings, click **Submit**.
-

To configure load balancing using the CLI, use the **port-channel load-balance** global configuration command.

```
CE(config)# port-channel load-balance
```

The **load-balance** command is effective globally. If two channel groups are configured, they must use the same load-balancing option.

Configuring Interfaces for DHCP



Note

Autoregistration must be disabled before you can manually configure an interface for DHCP.

An interface can be enabled for DHCP by using the **ip address dhcp** [*client_id* | *hostname*] interface configuration command. The client identifier is an ASCII value.

```
ContentEngine# configure
ContentEngine(config)# interface FastEthernet 0/0
ContentEngine(config-if)#
ContentEngine(config-if)# ip address ?
  A.B.C.D  IP address of the interface
  dhcp    IP address negotiated via DHCP
ContentEngine(config-if)# ip address dhcp ?
  client-id Specify client-id to use
  hostname  Specify value for hostname option
<cr>
```

The Content Engine sends its configured client identifier and host name to the DHCP server when requesting network information. DHCP servers can be configured to identify the client identifier information and the host name information that the Content Engine is sending and then send back the specific network settings that are assigned to the Content Engine.

To enable an interface for DHCP using the Content Distribution Manager GUI, follow these steps:

-
- Step 1** Choose **Devices > Devices**. The Devices window appears.
 - Step 2** Click the **Edit** icon next to the name of the device for which you want to configure interface settings. The Device Home window appears.
 - Step 3** In the Contents pane, choose **General Settings > Network > Network Interfaces**. The Network Interfaces listing window appears.
 - Step 4** Click the **Edit** icon for the physical interface FastEthernet or GigabitEthernet) that you want to modify. The Modifying Network Interface window appears. (See [Figure 14-1](#).)



Note Do not choose a logical interface (standby or portchannel) in this step. You cannot configure DHCP on a logical interface.

- Step 5** Check the **Use DHCP** check box. When checked, the secondary IP address and netmask fields are disabled.
 - Step 6** In the Hostname field, specify the hostname for the Content Engine or other device.
 - Step 7** In the Client Id field, specify the configured client identifier for the device. The DHCP server will use that identifier when the Content Engine requests the network information for the device.
 - Step 8** To save the settings, click **Submit**.
-

Configuring IP Addresses for Network Address Translation

ACNS software allows you to configure up to eight Network Address Translation (NAT) IP addresses to allow the router to translate up to eight internal addresses to registered unique addresses and translate external registered addresses to addresses that are unique to the private network.

To configure NAT IP addresses using the Content Distribution Manager (available as of the ACNS 5.3 software release), follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Device > Devices** (or **Devices > Device Groups**).
 - Step 2** Click the **Edit** icon next to the device or device group for which you want to configure external IP addresses.
 - Step 3** In the Contents pane, choose **General Settings > Network > External IP**. The External IP Settings window appears.
 - Step 4** To enable the window for configuration, check the **Enable** check box.
 - Step 5** In the External IP Address fields (1–8), enter up to eight IP addresses.
 - Step 6** To save the settings, click **Submit**.
-

To configure NAT IP addresses in the CLI, use the **external-ip** global configuration command. For example:

```
ContentEngine(config)# external-ip 192.168.43.1 192.168.43.2 192.168.43.3 192.168.43.4
```

Configuring Static IP Routes

ACNS software allows you to configure a static route for a network or host. Any IP packet designated for the specified destination uses the configured route.

To configure a static IP route, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. (Alternatively, choose **Devices > Device Groups**.)
 - Step 2** Click the **Edit** icon next to the device (or device group) that you want to configure.
 - Step 3** In the Contents pane, choose **General Settings > Network > IP Routes**. The IP Route Entries window appears.
 - Step 4** In the taskbar, click the **Create New IP Route Entry** icon. The Creating New IP Route window appears.
 - Step 5** In the Destination Network Address field, enter the destination network IP address.
 - Step 6** In the Netmask field, enter the destination host netmask.
 - Step 7** In the Gateway's IP Address field, enter the IP address of the gateway interface.
The gateway interface IP address should be in the same network as that of one of the device's network interface.
 - Step 8** To save the settings, click **Submit**.
-

**Note**

Do not use the **ip route** command to configure the default gateway; use the **ip default-gateway** command instead.



PART 5

Managing Website Access





Configuring Request Authentication and Authorization

Request authentication and authorization is a means to manage employee use of the Internet and restrict access to online content.

This chapter explains how to configure request authentication and authorization for serving content from Content Engines in a centrally managed ACNS network. It contains the following sections:

- [Understanding Request Authentication and Authorization for Content Engines, page 15-1](#)
- [Configuring Request Authentication on Centrally Managed Content Engines, page 15-17](#)
- [Configuring Authentication Server Settings, page 15-17](#)
- [Setting the Authentication Scheme for Request Authentication, page 15-36](#)



Note

In the ACNS network, Content Engines authenticate themselves as clients when acquiring content from origin servers and must be prepared to authenticate others when serving content. Authentication for acquiring content is discussed in [Chapter 6, “Configuring the ACNS Network for Content Acquisition.”](#) (See the [“Authentication Support”](#) section on page 6-42.)



Note

Content request authentication and authorization is independent of login (user) authentication and authorization. For information about login authentication and authorization, see [Chapter 12, “Configuring Login Authentication, Configuration Authorization, and Accounting.”](#)

Understanding Request Authentication and Authorization for Content Engines

Request authentication and authorization controls the serving of content through the Content Engine. When a Content Engine receives a request to serve content, it must determine whether to allow or deny that request. The determination to allow or deny the request might be made by the Content Engine itself, or it might be made by an external authentication service. If the request is allowed, the Content Engine begins to serve that content.

**Note**

In this chapter, *HTTP request* is used to refer collectively to requests over HTTP that include HTTP, MMS-over-HTTP, FTP-over-HTTP, and HTTPS-over-HTTP requests. (See the [“About Authenticating and Authorizing HTTP Requests”](#) section on page 15-6.)

**Note**

Native FTP requests use the FTP protocol (FTP-over-FTP) rather than FTP-over-HTTP. Native FTP can be used in proxied authentication mode only. (See the [“About Authenticating Native FTP Requests”](#) section on page 15-15.)

About Proxied Authentication and Pass-Through Authentication Modes

ACNS software supports two modes of operation when making content request authentication decisions: proxied authentication and pass-through authentication. Proxied authentication controls web use through a proxy. In the Cisco ACNS network, Content Engines can operate as proxy servers to collect user credentials from web browser clients. In proxied authentication mode, the Content Engine requests client credentials, collects the response, and then passes that information to an authentication service, which decides whether to allow or deny the request. The authentication service queries the configured user database to validate the credentials. The Content Engine can collect user credentials by using HTTP basic authentication or by using NTLM authentication as challenge mechanisms.

When configured for direct proxy routing, Cisco Content Engines use HTTP 407 headers as the default authentication response header for collecting credentials. When configured for transparent WCCP router redirection, Cisco Content Engines use HTTP 401 headers as the default authentication response header for collecting credentials. You can change the default behavior by using the CLI.

For more information about configuring 401 and 407 response messages in the proxy authentication header, see the [“Understanding Direct Proxy Routing and HTTP Request Authentication”](#) section on page 15-8 and the [“Understanding Transparent WCCP-Enabled Router Redirection and HTTP Request Authentication”](#) section on page 15-9.

In pass-through authentication mode, the origin server (rather than the Content Engine) challenges the client for credentials, and the Content Engine passes the authentication challenge to the client. If the origin server issues an HTTP basic authentication challenge, the client will prompt the end user for the credentials. If the origin server issues an NTLM challenge, the client will automatically respond with the network login credentials. The Content Engine then forwards the client challenge response to the upstream origin server. The origin server validates the credentials against the user database (or some other service). The Content Engine learns the net result of the decision (whether to allow or deny) based on the response code of the origin server. The Content Engine is not involved in the credential collection, evaluation, or validation.

Occasionally a website (origin server) requires the client’s IP address to authenticate a client. This authentication method is typically used in older web servers and is not the preferred solution for client authentication. However, for such older websites there must be a way for the Content Engine “to get out of the way” between the client and the web server so that the client can be authenticated. The authentication traffic bypass feature can be used in these cases. For information on this topic, see the [“Configuring Authentication Traffic Bypass”](#) section on page 4-22.

Content Delivery Protocols Supported

In ACNS software, three different content delivery protocols support authentication: HTTP, native FTP, and RTSP. These protocols are compatible with the following challenge and response authentication mechanisms:

- HTTP is compatible with both basic and NTLM authentication mechanisms. (See the [“About NTLM Authentication” section on page 15-3.](#))
- RTSP is compatible with basic authentication only. (See the [“About Basic Authentication” section on page 15-3.](#))
- Native FTP supports neither basic nor NTLM challenge and response authentication mechanisms.

The following list provides release-specific information about supported protocols:

- ACNS 5.2 software and later releases support HTTP request authentication. Throughout this chapter, the term *HTTP request* refers collectively to requests over HTTP that include HTTP, MMS-over-HTTP, FTP-over-HTTP, and HTTPS-over-HTTP requests. (For more information, see the [“About Authenticating and Authorizing HTTP Requests” section on page 15-6.](#))
- ACNS 5.4 software and later releases include proxy authentication of nontransparent native FTP requests. (That is, authentication of FTP-over-FTP requests from such FTP clients as a Reflection X client or a UNIX command line program by the Content Engine that is acting as the FTP proxy.) (For more information, see the [“About Authenticating Native FTP Requests” section on page 15-15.](#))
- ACNS 5.4 software and later releases enable you to use IP access control lists (ACLs) to control access for FTP native request (for example, enable the Content Engine that is acting as an FTP proxy to grant or deny access requests for incoming FTP connections from such FTP clients as a Reflection X client or a UNIX command line program). (For more information, see the [“Using IP ACLs to Control FTP Access” section on page 19-19.](#))
- ACNS 5.5 software does not support the MMS protocol, including MMST and MMSU protocols.

Authentication Mechanisms Supported

The Content Engine in proxied authentication mode supports both basic and NTLM authentication challenges.

About Basic Authentication

Basic authentication is a simple challenge and response authentication mechanism. When basic authentication is used, credentials, such as username and password, are transmitted to the origin server in clear text format. For this reason, basic authentication poses some security risks.

About NTLM Authentication

Windows NT LAN Manager (NTLM) is the challenge and response authentication mechanism used by Microsoft browsers (Internet Explorer), proxies, and web servers (IIS). When NTLM is used, clients encrypt the server challenge with their password hash and send the response to the server for validation. The main advantage of using NTLM for HTTP request authentication is that NTLM sends the password in an encrypted format to the server, which is more secure than transmitting in clear text across the Internet as is the case with basic authentication.

NTLM support on the Content Engine includes the following three types of support: (1) NTLM to NTLM pass-through authentication support, (2) NTLM proxied authentication of HTTP requests, and (3) NTLM group information query for authorization purposes. (See [Table 15-1.](#))

When the Content Engine is configured to use the NTLM challenge response mechanism, a preconfigured primary domain controller (PDC) validates the user's domain, username, and password before it allows the Content Engine to serve the request. In ACNS 5.2 software and later releases you can specify the list of domains that are allowed to perform NTLM authentication through the Content Engine.

Table 15-1 summarizes the relationship between protocols, challenge and response authentication mechanisms, authentication modes, and authentication decision services for the Content Engine as a server. It shows that in pass-through authentication mode, each authentication challenge mechanism that is supported can be passed only to an authentication server with a matching authentication service. For proxied authentication, the table shows that basic authentication of HTTP requests can be used with any of the supported decision services: NTLM authentication of HTTP can be proxied only to an NTLM service, and the streaming protocol, RTSP, does not support proxied authentication. Native FTP requests can be proxied to LDAP, RADIUS, and TACAS+ decision services.

Table 15-1 Authentication for Content Delivery—Content Engine as Server

Protocol	Pass-Through Authentication Mode		Proxied Authentication Mode			
	Auth. Mechanism	Decision Service	to NTLM Decision Service	to LDAP Decision Service	to RADIUS Decision Service	to TACACS+ Decision Service
HTTP ¹	Basic	to basic	Yes	Yes	Yes	Yes
	NTLM	to NTLM	Yes	No	No	No
FTP native	—	—	No	Yes	Yes	Yes
RTSP	Basic	to basic	No	No	No	No

1. HTTP refers collectively to HTTP, MMS-over-HTTP, FTP-over-HTTP, and HTTPS-over-HTTP.

NTLM Version 1 and 2 Support on the Content Engine

ACNS 5.4 software supports both NTLM Version 1 and Version 2 for both pass-through and proxied authentication of HTTP requests. (Earlier ACNS releases support NTLM Version 1 only.)

NTLM Version 2, the successor to NTLM Version 1, is more secure than NTLM Version 1. The algorithm that is used by NTLM Version 2 to encrypt the challenge is more complex than the algorithm used by NTLM Version 1, and the NTLM Version 2-encrypted response includes a timestamp to prevent replay attacks.

Unlike NTLM Version 1, NTLM Version 2 is not negotiated between the client and the server (or proxy), but it is configured separately on Windows-based clients and servers by modifying a registry variable. (For information about how to change a Microsoft Windows client registry or a Windows server registry, see your Microsoft Windows documentation.)

When an HTTP response to the client indicates that the server (or the proxy) requires authentication, and the NTLM credentials are accepted, the client generates the appropriate NTLM response based on the client configuration:

- If NTLM Version 1 is configured on the client, the client generates an NTLM Version 1 response when it receives the server challenge.
- If NTLM Version 2 is configured on the client, the client generates an NTLM Version 2 response when it receives the server challenge.

In the ACNS network, it is important that the NTLM version configurations be synchronized among the client, proxy, and server. By default, NTLM Version 2 is disabled on the Content Engine, and the Content Engine automatically uses NTLM Version 1 unless you enable NTLM Version 2.

To enable the Content Engine to use NTLM Version 2 to communicate with a specific NTLM server, use the **v2** option of the **ntlm server host** global configuration command or choose the **v2** option in the Content Distribution Manager GUI. When you enter the **v2** command option, you are specifying that the Content Engine should use NTLM Version 2 for request authentication with that particular NTLM server.

Specifying that the Content Engine use NTLM Version 2 when communicating with a specific NTLM host server is especially important for basic authentication because during basic authentication the Content Engine generates the NTLM response and communicates directly with the NTLM server.

In the following example, the Content Engine is configured to use a host list that consists of eight NTLM servers:

```
ContentEngine(config)# ntlm server host 172.16.10.10
ContentEngine(config)# ntlm server host 172.16.10.12
ContentEngine(config)# ntlm server host 172.16.10.14
ContentEngine(config)# ntlm server host 172.16.10.16
ContentEngine(config)# ntlm server host 172.16.10.18
ContentEngine(config)# ntlm server host 172.16.10.20
ContentEngine(config)# ntlm server host 172.16.10.22
ContentEngine(config)# ntlm server host 172.16.10.24
```

**Note**

In software releases that do not support NTLM Version 2, the Content Engine will fail to authenticate the client if the client attempts to use NTLM Version 2.

**Note**

In ACNS 5.4 software and later releases, when NTLM Version 2 is enabled on the Content Engine, if the domain controller accepts an NTLM Version 1 response, the Content Engine will successfully authenticate the client if the client is using NTLM Version 1.

NTLM Version 2 Pass-Through Authentication Support for NTLM-Enabled Browsers

If the Content Engine is using ACNS 5.4 software or a later release and the client browser is configured for NTLM Version 2, then pass-through authentication and proxied authentication are both supported.

For pass-through authentication, the following process occurs:

1. After the client browser receives an NTLM challenge from the Content Engine, the client browser generates an NTLM Version 2 response and sends it to the Content Engine.
2. The Content Engine performs the pass-through authentication and forwards the NTLM Version 2 response to the domain controller that sent the Content Engine the authentication challenge.
3. The Content Engine grants or denies the client access to the requested content based on the domain controller decision.

For proxied authentication, after the user is validated by the domain controller, the Content Engine saves the client username and domain name information in its authentication cache and does not challenge future requests from the same IP address. (See the [“Understanding the Content Engine Authentication Cache”](#) section on page 15-7.)

NTLM Version 2 Proxied Authentication Support for Non-NTLM Enabled Browsers

With non-NTLM-enabled browsers (browsers that must rely on basic authentication) you can configure the Content Engine to communicate with the domain controller by using NTLM Version 2 preventing the Content Engine from passing clear-text passwords over the network.

For the Content Engine to be able to use NTLM Version 2 to communicate with the domain controller, the following conditions must be met:

- The Content Engine is configured to use NTLM Version 2.
- The Content Engine is using ACNS 5.4 software or a later release.
- Basic authentication is not disabled on the Content Engine.

If the above conditions are met, when the client browser receives an HTTP response from the Content Engine indicating that client authentication is required, the client browser performs basic authentication instead of NTLM authentication. The client browser sends the user credentials to the Content Engine in clear text over the network. The Content Engine requests the challenge from the domain controller, generates an NTLM Version 2 response using the password, and then sends the NTLM Version 2 response to the domain controller for validation. (See the [“NTLM Version 1 and 2 Support on the Content Engine” section on page 15-4.](#))

**Note**

Even though the client browser sends its password to the Content Engine in clear text, the Content Engine does not forward the clear-text password over the network to the domain controller.

About Authenticating Windows Media Requests

ACNS 5.2 software and later releases up to ACNS 5.4 include pass-through basic and NTLM authentication support for MMS requests from Windows Media clients. ACNS 5.3 software and later releases include pass-through basic authentication support for Windows Media RTSP requests from Windows Media 9 players. With this support, the Content Engine establishes a tunnel between the client and the origin server for the origin server to authenticate the client. The Content Engine performs a pass through and does not function as the proxy authentication server.

**Note**

In the ACNS 5.5 release, the software no longer supports the MMS protocol, including the MMST and MMSU protocols.

When the Content Engine is operating in direct proxy routing mode, the Content Engine Windows Media 9 server supports two authentication mechanisms for pass-through authentication of WMT requests:

- Anonymous authentication
- Network authentication
 - Negotiates plug-in (NTLM or Kerberos) authentication
 - Digests plug-in authentication

About Authenticating and Authorizing HTTP Requests

For HTTP requests, proxies typically use basic authentication to collect client credentials and authenticate through a TACACS, RADIUS, or LDAP authentication database service. When the client is authenticated using basic authentication, the client browser automatically includes the client credentials with all subsequent HTTP requests so that any upstream proxy can validate the user credentials.

Organizations can use HTTP request authentication to restrict access to online content. For example, the BigCorp company might want to limit Internet access to its regular employees and not allow temporary employees to access the web at all. In this scenario, BigCorp places a Content Engine at its Internet gateway. The Content Engine is used to enforce the company's employee Internet access policy. When the Content Engine receives a client request to access content that is served through the Content Engine, the following occurs:

1. The Content Engine sends an authentication challenge to the client and prompts the client to enter authentication information, such as the username and password.
2. The Content Engine communicates with the AAA server to determine if the supplied authentication information is valid.
3. Based on the responses from the AAA server, the following occurs:
 - a. If the AAA server validates the user, the Content Engine grants the request and allows the client to access the requested content.
 - b. If the AAA server does not validate the user, the Content Engine denies the request and sends an authentication failed message to the client.

For NTLM deployments, Content Engines can be configured to challenge the client with NTLM. The Content Engine uses NTLM to collect the credentials and then uses Server Message Block (SMB) to validate the user's credentials against a domain controller. When the Content Engine sends a 401 or 407 response, the Microsoft Internet Explorer (IE) browser either prompts for a username, password, and domain, or it automatically responds with the known network login credentials. After collecting the credentials, the Content Engine validates the credentials with a query to the NT domain controller (or Active Directory server) using SMB.

Because NTLM can be implemented without a prompt, and the user database credentials are the same as the client's network login credentials, you can use promptless NTLM to implement single sign-on or to transparently log the domain and username in the transaction logs. The NTLM challenge is proprietary and is only understood by Microsoft IE browsers. If the Content Engine challenges a Netscape browser with NTLM, the Netscape browser will simply ask for a username and password.

**Note**

NTLM support on the Content Engine includes the following three types of support: NTLM pass-through authentication support, NTLM authentication of HTTP requests, and NTLM group information query for authorization purposes.

Understanding the Content Engine Authentication Cache

To minimize the impact of many authentication database queries per second on the Content Engine and on the authentication database, the Content Engine has an authentication cache. The authentication cache is maintained independently from other caches on the Content Engine. When a new request is routed to the Content Engine for the first time, the Content Engine collects a set of credentials from the user (user ID and password) and compares them against those in the authentication server database. If the credentials are valid, then an entry is created in the authentication cache with the user's credentials, the client IP address, and so forth. As long as the authentication entry is kept, subsequent attempts to access restricted Internet content by that particular user do not require authentication server lookups.

**Note**

The Content Engine must also guarantee that if it caches any objects that require client authentication (that is, if it caches authenticated content), then it will not deliver cached authenticated content to other clients unless those clients are authenticated to access the cached content.

The Content Engine routing method that is being used affects the way authenticated content is entered and indexed in the authentication cache. The routing methods supported for HTTP request authentication are direct proxy routing and transparent WCCP-enabled router redirection. The Content Engine stores and indexes authentication records differently, depending on the routing method being used, as follows:

- When the Content Engine is operating in direct proxy routing mode, it uses the client's user ID as a key for the its authentication cache.

When LDAP, RADIUS, and TACACS+ are used in direct proxy routing mode, the authentication record kept in the authentication cache is indexed by the username and the password entered. The username is entered with each client GET request.

- When the Content Engine is operating in transparent WCCP-enabled router redirection mode, it uses the client's IP address as a key for its authentication cache.

When LDAP, RADIUS, and TACACS+ are used in transparent WCCP-enabled router redirection mode, the authentication record is indexed by the IP address of the Content Engine that is sending the request in transparent mode. The client credentials are submitted only upon challenge, not with each request.


Note

When an NTLM server is used in either direct proxy routing mode or in transparent WCCP-enabled router redirection mode, all authentication records are indexed by using the IP address of the requesting Content Engine.

Understanding Direct Proxy Routing and HTTP Request Authentication

When the Content Engine is operating in direct proxy routing mode and is configured for HTTP request authentication, the following events occur if one of the following two situations is true: (1) the Content Engine receives a proxy-style request directly from a client, or (2) the Content Engine receives a WCCP-redirection request and the Content Engine **http authentication header** global configuration command option is set to 407 (Proxy Authorization Required) because there is an upstream proxy.

1. The Content Engine examines the HTTP headers of the client request to find user information (contained in the Proxy-Authorization header).
2. If no user information is provided, the Content Engine returns a 407 message (Proxy Authorization Required) to the client.
3. The client resends the request, including the user information.
4. The Content Engine searches its authentication cache (based on user ID and password) to determine whether or not the client has been authenticated previously.
5. If a match is found, the request is serviced normally.
6. If no match is found, the Content Engine sends a request to the authentication server to find an entry for this client.
7. If the server finds a match, the Content Engine allows the request to be serviced normally and stores the client user ID and password in the authentication cache.
8. If no match is found, the Content Engine again returns a 407 message to the client.

In some cases, the primary Content Engine (CE1) resides in the branch office and is configured in proxy mode, and another Content Engine (CE2) configured in proxy mode or another HTTP-compatible proxy device resides upstream with a TACACS+, RADIUS, NTLM, or LDAP server available to both Content Engines or proxy devices for login authentication.

**Note**

The **http append proxy-auth-header** global configuration command must be configured on the downstream Content Engines to ensure that proxy authorization information, required by upstream Content Engines, is not stripped from the HTTP request by the downstream Content Engines. Up to eight upstream IP addresses can be configured on each downstream Content Engine.

If branch office user 1 accesses the Internet and content is cached at CE1, then this content cannot be served to any other branch office user unless that user is authenticated. CE1 must authenticate each user before it serves the content.

Assuming that both CE1 and CE2 are connected to the server and authenticate the users, when branch office user 2 first requests Internet content, CE1 responds to the request with an authentication failure response (either HTTP 407 if in proxy mode or HTTP 401 if in transparent mode). User 2 enters the user ID and password, and the original request is repeated with the credentials included. CE1 contacts the HTTP request authentication server to authenticate user 2.

Assuming authentication success and a cache miss, the request and the credentials are forwarded to CE2. CE2 also contacts the authentication server to authenticate user 2. Assuming authentication success, CE2 either serves the request out of its cache or forwards the request to the origin server. (This credential forwarding capability is not configured by default. If you want to enable credential forwarding, you must explicitly configure it through the **http append proxy-auth-header host CE2ipaddress** global configuration command.)

User 2 authentication information is now stored in the authentication cache in both CE1 and CE2. Neither CE1 nor CE2 needs to contact the authentication server for user 2's subsequent requests (unless user 2's entry expires and is removed from the authentication cache).

This scenario assumes that CE1 and CE2 use the same method for authenticating users. Specifically, both Content Engines must expect the user credentials (user ID and password) to be encoded in the same way.

**Tip**

If you want to avoid authentication on an upstream Content Engine after authentication is performed downstream, you can use the **rule no-auth** global configuration command to exclude the downstream Content Engine IP address.

Understanding Transparent WCCP-Enabled Router Redirection and HTTP Request Authentication

When the Content Engine is operating in transparent mode (using WCCP-enabled router redirection), the user IP address is used as a key to the authentication cache. The Content Engine will always look for an X-Forwarded-For header first, then the source IP address.

If there are two levels of Content Engines in a proxy chain (CE1 at the first level [the Content Engine that is nearest to the client] and CE2 at the second level), and CE1 and CE2 both have the **http append x-forwarded-for-header multiple-ip-address** global configuration command configured on them, then the following events occur:

1. After receiving a request from the client, CE1 will append the default client's IP address to the X-Forwarded-for header and then forward the request to CE2. For example, if the client's IP address is 10.1.1.20, the X-Forwarded-For header would be "X-Forwarded-for: 10.1.1.20".
2. After CE2 receives the request from CE1, CE2 will append the IP address of CE1 to the X-Forwarded-for header. Consequently, the X-Forwarded-For header will contain the client's IP address (which is already present in the header) and the CE1's IP address, separated by comma. For example, if the IP address of CE1 is 10.40.1.40, the X-Forwarded-For header would be "X-Forwarded-for: 10.1.1.20, 10.40.1.40".

In the ACNS 5.4.1 software and later releases, multiple IP addresses in the X-Forwarded-For header are supported. Enter the **http append x-forwarded-for-header multiple-ip-address** global configuration command to enable support for appending multiple IP addresses to the X-Forwarded-for header. If you specify this command and if the request arrives at CE2 from CE1 with an X-Forwarded-For header, then CE1's IP address is appended to the existing client's IP address in the X-Forwarded-For header separated by a comma.

**Note**

If CE1 does not create an X-Forwarded-For header (for example, if it is not a Cisco Content Engine and does not support this header), then authentication on CE2 will not work.

In a topology with two Content Engines, assume that CE1 is operating in transparent mode and CE2 is operating in proxy mode, with the browsers of all users pointing to CE2 as a proxy.

Because the browsers are set up to send requests to a proxy, an HTTP 407 message (Proxy Authorization Required) is sent from CE1 back to each user to prompt for credentials. By using the 407 message, the problem of authenticating based on source IP address is avoided. The username and password can be used instead.

This mode provides better security than using the HTTP 401 message. The Content Engine examines the style of the address to determine whether or not there is an upstream proxy. If a proxy exists, the Content Engine uses an HTTP 407 message to prompt the user for credentials, even when operating in transparent mode.

When the Content Engine is operating in transparent mode and is configured for HTTP request authentication, the following events occur if either of the following is true: (1) the Content Engine receives a redirected request from a client, or (2) the **http authentication header** global configuration command option is set to 401 (Unauthorized) because there is no upstream proxy.

1. The Content Engine searches its authentication cache to see whether or not the user's IP address has been authenticated previously.
2. If a match is found, the Content Engine allows the request to be serviced normally.
3. If no match is found in the first step, the Content Engine examines the HTTP headers to find user information (contained in the Authorization header).
4. If no user information is provided, the Content Engine returns a 401 (Unauthorized) message to the client.
5. The client resends the request, including the user information.
6. The Content Engine sends a request to the authentication server to find an entry for this user.
7. If the server finds a match, the Content Engine allows the request to be serviced normally and stores the client IP address in the authentication cache.
8. If no match is found, the Content Engine again returns a 401 (Unauthorized) message to the client.

In transparent mode, the Content Engine uses the client IP address as a key for the Content Engine authentication cache.

In the ACNS 5.4.1 software and later releases, multiple IP addresses in the X-Forwarded-For header are supported.

**Note**

See the [“Configuring Authenticated HTTP Cache Settings”](#) section on page 8-13 for information about how to configure HTTP headers using the Content Distribution Manager GUI.)

Authentication Cache Adjustments

If the authentication cache is not large enough to accommodate all authenticated users at the same time, the Content Engine purges older entries that have not yet timed out. The default time interval between the user's last Internet access and the removal of that user's entry from the authentication cache is 480 minutes.

If 480 minutes passes without another HTTP request from the client, the Content Engine removes the client information from the authentication cache. If there is no entry in the authentication cache, the Content Engine challenges the next HTTP request from the client.

You can make adjustments to the authentication cache timeout value by using the **http authentication cache timeout** global configuration command, either through the CLI or through the Content Distribution Manager GUI. The minimum time interval is 1 minute, and the maximum is 1440 minutes (24 hours). The Content Engine forces reauthentication with the authentication server as soon as this time interval expires.

If you are using HTTP request authentication in transparent mode, we recommend that you configure the authentication cache timeout interval to be short. IP addresses can be reallocated, or different users can access the Internet through an already authenticated device (PC, workstation, and so on). Shorter timeout values help reduce the possibility that individuals can gain access by using previously authenticated devices. (When the Content Engine operates in proxy mode, however, unauthorized users cannot gain Internet access as easily because in proxy mode the end user must supply a valid user ID and password.)

In NTLM deployments, (unlike basic HTTP proxy authentication, which delivers client credentials with every HTTP request) the IE browser does not send credentials with subsequent web requests.

Consequently, the Content Engine authentication cache might rely on the client's source IP address rather than the actual user credentials to maintain an authenticated state on the client. When using NTLM instead of basic HTTP authentication, we recommend that you set the HTTP authentication cache timeout to 20 minutes (instead of using the default 480 minutes) to ensure that the Content Engine does not log in the incorrect user in circumstances where client IP addresses change frequently, for example, when dial modem access, VPN, or short term DHCP leases are used. When you use promptless NTLM authentication, you can set the authentication cache timeout even lower, since the authentication challenges will be transparent to the client. (See the [“Configuring the Microsoft Internet Explorer Browser for Promptless Authentication”](#) section on page 15-11.)

The ACNS 5.2 software release introduced an absolute timeout configuration option with the **http authentication cache ttl** global configuration command. This absolute timeout can also be configured to help reduce the possibility of individuals gaining access by using previously authenticated browsers. For more information, see the [“Specifying a Reauthentication Interval”](#) section on page 8-19.

Configuring the Microsoft Internet Explorer Browser for Promptless Authentication

To enable promptless authentication in the IE browser, follow these steps:

-
Step 1	From a Windows PC, choose Start > Settings > Control Panel > Internet Options . The Internet Properties dialog box appears.
Step 2	In the Internet Properties dialog box, choose the Security tab, and click the Custom Level button. The Security Settings dialog box appears.
Step 3	In the Security Settings dialog box, scroll to User Authentication > Logon, and click the Automatic logon only in Intranet zone radio button.
Step 4	To save the promptless authentication setting, click OK .
-

Using ACLs for Group-Based Authorization

Organizations that want to exercise more control can use group-based authorization (for NTLM and LDAP users) in addition to HTTP authentication. In ACNS 5.x software you can specify which groups of users are allowed to access the Internet and which groups are not allowed by creating access control lists (ACLs).

When you configure the Content Engine for group-based authorization, the Content Engine asks the AAA server if the clients are who they claim to be and to which groups they belong. The Content Engine then checks its ACLs to determine if the group to which the user belongs should be granted or denied access to the requested content.

Using Rules for Group-Based Authorization

The Content Engine also supports group-based rules for group-based authorization. With this feature, you can require that end users who are making an HTTP request be authenticated by an external AAA server and then authorized by the configured Rules Template.


**Note**

Group-based authorization occurs only after HTTP request authentication has occurred.

The Rules Template feature allows you to configure a set of rules, each clearly identified by an action and a pattern or a group of patterns that the Content Engine uses to filter HTTP requests, as well as HTTPS, and RTSP requests. If you have enabled the Rules Template feature on the Content Engine and configured rules for the Content Engine, the Content Engine checks every incoming client request to determine if a rule pattern matches the requested content. If a rule pattern matches the given request, the Content Engine uses the specified action (or policy) to filter this incoming content.

In the ACNS 5.2 software release, three new rule patterns were added (**groupname**, **username**, and **groupname-regex**). These new rule patterns support access control policies that are based on the group name and username of the authenticated NTLM and LDAP users. Rules based on group names apply to users who have been authenticated through NTLM and LDAP. Rules based on usernames apply to users who are authenticated through LDAP, NTLM, RADIUS, and TACACS+, which are request authentication methods that involve a username for authentication.

[Figure 15-1](#) shows how HTTP request authentication and group-based authorization can be used as mechanisms to control access to content.

Figure 15-1 HTTP Request Authentication and Group-Based Authorization

Using Microsoft Active Directory for Group Searches

Microsoft Active Directory, a feature used with group-based authorization, is a software application that runs on a Windows 2000 server. An Active Directory (AD) database is a user database that resides on a Windows 2000 server that is running the Microsoft Active Directory program. The LDAP client on a Content Engine provides LDAP support for Active Directory group searches.



Note

Microsoft Active Directory supports LDAP version 3 only. By default, the Content Engine uses LDAP Version 2, so you must configure the Content Engine to use LDAP Version 3 before enabling the LDAP Active Directory feature on the Content Engine.

In ACNS software Release 5.1 and later releases you can retrieve nested group names using an LDAP recursive search and apply all access lists configured for the nested groups. When you use nested groups with Active Directory servers, the policies configured for parent groups are automatically applied to members in subgroups.

The following items can be used to trigger an Active Directory group search:

- Group name-based access lists are configured.
- Group-based rules are configured in the Rules Template.
- ICAP is configured to append the authenticated-group header.
- SmartFilter is enabled on the Content Engine.

Using Digest Authentication

Digest Authentication was introduced in the Windows XP operating system. The Digest Authentication protocol is used with HTTP and Simple Authentication Security Layer (SASL) exchanges, as documented in RFCs 2617 and 2831. These exchanges require that parties who seek to authenticate must demonstrate their knowledge of secret keys. This process improves upon earlier versions of HTTP authentication, in which users provide passwords that are not encrypted when they are sent to a server, leaving them vulnerable to capture by attackers, or users must provide passwords that are encrypted but sent in an expensive, ongoing, Secure Sockets Layer (SSL) session.

Digest Authentication has similar security characteristics to the proprietary NTLM protocol. Both Digest Authentication and NTLM are challenge and response protocols. Challenge and response protocols require an authenticating server to generate a challenge containing some amount of unpredictable data. A client then uses a key derived from the user's password to encrypt the challenge and form a response. The server, or a trusted service such as Active Directory, can verify that the user possesses the correct password by comparing the client's encrypted response to a stored response based on the credential associated with the user in Active Directory or in the server account database for local users. If the responses match, the user is authenticated.

Limitations of Digest Authentication

SSL and Transport Layer Security (TLS) are often used to protect Digest Authentication from an offline attack against the Digest Authentication challenge and response.

Digest Authentication offers single sign-on to a single web URL protection space only. If users navigate to a different website or to a different server in the same site, they will usually be prompted to enter credentials again.

Common Digest Authentication Scenarios

Digest Authentication is a protocol that is used with web browsers for authenticating users browsing the Internet. Digest Authentication is also a general purpose protocol that can be used for authentication, and, by using SASL, it can provide integrity protection. For example, you can use Digest Authentication for the following purposes:

- Authenticated client access to a website.

Digest Authentication can be used to provide user authentication when users access pages on a web server.

- Authenticated client access using SASL.

Digest Authentication can be used as a SASL mechanism for any protocol that has a SASL profile. Using it as a SASL mechanism is a convenient way to support a single authentication mechanism for web, e-mail, LDAP, and other protocols.

- Authenticated client access with integrity protection to a directory service using LDAP.

When using SASL, integrity protection can be added.

Because Digest Authentication uses a message digest function to enable authentication, it works well with devices such as personal digital assistants (PDAs) that have little processor power and that need to have authenticated access to resources to read e-mail, access appointment information on an Exchange server, or view web pages on a server running services such as Internet Information Services (IIS).

Digest Authentication Dependencies

Digest Authentication depends on several related technologies and resources to function properly. The following section describes these technologies and resources and summarizes how they relate to Digest Authentication.

Active Directory Domains

Digest Authentication is not supported in earlier operating systems, such as Windows NT.

Users and services must have a valid Active Directory domain account.

The web server must be a member of the same forest as the user accounts.

Operating Systems

Digest Authentication is supported in Windows Server 2003 and Windows XP.

Windows Server 2003 Domain Controllers

All domain controllers for the domains of the users and services using Digest Authentication must be running on Windows Server 2003 to use the latest implementation of Digest Authentication. This requirement is necessary because password hash storage in Active Directory makes it necessary. The domains do not need to be configured for the Windows Server 2003 domain functional level.

If any of the domain controllers in the account domains are running Windows 2000 Server, then subauthentication, which requires reversible encryption, is required for Digest Authentication to work.

About Authenticating Native FTP Requests

ACNS 5.4 software and later releases support proxied authentication for nontransparent connections between FTP clients and FTP proxies. The FTP proxy communicates with the authentication daemon on the Content Engine to provide proxy authentication services for RADIUS, TACACS, LDAP, and NTLM protocols.



Note

In the case of NTLM, the FTP proxy uses the default NTLM domain name as configured on the Content Engine. The FTP client authenticates with the FTP server (origin server) using basic authentication. No NTLM authentication credentials are passed to the server.

The following is a sample FTP client session with NTLM proxy authentication enabled.

```
bash-2.04$ ftp -d 10.19.228.108 8021
Connected to 10.19.228.108.
220 Welcome to FTP-proxy. Login to the proxy using username and password.
Name (10.19.228.108:cisco_user): cnbul\cisco_user
---> USER cnbul\cisco_user
331 Password required for cnbul\cisco_user.
Password:
---> PASS XXXX
220-Welcome to FTP-proxy.
220-Login to origin server using the 'USER username@server-hostname' command, or
```



```

220 Login to origin server using the 'SITE server-hostname' followed by the 'USER
username' command.
ftp>
ftp> user admin@22.9.192.10
---> USER admin@22.9.192.10
331 Password required for admin.
Password:
---> PASS XXXX
230 User admin logged in. Access restrictions apply.
ftp> pwd
---> PWD
257 "/" is current directory.
ftp>

```

Considerations When Configuring Native FTP Proxied Authentication

When you configure request authentication for nontransparent native FTP requests on a Content Engine, remember the following important points:

- For security reasons, native FTP proxy authentication is disabled on the Content Engine by default. The FTP protocol is inherently insecure, and user credentials that would otherwise have been provided over a secure channel (such as, when using HTTP proxy authentication) can be exposed when native FTP proxy authentication is being used.
- If you configure native FTP proxy authentication and you have not already configured an authentication service (such as, RADIUS, LDAP, TACACS+, or NTLM) on the Content Engine, a warning message is displayed. This message indicates that you must configure an authentication service on the Content Engine before you can enable the FTP proxy authentication feature on the Content Engine.
- For request authentication for nontransparent native FTP requests, ACNS supports TACACS+, RADIUS, and LDAP as authentication services. The FTP client that is sending the native FTP request is queried for proxy authentication by the Content Engine only when one of the supported authentication services (TACACS+, RADIUS, and LDAP) is enabled on that Content Engine.



Note If NTLM is configured on the Content Engine, the Content Engine does not query the FTP client (that has sent a native FTP request) for proxied authentication.

In ACNS 5.4 software and later releases, you can use the same process to enable and configure an authentication service for HTTP requests and nontransparent FTP native requests. (For example, you use the same process to enable and configure RADIUS as an authentication service for HTTP requests and nontransparent native FTP requests.)

However, the following restrictions apply to native FTP caching support:

- No support for FTP request proxy rules
- No support for any URL filtering schemes (good list, bad list, N2H2, Websense, and SmartFilter)
- In ACNS 5.4 software and later releases, you can use IP access control lists (ACLs) to control access for nontransparent and transparent FTP native requests that are sent to the Content Engine that is acting as an FTP proxy for the user (an FTP client). (For more information, see [Chapter 17, “Creating and Managing IP Access Control Lists.”](#))

- In ACNS 5.4 software and later releases you can configure customized messages, which the Content Engine sends to an FTP client in response to an incoming proxy mode connection. (For more information, see the [“Configuring Native FTP Custom Messages”](#) section on page 8-62.)

Configuring Request Authentication on Centrally Managed Content Engines

To configure request authentication on the Content Engine, you must complete the following tasks:

1. Determine which one of the supported external authentication servers that you want the device to use when authenticating content requests.
2. Configure the authentication server settings that you wish to use on the Content Engine. (See the next section, [“Configuring Authentication Server Settings.”](#))
3. Specify which authentication database the device should check to process the request authentication. (See the [“Setting the Authentication Scheme for Request Authentication”](#) section on page 15-36.)

Configuring Authentication Server Settings

Some Cisco ACNS network users use an external access server as a centralized location for controlling the authentication, authorization, and accounting of user accounts and activities. External authentication servers are implemented at the protocol and application level with TACACS+, RADIUS, LDAP, and NTLM.

Only one type of request authentication can be enabled at a time. For example, you cannot enable both LDAP authentication and NTLM authentication at the same time.

The following sections describe how to configure LDAP, NTLM, RADIUS, and TACACS+ server settings for the Content Engine:

- [Configuring LDAP Server Settings, page 15-17](#)
- [Configuring NTLM Server Settings, page 15-23](#)
- [Configuring RADIUS Server Settings, page 15-32](#)
- [Configuring TACACS+ Server Settings, page 15-34](#)

Configuring LDAP Server Settings

System administrators can use the Content Engine to restrict user Internet access using an LDAP server for authentication purposes. ACNS 5.x software supports LDAP version 2 and version 3 and supports all LDAP features except for Secure Authentication and Security Layer (SASL).

To configure LDAP server settings using the Content Distribution Manager GUI, follow these steps:

-
Step 1	From the Content Distribution Manager GUI, choose Devices > Devices .
Step 2	Click the Edit icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.
Step 3	To display the entire table of contents, click the Show All button above the Contents pane.

- Step 4** In the Contents pane, choose **General Settings > Authentication > LDAP Server**. The LDAP Server Settings window appears. (See [Figure 15-2](#).) [Table 15-2](#) describes the fields shown in this figure.

Figure 15-2 LDAP Server Settings Window

The screenshot shows the 'LDAP Server Settings' window for the 'CONTENTENGINE' system. The left pane shows the 'Contents' tree with 'Authentication' > 'LDAP Server' selected. The main pane displays the following settings:

- Current settings: None (Using Factory Defaults)
- Enable LDAP Servers *: ☐
- LDAP Version: version 2
- Time to wait: 2
- Number of Retransmits: 2
- User-id Attribute: uid
- Filter:
- Base Distinguished Name:
- Administrative DN:
- Administrative DN Password:
- Confirm Password:
- Allow-Mode: ☐
- Server Port: 389
- Primary Host:
- Secondary Host:
- Enable Password Expiry Support: ☐

Buttons at the bottom right: Submit, Cancel.

- Step 5** To enable LDAP request authentication, check the **Enable LDAP Servers** check box.
- Step 6** From the LDAP Version drop-down list, choose the LDAP protocol version to be used.
- Step 7** In the Time to wait field, specify the number of seconds that the Content Engine waits before timing out.
- Step 8** In the Number of Retransmits field, specify the number of times that the Content Engine can try to reestablish its connection to the LDAP server if the timeout value is exceeded while it tries to contact the LDAP server. The number of transmission attempts can be from 1 to 3. The default is two attempts.
- Step 9** In the User-id Attribute field, enter the user ID attribute.
- Step 10** In the Filter field, enter the filter string to be used by the LDAP server.
- Step 11** In the Base Distinguished Name field, enter the base distinguished name string for the search in the LDAP server.
- Step 12** In the Administrative DN field, enter the administrative distinguished name.
- Step 13** In the Administrative DN password field, enter the administrative distinguished name password.
- Step 14** To enable access to users when the LDAP server is unavailable, check the **Allow-Mode** check box.
- Step 15** To allow the use of Windows Active Directory groups, check the **Active Directory Groups** check box.
- Step 16** In the Server Port field, specify a TCP port number for LDAP server authentication. We suggest using the default port value of 389.
- Step 17** In the Primary Host field, enter the IP address of the primary LDAP server.

Step 18 In the Secondary Host field, enter the IP address of the secondary LDAP server.

Step 19 To save the settings, click **Submit**.

Table 15-2 LDAP Server Settings

GUI Parameter	Function	CLI Command
Enable LDAP Servers	Enables HTTP authentication using LDAP servers.	ldap server enable
LDAP Version	LDAP protocol version (version 2 or version 3) to be used.	ldap server version
Time to wait	Number of seconds that the Content Engine waits for a response before timing out on a connection to a particular LDAP server. The default value is 5 seconds.	ldap server timeout
Number of Retransmits	Number of attempts allowed to connect to an LDAP server. The default value is 2 times.	ldap server retransmit
User-id Attribute	Name of the user ID attribute on the LDAP server. The default is “uid.”	ldap server userid-attribute
Filter	LDAP filter string. There is no default value.	ldap server filter
Base Distinguished Name	Base distinguished name of the starting point for the search of the LDAP server. This allows for a search on a particular string, such as the domain “com.”	ldap server base
Administrative DN	Administrative distinguished name. Allows a search for a particular user associated with the base distinguished name.	ldap server administrative-dn
Administrative DN Password	Password for the administrative distinguished name.	ldap server administrative-passwd
Allow-Mode	Allows access to users when the LDAP server is unavailable.	ldap server allow-mode
Active Directory Group	Allows access to Windows Active Directory groups.	ldap server active-directory-group
Server Port	Port number on which the LDAP server is listening. The default port number is 389.	ldap server port
Primary Host	IP address of the primary LDAP server.	ldap server host
Secondary Host	IP address of the secondary LDAP server.	ldap server host

In ACNS 5.4 software, the LDAP server password expiration, policy redirection, and group settings can be configured from the Content Distribution Manager GUI.

To configure the settings, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** To display the entire table of contents, click the **Show All** button above the Contents pane.
- Step 4** In the Contents pane, choose **General Settings > Authentication > LDAP Server**. The LDAP Server Settings window appears. (See [Figure 15-3](#).) [Table 15-3](#) describes the fields shown in this figure.

Figure 15-3 Additional LDAP Server Settings

- Step 5** Under the LDAP Password Expiry Settings heading, configure the following:
 - a. To enable support for the expiration of authorization passwords, check the **Enable Password Expiry Support** check box.
 - b. To direct users to a URL where they can reset their expired passwords, enter the URL in the Redirect URL to Reset Expired Password field.
- Step 6** Under the Policy Redirection Settings heading, configure the following:
 - a. To enable policy redirection support, check the **Enable Policy Redirection** check box.
 - b. To define the LDAP attribute under which the policy version number is stored, enter a name in the Attribute Name field.
 - c. To define the policy version number, enter a value from 1 to 999999999 in the Current Policy Version Number field.

- d. To redirect the user to a Usage Policy acceptance web page, enter the URL of the web page in the URL to Redirect field.
- e. Upon acceptance of the Usage Policy, if you want the browser to go directly to the requested URL, check the **Append Request URL to Redirect URL** check box.

Step 7 Under the Group Settings heading, configure the following:

- a. To obtain the group name from the *memberOf* attribute of a user's account and enable group membership based on that attribute, check the Active Directory Groups check box. This check box can be checked only if the LDAP version chosen from the LDAP Version drop-down list is version 3. (See [Figure 15-2](#).)
- b. To obtain the group name from the *custom* attribute of a user's account and enable group membership based on that attribute, enter a name in the Custom Group Name field. The maximum number of characters that you can enter for a name is 256 characters.
- c. To obtain the group name from the *organizationUnit* attribute of a user's account and enable group membership based on that attribute, check the Enable Organization Unit Group check box.
- d. To enable static group queries for group membership, check the **Enable Static Group** check box.
- e. To configure the static group attribute name, enter a name in the Static Group Attribute field.
- f. Choose one of the following static group member attributes by clicking the corresponding radio button:
 - **Do Not Set**
 - **Member**
 - **Unique Member**
 - **Custom Member**If you choose **Custom Member**, enter the custom member name of the static group in the text field.
- g. To enable nested static group queries for a user's membership, check the **Enable Nested Static Group** check box.
- h. To configure the level of nested static groups to query, enter a number from 1 to 100 in the Level of Nested Static Group field.

Step 8 To save the settings, click **Submit**.

Table 15-3 Additional LDAP Server Settings

GUI Parameter	Function	CLI Command
LDAP Password Expiry Settings		
Enable Password Expiry Support	Enables support for the expiration of authorization passwords.	ldap server password-expiry enable
Redirect URL to Reset Expired Password	Directs users to a URL where they can reset their expired passwords. This parameter accepts any URL with a valid format, such as the following: <code>http://www.someserver.com/file.html</code>	ldap server password-expiry redirect-url <i>url</i>
Policy Redirection Settings		
Enable Policy Redirection	Enables policy redirection support.	ldap server policy-redirect enable
Attribute Name	Defines the LDAP attribute under which the policy version number is stored.	ldap server policy-redirect attribute <i>name</i>
Current Policy Version Number	Defines the policy version number. The range is from 1 to 99999999.	ldap server policy-redirect version-number <i>number</i>
URL to Redirect	Redirects the user to a Usage Policy acceptance web page.	ldap server policy-redirect redirect-url <i>url</i>
Append Request URL to Redirect URL	Directs the user to the requested URL upon acceptance of the Usage Policy.	ldap server policy-redirect append-request-url
Group Settings		
Active Directory Groups	Queries the <i>memberOf</i> attribute of a user's account and enables group membership based on that attribute. To enable active directory query, the LDAP version must be configured, and the version must be version 3.	ldap server group active-directory enable
Custom Group Name	Queries the <i>custom</i> attribute of a user's account and enables group membership based on that attribute. Maximum characters for the custom group name is 256.	ldap server group custom <i>name</i> enable
Enable Organization Unit Group	Queries the <i>organizationUnit</i> attribute of a user's account and enables group membership based on that attribute.	ldap server group organizationUnit enable
Enable Static Group	Enables static group queries for group membership.	ldap server group static enable
Static Group Attribute	Configures the static group attribute name. Maximum characters for the static group attribute is 256. Cannot be configured if active directory group is enabled.	ldap server group static group-attribute <i>name</i>

Table 15-3 Additional LDAP Server Settings (continued)

GUI Parameter	Function	CLI Command
Group Settings		
Static Member Attribute:	Configures a static group member attribute. Maximum characters for the custom static member attribute is 256. Cannot be configured if active directory group is enabled or if the LDAP version is other than version 3.	ldap server group static member-attribute {custom-member <i>name</i> member uniquemember}
Do Not Set		
Member		
Unique Member		
Custom Member		
Enable Nested Static Group	Enables nested queries for static group membership.	ldap server group static nested enable
Level of Nested Static Group	Configures the level of nested static groups to query. Number of levels must be from 1 to 100. The default is 1.	ldap server group static nested level <i>number</i>

Configuring NTLM Server Settings

To configure NTLM server settings using the Content Distribution Manager GUI, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** In the Contents pane, choose **General Settings > Authentication > NTLM Server**. The NTLM Server Settings window appears. (See [Figure 15-4](#).) [Table 15-4](#) describes the fields shown in this window.

Figure 15-4 NTLM Server Settings Window—Top of Window

- Step 4** To enable NTLM authentication, configure the NTLM server domain name, NT primary domain controller name or IP address, and optionally set the host name or address as primary or secondary, check the **Enable NTLM Servers** check box.

Before making an NTLM authentication request, make sure that the following conditions are met:

- The NTLM primary domain controller has an entry in the Domain Name System (DNS) that matches its NetBIOS-named computer account.
- The primary domain controller is both forward and reverse DNS-resolvable.
- The domain name configured on the Content Engine matches the domain of which the primary domain controller is a part. This domain can be either a domain hosted by the PDC, or a trusted domain that the PDC can authenticate by contacting other PDCs.



Note

This domain can be either a domain hosted by the PDC or a trusted domain that the PDC can authenticate by contacting other PDCs.

- Step 5** In the Domain Name field, specify the domain name in which the user should be authenticated. (See the [“Support for No Default NTLM Domain”](#) section on page 15-30.)
- Step 6** From the Select NTLM scheme drop-down list, choose an option to specify the scheme for the NTLM servers for HTTP request authentication. This option allows you to specify the scheme (load balancing or failover) that is to be used among the configured NTLM servers. The default scheme is **fail-over**; the other supported scheme is **load-balanced**.

When the load balancing scheme is enabled, only the first request is sent to the first configured server, and then round robin is used among the configured servers. (See the [“About NTLM Load Balancing for HTTP Request Authentication”](#) section on page 15-30.) When the failover scheme is enabled, the Content Engine sends all the requests to the first configured server. (See the [“About NTLM Failover for HTTP Request Authentication”](#) section on page 15-31.)

- Step 7** In the Domain Server1 field, enter the host name or IP address for the domain server. The server configured as Server1 becomes the primary NTLM domain server when multiple servers are configured for failover purposes.



Note Only if the fail-over scheme is enabled, the first server configured is the “primary” server and is sent all of the requests first. If the load-balancing scheme is enabled, the definition of a “primary” server is not applicable.

- Step 8** In the Domain Server fields 2 through 8, enter the host names or IP addresses of each NTLM server that you want the Content Engine to use for HTTP request authentication. This list of configured NTLM servers is referred to as the “host list.”

In ACNS software, you can configure a Content Engine to use up to eight NTLM servers for HTTP request authentication. The order of the server configuration determines the order of load balancing or failover.

- Step 9** To configure the number of connection attempts, enter a number in the Maximum attempts to connect to server field.
- Step 10** To configure the wait time before a server connection fails, enter the number of seconds in the Time to wait when connecting to server field.
- Step 11** To enable support for active directory group search, check the **Enable Active Directory group search** check box. (See [Figure 15-5](#).)

Figure 15-5 NTLM Server Settings Window—Bottom of Window

The screenshot shows the 'NTLM Server Settings for Content Engine, CONTENTENGINE' window. The left sidebar contains a 'Contents' menu with options like Device Home, Device Activation, Assignments, Prepositioning, Request Routing, Request Processing, Applications, General Settings, Login/Access Control, Authentication Scheme, LDAP Server, NTLM Server, NTLM Allow Domain, RADIUS Server, TACACS+ Server, AAA Accounting, Access List, Services, Storage, Network, Notification and Tracking, Troubleshooting, and Device Monitoring. The main area is titled 'NTLM Server Settings for Content Engine, CONTENTENGINE'. It contains several configuration fields: 'Enable Referral Chasing' (checkbox), 'No. of Nested Referrals' (5), 'LDAP search server port' (506), 'Groupname attribute' (cn), 'Membership attribute' (memberOf), 'Object class for user object' (user), 'Username attribute' (sAMAccountName), 'Select scheme to user for the host list' (fail-over), and 'Port for global catalog server' (3268). Below these is a 'Global catalog server' section with a table of 8 rows for Host and Domain. The first row is filled with '172.16.30.45' and 'gcdom'. The bottom of the window has a note: '* To use NTLM for Request Authentication, please go to the Authentication Scheme Settings page.' and 'Submit' and 'Cancel' buttons.

Host	Domain
Host 1: 172.16.30.45	Domain 1: gcdom
Host 2:	Domain 2:
Host 3:	Domain 3:
Host 4:	Domain 4:
Host 5:	Domain 5:
Host 6:	Domain 6:
Host 7:	Domain 7:
Host 8:	Domain 8:

Active directory group search enables the ACNS software to interoperate with the Microsoft Active Directory database for NTLM group authentication.

Step 12 In the Domain Name field, enter the enumeration user's domain.



Note

An enumeration user is an account defined on the Content Engine to allow the Content Engine to perform a search in the Microsoft Active Directory database. This enumeration user needs to have read privileges throughout the whole directory.

Step 13 In the Username field, enter the enumeration username.

Step 14 In the Password field, enter the password of the enumeration user.

Step 15 In the Confirm Password field, reenter the enumeration user password.

Step 16 To allow searching in the Microsoft Active Directory database, check the **Enable Referral Chasing** check box.

Step 17 In the No. of Nested Referrals field, enter the directory level to which you want the search made.

Step 18 In the LDAP search server port field, enter the port number for the LDAP search server. The default is 389.

Step 19 In the Groupname attribute field, enter the groupname attribute in the Microsoft Active Directory database. The default is cn.

Step 20 In the Membership attribute field, enter the membership attribute in the Microsoft Active Directory database. The default is MemberOf.

- Step 21** In the Object class for user object field, enter the object class for the user object in the Microsoft Active Directory database. The default is user.
- Step 22** In the Username attribute field, enter the username attribute in the Microsoft Active Directory database. The default is sAMAccountName.
- Step 23** From the **Select scheme to use for the host list** drop-down list, choose a scheme:
- If you want failover to occur between the hosts, choose **failover**.
 - If you want round-robin load balancing to occur among the hosts, choose **load-balanced**.
- Step 24** In the Port for global catalog server field, enter the port number for the global catalog server. The default port number is 3268.
- Step 25** In the Host field, enter the global catalog server host name or IP address. You can add up to eight global catalog server host names.
- Step 26** In the Domain field that corresponds to the host name, enter the global catalog server domain name. You can enter up to eight global catalog server domain names.
- Step 27** To save the settings, click **Submit**.

Table 15-4 NTLM Server Settings

GUI Parameter	Function	CLI Command
Domain Name	Domain name in which the user should be authenticated.	ntlm server domain <i>name</i>
Select NTLM Scheme	Scheme for the NTLM servers for HTTP request authentication.	ntlm server scheme { <i>fail-over</i> <i>load-balanced</i> }
Domain Server1	NTLM server that serves as the primary host.	ntlm server host { <i>hostname</i> <i>ipaddress</i> } primary
Domain Server2–8	NTLM servers that serve as backup hosts.	ntlm server host { <i>hostname</i> <i>ipaddress</i> } secondary
Maximum attempts to connect to server	Maximum number of attempts (1–3) to connect to the server. The default is 2 attempts.	ntlm server connection-retry <i>number</i>
Time to wait when connecting to server	Number of seconds (1–20) to wait to connect to the server. The default is 5 seconds.	ntlm server connection-timeout <i>seconds</i>
Enable Active Directory group search	Enables the ACNS software to interoperate with the Microsoft Active Directory database for NTLM group authentication.	ntlm server ad-group-search enable
Domain Name	The enumeration user's domain.	ntlm server ad-group-search enum-user domain <i>domainname</i>
Username	The enumeration user's username.	ntlm server ad-group-search enum-user username <i>username</i>
Password/Confirm Password	The enumeration user's password.	ntlm server ad-group-search enum-user password <i>password</i>
Enable Referral Chasing	Enables LDAP referral. The default is disabled.	ntlm server ad-group-search ldap-referral enable

Table 15-4 NTLM Server Settings (continued)

GUI Parameter	Function	CLI Command
No. of Nested Referrals	the directory level to which you want the search made.	ntlm server ad-group-search ldap-referral limit <i>number</i>
LDAP search server port	Port number for the LDAP search server.	ntlm server ad-group-search ldap-search-port <i>portnum</i>
Groupname attribute	The groupname attribute in the Active Directory database. The default attribute is <i>cn</i> .	ntlm server ad-group-search groupname-attribute <i>attribute</i>
Membership attribute	The membership attribute in the Active Directory database. The default is <i>memberOf</i> .	ntlm server ad-group-search membership-attribute <i>attribute</i>
Object class for user object	The object class for the user object. The default is <i>user</i> .	ntlm server ad-group-search user-objectclass <i>class</i>
Username attribute	The username attribute in the Active Directory database. The default attribute is <i>sAMAccountName</i> .	ntlm server ad-group-search username-attribute <i>attribute</i>
Select scheme to use for the host list	Scheme to use for the global catalog host list. Options are fail-over or load-balanced.	ntlm server ad-group-search gc-server scheme { <i>fail-over</i> <i>load-balanced</i> }
Port for global catalog server	Global catalog server port. The default port is 3268.	ntlm server ad-group-search gc-server port <i>portnum</i>
Host	Global catalog server IP address or host name.	ntlm server ad-group-search gc-server host { <i>hostname</i> <i>ipaddress</i> }
Domain	Host domain name for the global catalog server being configured.	ntlm server ad-group-search gc-server host { <i>hostname</i> <i>ipaddress</i> } domain <i>domain</i>

Configuring the NTLM Server LDAP Memory Cache Settings

The ACNS 5.4 software release supports LDAP memory cache for nested group searches. This feature enables the Content Engine to store the results of a nested group search locally in its LDAP cache. By default, this feature is enabled on the Content Engine.

In previous ACNS releases, when the Content Engine performed an NTLM group search, it had to contact the LDAP Global Catalog server to get group information for every request. If the user belonged to a nested group with many levels, queries were sent from the Content Engine to the Active Directory server to search each parent level in the group, and the Content Engine needed to contact the LDAP Global Catalog server several times before obtaining all group information for the user.

To enhance performance for nested group searches, ACNS 5.4 software uses the LDAP memory cache on the Content Engine to cache the search results of LDAP queries. If the same search request is submitted again, the results are fetched from the LDAP memory cache on the Content Engine, reducing the number times the Content Engine must connect to the LDAP Global Catalog server.

To configure the NTLM server LDAP memory cache settings, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure.
- Step 3** In the Contents pane, choose **General Settings > Authentication > NTLM Server**. The NTLM Server Settings for Content Engine window appears.
- Step 4** To locate the LDAP memory cache configuration settings, scroll to the bottom of the window. (See [Figure 15-6](#).)

Figure 15-6 NTLM Server LDAP Memory Cache Settings for the Content Engine

The screenshot shows the 'NTLM Server Settings for Content Engine, CONTENTENGINE' window. The 'Contents' pane on the left is expanded to 'Authentication > NTLM Server'. The main area shows the following settings:

- Select scheme to user for the host list: **fail-over**
- Port for global catalog server: **3268**
- Global catalog server:

Host 1:	172.16.30.45	Domain 1:	gcdom
Host 2:		Domain 2:	
Host 3:		Domain 3:	
Host 4:		Domain 4:	
Host 5:		Domain 5:	
Host 6:		Domain 6:	
Host 7:		Domain 7:	
Host 8:		Domain 8:	
- Enable LDAP memory cache: ☐
- Size of memory cache: **140** (KB) (128-10240)
- Time to live for the cache: **400** (minutes) (1-1440)

At the bottom, there is a note: '* To use NTLM for Request Authentication, please go to the Authentication Scheme Settings page.' and a 'Note: * - Required Field'.

([Table 15-5](#) describes the fields in this window and lists the corresponding CLI commands.)

- Step 5** To enable caching of nested group query results, check the **Enable LDAP memory cache** check box.
- Step 6** To configure the size of the LDAP memory cache, enter a value (in kilobytes) in the Size of memory cache field.
- Step 7** To configure the time for an entry to remain in the cache, enter a value (in minutes) in the Time to live for the cache field.
- Step 8** To save the settings, click **Submit**.



Note All of the required fields in the NTLM Server Settings for Content Engine window must also be configured, or you will see an error message when you click Submit. (See the [“Configuring NTLM Server Settings”](#) section on page 15-23.)

Table 15-5 NTLM Server LDAP Memory Cache Settings for the Content Engine

GUI Parameter	Function	CLI Command
Enable LDAP memory cache	Enables the LDAP in-memory cache on the Content Engine. The default is enabled.	[no] ntlm server ad-group-search mem-cache enable
Size of memory cache	Number of kilobytes to allocate in the in-memory cache for the results of LDAP queries. The range is 128–10240 KB. The default is 140 KB.	ntlm server ad-group-search mem-cache size <i>size</i>
Time to live for the cache	Time (in minutes) that an entry remains valid in the LDAP in-memory cache. The range is 1–1440 minutes. The default is 400 minutes.	ntlm server ad-group-search mem-cache max-ttl <i>max-ttl</i>

Support for No Default NTLM Domain

ACNS software sends a “no domain configuration” error message to the client when the user does not supply a domain name in the request authentication credential and has not configured a default domain on the Content Engine. The error message contains text indicating the reason for the error.



Note

The no domain configuration feature is only supported with browsers that do not support NTLM (for example, Netscape browsers). For the Netscape browser, the user must specify the domain if the Content Engine does not have an NTLM default domain configured; otherwise, the client receives an error. Note that for the Netscape browser, the domain can only be supplied as part of the username in the format of “domain\username.” Browsers that do support NTLM, such as Internet Explorer, always include a domain name in the authentication credentials that originate from either the users being prompted to specify the credentials or from the domain that was used to logon users to their desktops.

About NTLM Load Balancing for HTTP Request Authentication

In large-scale networks where all network traffic passes through the Content Engine, even though the Content Engine authentication cache helps reduce the load on the domain controller, it might still be impractical to have a single domain controller handling all the authentication queries from end users. ACNS 5.4 software provides the ability to configure up to eight servers (domain controllers) for load balancing and failover purposes.

When you choose **load-balanced** as the authentication scheme, requests follow the round-robin process among the domain controllers. The order in which the domain controllers (servers) are configured determines the order of load balancing. For example, if you have n servers, the first request is sent to Server 0 and the second request is sent to Server 1, the n th request is sent to Server $n - 1$ and the $n + 1$ request is sent to Server 0. If Server 0 fails, the Content Engine attempts to send the request to the next alive server (in this case, Server 1). However, failover to the next alive server occurs a maximum of one time. For example, if Server 1 goes down when handling request 1, then request 1 does not fail over again.

If load balancing is enabled and the server information is changed during run time, the change is picked up at run-time without disrupting the service.

About NTLM Failover for HTTP Request Authentication

ACNS 5.4 software supports failover between domain controllers. The order in which the domain controllers (servers) are configured determines the order of failover. The first server configured (Server 0) is the first server to be contacted. The last server configured (Server 7) is the last server to be contacted.

If the timeout period for one connection attempt is exceeded, the Content Engine drops the connection and attempts to reconnect to the same server. The Content Engine retries the connection up to the specified number of retries configured (**ntlm server connection-retry** global configuration command) before attempting to connect to the next configured server on the host list.

Configuring NTLM Allowed Domains for HTTP Request Authentication

The NTLM protocol can be used to authenticate and block user access to the Internet. When a user logs in to a Windows NT or a Windows 2000 domain and starts a browser, the authentication information is stored by the browser and later used as NTLM credentials to access the Internet. The browser sends the NTLM credentials with the domain name to the Content Engine, which in turns sends a request to the Windows NT domain controller to check the validity of the user in the domain. If the user is not a valid user in the domain, then the request to access the Internet is denied. If authentication succeeds, the source IP address is entered in the Content Engine authentication cache. Future requests from this IP address are not challenged until the authentication cache entry expires or is cleared.

To configure NTLM allowed domains for HTTP request authentication on Content Engines, follow these steps:

-
Step 1	Choose Devices > Device Groups . The Device Groups window appears.
Step 2	Click the Edit icon next to the device group for which you want to configure the domains for NTLM HTTP request authentication. The Modifying Device Group window appears.
Step 3	In the Contents pane, choose General Settings > Authentication > NTLM Allow Domain . The NTLM Allow Domain Settings window appears.
Step 4	To enable NTLM HTTP request authentication on the Content Engine, check the Enable Use of the List of Domains Allowed to Authenticate check box.
Step 5	In the Domain field, enter the name of domain for which you want to support NTLM HTTP request authentication.
Step 6	To add the domain name to the list of domains allowed for NTLM HTTP request authentication, click Add . You can add up to 32 domains in this list. The newly added domain name appears in the window.
Step 7	To edit a domain name, click the Edit icon next to the domain name in the list of domains that are allowed for NTLM HTTP authentication.
Step 8	Edit the domain name field, and click Edit to save the changes.
Step 9	To delete a domain name from the list of domains supported for NTLM HTTP authentication, click the Delete icon next to the domain name.
Step 10	To save the changes, click Submit .
-

To enable NTLM HTTP request authentication from the CLI, use the **ntlm allow-domain {enable | domain *domainname*}** global configuration command.

Configuring RADIUS Server Settings

RADIUS authentication clients reside on devices running ACNS 5.x software. When enabled, these clients send authentication requests to a central RADIUS server, which contains user authentication and network service access information.



Tip

The Content Distribution Manager does not cache the user authentication information, so the user is reauthenticated against the RADIUS server for every request. To prevent performance degradation caused by many authentication requests, install the Content Distribution Manager in the same location as the RADIUS server, or as close as possible to it, to ensure that authentication requests can occur as quickly as possible.

To configure RADIUS server settings using the Content Distribution Manager GUI, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** In the Contents pane, choose **General Settings > Authentication > RADIUS Server**. The RADIUS Server Settings window appears. (See [Figure 15-7](#).) [Table 15-6](#) describes the fields in this window.

Figure 15-7 RADIUS Server Settings Window

RADIUS Server Settings

Current applied settings from Content Engine, **CONTENTENGINE**

Enable RADIUS Servers*: ☒

Time to wait: Number of retransmits:

Enable Redirect: ☐

Redirect Message 1: Location 1:

Redirect Message 2: Location 2:

Redirect Message 3: Location 3:

Shared Encryption Key*:

Server 1 Name:	Server 1 Port: 0
Server 2 Name:	Server 2 Port: 0
Server 3 Name:	Server 3 Port: 0
Server 4 Name:	Server 4 Port: 0
Server 5 Name:	Server 5 Port: 0

* To use RADIUS for Request Authentication, please go to the Authentication Scheme Settings page.

Submit Cancel

- Step 4** To enable RADIUS authentication, check the **Enable RADIUS Servers** check box.
- Step 5** In the Time to wait field, specify how long the Content Engine should wait before timing out. The default value is 5 seconds.
- Step 6** In the Number of Retransmits field, specify the number of attempts allowed to connect to a RADIUS server.
- Step 7** To enable RADIUS redirection, check the **Enable Redirect** check box.
- Step 8** In the Redirect Message field, enter a redirect message for the user. Three redirect messages are allowed.
- Step 9** In the Location field, enter a location where the redirect message should be sent. Three separate locations are allowed.
- Step 10** In the Shared Encryption Key field, enter the secret key that is used to communicate with the RADIUS server.
- Step 11** In the Server Name field, enter an IP address or host name information. Five different hosts are allowed.
- Step 12** In the Server Port field, enter the port number on which the RADIUS server is listening. Five different ports are allowed.
- Step 13** To save the settings, click **Submit**.

Table 15-6 RADIUS Server Settings

GUI Parameter	Function	CLI Command
Enable RADIUS Servers	Enables HTTP authentication using RADIUS servers.	radius-server enable
Time to wait	Number of seconds to wait for a response before timing out on a connection to a particular RADIUS server. The range is from 1 to 20 seconds. The default value is 5 seconds.	radius-server timeout
Number of retransmits	Number of attempts allowed to connect to a RADIUS server. The default value is 2 times.	radius-server retransmit
Enable redirect	Redirects an authentication response to a different authentication server if an authentication request using the RADIUS server fails.	radius-server redirect enable
Redirect Message	Message sent to the user if redirection occurs.	radius-server redirect message
Location	Sets an HTML page location. This is the URL destination of the redirect message that is sent when authentication fails.	radius-server redirect message reply location url
Shared Encryption Key	Encryption key shared with the RADIUS server.	radius-server key keyword
Server Name	IP address or host name of the RADIUS server.	radius-server host {hostname ipaddress}
Server Port	Port number on which the RADIUS server is listening.	radius-server host auth-port port

Configuring TACACS+ Server Settings

The TACACS+ database validates users before they gain access to a Content Engine. TACACS+ is derived from the United States Department of Defense (RFC 1492) and is used by Cisco Systems as an additional control of nonprivileged and privileged mode access. ACNS 5.x software supports TACACS+ only and not TACACS or Extended TACACS.

To enable TACACS+ for HTTP request authentication, use the **tacacs enable** global configuration command.



Tip

The Content Distribution Manager does not cache user authentication information. Therefore, the user is reauthenticated against the TACACS+ server for every request. To prevent performance degradation caused by many authentication requests, install the Content Distribution Manager in the same location as the TACACS+ server, or as close as possible to it, to ensure that authentication requests can occur as quickly as possible.

To configure TACACS+ server settings using the Content Distribution Manager GUI, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** In the Contents pane, choose **General Settings > Authentication > TACACS+ Server**. The TACACS+ Server Settings window appears. (See [Figure 15-8](#).) [Table 15-7](#) describes the fields in this window.

Figure 15-8 TACACS+ Server Settings Window

The screenshot shows the 'TACACS+ Server Settings' window for Content Engine, CONTENTENGINE. The left pane shows the 'Contents' tree with 'Authentication' expanded. The main area contains the following settings:

TACACS+ Server Settings	
Current applied settings from Content Engine, CONTENTENGINE	
Enable TACACS+ Servers *	☒
Use ASCII Password Authentication:	☐
Time to wait: *	0
Number of retransmits: *	0
Security Word:	***
Primary Server: *	172.9.26.12
Secondary Server:	10.83.133.65
Tertiary Server:	

Footnote: * To use TACACS+ for Request Authentication, please go to the Authentication Scheme Settings page.
 * To use TACACS+ for Login or Configuration Authentication, please go to the Login Authentication page.

Note: * - Required Field

Buttons: Submit, Cancel

- Step 4** To enable TACACS+ authentication, check the **Enable TACACS+ Servers** check box.
- Step 5** To use the ASCII password type for authentication, check the **Use ASCII Password Authentication** check box. The default password type is PAP (Password Authentication Protocol). However, you can change the password type to ASCII when the authentication packets are to be sent in ASCII clear text format.
- Step 6** In the Time to wait field, specify how long the Content Engine should wait before timing out. The default value is 5 seconds.
- Step 7** In the Number of Retransmits field, specify the number of attempts allowed to connect to a TACACS+ server. The default value is 2.
- Step 8** In the Security Word field, enter the secret key that is used to communicate with the TACACS+ server.
- Step 9** In the Primary Server field, enter an IP address or host name information for the primary TACACS+ server.

- Step 10** In the Secondary Server field, enter an IP address or host name information for a secondary TACACS+ server.
- Step 11** In the Tertiary Server field, enter an IP address or host name information for a tertiary TACACS+ server.
- Step 12** To save the settings, click **Submit**.

Table 15-7 TACACS+ Server Key Parameter Settings

GUI Parameter	Function	CLI Command
Enable TACACS+ Servers	Enables TACACS+ authentication.	tacacs enable
Use ASCII Password Authentication	Changes the default password type from PAP (Password Authentication Protocol) to ASCII clear text format.	tacacs password ascii
Time to wait	Number of seconds to wait for a response before timing out on a connection to a particular TACACS+ server. The range is from 1 to 20 seconds. The default value is 5 seconds.	tacacs timeout
Number of retransmits	Number of attempts allowed to connect to a TACACS+ server. The default value is 2 times.	tacacs retransmit
Security Word	Encryption key shared with the TACACS+ server.	tacacs key
Primary Server	IP address or host name of the primary TACACS+ server.	tacacs host {hostname ipaddress} [primary]
Secondary Server Tertiary Server	IP address or host name of the backup TACACS+ server. Up to 2 backup servers are allowed.	tacacs host {hostname ipaddress}

Setting the Authentication Scheme for Request Authentication

To enable an authentication scheme for request authentication, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** In the Contents pane, choose **General Settings > Authentication > Authentication Scheme**.
- Step 4** From the Authentication Scheme drop-down list, choose one of the following authentication schemes:
- Disable Authentication
 - RADIUS
 - LDAP

- NTLM
- TACACS+



Note You must configure and enable an authentication server before you can save the authentication scheme settings from this window.

Step 5 To save the settings, click **Submit**.



Configuring Request Processing Services

You can configure various content filtering services for the Content Engine, such as ICAP services, rules, and URL filtering. The following sections describe how to define and configure these various features:

- [Configuring ICAP, page 16-1](#)
- [Configuring Service Rules, page 16-13](#)
- [Configuring URL Filtering, page 16-25](#)
- [Configuring URL Filtering with the N2H2 Server, page 16-33](#)
- [Configuring URL Filtering with Websense Enterprise Software, page 16-39](#)
- [Configuring URL Filtering with SmartFilter Software, page 16-55](#)
- [Configuring Content Engines to Bypass URL Filtering for Specific HTTP and HTTPS Requests, page 16-57](#)
- [Configuring the Content Engine GUI for Secure or Nonsecure Access, page 16-59](#)
- [Setting Up the Content Engine to Interoperate with Third-Party Policy Servers, page 16-60](#)

Configuring ICAP

This section describes how to configure the Internet Content Adaptation Protocol (ICAP) on centrally managed Content Engines. The ACNS 5.2.1 software and later releases support ICAP for HTTP requests. The ACNS 5.4.1 software and later releases support ICAP for FTP-over-HTTP requests. Support for native FTP requests is not supported.

Overview of ICAP

ICAP is an open standards protocol that can be used for content adaptation, typically at the network edge. Content adaptation includes virus scanning, content translation, content filtering, content insertion, and other ways of improving the value of content to end users. ICAP specifies how a Content Engine, acting as an HTTP proxy server, can communicate with an external device that is acting as an ICAP server, which filters and adapts the requested content.

ICAP provides two content-processing modes for HTTP services. These modes define the transactions that can occur between a Content Engine acting as an ICAP client and an ICAP server. The two modes are as follows:

- Request modification (**reqmod**)—Allows modification of requests as they are sent from the Content Engine to the ICAP server on their way to the origin server. The ICAP server can modify these requests depending on the services requested.
- Response modification (**respmod**)—Allows modification of requests after they return from the origin server. The ICAP server only acts on requested objects after they return from the origin server.

About ICAP Services

An ICAP service is a collection of attributes that define the service and one or more ICAP servers that provide the ICAP services. You can configure a maximum of ten ICAP services per Content Engine, with an upper limit of five ICAP servers per ICAP service. Also, you can choose to apply ICAP services on all HTTP requests processed by the Content Engine or apply ICAP processing only to requests that match the Rules Template.



Tip

To set the type of load balancing to use among a cluster of ICAP servers, use the **icap service load balancing** global configuration command.

About ICAP Services and Vectoring Points

The point at which ICAP services are applied to content is called the *vectoring point*, specified using the **vector-point** option. The following three vectoring points are supported:

- Client request vectoring point (**reqmod-postcache**)—The ICAP server performs one of the following actions in response to the client request:
 - Terminates the connection
 - Sends a modified error response
 - Searches the cache using the URL in the request
 - Searches the cache using a modified URL
 - Modifies the request header or request body in the case of a cache miss
- Cache miss vectoring point (**reqmod-precache**)—The ICAP server performs one of the following actions before forwarding the request to the origin server:
 - Terminates the connection
 - Sends a modified error response
 - Sends the request to the origin server using the original URL
 - Sends the request to the origin server using an alternative URL
 - Modifies the request header or request body
- Server response vectoring point (**respmod-precache**)—The ICAP server performs one of the following actions after receiving the response from the origin server:
 - Returns the response to the client
 - Modifies the request header or request body

- Caches the response using the original URL
- Caches the response using an alternative URL

**Note**

Different ICAP services assigned to the same vectoring point can use different load-balancing options.

The following commands show a typical configuration for a virus-scanning service that requires processing on two vectoring points: **reqmod-precache** and **respmod-precache**:

```
ContentEngine(config)# icap apply all
ContentEngine(config)# icap service trend-reqmod
ContentEngine(config-icap-service)# enable
ContentEngine(config-icap-service)# vector-point reqmod-precache
ContentEngine(config-icap-service)# server icap://172.19.227.150/REQ-Service
ContentEngine# exit
ContentEngine# icap service trend-respmod
ContentEngine(config-icap-service)# enable
ContentEngine(config-icap-service)# vector-point respmod-precache
ContentEngine(config-icap-service)# server icap://172.19.227.150/interscan
ContentEngine# exit
```

If an ICAP vendor supports the same service name for more than one vectoring point, you can configure a single service and add the supported vectoring points, as in the following example:

```
ContentEngine(config)# icap service myicap-service
ContentEngine(config-icap-service)# enable
ContentEngine(config-icap-service)# vector-point reqmod-precache
ContentEngine(config-icap-service)# vector-point respmod-precache
ContentEngine(config-icap-service)# server icap://172.19.227.150/icap-service-name
ContentEngine(config-icap-service)# exit
ContentEngine(config)#
```

About ICAP Performance

With the response modification (respmod) vectoring point, which is used by virus-scanning ICAP vendors, the performance of the Content Engine model CE-7305 will be 300 transactions per second.

With the request modification (reqmod)-precache vectoring point, which is used by URL filtering ICAP vendors, the performance of the Content Engine model CE-7305 will drop 20 percent from the rated performance.

**Note**

The performance of the Content Engine will be limited by the performance of the ICAP server.

ICAP Vendors Supported

The following is a complete list of the ICAP vendors that have been certified to interoperate with the Content Engine:

- TrendMicro for reqmod and respmod
- Symantec for respmod

Maximum File Size Supported

For ACNS 5.4.x software and later, the maximum file size that is supported in the ACNS software is 2 GB. Files that exceed this size limit are not supported for ICAP processing.

For releases prior to ACNS 5.4.x software, the maximum file size that is supported in the ACNS software in pass-through mode is 2 GB. Files that exceed this size limit are not supported for ICAP processing.

Enhancements to ICAP Support in ACNS 5.5 Software

The ACNS 5.5.1 software supports the following enhancements for ICAP support:

- ICAP request modification support for HTTPS requests

The ACNS 5.5.1 software supports an ICAP request modification (reqmod) for the HTTPS proxy-style requests. Support for ICAP reqmod processing allows proxy-style requests that use the HTTPS protocol to be modified as they are sent from the Content Engine to the ICAP server on their way to the origin server.

- Compatibility with the data trickling feature of the WebWasher ICAP server

The ACNS 5.5.1 software allows you to download large files over an ICAP-enabled network. Because of some compatibility issues with the data trickling feature of the WebWasher ICAP server, downloading large files over ICAP-enabled ACNS networks used to fail before the download was completed. However, starting with release 5.5.1, the ACNS software has become compatible with the data trickling feature of the WebWasher ICAP server.

- Support for accessing blank pages over ICAP-enabled networks

The ACNS 5.5.1 software allows you to access blank pages (null body pages, where there is no content) over an ICAP-enabled ACNS network. When you try to access blank pages, earlier versions of the ACNS software used to return error messages saying that the page could not be retrieved because the server was either unreachable or temporarily busy.

- New CLI command and GUI support for configuring the ICAP connection timeout


The ACNS 5.5.1 software allows you to configure a timeout value for ICAP connections. You can configure an ICAP connection timeout value using the **icap connection-timeout** *minutes* global configuration command.

Configuring ICAP Settings

To configure ICAP settings for the Content Engine, follow these steps:

-
Step 1	Choose Devices > Devices . The Devices window appears.
Step 2	Click the Edit icon next to the Content Engine that you want to configure. The Device Home window appears.
Step 3	To display the entire table of contents, click the Show All button above the Contents pane.
Step 4	In the Contents pane, choose Request Processing > ICAP . The ICAP Settings window appears. (See Figure 16-1 .) Table 16-1 describes the fields in this window and provides the corresponding CLI global configuration commands.

Figure 16-1 ICAP Settings Window



Step 5 From the Apply to drop-down list, choose the type of ICAP processing that you wish to apply to requests: The **all** option applies ICAP processing to all HTTP requests. The **rules-template** option applies ICAP processing only to requests that match the **rule action use-icap-service** rule defined in the Rules Template.

Step 6 If you do not want to bypass ICAP processing for streaming media requests (that is, requests from Windows, RealMedia, and QuickTime media players) that enter the Content Engine, uncheck the **Bypass Requests Streaming Media** check box. This check box is checked by default.

Requests from streaming media clients are by default bypassed for ICAP processing. Disable the **icap bypass streaming-media** option to force strict rechecking of the cached content every time the ISTag changes. The ISTag is a field in the HTTP response header that allows ICAP servers to send a service-specific cookie to an ICAP client, representing the current state of the service. The ISTag might change as a result of an update to the server version, to a virus-pattern-file, or to the policy.

You can configure the Content Engine to append the client and server IP address headers to requests that are passed to the ICAP server. The **icap append x-headers** option specifies that the ICAP extension headers are passed to the ICAP server during the session negotiation between the Content Engine and the ICAP server. This capability allows you to use your ICAP server to perform URL filtering based on the client IP address and server IP address.

Step 7 To add an X-client-IP header for ICAP processing of HTTP requests, check the **Append X-Client-IP headers** check box.

Step 8 To add an X-server-IP header for ICAP processing, check the **Append X-Server-IP headers** check box.

You can configure the Content Engine to append the username and groupname headers to the request that is passed to the ICAP server was added. This capability allows you to use your ICAP server to perform URL filtering based on username and groupname. The currently supported authentication schemes include LDAP, NTLM, RADIUS, and TACACS+.

- Step 9** To allow the username information to be passed to the ICAP server for global services, check the **Append X-Authenticated-User headers** check box. This check box is unchecked by default. When this check box is checked, the x-authenticated-user information is inserted into the ICAP request that is sent to the ICAP server.
- Step 10** To allow the groupname information to be passed to the ICAP server for global services, check the **Append X-Authenticated-User-Groups headers** check box. This check box is unchecked by default. When check box is checked, the x-authenticated-groups information is inserted into the ICAP request that is sent to the ICAP server.
- Step 11** To rescan cached objects, check the **Rescan Cache Service Update (ISTag)** check box. The ISTag is a field in the HTTP response header that can change as a result of such things as updates to the ICAP server or ICAP policy. A change in the ISTag triggers the ICAP server to send a cookie to the ICAP client, indicating that a change has occurred and that the client should not use any previously cached material and should instead rescan the cache.
- Step 12** To configure the ICAP connection timeout interval, enter a number between 1 and 480 in the Connection-Timeout field. The default is 20 minutes.
- The **icap connection timeout** option allows you to stop an ICAP transaction that might be blocking traffic for longer than the specified time. This feature allows you to control the use of system resources and any additional overhead to the system that might be caused by a blocked transaction.
- Step 13** To configure and enable logging in ICAP standard logging format for ICAP exchanges between ICAP servers and Content Engines, check the **Enable Logging** check box.
- Step 14** To save your settings, click **Submit**. A “Click Submit to Save” message appears in red next to the current settings line when there are pending changes to be saved after you have applied default and device group settings. To revert to the previously configured window settings, click **Reset**. The Reset button appears only when you have applied default or group settings to change the current device settings but the settings have not yet been submitted.

Table 16-1 ICAP Settings

GUI Parameter	Function	CLI Command
Apply to	The all option applies ICAP processing to all HTTP requests. The rules-template option applies ICAP processing only to requests that match the rule action use-icap-service rule defined in the Rules Template.	icap apply {all rules-template}
Bypass Requests Streaming Media	When checked, bypasses ICAP processing for streaming media requests.	icap bypass streaming-media
Append X-Client-IP headers	Appends the client IP address header to requests that are passed to the ICAP server.	icap append-x-headers x-client-ip
Append X-Server-IP headers	Appends the server IP address header to requests that are passed to the ICAP server.	icap append-x-headers x-server-ip

Table 16-1 *ICAP Settings (continued)*

GUI Parameter	Function	CLI Command
Append X-Authenticated-User headers	Allows the username information to be passed to the ICAP server for global services.	icap append x-authenticated-user
Append X-Authenticated-Groups headers	Allows the groupname information to be passed to the ICAP server for global services.	icap append x-authenticated-groups
Rescan Cache	Rescans cached objects.	icap rescan-cache ISTag-change
Connection-Timeout	Adjusts the timeout interval (in minutes) for ICAP server connections.	icap connection-timeout <i>minutes</i>
Enable Logging	Enables logging in ICAP standard logging format for ICAP exchanges between ICAP servers and Content Engines.	icap logging {format {standard} enable}

Configuring ICAP Services

ACNS 5.5 software supports the following three vectoring points:

- Request modification-precache

This vectoring point applies to requests from clients. In the case of a cache miss, the ICAP server specifies whether to terminate the connection, send a modified error response to the client, search the cache using the requested URL, use a modified URL before looking up the cache, or modify the request headers or bodies.

- Request modification-postcache

This vectoring point applies only to requests after a cache miss and before the request is forwarded to the origin server for content retrieval. The ICAP server determines whether to terminate the client connection, send a modified error response to the client, send the request to the origin server using the client-specified URL or an alternative URL, or modify the request headers or bodies.

- Response modification-precache

This vectoring point applies to responses that are received from the origin server. The ICAP server specifies whether to return the response to the client, modify the response headers or bodies before sending them to the client, or cache the response using the same or an alternative URL.

ICAP servers configured with various vectoring points (especially request modification-precache) may become overloaded with HTTP requests because all requests pass through this point. We recommend that you use a cluster of ICAP servers to load balance requests based on various parameters, such as weighted load, client IP and server IP address-hash based format, or round-robin format.

More than one ICAP service can be associated with a vectoring point. An ICAP service configured at a vectoring point can have only one load-balancing scheme, irrespective of the number of servers. However, multiple ICAP services that are configured at one or all of the vectoring points can have different load-balancing schemes.

**Note**

When you use the Aggregate Settings option, you cannot use the ICAP Services for Content Engine window to modify or delete ICAP services that have been previously configured for device groups to which the Content Engine belongs. In other words, you can only view the ICAP services created for the device groups.

To configure an ICAP service for a Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure an ICAP service. The Device Home window appears.
- Step 3** In the Contents pane, choose **Request Processing > ICAP Services**. The ICAP Services for Content Engine window appears.
- Step 4** The **Yes** radio button for aggregate settings is selected by default, meaning that the ICAP services defined for the Content Engine and the device groups to which the Content Engine belongs are displayed. To display only the settings defined for the Content Engine, click the **No** radio button.
- Step 5** In the taskbar, click the **Create New ICAP Service** icon. The Creating New ICAP Service for Content Engine window appears. (See [Figure 16-2](#).) [Table 16-3](#) describes the fields in this window and provides the corresponding CLI global configuration commands.

Figure 16-2 Creating New ICAP Service

Cisco Application and Content Networking System - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Cisco Systems Application and Content Networking System

System Status Devices: 3 Devices, Critical Content: 17 Channels, Minor

Home | Help | Logout

Devices Device Groups Locations Statistics

Expand All Show Basic

Contents

- Device Home
- Device Activation
- Assignments
- Prepositioning
- Request Routing
- Request Processing
 - URL Filter
 - WebSense Server
 - Enable Rules
 - Service Rules
 - ICAP
- ICAP Services
 - Transaction Logs
 - Applications
 - General Settings
 - Device Monitoring

Creating New ICAP Service for Content Engine, CONTENTENGINE

ICAP Service Configuration		ICAP Service Vectoring Points	
Name:		Request Modification:	☐ Pre-cache ☐ Post-cache
Enable:	☐	Response Modification:	☐ Pre-cache
Load Balancing:	Round Robin		
Error Handling:	Bypass		

ICAP Servers for ICAP Service			
Server URL	Maximum Connections	Weight	Keepalive Interval
ICAP Servers may be created after this new ICAP Service has been Submitted.			

Submit Cancel

Step 6 Configure the ICAP service, as follows:

- a. In the Name field, enter a string that identifies the ICAP service to be configured.
- b. To enable ICAP service, check the **Enable** check box.
- c. From the Load Balancing drop-down list, choose the type of load-balancing mechanism for ICAP processing. Table 16-2 describes the load-balancing options.

Table 16-2 ICAP Load-Balancing Options

Load-Balancing Type	Description
Client IP Hash	Uses a hash-based algorithm based on the client IP address for load-balancing the ICAP servers in the cluster.
Round Robin	Uses the round-robin method, in which ICAP servers take turns processing HTTP requests.
Server IP Hash	Uses a hash-based algorithm based on the server IP address for load balancing the ICAP servers in the cluster.
Weighted	Uses a farm of ICAP servers with different load capacities.

Step 7 From the Error Handling drop-down list, choose the type of error handling mechanism for ICAP processing. To bypass this ICAP service, choose **Bypass**. Otherwise, choose **Return Error** if you want errors to be returned for client requests. These errors are also entered in the transaction log to show the status of the action performed by the ICAP services.

Step 8 Under the ICAP Service Vectoring Points heading, configure the following options:

- a. To enable the vectoring point that is to be invoked when a Content Engine receives a request from a client, check the **Pre-cache** check box for request modification.
- b. To enable the vectoring point that is to be invoked if the request is a cache miss and must be sent to the origin server for the content, check the **Post-cache** check box for request modification.
- c. To enable the vectoring point that is to be invoked only when the response is from the origin server, check the **Pre-cache** check box for response modification.

Step 9 To save your settings, click **Submit**. The configured settings are saved to the database, and the ICAP Services for Content Engine window appears, listing the configured ICAP service.

Step 10 To display a subset of the entire list of ICAP services, click the **Filter Table** icon in the taskbar.

Step 11 To revert to the display of all configured ICAP services, click the **View All ICAP Services** icon in the taskbar.

Table 16-3 ICAP Service Settings

GUI Parameter	Function	CLI Command
ICAP Service Configuration		
Name	Identifies the ICAP service to be configured.	icap service <i>service-id</i>
Enable	Enables the ICAP service.	icap service <i>service-id</i> enable
Load Balancing	Configures the type of load-balancing mechanism for ICAP processing.	icap service <i>service-id</i> load-balancing { round-robin client-ip-hash server-ip-hash weighted-load }

Table 16-3 *ICAP Service Settings (continued)*

GUI Parameter	Function	CLI Command
Error Handling	Configures the type of error handling mechanism for ICAP processing.	icap service <i>service_id</i> error-handling {bypass return-error}
ICAP Service Vectoring Points		
Request Modification: Pre-cache	When checked, enables the vectoring point that is to be invoked when a Content Engine receives a request from a client.	icap service <i>service_id</i> vector-point reqmod-precache
Post-cache	When checked, enables the vectoring point that is to be invoked if the request is a cache miss and must be sent to the origin server for the content.	icap service <i>service_id</i> vector-point reqmod-postcache
Response Modification: Pre-cache	When checked, enables the vectoring point that is to be invoked only when the response is from the origin server.	icap service <i>service_id</i> vector-point respmod-precache

Configuring an ICAP Server for ICAP Service

ICAP servers process HTTP requests from clients based on the ICAP services configured using various vectoring points. ICAP servers perform content adaptation, such as request or response modification and filtering of requests or responses based on the configured vectoring points.

You can configure the maximum number of connections and the weight that can be handled by an ICAP server in a cluster of servers. The weight parameter represents the percentage of load that can be redirected to the ICAP server. An ICAP server with a weight of 40 means that this server handles 40 percent of the load. If the total weight of all ICAP servers in a load-balanced cluster exceeds 100, the weight parameter for each ICAP server is recalculated.



Note

Always locate the ICAP server on a public LAN and configure its public IP address on the Content Engine. The ICAP server should not be located behind a NAT device.

To configure an ICAP server for a previously configured ICAP service on a Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the desired Content Engine. The Content Engine Device Home window appears.
- Step 3** In the Contents pane, choose **Request Processing > ICAP Services**. The ICAP Services for Content Engine window appears.

- Step 4** The **Yes** radio button for aggregate settings is selected by default, so the ICAP services defined for the Content Engine and the device groups to which the Content Engine belongs are displayed. To display only the settings defined for the Content Engine, click the **No** radio button.
- Step 5** To configure an ICAP service, click the **Edit ICAP Service** icon next to the ICAP server for which you wish to configure an ICAP service. The Modifying ICAP Service for Content Engine window appears.
- Step 6** In the ICAP Servers for ICAP Service section, click the **Create new ICAP Server** icon in the taskbar. The Creating New ICAP Server for ICAP Service window appears. Table 16-4 describes the fields in this window and provides the corresponding CLI global configuration commands.
- Step 7** In the Server Host field, enter the host name or IP address of the ICAP server.
- Step 8** In the Service Port field, enter the port number on which the ICAP server is to be configured to process HTTP requests. This step is optional. The default port number for the ICAP server is 1344. Any valid port number can be set. If no port number is specified, the default port number is used for the ICAP server.
- Step 9** In the Server Service Name field, enter the path to the ICAP service configured on the Content Engine that will be used for the ICAP service by using the URL format `icap://ICAPserverIPAddress/service-name`. The service name entered here must be supported by the ICAP vendor.
- Step 10** In the Maximum Connections field, enter the maximum number of simultaneous connections that can be made to this server. This step is optional. The valid range is 1 to 5000.
- Step 11** In the Weight field, enter the percentage of load that can be redirected to this ICAP server. This step is optional. The valid range is 1 to 100. Do not leave this field blank if you chose the **Weighted** load-balancing method.
- Step 12** In the Keepalive Interval field, specify the keepalive probe interval (in seconds) for this ICAP server. The default is 60 seconds.
- Step 13** To save your settings, click **Submit**.

Table 16-4 ICAP Server for ICAP Service Settings

GUI Parameter	Function	CLI Command
Server Host	Host name or IP address of the ICAP server.	—
Server Port	Port number on which the ICAP server is to be configured to process HTTP requests.	—
Server Service Name	Path to the ICAP server configured on the Content Engine. Use the URL format: <code>icap://ICAPserverIPAddress:port/service-name</code> . The service name entered here must be supported by the ICAP vendor.	icap service server server <i>server_url</i>
Maximum Connections	Maximum number of simultaneous connections that can be made to this server.	icap service server server <i>server_url max-connections</i> <i>num</i>

Table 16-4 ICAP Server for ICAP Service Settings (continued)

GUI Parameter	Function	CLI Command
Weight	Percentage of load that can be redirected to this ICAP server.	icap service server server <i>server-url weight number</i>
Keepalive Interval	Keepalive probe interval (in seconds) for this ICAP server	icap service server server <i>server_url keepalive-interval</i> <i>seconds</i>

About ICAP Processing and Java Applets

When ICAP processing is enabled and an HTTP browser with a streaming Java applet is opened, several undesirable things occur:

- The data for the Java applet is not updated in the browser. For example, when viewing a stock investment website, a user would not see any streaming stock updates.
- The ICAP daemon on the Content Engine continues to send updates (from the HTTP response) to the ICAP server, and it overloads the ICAP server.

These conditions occur because the ICAP server is set up to inspect the entire data packet before it delivers a response to the client. However, because there is a streaming request, the data continues to flow to the ICAP server indefinitely, deadlocking any response to the requesting client.

Two workarounds are available. You can configure the ICAP server to bypass the scanning process, or you can configure rules on the Content Engine to skip ICAP processing on websites that are known to contain streaming Java applets.

To configure the ICAP server to bypass scanning, use rules such as `client_skip_content` or `server_skip_content`.

- The `client_skip_content` rule bypasses scanning on the basis of an HTTP request. The software looks for patterns in the HTTP header and bypasses all requests that exactly match the patterns specified in the `intscan.ini` file, as shown in this example:

```
client_skip_content=User Agent: Windows Media Player 9.0.1
```

- The `server_skip_content` rule bypasses scanning on the basis of an HTTP response. The software looks for patterns in the HTTP header and bypasses all responses that exactly match the patterns specified in the `intscan.ini` file, as shown in this example:

```
server_skip_content=Content-Type: X-Dave_Content
```

Alternatively, you can configure the Content Engine to bypass ICAP processing based on user agents or any of the patterns available in the Rules Template by using the **rule** command. In the following example, the Content Engine is configured to bypass ICAP processing on the intranet site `cisco.com` and on the trusted Internet site `datek.com`:

```
CE(config)# rule enable
CE(config)# rule action use-icap-service trend-reqmod pattern-list 1 protocol all
CE(config)# rule action use-icap-service trend-respmod pattern-list 1 protocol all
CE(config)# rule pattern-list 1 domain "!(.*cisco\.com|.*datek\.com)"
!
CE(config)# icap apply rules-template
CE(config)# icap service trend-reqmod
enable
```

```

vector-point reqmod-precache
server icap://172.19.227.150/REQ-Service
exit
CE(config)# icap service trend-respmod
enable
vector-point respmod-precache
server icap://172.19.227.150/interscan
exit

```

Displaying Information About an ICAP Configuration

To display the current ICAP configuration for a Content Engine, use the **show icap** EXEC command. The command output shows the status of the enabled ICAP features, the service definitions, a list of vectoring points, and an ordered list of ICAP services.

To display the definition and status of a specific ICAP service that is configured on the Content Engine, use the **show icap service service-name** EXEC command.

To display an ordered list of configured ICAP services and their status, use the **show icap vector-point vector-point-name** EXEC command.

```

ContentEngine# show icap vector-point ?
reqmod-postcache  Display reqmod-postcache information
reqmod-precache   Display reqmod-precache information
respmod-precache  Display respmod-precache information

```

Displaying Statistics for ICAP Services

To display ICAP statistics for all of the configured ICAP services, enter the **show statistics icap** EXEC command. This command has no arguments or keywords. There is no default behavior or values.

Configuring Service Rules

The Rules Template feature provides a flexible mechanism to specify configurable caching requests by allowing these requests to be *matched* against an arbitrary number of parameters, with an arbitrary number of *policies* applied against the matches. You can specify a set of rules, each clearly identified by an action and a pattern. Subsequently, for every incoming request, if a pattern for a rule matches the given request, the corresponding action for that rule is taken.

Requests can be matched against regular expressions symbolizing domain names, source IP addresses and network masks, destination IP addresses and network masks, destination port numbers, MIME types, or regular expressions symbolizing a URL.



Note

The processing time on the Content Engine is directly related to the number of service rules configured. Processing times increase with an increase in the total number of rules configured. If the Content Engine processing time is greater than twice the datafeed poll rate, then the device goes offline until the processing is completed. You can avoid having the device go temporarily offline during processing by configuring a higher datafeed poll rate. The recommended datafeed poll rate for 750 service rules is 300 seconds. To configure the datafeed poll rate, see the [“Modifying System Default Properties” section on page 20-18](#).

You can configure service rules by using the Content Distribution Manager GUI or CLI commands. This section provides instructions for using the Content Distribution Manager GUI. For more information on the types of policies that can be applied, actions and patterns, Rules Template processing, and using the CLI to configure service rules, refer to the *Cisco ACNS Software Command Reference, Release 5.5*.

To configure or modify service rule settings, you need to do the following:

- [Enabling Rule Settings, page 16-14](#)
- [Configuring Service Rules, page 16-14](#)

Enabling Rule Settings

Before you configure service rules, you need to enable rule settings for the Content Engine. To enable rule settings, follow these steps:

-
Step 1	From the Content Distribution Manager GUI, choose Devices > Devices . The Devices window appears.
Step 2	Click the Edit icon next to the Content Engine for which you want to enable rule settings. The Content Engine Device Home window appears.
Step 3	In the Contents pane, choose Request Processing > Enable Rules . The Service Rule Settings window appears. Check the Enable check box to enable the use of rule settings.
Step 4	To save the settings, click Submit .
-

To enable rules using the CLI, use the **rule enable** global configuration command.

Configuring Service Rules

Configuring a service rule consists of the following tasks:

1. Configuring a pattern list
2. Adding a pattern to an existing pattern list
3. Associating an action with an existing pattern list

An action is a process that the Content Engine performs while processing the request, for example, blocking the request, redirecting the request, and so on. A pattern defines the limits of the request, for example, a pattern may specify that the IP address must fall within the subnet range 10.0.*.*.

To configure service rule parameters, follow these steps:

-
Step 1	From the Content Distribution Manager GUI, choose Devices > Device Groups . If you have created device groups, the Device Group window appears.
Step 2	Click the Edit icon next to the name of the device group for which you want to configure pattern lists. The Contents pane appears on the left.
Step 3	In the Contents pane, choose Request Processing > Service Rules .
Step 4	In the taskbar, click the Create New Service Rules icon. The Creating New Service Rules window appears.

Step 5 Configure a pattern list and add a pattern to it.

- From the Rule Type drop-down list, choose **pattern-list**.
- In the Rule Parameters field, configure the pattern list number and the pattern type, following the rules usage guidelines shown in the GUI. (See [Table 16-5](#) for a description of pattern types.)

For example, to create pattern list number 72 with the pattern type *domain* and the yahoo.com domain as the domain to be acted on by an action, enter **72 domain \.yahoo.com** in the Rule Parameters field. (See [Figure 16-3](#).)

Figure 16-3 Creating a New Pattern List and Defining a Pattern




Table 16-5 Pattern Types

Pattern Type	Description	CLI Command
domain	Matches the domain name in the URL or the Host header against a regular expression. For example, “*.ibm.*” matches any domain name that contains the “ibm” substring. “\foo\com\$” matches any domain name that ends with the “foo.com” substring. In regular expression syntax, the dollar sign “\$” metacharacter directs that a match is made only when the pattern is found at the end of a line.	rule pattern-list <i>list_num</i> domain <i>dn_regex</i>
dst-ip	Matches the request’s destination IP address and netmask against the specified destination IP address and netmask. Specify an IP address and a netmask. In proxy mode, the Content Engine does a DNS lookup to resolve the destination IP address of the HTTP request, making the response time longer and possibly negating the benefit of setting a dst-ip rule. When an outgoing proxy is configured, cache miss requests are forwarded by the Content Engine to the outgoing proxy without examination of the destination server IP address, making the dst-ip rule unenforceable on the first Content Engine.	rule pattern-list <i>list_num</i> dst-ip <i>d_ipaddress d_subnet</i>
dst-port	Matches the request’s destination port number against the specified destination port number. Specify a port number.	rule pattern-list <i>list_num</i> dst-port <i>port</i>
groupname	Matches the groupname of the end user (web client that is requesting content) who was authenticated through LDAP or NTLM. This pattern can be applied only to request authentication for users who have been authenticated through LDAP or NTLM. Supports exact string comparison. The groupname comparison is case insensitive. Maximum length of groupname is 255 characters. Valid characters are an underscore and alphanumeric characters. If the groupname configuration in the Rules Template and the group name-based access list match, then the access list takes precedence. Tip If you intend to use the groupname pattern, make sure that you set the correct number of maximum group entries in the authentication group cache (the http authentication cache max-group-entries <i>number</i> global configuration command). This number should correspond to the maximum number of groups that could be returned during authorization queries (for example, the total number of groups defined on the AAA server.) The number can be from 500 to 12000. The number of entries in the authentication group cache is dependent on the physical resources available on the Content Engine.	rule pattern-list <i>list_num</i> groupname <i>name</i>

Table 16-5 Pattern Types (continued)

Pattern Type	Description	CLI Command
groupname-regex	Matches the group-name of the end user (web client that is requesting content) against a regular expression. Specify a regular expression. For example, use to configure a regular expression-based policy on groupnames or to OR multiple groupnames in the same line of in a single pattern list. For example, to OR three groupnames enter this command: <pre>ContentEngine(config)# rule pattern-list 1 groupname-regex Engineering Marketing Finance\</pre> The specified action takes effect if any of the above groupnames matched the one in the request.	<pre>rule pattern-list list_num groupname-regex group_name_regex</pre>
group-type	Specifies whether the pattern list is an AND or OR type. The default is OR.	<pre>rule pattern-list list_num group-type {and or}</pre>
header-field	Request header field pattern. Request header field patterns referer, request-line, and user-agent are supported for actions block, reset, redirect, and rewrite. The referer pattern is matched against the Referer header in the request, the request-line pattern is matched against the first line of the request, and the user-agent pattern is matched against the User-Agent header in the request. Request header field pattern x-forwarded-for is supported for the use-server rule action only. This pattern is matched against the x-forwarded-for IP address in the HTTP header. The use-server action sends server-style HTTP requests from the Content Engine to the specified IP address and port on a cache miss.	<pre>rule pattern-list list_num header-field {referer ref_regex request-line req_regex user-agent ua_regex x-forwarded-for ipaddress}</pre>
header-field-sub	Request header field subpattern and substitute replacement pattern.	<pre>rule pattern-list list_num header-field-sub {referer ref_regex ref_sub request-line req_regex req_sub user-agent ua_regex ua_sub}</pre>
icap-attribute	Specifies the attribute and value pair of the ICAP service.	<pre>rule pattern-list list_num icap-attribute icap_attribute icap_value</pre>
mime-type	Matches the MIME type of the response. Specify a MIME type string, for example, “image/gif,” as defined in RFC 2046 (http://www.faqs.org/rfcs/rfc2046.html). The administrator can specify a substring, for example, “java”, and have it apply to all MIME types with the “java” substring, such as “application/x-javascript.”	<pre>rule pattern-list list_num mime-type mt_regex</pre>
src-ip	Matches the request’s source IP address and netmask. Specify an IP address and a netmask.	<pre>rule pattern-list list_num src-ip s_ipaddress s_subnet</pre>

Table 16-5 Pattern Types (continued)

Pattern Type	Description	CLI Command
url-regex	Matches the URL against a regular expression. The match is case insensitive. Specify a regular expression that has syntax which can be found at the following URL: http://yenta.www.media.mit.edu/projects/Yenta/Releases/Documentation/regex-0.12/ .	rule pattern-list <i>list_num</i> url-regex <i>url_regex</i>
url-regsub	For the rewrite and redirect actions, matches the URL against a regular expression to form a new URL in accordance with the pattern substitution specification. The match is case insensitive. The valid substitution index range is from 1 to 9.	rule pattern-list <i>list_num</i> url-regsub <i>url_regex</i> <i>url_sub</i>
username	Matches the username of the end user (the web client that is requesting content) who was authenticated through LDAP, NTLM, RADIUS, or TACACS+. Specify the username or usernames. Maximum length of username is 255 characters for LDAP, RADIUS, or TACACS+ authentication. See below for information about the maximum length of usernames for NTLM authentication. Valid characters are an underscore and alphanumeric characters. The match supports exact string comparison. The username comparison is case insensitive. To specify multiple usernames in the same line for the same pattern list use a delimiter, as shown in this example: <pre>ContentEngine(config)# rule pattern-list 1 username jdoe8,dsmith7,jsmith50</pre> By default, the match does not consider the domain name, and matches for username only. To include domain name as well as username in the match, specify <i>domainname\username</i>, as shown in this example: <pre>ContentEngine(config)# rule pattern-list 1 username domain cisco\jdoe8</pre> For NTLM authentication, the domain\username:password:NTLM string must be 50 characters or less. If this string is greater than 50 characters, the domain name is truncated and the rule username pattern is not matched. An error message is generated in the system log in this situation. To match all users in a particular domain, enter this command: <pre>ContentEngine(config)# rule pattern-list 1 username domain domainname*</pre> where <i>domainname</i> is the name of the domain (for example, cisco).	rule pattern-list <i>list_num</i> username <i>user_name</i>

Step 6 To save the settings, click **Submit**.

Step 7 Next, associate an action with an existing pattern list.

- From the Creating New Service Rule window (see [Step 1](#) through [Step 4](#)), choose an action type from the Rule Type drop-down list. (See [Table 16-6](#) for a description of rule actions.)
- In the Rule Parameter field, enter the list number of the pattern list that you want associated with this action.

For example, if you want to block access by any protocol to yahoo.com, then choose **block** from the Rule Type drop-down list, and enter **pattern-list 72 protocol all** in the Rule Parameters field. (See Figure 16-4.)

Figure 16-4 Associating an Action with an Existing Pattern List




Table 16-6 Rule Actions

Action Type	Description	CLI Command
allow	Allows incoming requests that match the pattern list. This rule action can be used in combination with reset or block actions to allow selective types of requests. Allow does not carry any meaning as a standalone action.	rule action allow pattern-list <i>list_num</i> [protocol {all http https rtsp}]
append-username-header	Appends the username in the request sent to the origin server.	rule action append-username-header pattern-list <i>list_num</i> [protocol {all http https}]
block	Blocks this request and allows all others.	rule action block pattern-list <i>list_num</i> [protocol {all http https rtsp}]

Table 16-6 Rule Actions (continued)

Action Type	Description	CLI Command
cache-non-cacheable	Overrides the HTTP response headers and caches the objects. This rule action caches objects only if they are not authenticated. For authenticated objects, some origin servers do not send the Last-Modified and ETag entity headers, and revalidation of these objects cannot be performed by the Content Engine. These authenticated objects are served only from the origin server. If the server <i>does</i> send the Last-Modified and ETag headers, then these objects can be cached.	<pre>rule action cache-non-cacheable ttl {day days pattern-list list_num [protocol {all http https}] hours hours pattern-list list_num [protocol {all http https}] minutes minutes pattern-list list_num [protocol {all http https}] seconds seconds pattern-list list_num [protocol {all http https}]}</pre>
cache-only	Caches objects depending on the HTTP response headers. Caches this object only if it is a match and is allowed to be cached by HTTP. If one or more rules specify this action, an object is cached if it matches at least one of the selective-cache rules and passes every other caching restriction, such as the object-size check and the no-cache-on-authenticated-object check. If the object does not match any of the selective-cache rules, the object is not cached.	<pre>rule action cache-only pattern-list list_num [protocol {all http https}]</pre>

Table 16-6 Rule Actions (continued)

Action Type	Description	CLI Command
dscp	Configures the IP ToS or DSCP code point field. cache-miss—Sets the IP ToS/DSCP code point bits for the client-side connection to the configured value for cache miss responses to the client. cache-hit—Sets the IP ToS or DSCP code point bits for the client-side connection to the configured value for cache hit responses to the client. Setting the Type of Service (ToS) or differentiated services code point (DSCP) is called packet marking, which allows you to partition network data into multiple priority levels or types of service. With the ACNS 5.x releases, you can set the ToS or DSCP values in IP packets based on a URL match, a file type, a domain, a destination IP address, a source IP address, or a destination port. You can set specific ToS or DSCP values for the following: • Requests from the Content Engine to the server • Responses to the client on a cache hit • Responses to the client on a cache miss The ToS or DSCP can be set based on any of the policies matching the src-ip <i>s_ipaddress s_subnet</i>, dst-ip <i>d_ipaddress d_subnet</i>, dst-port <i>port</i>, domain <i>LINE</i>, url-regex <i>LINE</i>, or mime-type <i>LINE</i> options. In addition, you can now configure global ToS or DSCP settings with the ip dscp command. The Rules Template configuration takes precedence over the ip dscp command, and the url-filter command takes precedence over the rule command to the extent that even the rule no-block command is executed only if the url-filter command has not blocked the request.	rule action dscp client cache-hit {match-server pattern-list list_num [protocol {all http https}] set-dscp <i>dscpvalue</i> set-tos tovalue} rule action dscp client cache-miss {match-server pattern-list list_num [protocol {all http https}] set-dscp <i>dscpvalue</i> set-tos tovalue}
freshness-factor	Determines the Time To Live if the request URL matches a specified regular expression. The refresh configuration takes priority over freshness-factor configurations.	rule action freshness-factor exp_time pattern-list list_num [protocol {all http https}]
generate-url-signature	Generates the URL signatures in the Windows Media metafile response associated with pre-positioned content, based on the Content Engine configuration for the url-signature and this rule action.	rule action generate-url-signature [include-client-src-ip] key-id-owner <i>owner_num</i> key-id-number id_num pattern-list list_num [protocol {all http}]
insert-no-cache	Inserts a no-cache header in the response.	rule action insert-no-cache pattern-list <i>list_num</i> [protocol {all http https}]
no-auth	Does not authenticate.	rule action no-auth pattern-list list_num [protocol {all http https rtsp}]

Table 16-6 Rule Actions (continued)

Action Type	Description	CLI Command
no-cache	Does not cache this object. If both the no-cache and selective-cache actions are matched, no-cache takes precedence.	rule action no-cache pattern-list <i>list_num</i> [protocol { all http https }]
no-proxy	For a cache miss, does not use the configured upstream proxy but contacts the server directly.	rule action no-proxy pattern-list <i>list_num</i> [protocol { all http https }]
redirect	Redirects the original request to a specified URL. Redirect is relevant to the RADIUS server only if the RADIUS server has been configured for redirect.	rule action redirect <i>url</i> pattern-list <i>list_num</i> [protocol { all http https rtsp }]
redirect-url-for-cdn	Redirects the original request to a specified URL for the ACNS network.	rule action redirect-url-for-cdn pattern-list <i>list_num</i> [protocol { all http https rtsp }]
refresh	For a cache hit, forces an object freshness check with the server.	rule action refresh pattern-list <i>list_num</i> [protocol { all http https }]
reset	Issues a TCP RST. This reset request is useful when resetting Code Red or Nimda virus requests.	rule action reset pattern-list <i>list_num</i> [protocol { all http https rtsp }]
rewrite	Rewrites the original request as a specified URL.	rule action rewrite pattern-list <i>list_num</i> [protocol { all http https rtsp }]
use-dns-server	Caches this object only if it is a match and is allowed to be cached by HTTP. If one or more rules specify this action, an object is cached if it matches at least one of the selective-cache rules and passes every other caching restriction such as the object-size check and the no-cache-on-authenticated-object check. If the object does not match any of the selective-cache rules, the object is not cached.	rule action use-dns-server { <i>hostname</i> <i>ip-address</i> } pattern-list <i>list_num</i> [protocol { all http https }]
use-icap-service	Applies ICAP processing and uses a specific ICAP service only for those requests that match this Rules Template action. An ICAP service is a collection of attributes that defines the type of modification to be performed on HTTP requests and responses. If this action is configured, you can allow requests and responses to be processed by ICAP servers for content adaptation.	rule action use-icap-service <i>service-name</i> pattern-list <i>list_num</i> [protocol { all http https }]
use-proxy	For a cache miss, uses a specific upstream proxy. Specify the upstream proxy IP address (or domain name) and port number. If both no-proxy and use-proxy are matched, no-proxy takes precedence.	rule action use-proxy { <i>hostname</i> <i>ip-address</i> } <i>port</i> pattern-list <i>list_num</i> [protocol { all http https }]
use-server	Sends server-style HTTP requests from the Content Engine to the specified IP address and port on a cache miss.	rule action use-server { <i>hostname</i> <i>ip-address</i> } <i>port</i> pattern-list <i>list_num</i> [protocol { all http https }]

Table 16-6 Rule Actions (continued)

Action Type	Description	CLI Command
use-xforward-clt-ip	Uses client IP address in the forwarded header for filtering.	rule action use-xforward-clt-ip pattern-list <i>list_num</i> [protocol {all http https}]
validate-url-signature	Validates the URL signature for a request using the configuration on your Content Engine for the url-signature and allows the request processing to proceed for this request.	rule action validate-url-signature error-redirect-url <i>url</i> pattern-list <i>list_num</i> [protocol {all http rtsp}]

Step 8 To save the settings, click **Submit**.

Configuring the Rules Template Using CLI Commands

These sections describe how to configure pattern lists and actions for the Rules Template using CLI commands.

- [Configuring a Pattern List, page 16-23](#)
- [Adding a Pattern to an Existing Pattern List, page 16-24](#)
- [Associating an Action with an Existing Pattern List, page 16-24](#)

Configuring a Pattern List

To create a new pattern list, follow these steps:

	Command"	Purpose
Step 1	ContentEngine(config)# rule enable	Enables the Rules Template.
Step 2	ContentEngine(config)# rule pattern-list <i>1-512</i>	Creates a pattern list.
Step 3	ContentEngine# show rule pattern-list <i>1-512</i> pattern-type <i>pattern</i>	Displays the Rules Template configuration.

In the following example, the **rule pattern-list** command is configured to create a pattern list to block all domains that contain .foo.com in the URL request using the domain \.foo.com pattern.

```
ContentEngine(config)# rule pattern-list 10 domain foo.com
ContentEngine# show rule pattern-list 10 domain
Rules Template Configuration
-----
Rule Processing Enabled

Pattern-Lists :

rule pattern-list 10 domain foo.com
ContentEngine#
```

Adding a Pattern to an Existing Pattern List

To add a new pattern to an already existing pattern list, follow these steps:

	Command	Purpose
Step 1	<code>ContentEngine(config)# rule pattern-list 1-512 pattern-type pattern</code>	Adds a pattern to a pattern list.
Step 2	<code>ContentEngine# show rule pattern-list 1-512 pattern-type pattern</code>	Displays the Rules Template configuration.

In the following example, the **rule pattern-list** command is configured to add a pattern to an existing pattern list to perform an action yet to be defined on the destination IP address 172.16.25.25 using the **dst-ip** pattern.

```
ContentEngine(config)# rule pattern-list 10 dst-ip 172.16.25.25 255.255.255.0
ContentEngine# show rule pattern-list 10 all
Rules Template Configuration
-----
Rule Processing Enabled

Pattern-Lists :

rule pattern-list 11 dst-ip 172.16.25.25 255.255.255.0
rule pattern-list 11 domain foo.com
ContentEngine#
```

Associating an Action with an Existing Pattern List

To associate an action with an existing pattern list, follow these steps:

	Command	Purpose
Step 1	<code>ContentEngine(config)# rule action action_type pattern-list 1-512 protocol protocol_type all</code>	Associates an action with an existing pattern list.
Step 2	<code>ContentEngine# show rule action action_type protocol protocol_type all</code>	Displays the Rules Template configuration.

In the following example, the **rule action block** command is configured and associated with an existing pattern list.

```
ContentEngine(config)# rule action block pattern-list 10 protocol all
ContentEngine# show rule action block
Rules Template Configuration
-----
Rule Processing Enabled

Actions :

rule action block pattern-list 10 protocol all
ContentEngine#
```

The **rule action use-proxy proxy pattern-list number** global configuration command can be used to configure only one proxy for a particular pattern list. If the use-proxy feature is configured without failover (for example, you have entered the **rule action use-proxy 10.16.0.0 8080 pattern-list 1**

command), the Content Engine will send the request to the use-proxy (the server in this example with the IP address 10.16.0.0). If the Content Engine does not obtain a response from the use-proxy, then it will send an error message to the client without failing over to the HTTP outgoing proxy.

If the use-proxy feature is configured with failover (for example, you have entered the **rule action use-proxy 10.16.0.0 8080 failover pattern-list 1** command), the Content Engine will send the request to the use-proxy (in this example, the server with the IP address of 10.16.0.0). If the Content Engine does not obtain a response from the use-proxy, then it fails over to the specified HTTP outgoing proxy.

Verifying an Action Performed on a Pattern List

To verify the response sent by the Content Engine to confirm that a certain action is performed on a pattern list, follow these steps:

	Command	Purpose
Step 1	<code>ContentEngine(config)# rule action action_type pattern-list 1-512 protocol protocol_type all</code>	Associates an action with an existing pattern list.
Step 2	<code>ContentEngine# show rule action action_type protocol protocol_type all</code>	Displays the Rules Template configuration after a new action has been added.
Step 3	<code>ContentEngine# show statistics rule action action_type</code>	Displays the local Rules Template configuration statistics after a request is issued on which an action should be performed.

In the following example, the **rule action block** command is configured and associated with an existing pattern list, which lists as its pattern the domain yahoo.com.

```
ContentEngine(config)# rule pattern-list 30 domain yahoo.com
ContentEngine(config)# rule action block pattern-list 30 protocol all
ContentEngine# show statistics rule action block
Rules Template Statistics
-----
Rule hit count = 3    Rule: rule action block pattern-list 30 protocol all
ContentEngine#
```

In this example, the request to yahoo.com was denied three times.

Configuring URL Filtering

Some enterprises have a requirement to monitor, manage, and restrict employee access to nonbusiness and objectionable content on the Internet. Employees or students can be allowed or denied access to websites or can be coached with information about acceptable use of the Internet. By having a URL filtering scheme on Content Engines, organizations get an immediate return on investment by increasing productivity and recapturing network bandwidth, while reducing legal liability.

The URL filtering features presented in this section allow the Content Engine to control client access to websites in any of the following ways:

- Deny access to URLs specified in a list.
- Permit access only to URLs specified in a list.
- Direct traffic to an N2H2 server for filtering.
- Direct traffic to a Websense enterprise server for filtering.

For information about configuring the Websense software, go to the following website:
<http://www.websense.com>

- Filter traffic with Secure Computing Corporation SmartFilter Software, Release 3.2 (HTTP traffic only).

For information about configuring the SmartFilter software, go to the following website:
<http://www.securecomputing.com>

**Note**

Although only one form of URL filtering scheme per protocol can be active, many URL filtering schemes can be supported at one time. In other words, if an N2H2 filter is applied to HTTP URLs, no other URL filtering scheme, such as Websense, or SmartFilter, can be applied to this protocol. However, the use of good and bad lists can be applied to the streaming media protocol. The scheme enabled for a particular protocol is independent of that of other protocols.

Configuring URL Filter Settings Using the Content Distribution Manager GUI

To configure URL filter settings for the Content Engine, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** In the Contents pane, choose **Request Processing > URL Filter**. The URL Filter Settings window appears. (See [Figure 16-5](#).)

Figure 16-5 URL Filter Settings Window

URL Filter Settings for Content Engine, CONTENTENGINE

Current applied settings from Content Engine, **CONTENTENGINE**

General URL Filter Settings

FTP Server Hostname: FTP Server Username:

FTP Server Password: Confirm Password:

Remote Custom Msg File Pathname: Custom Msg Directory:

Protocol URL Filter Settings

Protocol	Filter Type	Filter Param
HTTP	None Enabled	
RTSP	None Enabled	
Windows Media	None Enabled	

Note: * - Required Field

Submit Cancel

- Step 4** Enter information for the general URL filter settings. (See [Table 16-7](#) for descriptions of the filter-setting parameters.)
- Step 5** To create custom blocking messages, do the following:
- In the Remote Custom Msg File Pathname field, enter the path name for the remote file that contains the custom message directory.
 - In the Custom Msg Directory field, enter the directory name that contains the *block.html* file. (For more information, see the “[Configuring Custom Blocking Messages Using the CLI](#)” section on [page 16-31](#).)
- Step 6** Under the Protocol URL Filter Settings heading, click the **Edit** icon next to the name of the filter type and enter information for HTTP URL filter settings, RTSP URL filter settings, or WMT filter settings in the appropriate fields. (See [Table 16-7](#).)
- Step 7** To save the settings, click **Submit**.

Table 16-7 URL Filter Settings Window

GUI Parameter	Function	CLI Command
General URL Filter Settings		
FTP Server Hostname	DNS name or IP address of the FTP server from which the URL filter files are downloaded.	—

Table 16-7 URL Filter Settings Window (continued)

GUI Parameter	Function	CLI Command
FTP Server Username	Name needed to access the FTP server from which the URL filter files are downloaded.	—
FTP Server Password	Password of the FTP server from which the URL filter files are downloaded.	—
Confirm Password	Confirms the FTP server password.	—
Remote Custom Msg File Pathname	Path name of the remote file that contains the custom message directory.	—
Custom Msg Directory	Creates a customized URL blocking message to display to the client. This custom message must be an administrator-created HTML file named <i>block.html</i> .	url-filter http custom-message <i>dirname</i>
Protocol URL Filter Settings for HTTP, RTSP, and WMT		
Enable Bad Site Filtering	Enables the use of local list filtering for bad sites.	url-filter http bad-sites-deny enable url-filter rtsp bad-sites-deny enable url-filter wmt bad-sites-deny enable
Remote Bad Site File Pathname	Path name of the remote bad site file.	—
Bad Site Filename	File containing URLs to which access is denied.	url-filter http bad-sites-deny file <i>filename</i> url-filter rtsp bad-sites-deny file <i>filename</i> url-filter wmt bad-sites-deny file <i>filename</i>
Enable Good Site Allow	Enables URL filtering of the local good sites list over HTTP.	url-filter http good-sites-allow enable url-filter rtsp good-sites-allow enable url-filter wmt good-sites-allow enable
Remote Good Site File Pathname	Path name of the remote good site file.	—
Good Site Filename	File containing URLs to which access is permitted.	url-filter http good-sites-allow file <i>filename</i> url-filter rtsp good-sites-allow file <i>filename</i> url-filter wmt good-sites-allow file <i>filename</i>

Table 16-7 URL Filter Settings Window (continued)

GUI Parameter	Function	CLI Command
FTP Server Hostname	DNS name or IP address of the FTP server from which the URL filter files are downloaded.	—
FTP Server Username	Name needed to access the FTP server from which the URL filter files are downloaded.	—
FTP Server password	Password of the FTP server from which the URL filter files are downloaded.	—
Confirm password	Confirms the FTP server password.	—
Enable SmartFilter Filtering	Enables the use of SmartFilter software.	url-filter http smartfilter enable
Enable N2H2 Filtering	Enables the use of an N2H2 server for URL filtering.	url-filter http N2H2 enable
N2H2 Server Hostname	Host name or IP address of the N2H2 server.	url-filter http N2H2 server
N2H2 Port	Port number on which the N2H2 server is accepting requests.	url-filter http N2H2 server <i>IPaddress or hostname</i> port
Enable N2H2 Allow Mode	Allows the request to be served if there is no response from the N2H2 server.	url-filter http N2H2 allowmode enable
N2H2 Request Timeout	Number of seconds that the Content Engine should wait for a response from the N2H2 server.	url-filter http N2H2 server {<i>ipaddress</i> <i>hostname</i>} port 1-65535 timeout
Enable Websense Filtering	Enables the use of a Websense server for URL filtering.	url-filter http websense enable
Embedded	Enables the Websense server on the Content Engine. Ensures that the URL filtering software uses the local Websense server and not a remote host as the Websense server.	websense-server enable
WebSense Server Hostname	Host name or IP address of an external Websense server	url-filter http websense server <i>hostname</i>
WebSense Port	Port number on which the Websense server is accepting requests.	url-filter http websense server {<i>ipaddress</i> <i>hostname</i>} port 1-65535
Enable WebSense Allow Mode	Allows the request to be served if there is no response from the Websense server.	url-filter http websense allowmode enable

Table 16-7 URL Filter Settings Window (continued)

GUI Parameter	Function	CLI Command
WebSense Request Timeout	Number of seconds the Content Engine should wait for a response from the Websense server.	url-filter http websense server {ipaddress hostname} port 1-65535 timeout
Websense Request Connections	Number of persistent connections to the Websense server. The range is 1–250 connections per CPU. The default is 40 connections. Do not change the default unless you are certain that a different value is required.	url-filter http websense server {ipaddress hostname} port 1-65535 timeout seconds connections 1-250

Configuring URL Filters Using the CLI



Note

Replace HTTP with either RTSP or WMT in the examples that follow if you want to use URL lists on these protocols. Make sure the URLs that are either accessed or denied match the protocol used in the command.

Use the **url-filter wmt** global configuration commands to configure local list URL filtering for WMT requests. If you configure an RTSP URL in a WMT bad sites list, then it blocks both the RTSP and RTSPU URLs, as well as the RTSP URL that is specified in the bad sites list.

To deny requests for specific HTTP URLs using the CLI, follow these steps:

Step 1 Create a plain text file named badurl.lst.

In this file, enter the HTTP URLs that you want to block. The list of URLs in the badurl.lst file must be written in the form www.domain.com and delimited with carriage returns.

Step 2 Copy the badurl.lst file to the /local1 sysfs directory of the Content Engine.



Note

We recommend creating a separate directory under local1 to hold the bad and good lists, for example, /local1/filtered_urls.

Step 3 Point to the badurl.lst file by entering the following command:

```
ContentEngine(config)# url-filter http bad-sites-deny local/local1/badurl.lst
```

Step 4 Deny the bad URLs by entering the following command:

```
ContentEngine(config)# url-filter http bad-sites-deny enable
```

Use the **no** form of this command to disable URL blocking as follows:

```
ContentEngine(config)# no url-filter http bad-sites-deny enable
```

To permit specific HTTP URLs to the exclusion of all other URLs using the CLI, follow these steps:

Step 1 Create a plain text file named `goodurl.lst`.

In this file, enter the HTTP URLs that you want to exclusively allow. The list of URLs in the `goodurl.lst` file must be written in the form `www.domain.com` and delimited with carriage returns.

Step 2 Copy the `goodurl.lst` file to the `/local1 sysfs` directory of the Content Engine.



Note We recommend creating a separate directory under `local1` to hold the bad and good sites lists, for example, `/local1/filtered_urls`.

Step 3 Point to the `goodurl.lst` file by entering the following command:

```
ContentEngine(config)# url-filter http good-sites-allow local/local1/goodurl.lst
```

Step 4 Permit only the good URLs by entering the following command:

```
ContentEngine(config)# url-filter http good-sites-allow enable
```

Use the **no** form of this command to disable the allowance of good URLs as follows:

```
ContentEngine(config)# no url-filter http good-sites-allow enable
```

Configuring Custom Blocking Messages Using the CLI

The Content Engine can be configured to return a customized blocking message to the client. The custom message must be an administrator-created HTML file named `block.html`. Make sure to copy all embedded graphics associated with the custom message HTML window to the same directory that contains the `block.html` file. To enable the customized blocking message, use the **url-filter http custom-message** command, and specify the directory name.

To disable the custom message, use the **no url-filter http custom-message** command.

The **url-filter http custom-message** command can be enabled and disabled without affecting the **good-sites-allow** and **bad-sites-deny** configuration.



Note Do not use `local1` or `local2` as directories for custom blocking messages. Create a separate directory under `local1` or `local2` for holding the custom message file.

In this example, a `block.html` file displays the following custom message when the Content Engine intercepts a request to the blocked site:

```
This page is blocked by the Content Engine
```

In the `block.html` file shown, objects (such as `.gif`, `.jpeg`, and so on) must be referenced within the custom message directory string `/content/engine/blocking/url`, as shown in the example below.



Note Contact your administrator if you have any questions concerning access to the blocked site that you requested.

```
<TITLE>Cisco Content Engine example customized message for url-filtering</TITLE>
<p>
```

```

<H1>
<CENTER><B><I><BLINK>
<FONT COLOR="#800000">P</FONT>
<FONT COLOR="#FF00FF">R</FONT>
<FONT COLOR="#00FFFF">A</FONT>
<FONT COLOR="#FFFF00">D</FONT>
<FONT COLOR="#800000">E</FONT>
<FONT COLOR="#FF00FF">E</FONT>
<FONT COLOR="#00FFFF">P</FONT>
<FONT COLOR="#FF8040">'</FONT>
<FONT COLOR="#FFFF00">S</FONT>
</BLINK>
<FONT COLOR="#0080FF">Blocked Page</FONT>
</I></B></CENTER>
</H1>
<p>
<p>
<IMG src="/content/engine/blocking/url/my.gif">
<p>
This page is blocked by the Content Engine.
<p>

```

If the block.html file is updated, it will automatically display its new message without your having to reenter the **url-filter http custom-message** command.

Creating a Text File URL List for URL Filtering

You can configure the Content Engine to deny client requests for URLs that are listed in a badurl.lst file, or you can configure it to fulfill only requests for URLs in a goodurl.lst file.

The use of URL lists applies to requests in HTTP, HTTPS, and FTP format as well as streaming media protocols such as RTSP.



Note

The local list file for each protocol should not contain URLs that belong to other protocols. For example, the HTTP local list file should contain only HTTP, HTTPS, or FTP URLs.



Caution

If the size of the local list file is too large, it can affect device performance because the file is loaded into memory when local list file filtering is enabled. If the size of the file is larger than 5 megabytes, the device issues a warning to notify you of its impact on performance.

To deny requests for specific HTTP URLs, follow these steps:

Step 1 Create a plain text file named badurl.lst.

In this file, enter the URLs that you want to block. You must write the list of URLs in the badurl.lst file in the form `http://www.domain.com/` delimited with carriage returns.

Step 2 Copy the badurl.lst file to the /local1 system file system (sysfs) directory of the Content Engine.



Note

We recommend creating a separate directory under local1 to hold the bad lists, for example, /local1/filtered_urls.

Step 3 Use the **url-filter http bad-sites-deny file** command to point to the bad URL list.

```
Console(config)# url-filter http bad-sites-deny file local/local1/badurl.lst
```

Step 4 Use the **url-filter http bad-sites-deny enable** command to actively deny the URLs.

```
Console(config)# url-filter http bad-sites-deny enable
```

To permit specific HTTP URLs to the exclusion of all other URLs, follow these steps:

Step 1 Create a plain text file named goodurl.lst.

In this file, enter the URLs that you want to exclusively allow. You must write the list of URLs in the goodurl.lst file in the form `http://www.domain.com/` delimited with carriage returns.

Step 2 Copy the goodurl.lst file to the /local1 sysfs directory of the Content Engine.



Note We recommend that you create a separate directory under local1 to hold the good lists, for example, /local1/filtered_urls.

Step 3 Use the **url-filter http good-sites-allow file** command to point to the goodurl.lst file.

```
Console(config)# url-filter http good-sites-allow file local/local1/goodurl.lst
```

Step 4 Use the **url-filter http good-sites-allow enable** command to actively permit only the good URLs.

```
Console(config)# url-filter http good-sites-allow enable
```



Note Only one good sites file or one bad sites file can be active at a time for each protocol.



Note When you update the badurl.lst or goodurl.lst file, use the **url-filter local-list-reload EXEC** command to recopy the URL list file from any protocol to the Content Engine.

Use the **no** form of the command to disable blocking, Websense, or N2H2 permission requests (for example, **no url-filter bad-sites-deny**).

Configuring URL Filtering with the N2H2 Server




Note URL filtering with the N2H2 server applies only to HTTP, FTP, or HTTPS protocols.

N2H2 is a globally deployed URL-filtering software that can filter HTTP, FTP, or HTTPS requests based on destination host name, destination IP address, and username and password. It relies on a sophisticated URL database exceeding 15 million sites and is organized into over 40 categories using both Internet technology and human review.

The Content Engine can perform URL filtering by using the N2H2 server. (See [Figure 16-6](#).) The Content Engine and the N2H2 server use the Internet Filtering Protocol (IFP) version 1 to communicate with each other. When the Content Engine receives a URL request, it sends an IFP request to the N2H2 server. The N2H2 server does some necessary lookups for the URL and sends back an IFP response. Based on the N2H2 server's IFP response, the Content Engine either blocks the HTTP request by redirecting the browser to a block page or proceeds with normal HTTP processing by sending the URL request to an origin server.

Figure 16-6 N2H2 Filtering



URL filtering that uses an N2H2 server is applied to HTTP, FTP, or HTTPS traffic before the service rule mechanism is applied, regardless of whether or not the requested object is in the cache. Filtering is applied to these traffic types:

- Proxy-style or transparent-style HTTP or HTTPS requests
- Proxy-style and transparent redirect proxy-style FTP over HTTP requests

N2H2 Features Supported

N2H2 supports three filtering methods. [Table 16-8](#) lists the N2H2 features supported by the Content Engine. One N2H2 server can support multiple Content Engines simultaneously.

Table 16-8 N2H2 Features Supported

N2H2 Feature Name	Description
Global filtering	Applies filtering to all HTTP, FTP, or HTTPS requests.
User-based filtering	Applies filtering to specific users or groups.
Client IP-based filtering	Applies filtering to specific client IP addresses.
Transparent Authentication	Performs transparent authentication by passing back the initial response header to the client using the HTML page in the IFP responses.

Enabling N2H2 Filtering Using the Content Distribution Manager GUI

To configure N2H2 filter settings for the Content Engine, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** Choose **Request Processing > URL Filter**. The URL Filter Settings window appears. (See [Figure 16-5](#).)
- Step 4** Under the Protocol URL Filter Settings heading, click the **Edit Filter** icon next to the HTTP protocol. The URL Filter Settings for HTTP Protocol window appears. (See [Figure 16-7](#).) [Table 16-9](#) describes the fields in this window and provides the corresponding CLI global configuration commands.

Figure 16-7 URL Filter Settings—HTTP URL Filter Settings

- Step 5** To save the N2H2 settings, click **Submit**.

Table 16-9 URL Filtering Using the N2H2 Server

GUI Parameter	Function	CLI Command
Enable N2H2 Filtering	Enables the use of an N2H2 server for URL filtering.	<code>url-filter http N2H2 enable</code>
N2H2 Server Hostname	Host name or IP address of the N2H2 server.	<code>url-filter http N2H2 server</code>

Table 16-9 URL Filtering Using the N2H2 Server (continued)

GUI Parameter	Function	CLI Command
N2H2 Port	Port number on which the N2H2 server is accepting requests.	url-filter http N2H2 server IPaddress or hostname port
Enable N2H2 Allow Mode	Allows the request to be served if there is no response from the N2H2 server.	url-filter http N2H2 allowmode enable
N2H2 Request Timeout	Number of seconds that the Content Engine should wait for a response from the N2H2 server.	url-filter http N2H2 server IPaddress or hostname port 1-65535 timeout

Enabling N2H2 Filtering Using the CLI

To configure a Content Engine to use an external N2H2 server for URL filtering from the CLI, follow these steps:

- Step 1** Display the URL filtering schemes that are currently enabled on this Content Engine for requests over HTTP by entering the following command:

```
Content Engine# show url-filter http
URL filtering is DISABLED

Local list configurations
=====
Good-list file name :
Bad-list file name : /local1/badfile.txt
Custom message directory :

Websense server configuration
=====
Websense server IP      : 2.42.0.144
Websense server port    : 15868
Websense server timeout: 20 (in seconds)
Websense server connections: 40
Websense server IP      : <none>
Websense server port    : 15868
Websense server timeout: 20 (in seconds)
Websense server connections: 40
Websense allow mode is ENABLED

N2H2 server configuration
=====
N2H2 server IP          : <none>
N2H2 server port        : 4005
N2H2 server timeout     : 5 (in seconds)
N2H2 allow mode is ENABLED
```

- Step 2** Make sure that no other URL filtering scheme (for example, Websense or SmartFilter software) is currently enabled for requests over HTTP.

Only one URL filtering scheme per protocol can be active at a time.

- Step 3** Configure the Content Engine to use an external N2H2 server for URL filtering by entering the following command:

```
ContentEngine(config)# url-filter http N2H2 server { [hostname | ip-address] } [port portnum  
[timeout seconds]]
```

where

- *hostname* is the hostname of the external N2H2 server.
- *IP address* is the IP address of the external N2H2 server.
- *portnum* is the port number (1–65535) to which the Content Engine sends the IFP requests to the specified N2H2 server. The default port number is 4005.
- *seconds* is the number of seconds (1–120) that the Content Engine is to wait for an IFP response from the N2H2 server before timing out the connection. The default timeout is 5 seconds.

In the following example, the Content Engine is configured to use an N2H2 server that has an IP address of 172.16.22.10. The Content Engine will send IFP requests to this N2H2 server on port 4008 and will wait for up to 20 seconds for an IFP response from this server before timing out the connection:

```
ContentEngine(config)# url-filter http N2H2 server 172.16.22.10 port 4008 timeout 20
```

The server IP address and port number configured on the Content Engine must match the IP address of the N2H2 server and the port that the N2H2 server listens to for IFP requests. If the configuration on the Content Engine does not match the configurations on the N2H2 server, the Content Engine will time out all HTTP requests (HTTP, FTP-over-HTTP, or HTTPS-over-HTTP requests) and either block or allow all HTTP traffic based on the **allowmode** option configuration.



Note

The **url-filter http N2H2 server** global configuration command does not verify whether or not an N2H2 server is accessible at the specified IP address in the current implementation. The configuration can be changed while N2H2 is enabled. The Content Engine will adopt the new configuration at run time.

- Step 4** Enable the N2H2 URL filtering scheme on this Content Engine by entering the following command:

```
ContentEngine(config)# url-filter http N2H2 enable
```

If the URL filter is already enabled with N2H2 or other filtering schemes, the **enable** command fails. Even if the server IP address is not configured, the command is accepted, but filtering does not take effect until the N2H2 server is properly configured. To ensure that all traffic is filtered by the N2H2 server with a cluster of Content Engines, make sure to use the **url-filter http N2H2 enable** command on each Content Engine in the cluster.

- Step 5** Allow HTTP requests (HTTP, FTP-over-HTTP, or HTTPS-over-HTTP requests) to pass through when the N2H2 server is enabled but the Content Engine has problems communicating with the N2H2 server by entering the following command:

```
ContentEngine(config)# url-filter http N2H2 allowmode enable
```

By default, **allowmode** is enabled.

When **allowmode** is enabled, the Content Engine allows all HTTP traffic to continue through it (it proceeds with normal traffic processing) even if it fails to receive responses from the N2H2 server.

When **allowmode** is disabled, the Content Engine blocks all HTTP traffic that is served through it if it fails to receive responses from the N2H2 server.

You can configure the **allowmode** option with or without N2H2 being enabled; it is independent of the N2H2 server configuration. The Content Engine adopts the new configuration for **allowmode** if N2H2 URL filtering is already being used.

- Step 6** Display the request-reply statistics for the communication between the Content Engine and the N2H2 server by entering the following command:

```
ContentEngine(config)# show statistics url-filter http N2H2
N2H2 URL Filtering Statistics:
    Lookup requests transmitted = 0
    Lookup response received = 0
    Requests timed out = 0
    Number of retransmits = 0

    Requests BLOCKed by N2H2 = 0
    Requests OKed by N2H2 = 0

Allow Mode Statistics:
    No available connection = 0
    Error sending lookup requests = 0
    Error recving lookup responses = 0
    Server error in responses = 0

Server Error in Responses:
    Error in Filter Server = 0
    Error in IFP server = 0
    Seq number mismatch = 0
    Multiple responses rcvd = 0

TCP error statistics:
    Bad network endpoint = 0
    Network unreachable = 0
    Underlying connection broken = 0
    Timeout specified is reached = 0
    Address already in use = 0
    Client connection broken = 0
    Client connection timeout = 0
    Server connection broken = 0
    Server connection timeout = 0
    Register read cancelled = 0
    Other errors = 0

Queue statistics:
    Number of xacts in Queue = 0

Overhead statistics:
    Avg total process time = 0
    Avg response time = 0
    Socket update count = 0
```

These statistics show the number of requests sent, replies received, pages blocked, pages allowed, and failure cases. Detailed URL filtering statistics are available on the N2H2 server.

You can clear the statistics by entering the **clear statistics url-filter http N2H2** and **clear statistics all EXEC** commands. The **clear statistics url-filter http N2H2 EXEC** command resets the statistics counters for the N2H2 server. All statistics counters are reset to 0.

Go to <http://www.n2h2.com> for more information about N2H2 filtering configuration and policies.

N2H2 Configuration and Restrictions

Only one URL filtering scheme can be active for each protocol. To enable N2H2 URL filtering, you should first make sure that no other URL filtering scheme is configured. You can then configure the server information for N2H2 using the **url-filter N2H2 server IP address [port 1-65535] [timeout 1-120]** command (or the GUI equivalent) and enabling the N2H2 server.

The server IP address and port number configured in the Content Engine must match the IP address of the N2H2 server and the port that the N2H2 server uses to listen for IFP requests. If the configuration on the Content Engine does not match the configurations on the N2H2 server, the Content Engine will time out all HTTP, FTP, or HTTPS requests and either block or allow all HTTP traffic based on the **allowmode** option configuration.

Configuring URL Filtering with Websense Enterprise Software

**Note**

URL filtering with the Websense server only applies to HTTP, FTP, or HTTPS protocols.

Websense Enterprise is an Internet filtering software that transparently monitors, reports on, and manages employee use of the Internet. Cisco ACNS 5.5 software supports Websense Enterprise software version 5.5 on all Cisco Content Engine platforms.

**Note**

The integrated Websense Enterprise software version 5.5 in ACNS 5.5 software requires a minimum of 512 MB of RAM. We recommend that you upgrade the RAM on your device to 512 MB or greater, or move your integrated Websense server to another device that has at least 512 MB of RAM.

You can configure your ACNS network to use an external enterprise server with Websense Enterprise software installed, or you can enable the Content Engine integrated Websense server, or you can configure both an external and an integrated Websense server for the Content Engine. An external Websense server communicates with the Content Engine over the network, whereas an integrated Websense server runs internally to the Content Engine. Communications between the Content Engine caching processes and an integrated Websense server happen internally to the Content Engine.

To access the set of documents on Websense product setup and implementation, use the following URL:

<http://ww2.websense.com/global/en/SupportAndKB/ProductDocumentation/?Section=All>


Websense integration documents are listed under the heading, “Integration-specific Installation Guides.”

Using a Websense Enterprise Server

Figure 16-8 shows a network configuration that uses an external Websense enterprise server as a filtering engine for the Content Engine. The Content Engine enforces the filtering policy configured on the remote Websense server. Refer to the Websense documentation for further information on Websense filtering policies.

**Note**

Clicking the **Save Changes** button from the Websense Enterprise Manager window in the Websense GUI does not save the Websense configuration across device reboots. You need to use the **write memory** command to save the Websense configuration changes across reboots.

Figure 16-8 URL Filtering with a Websense Server

68684

Configuring Websense URL Filtering with External Websense Servers

When configuring a Content Engine to use an external Websense server, you must specify an IP address and port number for that server. That specified IP address and port number must match the IP address of the external Websense server and the port that the external Websense server listens to for filtering requests. Otherwise, the Content Engine times out all HTTP requests (HTTP, FTP-over-HTTP, or HTTPS-over-HTTP requests) and blocks and allows all HTTP traffic based on the **allowmode** option configuration. By default, allow mode is enabled on the Content Engine. When allow mode is enabled, the Content Engine is permitted to fulfill an HTTP request from a client if the external Websense server does not respond. If allow mode has been disabled, use the **url-filter http websense allowmode enable** command to reenale it.

In ACNS software Release 5.2 and later releases you can configure up to two Websense servers for failover purposes. The order in which you configure Websense servers determines which server is the primary server. The first configured Websense server is automatically designated the primary Websense server, whereas the second configured server becomes the secondary Websense server. For a list of supported Websense server configurations, see [Table 16-10](#).

To configure a Content Engine to use an external Websense server for URL filtering, follow these steps:

- Step 1** Specify the necessary information about the external Websense server by entering the following command:

```
ContentEngine(config)# url-filter http websense server { [hostname | ip-address] } [port
portnum [timeout seconds [connections connection]]
```

where

- *hostname* is the hostname of the external Websense server.
- *IP address* is the IP address of the external Websense server.
- *portnum* is the port number (1–65535) of the external Websense server to which the Content Engine is to send HTTP requests. The default is port 15868.

- *seconds* is the number of seconds (0–240) that the Content Engine is to wait for an HTTP response from the external Websense server before timing out the connection. The default is 20 seconds.
- *connections* is the number of persistent connections (1–250) per CPU (the default is 40 per CPU). Use this option to configure the number of persistent connections to the external Websense server. Do not change the default number unless you know for certain that a different value is required.

The following example shows how to configure a Content Engine to point to an external Websense server that has the IP address 172.18.22.10 and is running on Host A. The Content Engine is configured to send requests to this external Websense server on port 4006 and to wait up to 90 seconds for a response from this server before timing out the connection:

```
ContentEngine(config)# url-filter http websense server 172.18.22.10 port 4006 timeout 90
```



Note To use an external Websense server for URL filtering with a cluster of Content Engines, make sure to use the **url-filter http websense server** global configuration command on each Content Engine in the Content Engine cluster to ensure that all traffic is filtered.

Step 2 (Optional) Configure a secondary Websense server for failover purposes as follows:

- To configure the integrated Websense server as the secondary Websense server, use the **url-filter http websense server local** global configuration command. (For more information about configuring the integrated Websense server, see the “Using the Integrated Websense Server” section on page 16-41.)
- To configure an external Websense server that is running on a different host (Host B) than the primary Websense server (Host A), enter another **url-filter http websense server** command. This time, the command should specify the parameters for the secondary Websense server (for example, the IP address, port number, timeout, and number of connections of Websense server running on Host B.)

Step 3 Enable Websense as the current URL filtering scheme for HTTP on this Content Engine by entering the following command:

```
Console(config)# url-filter http websense enable
```



Note For information about configuring an external Websense server, go to the following website: <http://www.websense.com>.

Use the **show websense-server EXEC** command to view the current Websense server configuration.

Using the Integrated Websense Server

You can also configure the Content Engines in your network to use the integrated Websense server. The integrated Websense server uses approximately 60 MB to 140 MB of RAM in the Content Engine. We recommend that you run the integrated Websense server on Content Engines with at least 512 MB of RAM for best results.

When the Websense server is enabled on the Content Engine and the Websense URL database is downloaded the first time, CPU usage can be very high, so we recommend that you enable the Websense server during off-peak times or times of low network traffic. Otherwise, other processes running on the Content Engine might be affected. If the Websense server stalls, it restarts automatically.

Websense provides an image of the Websense server that resides in the /local/local1/WebsenseEnterprise directory. All executables, as well as the configuration and logging files, are stored in the above directory. This package requires about 150 MB of disk space in the /local/local1/WebsenseEnterprise/EIM directory. An additional 140 MB of disk space is required at the time of downloading the Websense URL database, increasing the total disk space requirement to 290 MB. To ensure that you have enough disk space to properly download the Websense software, we recommend that you increase the amount of sysfs disk space to an amount that is greater than the default sysfs on your Content Engine.

Configuring Ports for the Integrated Websense Server

The integrated Websense process requires that four ports be opened for connections either from processes internal to the Content Engine or from external processes such as the PIX firewall. These four ports and default port numbers are as follows:

- 15868—Websense server port

This port is the TCP port that receives requests for content filtering according to the Websense protocol.

- 15871—Block message server port

If the Websense process blocks a URL, it sends a redirect URL to the user. The redirect URL is configured to print out the blocked page and policy for the user. The Websense process listens on this port to receive the blocked pages that are served by the Websense server in response to the redirected request.

- 55806—Configuration server port

This port is required by the Websense GUI to configure the Websense server.

- 15869—Diagnostics server port

The Websense server has an exhaustive set of diagnostics that the users can run remotely to diagnose problems in the Websense process. These diagnostics connect to this port.

You can configure these ports by modifying the websense.ini file that resides in the /local/local1/WebsenseEnterprise directory. The Websense server must be restarted so it can pick up the newly configured ports.

You can modify the ports by exporting a copy of the websense.ini file by using FTP from the /local/local1/WebsenseEnterprise directory on the Content Engine, by modifying the file, by deleting the websense.ini file on the Content Engine, and then by sending back the modified file to the Content Engine using FTP.



Note

The Websense server must to be disabled and then enabled to pick up the newly configured ports.

About Websense Server Failover

The Websense server failover feature allows you to configure a Content Engine to use up to two Websense servers for failover purposes (one primary and one secondary server) for URL filtering. [Table 16-10](#) lists the supported Websense server failover configurations.

Table 16-10 *Supported Websense Server Failover Configurations*

Supported Configurations	Integrated Websense Server	Remote Websense Server
Option A	The integrated Websense server is disabled on the Content Engine.	The primary Websense server is running on an external host (for example, Host A). The secondary Websense server is running on a second external host (for example, Host B).
Option B	The integrated Websense server is acting as the primary Websense server.	The secondary Websense server is running on an external host.
Option C	The integrated Websense server is acting as the secondary Websense server.	The primary Websense server is running on an external host.

The order in which you configure the Websense servers determines which server is designated the primary Websense server. The first configured Websense server is designated the primary server. Configuration of a secondary Websense server is optional. To configure Websense server failover using the Content Distribution Manager GUI, see the next section, “[Enabling Websense Filtering and Configuring a Failover Scheme](#).”

Enabling Websense Filtering and Configuring a Failover Scheme

To enable Websense filtering using the Content Distribution Manager GUI, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the Content Engine that you want to view. The Content Engine Device Home window appears.
- Step 3** In the Contents pane, choose **Request Processing > URL Filter**. The URL Filter Settings for Content Engine window appears. (See [Figure 16-5](#).)
- Step 4** Click the **Edit Filter** icon next to the HTTP protocol. The URL Filter Settings for HTTP Protocol window appears. (See [Figure 16-7](#).) [Table 16-11](#) describes the fields in this window and provides the corresponding CLI global configuration commands.
- Step 5** To enable URL filtering using a Websense server, check the **Enable WebSense Filtering** check box.



Note You can specify up to two Websense servers. The first server specified becomes the primary Websense server, and the second server specified becomes the secondary Websense server. You can configure the integrated Websense server to be the primary server and an external Websense server to be the secondary server, or the reverse.

- Step 6** To configure the Content Engine to use an external Websense server as the primary URL filtering server, enter the server IP address or hostname in the Websense Server 1 Hostname field.

Alternatively, to configure the Content Engine integrated Websense server as the primary URL filtering server, check the **Embedded** check box. The Websense Server Hostname field becomes unavailable. This option ensures that the URL filtering software uses the local Websense server and not a remote host as the Websense server.



Note When you configure the Content Engine to use the integrated Websense server, you must install the appropriate server components and enable the Websense server service on the Content Engine before any URL filtering can take place. (See the [“Installing Websense Server Components for the Content Engine”](#) section on page 16-50.)

- Step 7** In the Port and Request Timeout fields, leave the default settings as they appear unless different settings are required.
- Step 8** In the Request Connections field, leave the default setting at 40 connections for each CPU unless you are certain that a different value is required.
- Step 9** To configure a secondary Websense server for the Content Engine, choose either an external server or the integrated server, and follow the same steps outlined for configuring the primary Websense server options.
- Step 10** To enable HTTP access to a website if the Websense server does not respond, check the **Enable WebSense Allow Mode** check box.
- Step 11** To confirm your settings, click **Submit**.

Table 16-11 Enabling Websense URL Filtering

GUI Parameter	Function	CLI Command
Enable WebSense Filtering	Makes the WebSense Server 1 and 2 Hostname and Port fields available.	url-filter http websense server
WebSense Server 1 Hostname	Enables the external Websense server specified by the IP address or name of the external Websense server.	url-filter http websense server <i>hostname</i>
Port	Port number on which the Websense server will accept requests from the Content Engine. The default port is 15868.	url-filter http websense server <i>hostname port port_num</i>
Request Timeout	Maximum amount of time that the Content Engine will wait for a response from the Websense server. The default is 20 seconds.	url-filter http websense server <i>hostname port port_num</i> timeout num connections num
Request Connections	Number of persistent connections to the Websense server per CPU.	

Table 16-11 Enabling Websense URL Filtering (continued)

GUI Parameter	Function	CLI Command
Embedded	Configures the Content Engine to send URL filtering requests to the integrated Websense server running on the Content Engine.	url-filter http websense server local port <i>port_num</i>
Enable WebSense Allow Mode	Allows HTTP access to a website if the Websense server does not respond. Permits the Content Engine to fulfill the client request after a Websense server timeout. The Websense server returns its own blocking message when a request is denied. With allowmode disabled, however, the Content Engine blocks all traffic through the Content Engine. By default, allowmode is enabled.	url-filter http websense allowmode enable

**Note**

To use Websense URL filtering with a cluster of Content Engines, enable Websense filtering with the **url-filter http websense enable** command (or the GUI equivalent), and configure the **url-filter http websense server** command (or the GUI equivalent) on each Content Engine in the cluster to ensure that all traffic is filtered.

Configuring Websense Server Failover and URL Filtering Using the CLI

In the following scenario, the Content Engine acts as the HTTP proxy for URL filtering. First, the Content Engine is configured to use either the local or the remote policy server, and then the local Websense server services (the local EIM server, the local user service, and the local network agent) are activated on the Content Engine.

Next, the Content Engine is configured to use the local (internal) Websense server as its primary Websense server and an external Websense server as the secondary Websense server. If the primary Websense server is unavailable, the Content Engine sends the filtering requests to this secondary server.

After allow mode is reenabled on the Content Engine, URL filtering is enabled on the Content Engine. Finally, the Websense manager GUI is used to configure the default policy for the local and the remote Websense servers, and then the HTTP proxy is enabled on the Content Engine.

Step 1 Specify whether the local or remote Websense policy server will be used to activate the individual Websense services on the Content Engine as follows:

- To use the local policy server, use the **websense-server service policy local activate** global configuration command to activate the local policy server on the Content Engine.

```
ContentEngine(config)# websense-server service policy local activate
```
- To use a remote policy server, use the **websense-server service policy remote host** global configuration command to configure the necessary information about the remote policy server (for example, its hostname or IP address, and its port number) on the Content Engine.

```
ContentEngine(config)# websense-server service policy remote host {hostname|
IP address} [port policy-server-port]
```

where

- *hostname* or *IP address* is the hostname or IP address of the remote policy server.
- The port number is optional. The default port number is 55806.

Either the local or the remote policy server must be running before you can activate any of the services of the local Websense server (the local EIM server, the local network agent, the local eDirectory agent, the local RADIUS agent, and the local user service) on the Content Engine. Local and remote policy server configuration are mutually exclusive.

Step 2 Activate the local EIM server on the Content Engine by entering the following command:

```
ContentEngine(config)# websense-server service eim activate
```

Step 3 Activate the local user service on the Content Engine by entering the following command:

```
ContentEngine(config)# websense-server service user activate
```

Step 4 Activate the local network agent on the Content Engine by entering the following command:

```
ContentEngine(config)# websense-server service network-agent activate
```

Step 5 Activate the local eDirectory agent on the Content Engine by entering the following command:

```
ContentEngine(config)# websense-server service edir-agent activate
```

Step 6 Activate the local RADIUS agent on the Content Engine by entering the following command:

```
ContentEngine(config)# websense-server service radius-agent activate
```

Step 7 Enable all of the services of the local Websense server (the local EIM server, the local network agent, and the local user service) that have been activated on the Content Engine by entering the following command:

```
ContentEngine(config)# websense-server enable
```

By default, the local Websense server, which consists of the local EIM server, the local network agent, and the local user service, is disabled on a Content Engine.



Note If you are using the local Websense server with a cluster of Content Engines, make sure that you enable the local Websense server on each Content Engine (enter the **websense-server enable** global configuration command on each Content Engine in the Content Engine cluster).

Step 8 Configure the Content Engine to use the local Websense server as the primary Websense server by entering the following command:

```
ContentEngine(config)# url-filter http websense server local port 4005 timeout 60
connections 90
```



Note You can use the **url-filter http websense server** global configuration command to configure different settings (for example, the timeout, the port number, and the number of connections) for the primary and secondary Websense servers. By default, the Content Engine (that is acting as the HTTP proxy) sends filtering requests to Websense server on port 15868, waits 20 seconds for a response from Websense server before timing out the connection, and establishes 40 persistent connections per CPU.

In the previous example, the Content Engine (that is acting as the HTTP proxy) sends filtering requests to the local Websense server on port 4005, waits 60 seconds for a response from the local Websense server before timing out the connection, and establishes 90 persistent connections to this local Websense server. Because the local Websense server is configured first, it is designated as the primary Websense server for the Content Engine.



Note The IP address of the local Websense server cannot be configured and is set at 127.0.0.1.

- Step 9** Configure the Content Engine to use an external Websense server as the secondary Websense server by entering the following command:

```
ContentEngine(config)# url-filter http websense server 172.18.22.10 port 4006
timeout 90
```

Because the local Websense server is already the primary Websense server, you must specify an external Websense server as the secondary Websense server.

In the previous example, the external Websense server with an IP address of 172.18.22.10 is configured as the secondary Websense server. If the local Websense server is unavailable, the Content Engine sends the requests to this secondary Websense server on port 4006, waits up to 90 seconds for a response from this server before timing out the connection, and establishes 90 persistent connections per CPU.

- Step 10** By default, allow mode is enabled. To reenable allow mode, enter the following command:

```
ContentEngine(config)# no url-filter http websense allowmode enable
```

If the primary Websense server is unavailable, then the Content Engine sends the requests to the specified secondary Websense server. If both the primary and the secondary Websense servers are unavailable, then the requests are sent to allow mode as follows:

- When allow mode is enabled, the Content Engine allows all HTTP traffic to continue through it (it proceeds with normal traffic processing) even if it fails to receive responses from Websense server.
- When allow mode is disabled, the Content Engine blocks all HTTP traffic that is served through it if it fails to receive responses from Websense server.

You can configure the **allowmode** option with or without Websense server being enabled; it is independent of Websense server configuration. The Content Engine adopts the new configuration for **allowmode** if Websense URL filtering is already being used.

- Step 11** Enable URL filtering on the Content Engine by entering the following command:

```
ContentEngine(config)# url-filter http websense enable
```

- Step 12** Configure the default policy using the Websense Manager GUI. This step should be performed for both the local and the remote Websense server.

- Use the Websense Manager GUI to add a policy server:
 - Right-click the left pane of the Websense Manager main window.
 - Choose **Add Policy Server**.
 - In the displayed dialog box, enter the IP address of the Content Engine that is running the local (internal) Websense server.
- Connect to the Websense policy server that is running on the Content Engine:
 - In the left pane, double-click on the policy server (which could be the Content Engine IP address, for example).
 - Enter the username and password, and then click **OK**.

- c. Use the Websense Manager GUI to configure a Websense policy:
- Use the Websense Manager GUI to connect to the Websense policy server.
 - In the left pane, double-click **Filter Definition** and then **Policies**.
 - Choose **Global**.
 - In the right pane, click the **Edit** button.
 - In the displayed dialog box, apply such category sets as the default settings, basic settings, always block, and never block. The default policy is global, and the default category set is the default settings.

**Note**

For more information about how to use the Websense Manager GUI, go to the following website: <http://www.websense.com>.

**Note**

Clicking the **Save Changes** button from the Websense Enterprise Manager window does not save the Websense configuration modifications across device reboots. You need to use the **write memory** command to save the Websense configuration changes across reboots.

- Step 13** Configure the HTTP proxy on the Content Engine by entering the following command:

```
ContentEngine(config)# http proxy incoming 8080
```

- Step 14** Display statistics for both the primary and the secondary Websense servers by entering the following command:

```
ContentEngine# show statistics url-filter http websense
```

About Websense Server Components

ACNS 5.5 software supports the following Websense Enterprise software version 5.5.2 server components on the integrated Websense server:

- Policy Server
- EIM Server
- User Service
- Network Agent
- Radius Agent
- eDirectory Agent
- Logon Agent (ACNS support for this server component has been added in the ACNS 5.5 release.)

**Note**

When you install or activate additional Websense components for the integrated Websense server, the ACNS software requires a minimum of 1 GB of RAM.

Table 16-12 describes these Websense 5.5.2 services.

Table 16-12 **Integrated Websense 5.5.2 Server Components Supported in ACNS 5.5 Software**

Name	Description
Policy Server	Hosts all of the policy information that you have configured through the external Websense Manager GUI. Communicates the policy information to the other components of the integrated Websense server. Note The integrated Policy Server or the specified remote Policy Server must be already installed and running before you can activate the EIM Server, User Service, Network Agent, RADIUS Agent, and eDirectory components. All of these five components need the IP address of the Policy Server for installation. With Websense 5.2 software, you can use the integrated Websense Policy Server or a remote Websense Policy Server to activate the individual server components on the Content Engine.
EIM Server	Provides the URL filtering functionality when used with proxy servers, firewalls, and caching appliances.
User Service	Enables URL filtering based on user-based or group-based policies. If you are using a user service and you want to configure user-based or group-based URL filtering using a Windows NT directory, then you must use the external user service on a Windows machine.
Network Agent	Enables URL filtering of requests that use protocols other than HTTP, HTTPS-over-HTTP, and FTP-over HTTP. If the local network agent is activated on the Content Engine, the network agent can filter incoming requests from the following protocols and applications: • Database applications, such as SQL Net • File transfer applications, such as FTP and Gopher • Instant messaging and chat applications, such as Yahoo Messenger, and MSN Messenger • Mail and collaborative tools, such as POP3, SMTP, and NetMeeting • Network operating system applications, such as Daytime, finger, NTP, SSH, and Telnet • Remote access applications, such as VNC and pcANYWHERE • Streaming media applications, such as RTSP, Windows Media, and Liquid Audio • Other (for example, Network News Transfer Protocol [NNTP])
RADIUS Agent	Enables URL filtering based on user-based or group-based policies for users who are authenticated through an external RADIUS server. This agent transparently identifies the users who access the network and are authenticated through the RADIUS authentication scheme. When the Content Engine is supplied with this information, the Content Engine can apply policies to users and groups of the users who access the network remotely. This agent acts as a proxy that forwards the RADIUS messages between the RADIUS client and the external RADIUS server. For the integrated RADIUS agent to work properly, you must have configured the RADIUS settings (for example, the IP address of the external RADIUS server) on the Content Engine.

Table 16-12 Integrated Websense 5.5.2 Server Components Supported in ACNS 5.5 Software

Name	Description
eDirectory Agent	Enables URL filtering based on user-based or group-based policies for users who are authenticated through LDAP. This agent works in conjunction with the Novell eDirectory to transparently identify users who access the network and are authenticated through the LDAP authentication scheme. When the Content Engine is supplied with this information, the Websense filtering service can filter requests based on the policies applied to the users or groups. This agent uses LDAP to gather the user login session information from the Novell eDirectory, which authenticates users logging in to the network. This agent associates each authenticated user with the IP address. With the help of the Websense integrated user service, the integrated eDirectory agent supplies this information to the Websense filtering service. For this integrated eDirectory agent to work properly, you must have configured such settings as the administrative distinguished name.
Logon Agent	Enables URL filtering based on user-based or group-based policies for users as they log on to the network through a Windows client machine. This agent's associated logon application captures logon sessions as users log on to Windows domains in a network. The Logon Agent communicates with the Websense User Service to provide up-to-date user logon session information to Websense for filtering purposes. The Logon Agent identifies users in a real-time manner as they log on to a domain, whereas a Domain Controller (DC) Agent identifies users by periodically querying domain controllers and workstations. Because the Logon Agent identifies users in a real-time manner, the Websense Filtering Service can accurately filter Internet access based on the policies assigned to particular users, groups, workstations, or networks. The Logon Agent, like the eDirectory agent, transparently identifies users. Support for the Logon Agent was added in the ACNS 5.4.1 software release.

Installing Websense Server Components for the Content Engine

To install Websense server components using the Content Distribution Manager GUI, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the name of the Content Engine for which you want to configure URL filter settings. The Device Home for Content Engine window appears.
- Step 3** In the Contents pane, choose **Request Processing > Websense Server**. The Websense Server Settings for Content Engine window appears. [Table 16-13](#) describes the fields in this window and lists the corresponding CLI commands.
- Step 4** To install the integrated Websense Policy Server component on the Content Engine, check the **Install policy server component** check box. (See the [“About Websense Server Components”](#) section on [page 16-48](#).)

Alternatively, to specify a remote Websense server that is running the Websense Policy Server component, check the **Specify remote policy server** check box. You must specify either local Content Engine installation of Policy Server, or specify a remote Policy Server, but not both.

If you specified a remote Policy Server, enter the IP address of the remote host in the Remote Host field. The Policy Server port field displays port 55806 and needs to be changed only if you need to configure a port other than 55806.

- Step 5** To install the Websense server EIM Server component on the Content Engine, check the **Install EIM component** check box.
- Step 6** To install the Websense server User Service component on the Content Engine, check the **Install User component** check box.
- Step 7** To install the Websense server Network Agent component on the Content Engine, check the **Install Network Agent component** check box.
- Step 8** To install the Websense server Logon Agent component on the Content Engine, check the **Install Logon Agent component** check box.
- Step 9** To install the Websense server RADIUS Agent component on the Content Engine, check the **Install Radius Agent component** check box.
(For more information, see the [“Specifying RADIUS Agent Settings for the Content Engine Using the CLI”](#) section on page 16-52.)
- Step 10** To install the Websense server eDirectory Agent component, check the **Install eDirectory Agent component** check box.
(For more information, see the [“Specifying eDirectory Agent Settings for the Content Engine Using the CLI”](#) section on page 16-52.)
- Step 11** To configure an administrative password that the Content Engine should use to contact the external eDirectory Server to request a database search, enter a string in the Password field.
- Step 12** To confirm the password, reenter the same string in the Confirm Password field.
- Step 13** To enable the Websense server and activate the installed components, check the **Enable Websense Server** check box.
- Step 14** To confirm your settings, click **Submit**.

Table 16-13 Websense Server Settings for the Content Engine

GUI Parameter	Function	CLI Command
Install policy server locally	Installs the integrated policy server on the Content Engine.	websense-server service policy local activate
Specify remote policy server	Activates GUI fields for specifying the remote host IP address and port number.	—
Remote Host	Specifies the remote policy server to be used to activate the other integrated server components on the Content Engine.	websense-server service policy remote host <i>remote-policy-server IP-address</i> port <i>remote-policy-server port-number</i>
Port	Specifies the port number of the remote Policy Server. The default port number is 55806.	
Install EIM component	Installs the integrated EIM server on the Content Engine.	websense-server service eim activate
Install User component	Installs the integrated user service on the Content Engine. Use the no form of this command to deactivate it.	websense-server service user activate
Install Network Agent component	Installs the integrated network agent on the Content Engine.	websense-server service network-agent activate

Table 16-13 Websense Server Settings for the Content Engine (continued)

GUI Parameter	Function	CLI Command
Install Logon Agent component	Installs the integrated logon agent on the Content Engine. The Logon Agent uses port 15880.	websense-server service logon-agent activate
Install Radius Agent component	Installs the integrated RADIUS agent on the Content Engine.	websense-server service radius-agent activate
Install eDirectory Agent component	Activates the integrated eDirectory agent on the Content Engine.	websense-server service edir-agent activate
Password / Confirm Password	Administrative password that the Content Engine should use to contact the external eDirectory Server to request a database search.	websense-server service edirectory-agent edir-server administrative-passwd password
Enable Websense Server	Activates the Websense server components that you installed.	websense-server enable

Specifying RADIUS Agent Settings for the Content Engine Using the CLI

In the Websense GUI Manager versions 5.5 and 6.1 that are supported by the Websense 5.5.2 software, you can configure the RADIUS agent through the GUI, as well as through CLI commands. Consequently, in the ACNS 5.4.1 software release, all of the global configuration CLI commands that were related to configuring the RADIUS agent have been removed from the ACNS CLI command set. However, the CLI command that is used to active the RADIUS agent has been retained.

The following global configuration commands for configuring the RADIUS agent were removed in the ACNS 5.4.1 software release:

websense-server service radius-agent incoming [auth-port *port number*]
[acct-port *port number*]

websense-server service radius-agent outgoing [host *remote-RADIUS-server IP-address*]
[auth-port *port number*] [acct-port *port number*]



Note

If the Content Engine is running the ACNS 5.4.1 software or a later release, and you attempt to configure the RADIUS agent through the CLI by entering any of the removed global configuration commands, the command is nullified. (You will not receive an error message if you enter one of these nullified commands).

Specifying eDirectory Agent Settings for the Content Engine Using the CLI

In the Websense GUI Manager versions 5.5 and 6.1 that are supported by the Websense 5.5.2 software, you can configure the eDirectory agent through the GUI, as well as through CLI commands. Consequently, in the ACNS 5.4.1 software release, all of the global configuration CLI commands that were related to configuring the eDirectory agent have been removed from the ACNS CLI command set. However, the CLI command that is used to active the eDirectory agent has been retained.

The following global configuration command for configuring the eDirectory administrative password was also retained in the ACNS 5.4.1 software release:

websense-server service edirectory-agent edir-server administrative-passwd *password*

The following global configuration command for configuring the eDirectory agent was removed in the ACNS 5.4.1 software release:

websense-server service edir-agent edir-server [**administrative-dn** *administrative-distinguished-name*] [**host** *remote-eDirectory-server IP-address*] [**root-context** *root-context*]



Note

If the Content Engine is running ACNS 5.4.1 software or a later release, and you attempt to configure the eDirectory agent through the CLI by entering any of the removed global configuration commands, the command is nullified. (You will not receive an error message if you enter one of these nullified commands).

Websense Status and Statistics Commands

Additional CLI commands to use with Websense URL filtering are as follows:

- **show url-filter http**

This command shows the IP address of the local host in the Websense sever IP field when the local host is configured as the Websense server for Websense URL filtering.

The **show url-filter http** command also shows the URL filtering scheme enabled on the Content Engine for HTTP traffic and the configurations for each URL filtering scheme, such as the configuration data for Websense.

In this example, the **show url-filter http** command displays the status of all HTTP URL filtering schemes presently configured on the Content Engine:

```
ContentEngine# show url-filter http
URL filtering is set to use bad-list

Local list configurations
=====
Good-list file name :
Bad-list file name : /local1/url-filter/badlist.http
Custom message directory :

Websense server configuration
=====
Websense server IP      : 172.16.193.165
Websense server port    : 15868
Websense server timeout: 20 (in seconds)
Websense allow mode is ENABLED

N2H2 server configuration
=====
N2H2 server IP          : 172.16.193.165
N2H2 server port        : 4005
N2H2 server timeout     : 5 (in seconds)
N2H2 allow mode is DISABLED
ContentEngine#
```

- **show websense-server**

This command shows the configuration for the Websense server configured on the Content Engine. The output of the command includes the configured port numbers for the Websense server port, blocked message server port, configuration server port, and diagnostics server port; the Websense server version number; and the maximum number of connections.

- **show statistics url-filter http websense**

This command shows the request-reply statistics of the communication between the Content Engine and the Websense server. These statistics show the number of requests sent, replies received, pages blocked, pages allowed, and failure cases. More detailed URL filtering statistics are available on the Websense server.

```
ContentEngine# show statistics url-filter http websense
Websense URL Filtering Statistics:
Transmission statistics:
    Lookup requests transmitted = 1
    Lookup requests timed-out = 1
    Lookup responses received = 1
    Lookup responses received with error = 0
    Multiple response received = 0
    Sequence number mismatch = 0

TCP errors:
    Connection reset = 0
    Connection timeout = 3
    Other errors = 0

Filter results:
    Requests BLOCKed by Websense = 1
    Requests OKed by Websense = 0
    Sent to Allowmode ok = 1
    Sent to Allowmode block = 0
    desc_filtered_and_passed = 0
    desc_category_blocked = 1
    desc_category_not_blocked = 0
    desc_category_blocked_custom_deny = 0
    desc_category_not_blocked_custom_permit = 0

Websense log statistics:
    Logs sent successfully = 1
    Connection error = 0
    Error during log processing = 0
    Log not complete = 1
    Log not sent because Websense disabled = 0
    No available connection = 0

Congestion statistics:
    Pending requests = 0
    Pending log requests = 0

ContentEngine#
```

The statistics shown can be cleared by using the **clear statistics url-filter http websense**, and **clear statistics all** commands. All the statistics counters are then reset to 0.

- **write memory**

This EXEC command saves Websense configuration files (websense.init and ws.cfg) modified from the external Websense Manager GUI across disk reconfiguration and ACNS software release upgrades (which might erase disk content).

You must enter this command to have the most recent configuration modifications, including websense.ini file modifications and the Websense URL filtering configuration changes.

If you do not use the **write memory** command before a reboot but after a disk reconfiguration or an ACNS software upgrade that erases disk content, the Websense configurations that were saved when the **write memory** command was last used are retained. However, if the **write memory** command was never used before, then default configurations are applied when the content on /local/local1/WebsenseEnterprise directory is erased.

Websense Configuration and Restrictions

Only one URL filtering scheme can be active per protocol. To enable Websense URL filtering, first make sure that no other URL filtering scheme is enabled on the same protocol. You can then configure the information for the Websense server by using the **url-filter http websense server IP address [port 1-65535] [timeout 1-120]** command (or the GUI equivalent) and enabling the Websense server with the **url-filter http websense enable** command (or the GUI equivalent).

The server IP address and port number that are configured in the Content Engine must match the IP address of the Websense server and the port on which the Websense server listens to filter requests. If the configuration on the Content Engine does not match the configurations on the Websense server, the Content Engine will time out all HTTP, FTP, or HTTPS requests and either block or allow all HTTP traffic based on the **allowmode** option configuration.

Configuring URL Filtering with SmartFilter Software



Note

URL filtering with SmartFilter software applies only to HTTP, FTP, or HTTPS protocols.

SmartFilter software for the Content Engine provides employee Internet management (EIM) functionality with proxy servers, firewalls, and caching appliances. The integrated Content Engine and SmartFilter product preserves all functionality that is available in a regular Content Engine. The SmartFilter filtering capability is available as an add-on service on the Content Engine, and the service may be enabled or disabled as desired through the Content Engine CLI or GUI.

The integrated Content Engine and SmartFilter product provides a one-box solution for server functionality. The Content Engine uses a suite of plug-in APIs to allow the SmartFilter software to implement hooks at strategic points during an HTTP transaction and provide URL filtering.

The integrated Content Engine and SmartFilter product provides two end user management tools called the SmartFilter Administration Console and the SmartFilter Administration Server. These GUI components download configurations into the Content Engine to be used by the SmartFilter process.

To use SmartFilter URL filtering with a cluster of Content Engines, enter the **url-filter http smartfilter enable** command (or the GUI equivalent) on each Content Engine in the cluster to ensure that all traffic is filtered.

When you configure URL filtering with Smartfilter software, remember the following important points:

- When you upgrade or downgrade the Content Engine to a different release of the ACNS software, if there is a difference in the SmartFilter plug-in version, the SmartFilter database and configuration files are deleted and default configurations are loaded. This change occurs because the configuration details might be changed with each new version of the SmartFilter software. After each upgrade or downgrade of the SmartFilter plug-in, a fresh database has to be downloaded from the SmartFilter Administration Console to the Content Engine.

**Note**

The ACNS 5.2.x, ACNS 5.3.x, 5.4.x, and 5.5.1 software releases all support the SmartFilter software version 4.0.1.

- If the Content Engine is deployed behind a firewall, you must use the **external-ip external-ip-address** global configuration command to configure the Content Engine's external IP address; otherwise, the advanced block page feature may not work properly.
- The Smartfilter software and the cache process on the Content Engine will be restarted in the following situations:
 - If you configure the Content Engine's external IP address when the Smartfilter software is running on the Content Engine
 - If the Content Engine's external IP address is not configured and you change the Content Engine's interface IP address when the Smartfilter software is running on the Content Engine.
- Port 9014 is reserved for Smartfilter Advanced block pages.

**Note**

To obtain a copy of the Smartfilter authentication server software and information about configuring the SmartFilter software, go to the following website: <http://www.securecomputing.com>.

About the SmartFilter Control List

A SmartFilter Control List categorizes 2 million websites into content groups. There are 30 predefined SmartFilter Control List categories that encompass a wide variety of material. Some categories are focused on reducing legal liability of a company. These 30 categories are set to Deny in the default SmartFilter software policy. Some categories contain such sites as MP3 sites (sites with content that consumes excessive bandwidth). The remainder of these 30 categories are considered unproductive or inappropriate for business or educational environments.

The SmartFilter software also provides ten user-defined categories that allow you to tailor access further by defining and filtering sites that are not included in the SmartFilter Control List. Additionally, you can exempt any site that you would like specific groups or individuals to access quickly and easily. You can use the SmartFilter Administration Console to define a SmartFilter Control List download schedule. The Download Setup window tracks the download site, your username, and your password. If you do not download an updated SmartFilter Control List at least monthly, the SmartFilter software considers the Control List expired and invokes the action that you specified in the SmartFilter License window.

SmartFilter Software and the Action No-Auth Command Rule Interaction

The **rule action no-auth** global configuration command permits specific login and content requests to bypass authentication and authorization features, such as LDAP, RADIUS, NTLM, or TACACS+. In the following example, any requests from the source IP address (src-ip) 172.16.53.88 are not authenticated:

```
ContentEngine(config)# rule enable
ContentEngine(config)# rule action no-auth pattern-list 1 protocol all
ContentEngine(config)# rule pattern-list 1 src-ip 172.16.53.88 255.255.255.255
```

If the software is configured for authentication and SmartFilter URL filtering, requests that are allowed to bypass authentication will also bypass the SmartFilter URL filter.

Configuring Content Engines to Bypass URL Filtering for Specific HTTP and HTTPS Requests

In ACNS 5.2.3 software, the ability to configure a Content Engine to bypass URL filtering for certain HTTP and HTTPS requests was added. This feature is supported for local list URL filtering (good and bad site lists), as well as Websense, SmartFilter, or N2H2 URL filtering.

For example, if you enable local URL filtering on the Content Engine and enable the bad sites deny feature (for example, the badfile.txt file contains the URLs that should be blocked), and if the **rule no-url-filtering** action is a hit (a match), the Content Engine bypasses the URL filtering for that particular request; otherwise, it proceeds with URL filtering and blocks the URL request.

To support this new feature, the following CLI changes were made:

- The **rule** global configuration command supports a new action called the **no-url-filtering** action.

The **no-url-filtering** action supports the following rule patterns: src-ip, dst-ip, dst-port, domain, group-name, groupname-regex, header-field, url-regex, and username.



Note

Patterns can be ANDed or ORed by using the group-type pattern (for example, **rule pattern-list 1 group-type and**). The default is OR.

- The output of the **show run**, **show statistics rule all**, and **clear statistics rule all EXEC** commands now includes information about the new **no-url-filtering** action.
- The **show statistics rule http action no-url-filtering EXEC** command was added to enable you to display statistics for the **no-url-filtering** action.

The following is an example of how you can use this new bypass URL filtering feature with Websense URL filtering. First, the **rule action no-url-filtering** command is specified and then associated with a specific pattern list (pattern list 100). Next, the **domain** pattern type is added to pattern list 100 in order to configure the Content Engine to match requests that have “foo.com” as the domain. In this scenario, Websense URL filtering has already been configured and enabled on the Content Engine.

```
ContentEngine (config)# rule action no-url-filtering pattern-list 100
ContentEngine (config)# rule pattern-list 100 domain .*foo.com
ContentEngine (config)# rule enable
```

When the Content Engine receives an HTTP or HTTPS request that has “foo.com” as the domain, the **rule action no-url-filtering** rule is matched. The Content Engine bypasses URL filtering for that particular request (as shown in the partial output of the **debug http proxy** command below).

```
Oct 28 12:25:12 Content Engine 3: Rule action no-url-filtering match - Bypassing
urlfiltering
```

If the **rule action no-url-filtering** rule is matched and SmartFilter URL filtering is being used instead of Websense URL filtering, the output of the **debug http proxy** command would be as follows:

```
Oct 28 12:25:12 Content Engine 3: Rule action no-url-filtering match - Bypassing
SmartFilter processing
```

When the Content Engine receives an HTTP or HTTPS request for websites other than “foo.com” (for requests that have “www.abc.com” as the domain), the **rule action no-url-filtering** rule is not matched. The Content Engine proceeds with Websense URL filtering for that particular request (as shown in the partial output of the **debug http proxy** command below).

```
Oct 28 12:28:06 Content Engine 3: Rule action no-url-filtering not hit - Proceed with
urlfiltering
```


If the **rule action no-url-filtering** rule is not matched and SmartFilter URL filtering is being used instead of Websense URL filtering, the output of the **debug http proxy** command would be as follows:

```
Oct 28 12:25:12 Content Engine 3: Rule action no-url-filtering not hit- Proceed with
SmartFilter processing
```

Execution Order of Rule Actions

In ACNS 5.2.3 software or later, the order in which the rule actions are executed is as follows:

1. Redirect-url-for-cdn (this action is only applicable for Content Engines that are registered with a Content Distribution Manager and is not applicable for standalone Content Engines)
2. No-auth (before authentication using RADIUS, LDAP, or NTLM)
3. Reset
4. Block / allow (See the [Note](#) that follows item 29.)
5. Redirect (before cache lookup)
6. Rewrite (before cache lookup)
7. No-url-filtering
8. Refresh (after cache lookup, in the case of cache hit)
9. Freshness-factor (after cache lookup, in the case of a cache hit)
10. Use-server
11. No-proxy
12. Use-proxy-failover
13. Use-proxy
14. Use-dns-server
15. ToS/DSCP server (TOS bits on the connection to the server)
16. ToS/DSCP client (TOS bits on the connection that the server uses to send response to client)
17. DSCP client cache-miss
18. DSCP client cache-hit
19. Insert-no-cache
20. No-cache
21. Cache (when the response is received from the server)
22. Selective-cache (when the response is received from the server)
23. Append-username-header
24. Use-icap-service
25. Use-xforward-clt-ip
26. No-persistent-connection
27. Cache-cookie

28. No-selective-cache

29. Allow



Note

Allow and block carry the same precedence. The order of execution depends on the order of configuration between allow and block actions. Other actions always take precedence over allow. For example, a reset action always takes precedence over allow regardless of the order of configuration.

Configuring the Content Engine GUI for Secure or Nonsecure Access

You can configure the Content Engine GUI on a centrally deployed Content Engine for secure or nonsecure access. The secured Content Engine GUI is the default.

Either secure or nonsecure access to the Content Engine GUI is possible, but not both. For example, if the secured Content Engine GUI is enabled (for example, https:// access on port 8003), then nonsecure access to the Content Engine GUI (for example, http:// access on port 8001) is not allowed. The port number of the Content Engine GUI is determined when the ACNS software is installed on the Content Engine.

Before you log in to the Content Engine GUI, make sure that you have the following information:

- Name or IP address of the Content Engine that you want to log in to.
- User account (username and password) that you want to log in with. If you do not have a user account, your ACNS system administrator must create one for you.
- Type of access enabled on the Content Engine GUI (secure or nonsecure).

To access the Content Engine GUI, you must enter the URL or IP address of the Content Engine and the port number. The URL (location) of the Content Engine is determined during the installation of the ACNS software. If your network supports DNS and the IP address of the Content Engine has been added to your DNS table, you can access the Content Engine GUI by using the DNS name of the Content Engine.

To configure the Content Engine GUI for secure or nonsecure access, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to enable the graphical user interface. The Device Home for Content Engine window appears.
- Step 3** In the Contents pane, choose **Request Processing > GUI Server**. The GUI Server Settings for Content Engine window appears.
- Step 4** Under the GUI Server Settings section, follow these steps:
 - a. To enable nonsecure access to the Content Engine GUI, check the **GUI Server Enable** check box.
 - b. In the GUI Server Port field, specify a port number for the graphical user interface server port. The default port for nonsecure access to the GUI is 8001.

- Step 5** Under the Secured GUI Server Settings section, follow these steps:
- To enable secure access to the Content Engine GUI, check the **Secured GUI Server Enable** check box.
 - In the Secured GUI Server Port field, specify a port number for the graphical user interface server port. The default port for secure access to the GUI is 8003.
- Step 6** To save the settings, click **Submit**. A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default and device group settings. You can also revert to the previously configured settings by clicking **Reset**. The **Reset** button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

If you try to leave this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box appears only if you are using the Internet Explorer browser.



Note When secured access is enabled, unsecured access is automatically disabled.

Setting Up the Content Engine to Interoperate with Third-Party Policy Servers

Cisco ACNS 5.5 software allows you to set up interoperability with Camiant PCMM-compliant third-party policy servers to allocate guaranteed bandwidth for authorized requests of pre-positioned content.



Note The settings in this section are for a feature that is not released for general use. You should not set them unless you have been specifically instructed to do so by Cisco.

In a typical deployment scenario involving a third-party policy server for request authorization and allocation of guaranteed bandwidth, the request and response flow happens as follows:

1. A client logs in to a web portal, authenticates the user credentials, and requests a Windows Media file.
2. The portal inserts a cookie with necessary information about the user credentials for subsequent authorization of the request, and it directs the request to its back-end network, which is a content distribution network that runs Cisco ACNS software.
3. A Content Router in the back-end network picks up the request and redirects the message to the nearest available Content Engine that can serve the requested media file.
4. The Content Engine validates the requested URL based on the rule patterns defined on the Content Engine for authorization of the request. If the requested URL matches the rule pattern for domains that offer guaranteed bandwidth, the Content Engine passes the request to the third-party policy server, which handles authorization of the request and allocation of guaranteed bandwidth. (See the [“Creating Rule Patterns and Rule Actions for Directing Requests to a Policy Server”](#) section on page 16-64 for more information about rule patterns.)

**Note**

The Content Engine uses ICAP protocol for communicating with the policy server for authorization of the request. You can set up an ICAP service for the policy server from the ICAP Services interface of the Content Distribution Manager to enable interoperability between your ACNS network and the policy server.

5. The entitlement server that is associated with the policy server verifies the user credentials. Based on the user credentials, the policy server returns one of the following messages to the Content Engine along with the requested URL:

Message Type	Explanation	Content Engine Action
OK with a Request ID	Valid user with privilege for guaranteed bandwidth.	Directs the streaming server to serve the content with guaranteed bandwidth.
OK	Valid user without privilege for guaranteed bandwidth.	Directs the streaming server to serve the content without guaranteed bandwidth.
Error	Valid user without privilege for receiving the requested content.	Drops the request.

6. The Content Engine, on receiving the OK with a Request ID message from the policy server, generates a URL signature and appends it to the requested URL, embeds the requested URL in an .asx file, and sends the .asx file back to the client. (See the [“About URL Signature Keys”](#) section on page 16-62.)

The Content Engine provides two URLs to the client, one based on the RTSP protocol and another based on the MMS-over-HTTP protocol. The RTSP-based URL is the default URL, whereas the MMS-over-HTTP URL is for failover when the client is unable to access RTSP-based URLs. However, if both RTSP-based and MMS-over-HTTP URLs fail, the client has a third option to use progressive HTTP download. URL signatures are generated based on ACNS rule patterns set for specific domains and are used to ensure that the client is authorized to view the content.

7. After receiving the .asx file containing the URLs, the client contacts the Content Engine using one of the URLs based on the client settings. The Windows Media Server on the Content Engine validates the URL using the URL signature. If the signature passes validation, the Windows Media server begins streaming the content. If the client is using the MMS-over-HTTP protocol, then the HTTP proxy cache application on the Content Engine does the URL validation and content delivery.
8. Upon receiving the request from the client with the appropriate URL, the Content Engine issues an HTTP callout to the policy server to provide the IP address and port number of both the client and the server (Content Engine).
9. The policy server guarantees the required bandwidth for the stream between the two points, Content Engine port and client port, as specified in the HTTP callout. The Content Engine begins streaming the content.
10. When the stream ends, the Content Engine sends a teardown message to the policy server. Upon receiving the teardown request from the Content Engine, the policy server releases the allocated bandwidth for use by other streams.

About URL Signature Keys

URL signature keys are word values that Cisco ACNS software uses to ensure URL-level security. The URL signature key is a shared secret between the device that assigns the key and the device that decrypts the key. Based on your network settings, either the Content Engine itself or some other external device can assign the signature key to the URL, but the Content Engine decrypts the URL signature key.

Cisco ACNS software uses a combination of key owners, key ID numbers, and a word value to generate URL signature keys. To create request-specific URL signature keys, you can choose to append the IP address of the client that has made the request, to the URL signature key.

You can have a maximum of 32 key owners. Each key owner can have up to 16 key ID numbers. You can use a number between 1 and 32 for a key owner and a number between 1 and 16 for the key ID number. You can have only one key configured for a unique combination of key owner and key ID number at any given time. Key values can have a maximum of 16 characters (excluding double quotes at the beginning and end of the string).

The GUI allows quoted and unquoted strings. Double quotes (") are allowed at the beginning and end of the string only. If you do not surround the key string with double quotes, quotes will be added when you click Submit in the GUI.

**Note**

The key supports using all special characters in the string, including space, pipe (|), question mark (?), forward slash (/), and back quote (`).

URL signature key authentication is implemented by using the generate-url-signature and validate-url-signature rule actions that can be applied to specific rule patterns. (See the [“Creating Rule Patterns and Rule Actions for Directing Requests to a Policy Server”](#) section on page 16-64.)

Setting QoS Policy Server Settings for Content Engines

Cisco ACNS 5.5 software enables you to configure Content Engines to interoperate with Camiant PCMM-compliant third-party policy servers to allocate guaranteed bandwidth for authenticated requests of prepositioned content.

**Note**

These settings are for a feature that is not released for general use. You should not set them unless you have been specifically instructed to do so by Cisco.

**Note**

You must have a Camiant PCMM-compliant policy server in your network to use this feature.

To configure QoS Policy Server settings for a Content Engine or a device group, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices** (or **Devices > Device Groups**). The Devices (or Device Groups) listing window appears.
- Step 2** Click the **Edit** icon next to the device (or device group) for which you want to configure the QoS Policy Server settings. The Devices (or Device Group) home window appears.
- Step 3** In the Contents pane, choose **Request Processing > PCMM > QoS Policy Service**. The QoS Policy Service Settings for Content Engine (or Device Group) window appears. [Table 16-14](#) describes the fields in this window and provides the corresponding CLI commands.

- Step 4** To enable policy server settings for the device (or device group), check the **Enable** check box.
- Step 5** To specify the configuration file for the HTTP callout to the policy server, check the **Set Config File or URL** check box. The Config File or URL field becomes active.
- The config file is an XML file that contains information on the callout URLs, attributes, application type, protocol, etc, for streaming the content.
- Step 6** In the Config File or URL field, enter the filename or a URL from which the Content Engine can download the configuration file. This field supports only URLs that use the HTTP, HTTPS, or FTP protocols.
- Step 7** To save the settings, click **Submit**.

Table 16-14 QoS Policy Service Settings

GUI Parameter	Function	CLI Command
Enable	Enables Camiant policy service.	qos camiant-cdn-am-service enable
Set Config File or URL	Allows you to specify the configuration file for the HTTP callout to the policy server.	qos camiant-cdn-am-service config-file <i>filename</i> or <i>URL</i>
Config File or URL	Filename or URL from which the Content Engine can download the configuration file.	

Creating New URL Signature Keys for Content Engines

Cisco ACNS 5.5 software enables you to create URL signature keys for unique identification of URL requests if your network is configured to use URL-level security features. The system allows you to define 32 different key owners and 16 keys for each key owner.



Note

These settings are for a feature that is not released for general use. You should not set them unless you have been specifically instructed to do so by Cisco.

To create a URL signature key for a Content Engine or a device group, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices** (or **Devices > Device Groups**). The Devices (or Device Group) listing window appears.
- Step 2** Click the **Edit** icon next to the device (or device group) for which you want to configure the QoS Policy Server settings. The Device (or Device Group) home window appears.
- Step 3** In the Contents pane, choose **Request Processing > PCMM > URL Signature Key**. The URL Signature Keys for Device (or Device Group) window appears.
- Step 4** To create a new URL Signature Key, and click the **Create New URL Signature** icon in the taskbar. The Creating New URL Signature Key window appears. [Table 16-15](#) describes the fields in this window and provides the corresponding CLI commands.
- Alternatively, to go to the Edit URL Signature Key ID window, click the **Edit** icon corresponding to a URL Signature Key ID owner.

- Step 5** In the Key ID Owner field, enter a number between 1 and 32 as the Key ID owner.
- Step 6** In the Key ID field, enter a number between 1 and 16.
- Step 7** In the URL Signature Key field, enter a unique URL signature key. This field does not support a space or the following special characters: pipe (|), question mark (?), double quotes ("), and back quote (`).
- Step 8** To save changes, click **Submit**.

Table 16-15 URL Signature Key Settings

GUI Parameter	Function	CLI Command
Key ID Owner	Specifies the ID number for the owner of this encryption key.	<code>url-signature key-id-owner owner_num</code>
Key ID	Specifies the encryption key ID number.	<code>url-signature key-id-owner owner_num key-id-number id_num</code>
URL Signature Key	Text of the encryption key. Maximum length is 16 characters.	<code>url-signature key-id-owner owner_num key-id-number id_num key encryption_key</code>

Creating Rule Patterns and Rule Actions for Directing Requests to a Policy Server

If your network is configured to work with third-party Camiant PCMM-compliant policy servers for servicing requests that require guaranteed bandwidth, you can use the following rule patterns and rule actions to filter the requests and to direct them to the policy server. The rule patterns and rule actions also enable you to generate URL signatures in the response for a valid request for a Windows Media metafile (.asx file extension) and to validate the URL signature on incoming requests to the Content Engine. (For more information on creating rule patterns and configuring rule actions, see the [“Configuring Service Rules”](#) section on page 16-13.)

The following rule patterns support the rule action **use-icap-service** for directing requests that require guaranteed bandwidth to the third-party policy server:

Rule Pattern	Description
url-regex	Filters the request based on any regular expression in the URL.
domain	Filters the request based on the domain name specified.
dst-ip	Filters the request based on the destination IP address.
dst-port	Filters the request based on the destination port number.
src-ip	Filters the request based on the IP address of the source.
header-field user-agent	Filters the request based on the user agent specified in the request header.
header-field referer	Filters the request based on the referer in the request header.
header-field request-line	Filters the request based on the request line in the request header.

You can set the rule action **use-icap-service** for any of the rule patterns above. If the request matches the parameters that you have set for the rule pattern, then the Content Engine redirects the request to the third-party policy server using ICAP services. However, you must make sure that your network is configured to interoperate with the third-party policy server using ICAP service. You can set up the necessary ICAP configurations from the ICAP Server Settings window of Content Distribution Manager. Cisco ACNS software also provides a number of CLI commands to set up interoperability with a third-party policy server using ICAP services. For more information about CLI commands related to ICAP configuration for third-party policy servers, see the *Cisco ACNS Software Command Reference, Release 5.5* publication.

You can also use the rule pattern and rule action to generate URL signatures in the response for a valid request for a Windows Media metafile. You can use the following rule patterns to filter out requests for which you want to generate a URL signature key:

Rule Pattern	Description
url-regex	Filters the request based on any regular expression in the URL.
domain	Filters the request based on the domain name specified.
dst-ip	Filters the request based on the destination IP address.

For the rule patterns mentioned above, you can set the following rule actions:

Rule Action	Description
generate-url-signature	Generates the URL signatures in the Windows Media metafile response associated with pre-positioned content, based on the Content Engine configuration for the url-signature and this rule action.
validate-url-signature	Validates the URL signature for a request by using the configuration on your Content Engine for the url-signature and allows the request processing to proceed for this request



Creating and Managing IP Access Control Lists

Access control lists (ACLs) provide a means to filter packets by allowing a user to permit or deny IP packets from crossing specified interfaces. Packet filtering helps to control packet movement through the network. Such control can help limit network traffic and restrict network use by certain users or devices.

You can also apply ACLs to management services such as SNMP, SSH, HTTPS, Telnet, FTP, and TFTP. ACLs can be used to control the traffic that these applications provide by restricting the type of traffic that the applications will handle. (The Content Distribution Manager GUI allows you to associate standard ACLs with SNMP, TFTP, and WCCP applications only.)

This chapter describes the procedure for applying ACLs to devices that use the Content Distribution Manager GUI. It contains the following sections:

- [Introducing IP ACLs, page 17-1](#)
- [Creating or Modifying an IP ACL, page 17-2](#)
- [Associating an IP ACL with an Application, page 17-12](#)
- [Applying an IP ACL to an Interface, page 17-13](#)
- [Deleting an IP ACL, page 17-14](#)

Introducing IP ACLs

In a managed ACNS network environment, administrators need to be able to prevent unauthorized access to various devices and services. ACNS software supports standard and extended ACLs that allow administrators to restrict access to or through an ACNS network device, such as a Content Engine, Content Router, or Content Distribution Manager. Administrators can use ACLs to reduce the infiltration of hackers, worms, and viruses that can harm the corporate network.

ACNS software also provides controls that allow various services to be tied to a particular interface. For example, the administrator can use IP ACLs to define a public interface on the Content Engine for content serving and a private interface for management services (for example, Telnet, SSH, SNMP, HTTPS, and software upgrades). A device attempting to access one of the services must be on a list of trusted devices before it is allowed access. The implementation of ACLs for incoming traffic on certain ports for a particular protocol type is similar to the ACL support for the Cisco Global Site Selector and Cisco routers.

To use ACLs, the system administrator must first configure ACLs and then apply them to specific services or interfaces. The following are some examples of how IP ACLs can be used in various enterprise deployments:

- An application layer proxy firewall with a hardened outside interface has no ports exposed. (“Hardened” means that the interface carefully restricts which ports are available for access primarily for security reasons. Because the interface is outside, many types of attacks are possible.) The Content Engine’s outside address is globally accessible from the Internet, while its inside address is private. The inside interface has an ACL to limit Telnet, SSH, GUI, and Content Distribution Manager traffic.
- A Content Engine is deployed anywhere within the enterprise. Like routers and switches, the administrator wants to limit Telnet, SSH, Content Distribution Manager, and GUI access to the IT source subnets.
- A Content Engine is deployed as a reverse proxy in an untrusted environment, and the administrator wishes to allow only port 80 inbound traffic on the outside interface and outbound connections on the backend interface.
- A Content Engine using WCCP is positioned between a firewall and an Internet router or a subnet off the Internet router. Both the Content Engine and the router must have IP ACLs.

**Note**

To allow or deny access to the user through the TFTP protocol, you must explicitly configure the **ip access-list** command. Unless the ACL is explicitly configured for the TFTP service, the security of the content can be at risk, and TFTP will not work properly.

The **ip access-list** command replaces the **trusted-host** command. If the **trusted-host** command is used on Content Engines using ACNS releases earlier than 5.1, and the devices are subsequently upgraded to ACNS 5.1 or later, the command shows up in the CLI but has no effect on the TFTP protocol. You can delete trusted host configurations by using the **no trusted-host** command.

Creating or Modifying an IP ACL

**Note**

IP ACLs are defined for individual devices only. IP ACLs cannot be managed globally across the ACNS network or through device groups. IP ACLs cannot be defined on IP/TV Program Manager devices.

When you create an IP ACL, you should note the following constraints:

- IP ACL names must be unique within the device.
- IP ACL names must be limited to 30 characters and contain no white space or special characters.
- The Content Distribution Manager can manage up to 50 IP ACLs and a total of 500 conditions per device.
- When the IP ACL name is numeric, numbers 1 through 99 denote standard IP ACLs and numbers 100 through 199 denote extended IP ACLs. IP ACL names that begin with a number cannot contain nonnumeric characters.
- Extended IP ACLs cannot be used with SNMP, TFTP, or WCCP applications.

Creating a New IP ACL

To create a new IP ACL using the Content Distribution Manager GUI, follow these steps:

-
- Step 1** In the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon next to the name of the device for which you want to create an IP ACL.
 - Step 3** To display the entire table of contents, click the **Show All** button above the Contents pane.
 - Step 4** In the Contents pane, choose **General Settings > Network > IP ACL**. The IP ACL window appears.
 - Step 5** In the taskbar, click the **Create a new IP ACL** icon. The Creating New IP ACL window appears. In the Name field, enter a name, observing the naming rules for IP ACLs.
 - Step 6** From the drop-down list, choose an IP ACL type (**Standard** or **Extended**). The default is Standard.
 - Step 7** To save the IP ACL, click **Submit**. The window refreshes, and the Modifying IP ACL window for the newly created IP ACL appears.



Note Clicking **Submit** at this point merely saves the IP ACL; IP ACLs without any conditions defined do not appear on the individual devices.

Adding Conditions to an IP ACL

To add conditions to an IP ACL, follow these steps:

-
- Step 1** In the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon next to the name of the device for which you want to create an IP ACL.
 - Step 3** In the Contents pane, choose **General Settings > Network > IP ACL**. The IP ACL window appears.
 - Step 4** Click the **Edit** icon next to the name of the IP ACL for which you want to add conditions. The Modifying IP ACL window appears.
 - Step 5** In the taskbar, click the **Create New Condition** icon. The Creating New Condition window appears. (See [Figure 17-1](#).)



Note The number of available fields for creating IP ACL conditions depends on the type of IP ACL that you have created, either standard or extended.

Figure 17-1 Creating a New Condition for an Extended IP ACL Window

Step 6 Enter values for the properties that are enabled for the type of IP ACL that you are creating.

- To create a standard IP ACL, go to [Step 7](#).
- To create an extended IP ACL, go to [Step 8](#).

Step 7 To set up conditions for a standard IP ACL, follow these steps:

- a. From the drop-down list, choose a purpose (**Permit** or **Deny**).
- b. In the Source IP field, enter the source IP address.
- c. In the Source IP Wildcard field, enter a source IP wildcard address.
- d. To save the condition, click **Submit**. The Modifying IP ACL window reappears, displaying the condition and its configured parameters in tabular format.
- e. To add another condition to the IP ACL, repeat the steps.
- f. To reorder your list of conditions from the Modifying IP ACL window, use the Up or Down Arrows in the Move column, or click a column heading to sort by any configured parameter.



Note The order of the conditions listed in the Content Distribution Manager GUI becomes the order in which IP ACLs are applied to the device.

- g. When you have finished adding conditions to the IP ACL, and you are satisfied with all your entries and the order in which the conditions are listed, click **Submit** in the Modifying IP ACL window to commit the IP ACL to the device database.

A green “Change submitted” indicator appears in the lower right corner of the Modifying IP ACL window to indicate that the IP ACL is being submitted to the device database.

Table 17-1 describes the fields in a standard IP ACL.

Table 17-1 Standard IP ACL Conditions

Field	Default Value	Description
Purpose* ¹	Permit	Specifies whether a packet is to be passed (Permit) or dropped (Deny).
Source IP*	0.0.0.0	Number of the network or host from which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Source IP Wildcard*	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.

1. * = required field.

Step 8 To set up conditions for an extended IP ACL, follow these steps:

- a. From the drop-down list, choose a purpose (**Permit** or **Deny**).
- b. From the Extended Type drop-down list, choose **Generic**, **TCP**, **UDP**, or **ICMP**. (See Table 17-2.)
After you choose a type of extended IP ACL, various options become available in the GUI, depending on what type you choose.
- c. In the fields that are enabled for the chosen type, enter the data. (See Table 17-3 through Table 17-6.)
- d. To save the condition, click **Submit**. The Modifying IP ACL window reappears, displaying the condition and its configured parameters in tabular format.
- e. To add another condition to the IP ACL, repeat the steps.
- f. To reorder your list of conditions from the Modifying IP ACL window, use the Up or Down Arrows in the Move column, or click a column heading to sort by any configured parameter.



Note The order of the conditions listed in the Content Distribution Manager GUI becomes the order in which IP ACLs are applied to the device.

- g. When you have finished adding conditions to the IP ACL, and you are satisfied with all your entries and the order in which the conditions are listed, click **Submit** in the Modifying IP ACL window to commit the IP ACL to the device database.

A green “Change submitted” indicator appears in the lower right corner of the Modifying IP ACL window to indicate that the IP ACL is being submitted to the device database.

Table 17-2 Extended IP ACL Conditions

Field	Default Value	Description
Purpose* ¹	Permit	Specifies whether a packet is to be passed or dropped. Choices are Permit or Deny.
Extended Type*	Generic	Specifies the Internet protocol to be applied to the condition. When selected, the GUI window refreshes with applicable field options enabled. Choices are as follows: • Generic (See Table 17-3.) • TCP (See Table 17-4.) • UDP (See Table 17-5.) • ICMP (See Table 17-6.)

1. * = required field.

Table 17-3 Extended IP ACL Generic Condition

Field	Default Value	Description
Purpose* ¹	Permit	Specifies whether a packet is to be passed (Permit) or dropped (Deny).
Extended Type*	Generic	Matches any Internet protocol.
Protocol	ip	Internet protocol (gre , icmp , ip , tcp , or udp). To match any Internet protocol, use the keyword ip .
Source IP*	0.0.0.0	Number of the network or host from which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Source IP Wildcard*	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.
Destination IP	0.0.0.0	Number of the network or host to which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Destination IP Wildcard	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.

1. * = required field.

Table 17-4 Extended IP ACL TCP Condition

Field	Default Value	Description
Purpose* ¹	Permit	Specifies whether a packet is to be passed (Permit) or dropped (Deny).
Extended Type*	TCP	Matches the TCP Internet protocol.
Established	Unchecked (false)	When checked, a match with the ACL condition occurs if the TCP datagram has the ACK or RST bits set, indicating an established connection. Initial TCP datagrams used to form a connection are not matched.
Source IP*	0.0.0.0	Number of the network or host from which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Source IP Wildcard*	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.
Source Port 1	0	Decimal number or name of a TCP port. Valid port numbers are 0 to 65535. Valid TCP port names are as follows: - ftp - ftp-data - https - netbios-dgm - netbios-ns - netbios-ss - nfs - rtsp - ssh - telnet - www
Source Operator	range	Specifies how to compare the source ports against incoming packets. Choices are <, >, ==, !=, or range.
Source Port 2	65535	Decimal number or name of a TCP port. See Source Port 1.
Destination IP	0.0.0.0	Number of the network or host to which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Destination IP Wildcard	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.

Table 17-4 **Extended IP ACL TCP Condition (continued)**

Field	Default Value	Description
Destination Port 1	0	Decimal number or name of a TCP port. Valid port numbers are 0 to 65535. Valid TCP port names are as follows: - ftp - ftp-data - https - netbios-dgm - netbios-ns - netbios-ss - nfs - rtsp - ssh - telnet - www
Destination Operator	range	Specifies how to compare the destination ports against incoming packets. Choices are <, >, ==, !=, or range.
Destination Port 2	65535	Decimal number or name of a TCP port. See Destination Port 1.

1. * = required field.

Table 17-5 **Extended IP ACL UDP Condition**

Field	Default Value	Description
Purpose* ¹	Permit	Specifies whether a packet is to be passed (Permit) or dropped (Deny).
Extended Type*	UDP	Matches the UDP Internet protocol.
Established	—	Not available for UDP.
Source IP*	0.0.0.0	Number of the network or host from which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Source IP Wildcard*	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.

Table 17-5 **Extended IP ACL UDP Condition (continued)**

Field	Default Value	Description
Source Port 1	0	Decimal number or name of a UDP port. Valid port numbers are 0 to 65535. Valid UDP port names are as follows: • bootpc • bootps • domain • netbios-dgm • netbios-ns • netbios-ss • nfs • ntp • snmp • snmptrap • tacacs • tftp • wccp
Source Operator	range	Specifies how to compare the source ports against incoming packets. Choices are <, >, ==, !=, or range.
Source Port 2	65535	Decimal number or name of a UDP port. See Source Port 1.
Destination IP	0.0.0.0	Number of the network or host to which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Destination IP Wildcard	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.

Table 17-5 **Extended IP ACL UDP Condition (continued)**

Field	Default Value	Description
Destination Port 1	0	The decimal number or name of a UDP port. Valid port numbers are 0 to 65535. Valid UDP port names are as follows: • bootpc • bootps • domain • netbios-dgm • netbios-ns • netbios-ss • nfs • ntp • snmp • snmptrap • tacacs • tftp • wccp
Destination Operator	range	Specifies how to compare the destination ports against incoming packets. Choices are <, >, ==, !=, or range.
Destination Port 2	65535	Decimal number or name of a UDP port. See Destination Port 1.

1. * = required field.

Table 17-6 **Extended IP ACL ICMP Condition**

Field	Default Value	Description
Purpose* ¹	Permit	Specifies whether a packet is to be passed (Permit) or dropped (Deny).
Extended Type*	ICMP	Matches the ICMP Internet protocol.
Source IP*	0.0.0.0	Number of the network or host from which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.
Source IP Wildcard*	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.
Destination IP	0.0.0.0	Number of the network or host to which the packet is being sent, specified as a 32-bit quantity in 4-part dotted decimal format.

Table 17-6 **Extended IP ACL ICMP Condition (continued)**

Field	Default Value	Description
Destination IP Wildcard	255.255.255.255	Wildcard bits to be applied to the source, specified as a 32-bit quantity in 4-part dotted decimal format. Place a 1 in the bit positions that you want to ignore and identify bits of interest with a 0.
ICMP Param Type*	None	Choices are None, Type/Code, or Msg. None—Disables the ICMP Type, Code, and Message fields. Type/Code—Allows ICMP messages to be filtered by ICMP message type and code. Also enables the ability to set an ICMP message code number. Msg—Allows a combination of type and code to be specified using a keyword. Activates the ICMP message drop-down list. Disables the ICMP Type field. Note Refer to the <i>Cisco ACNS Software Command Reference, Release 5.5</i> for further explanation of related keywords. (See the ip access-list extended command.)
ICMP Message*	administratively-prohibited	Allows a combination of ICMP type and code to be specified using a keyword chosen from the drop-down list. Note Refer to the <i>Cisco ACNS Software Command Reference, Release 5.5</i> for further explanation of these keywords. (See the ip access-list extended command.)
ICMP Type*	0	Number from 0 to 255. This field is enabled when you choose Type/Code .
Use ICMP Code*	Unchecked	When checked, enables the ICMP Code field.
ICMP Code*	0	Number from 0 to 255. Message code option that allows ICMP messages of a particular type to be further filtered by an ICMP message code.

1. * = required field.

Modifying, Deleting, or Reordering a Condition

To modify or delete an individual condition from an IP ACL, follow these steps:

- Step 1** In the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the device whose IP ACL that you want to modify.
- Step 3** In the Contents pane, choose **General Settings > Network > IP ACL**.

- Step 4** Click the **Edit** icon next to the name of the IP ACL that you want to modify. The Modifying IP ACL window appears, listing all the conditions applied to the IP ACL.
- Step 5** Click the **Edit Condition** icon next the condition that you want to modify or delete. The Modifying Condition window appears.
- Step 6** To modify the condition, change any allowable field as necessary.
- Step 7** To delete the condition, click the **Trash (Delete IP ACL Condition)** icon in the taskbar.
- Step 8** To reorder your list of conditions, use the Up or Down arrows in the Move column, and click **Submit**.

Associating an IP ACL with an Application

The Content Distribution Manager GUI allows the association of standard IP ACLs with SNMP, TFTP, native FTP, and WCCP. Any device that attempts to access one of these applications associated with an access control list (ACL) must be on the list of trusted devices to be allowed access. You can associate any previously configured standard IP ACL with SNMP, TFTP, native FTP and WCCP; however, you can associate an extended IP ACL only with the WCCP application.

To associate a standard IP ACL with one of these applications, follow these steps:

- Step 1** In the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the device for which you have created an IP ACL.
- Step 3** In the Contents pane, choose **General Settings > Network > IP ACL Feature Usage**. The IP ACL Feature Settings window appears. [Table 17-7](#) describes the fields in this window and provides the corresponding CLI global configuration commands.
- Step 4** From the drop-down lists, choose the name of an IP ACL for SNMP, TFTP, FTP-IN, FTP-OUT, and WCCP. If you do not want to associate an IP ACL with one of the applications, choose **Do Not Set**.



Note The drop-down lists for SNMP, TFTP, and FTP-IN contain the names of all your standard IP ACLs. Only standard IP ACLs (not extended IP ACLs) can be associated with SNMP, TFTP, and native FTP-IN applications.

- Step 5** To save the settings, click **Submit**.

Table 17-7 IP ACL Feature Settings

GUI Parameter	Function	CLI Command
SNMP	Associates a standard IP ACL with SNMP.	snmp-server access-list { <i>std_acl_number</i> <i>acl_name</i> }
TFTP	Associates a standard IP ACL with TFTP.	tftp-server access-list { <i>std_acl_number</i> <i>acl_name</i> }
FTP-IN	Associates a standard IP ACL with inbound FTP native proxy mode and transparently redirected connections.	ftp-native access-list in { <i>std_acl_number</i> <i>acl_name</i> }

Table 17-7 IP ACL Feature Settings (continued)

GUI Parameter	Function	CLI Command
FTP-OUT	Associates an extended IP ACL with outbound FTP native proxy mode and transparently redirected connections.	ftp-native access-list out { <i>ext_acl_number</i> <i>acl_name</i> }
WCCP	Associates any IP ACL with WCCP.	wccp access-list { <i>std_acl_number</i> <i>ext_acl_number</i> <i>acl_name</i> }

Applying an IP ACL to an Interface

ACNS software allows SNMP, TFTP, native FTP, and WCCP applications to be applied to a particular interface (such as management services to a private IP address) so that the Content Engine can have one interface in a public IP address space that serves content and another interface in a private IP address space that the administrator uses for management purposes. This feature ensures that clients can access the Content Engine only in the public IP address space for serving content and not access it for management purposes. A device attempting to access one of these applications that is associated with an IP ACL must be on the list of trusted devices to be allowed access.

To apply an IP ACL to an interface, follow these steps:

- Step 1** In the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the device for which you want to apply an IP ACL to an interface.
- Step 3** In the Contents pane, choose **General Settings > Network > Network Interfaces**.

The Network Interfaces window for the device appears. This window displays all the interfaces available on that device.



Note The Port Type column may contain a PortChannel interface indicating an EtherChannel configuration. EtherChannel for ACNS 5.x software supports the grouping of up to four same-speed network interfaces into one virtual interface. (See the [“Configuring EtherChannel” section on page 14-11.](#))

- Step 4** Click the **Edit** icon next to the name of the interface to which you want to apply an IP ACL. The Modifying Network Interface window appears. (See [Figure 17-2.](#))

Figure 17-2 Modifying Network Interface—Applying an IP ACL

Step 5 Scroll to the bottom of the window, and from the Inbound ACL drop-down list, choose the name of an IP ACL. In the example in [Figure 17-2](#), **ext_acl_2** is chosen.

Step 6 From the Outbound ACL drop-down list, choose the name of an ACL.



Note The only network interface properties that can be altered from the Content Distribution Manager GUI are the inbound and outbound IP ACLs. All other property values are populated from the device database and are read-only in the Content Distribution Manager GUI.

Step 7 To save the settings, click **Submit**.

To apply an IP ACL to an interface from the CLI, use the following interface configuration command:

```
interface {FastEthernet | GigabitEthernet} slot/port ip access-group {accesslistnumber | accesslistname} {in | out}
```

Deleting an IP ACL

You can delete an IP ACL, including all conditions and associations with network interfaces and applications, or you can delete only the IP ACL conditions. Deleting all conditions allows you to change the IP ACL type if you choose to do so. The IP ACL entry continues to appear in the IP ACL listing; however, it is in effect nonexistent.

To delete an IP ACL, follow these steps:

-
- Step 1** In the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the device with the IP ACL configured that you want to delete.
- Step 3** In the Contents pane, choose **General Settings > Network > IP ACL**. The IP ACL for Content Engine window appears.
- Step 4** Click the **Edit** icon next to the name of the IP ACL that you want to delete. The Modifying IP ACL window appears. If you created conditions for the IP ACL, you have two options for deletion:
- **Delete ACL**
This option removes the IP ACL, including all conditions and associations with network interfaces and applications.
 - **Delete All Conditions**
This option removes all the conditions, while preserving the IP ACL name.
- Step 5** To delete the entire IP ACL, click the large **Trash (Delete ACL)** icon in the taskbar. You are prompted to confirm your action. Click **OK**. The record is deleted.
- Step 6** To delete only the conditions, click the small **Delete All Conditions** Trash/List icon in the taskbar. You are prompted to confirm your action. Click **OK**. The window refreshes, conditions are deleted, and the ACL Type field becomes available.
-



Configuring Group Authorization Using an Access Control List

This chapter explains how to configure group authorization using an access control list. It contains the following sections:

- [Configuring an Access Control List for Group Authorization, page 18-1](#)
- [Enabling the Access Control List, page 18-2](#)
- [Modifying an Access Control List, page 18-3](#)
- [Viewing Access Control List Settings, page 18-5](#)

In ACNS 5.x software, you can configure group authorization using an access control list (ACL) only after a user has been authenticated against an NTLM or LDAP HTTP-request authentication server. The use of this list configures a group privilege when members of the group are accessing content provided by the Content Engine. Using the ACL allows or prevents users belonging to certain groups from viewing specific content. This authorization feature offers more specific access control because that access is only allowed to certain groups.

In ACNS 5.x software, the access control list contains the following feature enhancements and limitations:

- A user can belong to several groups.
- A user can belong to an unlimited number of groups within groupname strings.
- A groupname string is a case-sensitive string with mixed-case alphanumeric characteristics.
- Each unique groupname string cannot exceed 128 characters.



Note If the unique groupname string is longer than 128 characters, the group is ignored.

Configuring an Access Control List for Group Authorization

To configure an ACL for group authorization, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure group authorization.
 - Step 3** To display the entire table of contents, click the **Show All** button above the Contents pane.

- Step 4** In the Contents pane, choose **General Settings > Authentication > Access List > Configure Access Control List**. The Access List Settings for Content Engine window appears listing all configured ACL settings.
- Step 5** To configure a new access list setting, click the Create New Access List Setting icon in the taskbar. The Creating New Access List for Content Engine window appears.
- Step 6** Configure the Action parameter by choosing either **Permit** or **Deny**.
- Step 7** Configure the Group Name parameter.
- Specify any user group by choosing **Any Group Name**.
 - Specify a particular user group by choosing **Enter Group Name**. Enter a group name in the text field, such as, Marketing or Sales.
- Step 8** Click **Submit** to save the configuration. The Access List Settings for Content Engine window reappears.
- Step 9** To configure additional ACL settings, repeat [Step 5](#) through [Step 8](#).
- Step 10** Review the order of your ACL settings and reposition them in the list, if necessary.
- a. To change the position of the settings in the list that you are creating, click the **Change Position** button. A dialog box pops up.
 - b. In the Create Access List at position field, enter a position number.
 - c. In the dialog box, click the **Submit** button. The dialog box closes and the Access List Settings for Content Engine window reappears.
- Alternatively, see the [“Repositioning Settings in the Access Control List”](#) section on page 18-5.
- Step 11** To save the ACL, click **Submit**, and then proceed to the next section, [“Enabling the Access Control List,”](#) to enable the ACL.

Enabling the Access Control List

To enable the use of the ACL, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to enable the ACL.
- Step 3** In the Contents pane, choose **General Settings > Authentication > Access List > Enable Access Control List**. The Enable Access Control List window appears.
- Step 4** Check the **Enable Access Control List** check box.
- Step 5** Click **Submit**.

Using the CLI to Configure and Enable an ACL for Group Authorization

To configure an ACL from the CLI, use the **access-lists 300** global configuration command. This command allows you to permit or deny a group from accessing the Internet by using the Content Engine. For instance, use the **access-lists 300 deny groupname marketing** command to prevent any user from the marketing group from accessing content through the Content Engine.

Use the **access-lists enable** global configuration command to enable the use of the ACL.

You can display the configuration of the ACL by using the **show access-lists 300 EXEC** command, as shown in this example:

```
ContentEngine# show access-lists 300
Access Control List Configuration
-----
Access Control List is enabled
Groupname-based List (300)
1. permit groupname techpubs
2. permit groupname acme1
3. permit groupname engineering
4. permit groupname sales
5. permit groupname marketing
6. deny groupname any
```

To display statistical information for the access control list, use the **show statistics access-lists 300** command, as shown in this example:

```
ContentEngine# show statistics access-lists 300
Access Control Lists Statistics
-----
Groupname and username-based List (300)
Number of requests:      1
Number of deny responses: 0
Number of permit responses: 1
```

To reset the statistical information for the access control list, use the **clear statistics access-lists 300** command.

```
ContentEngine# clear statistics access-lists 300
```

```
ContentEngine(config)# access-lists 300 permit groupname acme1 position 2
```

Modifying an Access Control List

To modify an ACL, follow these steps:

-
Step 1	From the Content Distribution Manager GUI, choose Devices > Devices .
Step 2	Click the Edit icon next to the Content Engine for which you want to modify the group authorization.
Step 3	In the Contents pane, choose General Settings > Authentication > Access List > Configure Access Control List . The Access List Settings for Content Engine window appears listing all configured settings.
Step 4	Click the Edit Access List Setting icon next to the ACL that you want to change. The Modifying Access List window appears.
Step 5	Change any of the configurable parameters that you want to modify.
Step 6	To save the new ACL setting, click Submit .
-

Changing the Currently Applied Settings


Access list settings can be applied from the Content Engine or from a device group. The source of the currently applied settings is shown in the Access List Settings window (**Devices > General Settings > Authentication > Access List > Configure Access Control List**) for any chosen Content Engine. You can change the source of the currently applied settings from the Content Engine to a device group, provided that the following conditions are met:

- You have configured a device group.
- The Content Engine is assigned to the device group.
- The access list is configured in the device group.

To change the current applied settings, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to enable the ACL.
- Step 3** In the Contents pane, choose **General Settings > Authentication > Access List > Enable Access Control List**. The Enable Access Control List for Content Engine window appears. (See [Figure 18-1](#).)

Figure 18-1 Enable Access Control List for Content Engine Window



- Step 4** From the drop-down list in the taskbar, chose a device group. (In the example in [Figure 18-1](#), the device group *dg1* has been chosen.
- Step 5** To save the changes, click **Submit**.

Repositioning Settings in the Access Control List

To reposition the settings in the ACL, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon next to the Content Engine for which you want to reposition the ACL settings.
 - Step 3** In the Contents pane, choose **General Settings > Authentication > Access List > Configure Access Control List**. The Access List Settings for Content Engine window appears listing all configured settings.
 - Step 4** To move a setting up or down in the list order, click the **Up** or **Down** arrow in the Move column. The setting shifts one position up or down. The list is saved in the new order.
-

Viewing Access Control List Settings

To view the ACL settings, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon next to the Content Engine for which you want to reposition the ACL settings.
 - Step 3** In the Contents pane, choose **General Settings > Authentication > Access List > Configure Access Control List**. The Access List Settings for Content Engine window appears listing all configured settings.

Columns display the position, action, and group name of each setting configured in the ACL.



Using the Transaction Logs

This chapter explains how to use the transaction logs and contains the following sections:

- [Understanding Transaction Log Formats, page 19-1](#)
- [Transaction Logging and NTLM Authentication, page 19-7](#)
- [Usage Guidelines for Log Files, page 19-8](#)
- [Enabling Transaction Logging with the Content Distribution Manager GUI, page 19-10](#)
- [Using WMT Transaction Logging, page 19-15](#)
- [Using Real-Time Transaction Logging, page 19-19](#)

Transaction logs allow administrators to view the traffic that has passed through the Content Engine. Typical fields in the transaction log are the date and time when a request was made, the URL that was requested, whether it was a cache hit or a cache miss, the type of request, the number of bytes transferred, and the source IP address.

Understanding Transaction Log Formats

In ACNS 5.x software, the user can choose among Squid, Extended Squid, Apache, or customized log formats for the transaction log.

Squid-Style Transaction Logging

The Squid-style log format is the default format for transaction logging in the Content Engine. The Squid log file format used is the native log file format associated with the Squid-1.1 *access.log* file format. For details on the Squid-1.1 native log file format, refer to the Squid documentation “Frequently Asked Questions,” section 6.6, *access.log* heading at the following URL:

<http://wiki.squid-cache.org/SquidFaq/SquidLogs>

The Squid log file format is as follows:

time elapsed remotehost code/status bytes method URL rfc931 peerstatus/peerhost type

A Squid log format example looks like this:

```
1012429341.115 100 172.16.100.152 TCP_REFRESH_MISS/304 1100 GET
http://www.cisco.com/images/homepage/news.gif - DIRECT/www.cisco.com -
```


Extended Squid Log Format

The Extended Squid format logs the associated username for each record in the log file, in addition to the fields logged by the Squid-style format, and is used for billing purposes. In this format the Rfc931 field associated with the Squid format is used to log the authorized user. This field always contains a “-” (dash) if no user information is available.

An Extended Squid-style log format example looks like this:

```
1012429341.115 100 172.16.100.152 TCP_MISS/302 184 GET http://www.cisco.com/cgi-bin/login
myloginname DIRECT/www.cisco.com
```

Apache-Style Transaction Logging

The Apache format is the Common Log File (CLF) format defined by the World Wide Web Consortium (W3C) working group. This format is compatible with many industry-standard log tools. For more information, see the W3C Common Log Format website at the following URL:

<http://www.w3.org/Daemon/User/Config/Logging.html>.

The Apache-style log file format is as follows:

```
remotehost rfc931 authuser date request status bytes
```

An Apache-style log file format example looks like this:

```
172.16.100.152 - - [Wed Jan 30 15:26:26 2002]
“GET/http://www.cisco.com/images/homepage/support.gif HTTP/1.0” 200 632
```

Custom Format Transaction Logging

The **transaction-logs format custom** command allows you to use a log format string to log additional fields that are not included in the predefined native Squid format, the Extended Squid format, or the Apache CLF format. The log format string is a string that can contain the tokens listed in [Table 19-1](#) and that mimics the Apache log format string. The log format string can contain literal characters that are copied into the log file. Double backslashes (\) can be used to represent a literal backslash, and a backslash followed by a single quote (\') can be used to represent a literal single quote. A literal double quote cannot be represented as part of the log format string. The control characters \t and \n can be used to represent a tab and a new line character, respectively.

[Table 19-1](#) lists the acceptable format tokens for the log format string. The “...” portion of the format tokens shown in this table represents an optional condition. This portion of the format token can be left blank, as in %a. If an optional condition is included in the format token and the condition is met, then what is shown in the Value column of [Table 19-1](#) is included in the transaction log output. If an optional condition is included in the format token but the condition is not met, the resulting transaction log output is replaced with a hyphen (-). The form of the condition is a list of HTTP status codes, which may or may not be preceded by an exclamation point (!). The exclamation point is used to negate all of the status codes that follow it, meaning that the value associated with the format token is logged if none of the status codes listed after the ! match the HTTP status code of the request. If any of the status codes listed after the ! match the HTTP status code of the request, then a hyphen (-) is logged.

For example, “%400,501{User-Agent}i” logs the User-Agent header value on 400 errors and 501 errors (Bad Request, Not Implemented) only, whereas “%!200,304,302{Referer}i” logs the Referer header value on all requests that did not return a normal status.

The custom format currently supports the following request headers:

- User-Agent
- Referer
- Host
- Cookie

The output of each of the following Request, Referer, and User-Agent format tokens specified in the custom log format string is always enclosed in double quotation marks in the transaction log entry:

%r

%{Referer}i

%{User-Agent}i

The %{Cookie}i format token is generated without the surrounding double quotation marks because the Cookie value itself can contain double quotes. The Cookie value can contain multiple attribute-value pairs that are separated by spaces. When you use the Cookie format token in a custom format string, we recommend that you position it in the last field of the format string so that the Cookie format token can be more easily parsed by transaction log reporting tools. Alternatively, if you use the format token string “\%{Cookie}i\”, the Cookie header can be surrounded by single quotes.

The following command can be entered to generate the well-known Apache Combined Log Format:

```
transaction-logs format custom "[%d]t/%b)t/%Y}t:%H}t:%M}t:%S}t
%{z}t] %r %s %b %{Referer}i %{User-Agent}i"
```

The following transaction log entry example in the Apache Combined Format is configured by using the preceding custom format string:

```
[11/Jan/2003:02:12:44 -0800] "GET http://www.cisco.com/swa/i/site_tour_link.gif HTTP/1.1" 200
3436 "http://www.cisco.com/" "Mozilla/4.0 (compatible; MSIE 5.5; Windows NT 5.0)"
```

Table 19-1 Custom Format Log Format String Values

Format Token	Value
%...a	IP address of the requesting client.
%...A	IP address of the Content Engine.
%...B %...b	Bytes sent, excluding HTTP headers.
%...c	Connection status when response is completed, where: X = Connection was aborted before the response was completed. + = Connection can be kept alive after the response is sent. – = Connection is closed after the response is sent.
%...f	Filename.
%...h	Remote host (IP address of the requesting client is logged).
%...H	Request protocol.
%...{Foobar}i	Contents of Foobar: header lines in the request that is sent to the server. The value of Foobar can be one of the following headers: User-Agent, Referer, Host, or Cookie.
%...l	Remote log name. Not implemented on the Content Engine, so a hyphen (-) is logged.
%...m	Request method.

Table 19-1 Custom Format Log Format String Values (continued)

Format Token	Value
%...p	Canonical port of the server servicing the request. Not applicable on the Content Engine, so a hyphen (-) is logged.
%...P	Process ID of the child that serviced the request.
%...q	Query string (which is preceded by a ? if a query string exists; otherwise, it is an empty string).
%...r	First line of the request.
%...s	Status. The translog code always returns the HTTP response code for the request.
%...t	Time in common log time format (or standard English format).
%...{format}t	Time in the form given by the format token specified in Table 19-2 .
%...T	Time consumed to serve the request in seconds (a floating point number with 3 decimal places).
%...u	Remote user.
%...U	URL path requested, not including query strings.
%...v %...V	Value of the host request header field reported if the host appeared in the request. If the host did not appear in the host request header, the IP address of the server specified in the URL is reported.

[Table 19-2](#) specifies the format token for the date and time of the format token %...{format}t listed in [Table 19-1](#).

Table 19-2 Format Token for Date and Time

Format Token	Value
%a	Abbreviated weekday name.
%A	Full weekday name.
%b	Abbreviated month name.
%B	Full month name.
%c	Date and time representation.
%C	Century number (year/100) as a 2-digit integer.
%d	Day of the month as a decimal number (01–31).
%D	Equivalent to %m/%d/%y. (Note that in countries other than the USA, the form %d/%m/%y is commonly used. This means that in international context this format is ambiguous and should not be used.)
%e	Like %d, the day of the month as a decimal number, but a leading zero is replaced by a space.
%G	ISO 8601 year with the century as a decimal number. The 4-digit year corresponding to the ISO week number. (See %V.) This has the same format and value as %y, except that if the ISO week number belongs to the previous or next year, that year is used instead.
%g	Like %G, but without century; that is, with a 2-digit year (00–99).
%h	Equivalent to %b.

Table 19-2 *Format Token for Date and Time (continued)*

Format Token	Value
%H	Hour as a decimal number using a 24-hour clock (00–23).
%I	Hour as a decimal number using a 12-hour clock (01–12).
%j	Day of the year as a decimal number (001–366).
%k	Hour (24-hour clock) as a decimal number (0–23); single digits are preceded by a blank. (See also %H.)
%l	Hour (12-hour clock) as a decimal number (1–12); single digits are preceded by a blank. (See also %I.)
%m	Month as a decimal number (01–12).
%M	Minute as a decimal number (00–59).
%n	New line character.
%p	Either AM or PM according to the given time value, or the corresponding strings for the current locale. Noon is treated as pm and midnight as am.
%P	Like %p but in lowercase: am or pm or a corresponding string for the current locale.
%r	Time in a.m. or p.m. notation. This is equivalent to ‘%I:%M:%S %p.’
%R	Time in 24-hour notation (%H:%M). For a version including the seconds, see %T below.
%s	Number of seconds since the epoch; that is, since 1970-01-01 00:00:00 UTC.
%S	Second as a decimal number (00–61).
%t	Tab character.
%T	Time in 24-hour notation (%H:%M:%S).
%u	Day of the week as a decimal (1–7), Monday being 1. (See also %w.)
%U	Week number of the current year as a decimal number (00–53), starting with the first Sunday as the first day of week 01. (See also %V and %W.)
%V	ISO 8601:1988 week number of the current year as a decimal number (01–53), where week 1 is the first week that has at least 4 days in the current year, and with Monday as the first day of the week. (See also %U and %W.)
%w	Day of the week as a decimal (0–6), Sunday being 0. See also %u.
%W	Week number of the current year as a decimal number (00–53), starting with the first Monday as the first day of week 01.
%x	Date representation without the time.
%X	Time representation without the date.
%y	Year as a decimal number without a century (00–99).
%Y	Year as a decimal number, including the century.
%z	Time zone as hour offset from GMT. Required to emit RFC 822-conformant dates (using “%a, %d %b %Y %H:%M:%S %z”).
%Z	Time zone or name or abbreviation.
%%	Literal % character.

W3C Customizable Logging Format

To provide more flexibility in logging, ACNS software supports W3C Customizable Logging Format apart from the fixed formats like Apache Common Log Format (CLF), Squid, and Extended Squid formats.

W3C customizable Logging Format supports a set of format tokens that exposes the underlying Translog (Transaction Logs) tokens. The W3C Customizable Logging Format is limited in that it was defined from the HTTP web server perspective and does not offer certain web cache-specific custom options such as those supplied by the fixed Squid format. Consequently, additional format tokens that are extensions to the W3C Customized Logging Format were added (in the ACNS 5.3 software release) to support additional Cisco and Squid-like customized logging fields. These format tokens provide support for Squid-like logging format from within the W3C customizable token set.

ACNS 5.5 software supports the following transaction logging formats:

- Support for the Extended Squid-equivalent internal tokens that were not supported by the W3C format
- Support for an additional hierarchy token that treats a configured HTTP outgoing proxy ("http outgoing-proxy") as a Squid-style "DEFAULT_PARENT" hierarchy event

ACNS 5.5 software includes the following special token sequence for the W3C Customizable Logging Format:

%...{}C

The “...” is optional. If specified, it can be a sequence of conditional HTTP response codes separated by commas. The “C” is an uppercase “C” and defines the extended customizable behavior token set, for which tokens are defined by the directive, which is a two-character token directive.

See [Table 19-3](#) for a list of existing and new directives from the Extended Squid format, which are not currently supported by the W3C definitions, but are supported in ACNS 5.5 software.

Table 19-3 Translog Token Directives

Format Token	Value
%...{es}C	Current time presented as the number of seconds that have elapsed since the Epox (Jan. 1st. 1970).
%...{em}C	Current number of milliseconds that have elapsed since the Epox (Jan. 1st. 1970).
%...{te}C	Number of milliseconds that have elapsed until the request was completed.
%...{rd}C	Squid-like cache-status code string (for example, TCP_HIT and TCP_CLIENT_REFRESH_MISS).
%..{cs}C	Number of bytes sent to the client (including the protocol headers).
%...{rh}C	Strict Squid-style hierarchy as it applies to the Content Engine.
%...{rH}CE	Extended Squid-style hierarchy. Same as “%...{rh}C” except when an outgoing-proxy is explicitly defined and is used to satisfy a request, and then the “DEFAULT_PARENT/proxy_ip_address” is logged instead of the “DIRECT/origin_server_ip_address.”
%...{rt}C	Mime-Type of the object in the response, as specified by any protocol headers which define such.

Table 19-3 *Translog Token Directives (continued)*

Format Token	Value
%...{ru}C	URL being requested, including any additional query strings.
%...{as}C	Application specific information. Certain request handling applications might attempt to log a certain string here, which is supported as part of the Squid format specification. For example, SmartFilter URL filtering will log information where this token sequence is used.

In addition to the tokens listed in [Table 19-3](#), you can condense multiple “%...{xx}C” style tokens into a single embedded token sequence within the %...{xx}C style. To condense multiple style tokens into a single embedded token sequence, you must specify multiple tokens within the { } braces and prefixing each token with the `%' symbol. For example,

```
%{rh}C %{rt}C %{as}C
```

can be re-expressed in a condensed embedded token format as the following:

```
%{ %rh %rt %as}C.
```

The command line syntax will accept single tokens represented as

```
%{ %rh}C
```

and

```
%{rh}C
```

as equivalents.

Any character that is not part of an embedded token sequence (for example, the space character) is repeated verbatim in the output file.

The above set of tokens allow you to configure an extended Squid-like format line within the W3C Customizable Logging format specification. For example:

```
“%{es}C.%{em}C %{te}C %a %{rd}C/%s %{cs}C %m %ru}C %u %{rh}C %rt}C %as}C”
```

The following is an example of a Extended Squid-like format that specifies that user-readable time-stamps are used instead of the Squid “seconds-since-epoch” time-stamp format, and that a configured out-going proxy (as specified by “%...{rH}C”) is logged:

```
“[%{ %d/%b/%Y:%H:%M:%S %z}t] %{te}C %a %{rd}C/%s %{cs}C %m %ru}C %u %rH}C  
%rt}C %as}C”
```

Unknown or unsupported translog tokens are logged within the log file as the characters that made up the token. For example, “%{xy}C” is logged into the log file as “xy.” All characters outside of a token specification sequence are repeated verbatim within the log file.

Transaction Logging and NTLM Authentication

If your device is configured for NT LAN Manager (NTLM) authentication and uses Apache-style or Extended Squid-style format, you can record the Windows domain name and username in the “authenticated username” field of the transaction log. If the domain name is available, both the domain name and the username are recorded in the “authenticated username” field, in the form domain\username. If only the username is available, only the username is recorded in the “authenticated username” field. If neither a domain name nor a username is available, a “-” (hyphen) is recorded in the field.

Usage Guidelines for Log Files

This section provides some guidelines for working with log files.

Understanding Working Logs

Depending upon where the sysfs is mounted, transactions are logged to a working log on the local disk in one of these files:

- /local1/logs/working.log
- /local2/logs/working.log

Depending upon where the sysfs is mounted, the following log files are logged to a working log on the local disk as follows:

- WMT logs are logged to a working log on the local disk in one of these files:
 - /local1/logs/export/working.log
 - /local2/logs/export/working.log
- RealSubscriber logs are logged to a working log on the local disk in one of these files:
 - /local1/logs/real-subscriber-logs/working.log
 - /local2/logs/real-subscriber-logs/working.log
- Cisco Streaming Engine logs are logged to a working log on the local disk in one of these files:
 - /local1/logs/cisco-streaming-engine/working.log
 - /local2/logs/cisco-streaming-engine/working.log
- RealProxy logs are logged to a working log on the local disk in one of these files:
 - /local1/logs/real-proxy/working.log
 - /local2/logs/real-proxy/working.log

Archive Working Log

You can specify the interval at which the working log should be cleared by moving the data to an archive log. The archive log files are located on the local disk in the directory /local1/logs/ or /local2/logs/, depending upon where the sysfs is mounted.

Because multiple archive files are saved, the filename includes the time stamp when the file was archived. Because the files can be exported to an FTP/SFTP server, the filename also contains the IP address of this Content Engine.

The archive file name use this format:

celog_IPADDRESS_YYYYMMDD_HHMMSS.txt.

Sanitizing Transaction Logs

You can disguise the IP address and usernames of clients in the transaction log file. The default is that transaction logs are not sanitized. A sanitized transaction log disguises the network identity of a client by changing the IP address in the transaction logs to 0.0.0.0.

Exporting Log Files

To facilitate the postprocessing of cache log files, you can export transaction logs to an external host. This feature allows log files to be automatically exported by FTP to an external host at configurable intervals. The username and password used for FTP are configurable, as is the directory to which the log files are uploaded.

The log files automatically have a filename that uses this format:

```
<type>_<ipaddr>_yyyymmdd_hhmmss.txt
```

where

- *<type>* represents the type of log file with *celog* for cache logs such as HTTP, HTTPS, and FTP, and *mms_export* for Windows Media Technologies (WMT) logs.
- *<ipaddr>* represents the Content Engine IP address.
- *yyyymmdd_hhmmss* represents the date and time when the log was archived for export.



Note

For WMT logs, there is no .txt extension in the filename.

Exporting Transaction Logs to External FTP Servers

To export transaction logs to an FTP server, you must first enable exporting of transaction logs and then configure the FTP or secure FTP (SFTP) server parameters. This feature can support up to four FTP servers. The following information is required for each target FTP server:

- Server IP address or the host name
The Content Engine translates the host name with a DNS lookup and then stores the IP address in the configuration.
- FTP user login and user password
- Path of the directory where transferred files are written
Use a fully qualified path or a relative path for the user login. The user must have write permission to the directory.

You can also compress archived log files into gzip format before exporting them to external FTP servers. The compressed filename has a .gz extension in the filename. This compression feature uses less disk space than that required for noncompressed archived files on both the Content Engine and the FTP export server and also requires less bandwidth during export because of the smaller size of the files to be exported.

Restarting Export After Receiving a Permanent Error from the External FTP Server

When an FTP server returns a permanent error to the Content Engine, the archive transaction logs are no longer exported to that server. You must reenter the Content Engine transaction log export parameters for the misconfigured server to clear the error condition.

A permanent error (Permanent Negative Completion Reply, RFC 959) occurs when the FTP command to the server cannot be accepted, and the action does not take place. Permanent errors can be caused by invalid user logins, invalid user passwords, and attempts to access directories with insufficient permissions or directories that do not exist.

Exporting Transaction Logs to External SFTP Servers

You can also export transaction logs to a Secure File Transfer Protocol (SFTP) server. You must first enable the feature and configure the SFTP server parameters. The following information is required for each target SFTP server:

- SFTP server IP address or the host name

The Content Engine translates the host name with a DNS lookup and then stores the IP address in the configuration.

- SFTP user login and user password
- Path of the directory where transferred files are written

Use a fully qualified path or a relative path for the user login. The user must have write permission to the directory.

Secure File Transfer Protocol Access for Nonadministrative Users

In the ACNS 5.3.5 software release, the SFTP server on the Content Engine was enhanced to allow nonadministrative users (that is, a user with a nonzero UID) to use SFTP to access the Content Engine. In the ACNS 5.3.5 software release, the **sshd allow-non-admin-users** and **no sshd allow-non-admin-users** global configuration commands were added to enable and disable this new feature. By default, this feature is disabled on the Content Engine, and nonadministrative users cannot use SFTP to access the Content Engine. To enable this feature, enter the **sshd allow-non-admin-users** command on the Content Engine. After enabling this feature, you can disable it again by entering the **no sshd allow-non-admin-users** command on the Content Engine.

If this feature is enabled, the output of the **show running-config EXEC** command shows that this feature is enabled on the Content Engine.


Enabling Transaction Logging with the Content Distribution Manager GUI

To enable transaction logging, follow these steps:

Step 1 From the Content Distribution Manager GUI, choose **Devices > Devices**.

Step 2 Click the **Edit** icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.

From the Contents pane, choose **Request Processing > Transaction Logs**. The Transaction Logs settings window appears. (See [Figure 19-1](#).) [Table 19-4](#) describes the fields in this window and provides the corresponding CLI global configuration commands.

Figure 19-1 Transaction Log Settings Window—General Settings

- Step 3** Under the General Settings heading, to activate transaction logging on your device, check the **Transaction Log Enable** check box.
- Step 4** To disguise the IP address and username of clients, check the **Enable Sanitize Transaction Logs** check box.
- Step 5** To add markers to the transaction logs that indicate where the files begin and end, check the **Enable File-Marker** check box.
- Step 6** To record the Windows domain name and username in the “authenticated username” field of the transaction log, check the **Log Windows Domain** check box.



Note This option is operational if your device is configured for NT LAN Manager (NTLM) authentication and uses either the Apache-style or the Extended Squid-style transaction log format.

- Step 7** To enable compression of archived log files into gzip format before exporting them to external FTP servers, check the **Compress Files before Export** check box.
- Step 8** To choose a log file format, click the **Log File Format** radio button, and choose a log file format from the drop-down list. Choose **apache**, **extended-squid**, or **squid**.

Alternatively, if you want to use a custom format for the transaction log, click the **Log Format Custom** radio button, and then enter a custom format string in the field provided. (See the [“Custom Format Transaction Logging”](#) section on page 19-2.)

- Step 9** Under the Archive Settings heading (see [Figure 19-2](#)), in the Max Size of Archive File field, specify a value for the maximum size of the archive file in kilobytes. [Table 19-4](#) describes the fields in this window and provides the corresponding CLI global configuration commands.

This value is the maximum size of the archived file to be maintained on the local disk.

Figure 19-2 Transaction Log Settings Window—Archive Settings

The screenshot shows the Cisco Application and Content Networking System GUI. The 'Transaction Log Settings for Content Engine, CONTENTENGINE' window is open. The 'Archive Settings' section is expanded, showing the following fields:

- Max size of Archive File: 2000000 (KB) (1000-2000000)
- Archive occurs:
 - ☐ every (seconds) 120-604800
 - ☐ every hour ☐ at (minutes after the hour) 0-59
 - ☐ every 2 (minutes)
 - ☒ every day ☐ at (hh:mm) 0:0-23:59
 - ☒ every 1 (hours)
 - ☐ every week on ☐ Sun ☐ Mon ☐ Tue ☐ Wed ☐ Thu ☐ Fri ☐ Sat
 - at: (hh:mm) 0:0-23:59
- Export Settings:
 - Enable Export: ☒
 - Export occurs:
 - ☐ every (minutes) 1-10080
 - ☐ every hour ☐ at (minutes after the hour) 0-59

At the bottom right, there are 'Submit' and 'Cancel' buttons.

- Step 10** To schedule the interval when the working log should be cleared by moving data into the archive log, choose a time option from among the options given in the **Archive occurs** section.
- Step 11** To save the settings, click **Submit**.

Table 19-4 Transaction Log General and Archive Settings

GUI Parameter	Function	CLI Command
General Settings		
Transaction Log Enable	Enables transaction logging on the Content Engine.	transaction-logs enable
Enable Sanitize Transaction Logs	Disguises the IP address and username of clients.	transaction-logs sanitize
Enable File-Marker	Adds markers to the transaction logs to indicate where files begin and end.	transaction-logs file-marker
Log Windows Domain	Records the Windows domain name and username in the “authenticated username” field of the transaction log.	transaction-logs log-windows-domain

Table 19-4 Transaction Log General and Archive Settings (continued)

GUI Parameter	Function	CLI Command
Compress Files before Export	Enables compression of archived log files into gzip format before exporting them to external FTP servers.	transaction-logs export compress
Log File Format	Configures the log file format (apache , extended-squid , or squid).	transaction-logs format { squid extended-squid apache }
Log Format Custom	Configures a custom log file format.	transaction-logs format custom <i>string</i>
Archive Settings		
Max Size of Archive File	Maximum size (in kilobytes) of the archive file to be maintained on the local disk.	transaction-logs archive max-file-size <i>kilobytes</i>
Archive occurs every (interval)	Interval for the working log to be cleared by moving data into the archive log.	transaction-logs archive interval { <i>seconds</i> every-week [on weekdays at <i>hour:minute</i>] every-day { at <i>hour:minute</i> every hours } every-hour { at <i>minute</i> every minutes }}

If you want to enable exporting to an FTP server, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **Request Processing > Transaction Logs**. The Transaction Logs settings window appears.
- Step 4** Under the Export Settings heading, check the **Enable Export** check box. (See [Figure 19-3](#).) [Table 19-5](#) describes the fields in this window and provides the corresponding CLI global configuration commands.

Figure 19-3 Transaction Log Settings Window—Export Settings

The screenshot shows the 'Transaction Log Settings for Content Engine, CONTENTENGINE' window. The 'Export Settings' tab is selected. The 'Enable Export' checkbox is checked. Under 'Export occurs', the 'every day' option is selected with a frequency of 1 hour. Below this, there is a table for 'Export Server' with columns for Name, Password, Confirm password, and Directory. The table has four empty rows for configuration. At the bottom right, there are 'Submit' and 'Cancel' buttons.

Step 5 Using the time options shown, specify the interval when the working log should be cleared by moving data into the FTP server.

Step 6 Enter an IP address or host name information for the FTP server in the Export Server field.



Note The FTP export feature can support up to four servers. Each server must be configured with a username, password, and directory that are valid for that server.

Step 7 In the Name field, enter a user ID.

Step 8 In the Password and Confirm Password fields, enter a password for the user specified in [Step 7](#).

Step 9 In the Directory field, enter the name of a working directory that will contain the transaction logs.



Note The user specified in the Name field must have write permission to the specified directory.

Step 10 If the server chosen is a secure FTP server, check the **SFTP** check box.

Step 11 To save the settings, click **Submit**.

Table 19-5 Transaction Log Export Settings

GUI Parameter	Function	CLI Command
Enable Export	Enables transaction log export to an external FTP server.	transaction-logs export enable
Export occurs (interval)	Interval for the working log to be cleared by moving data to the FTP or SFTP server.	transaction-logs export interval {minutes every-week [on weekdays at hour:minute] every-day {at hour:minute every hours } every-hour {at minute every minutes }}
Export Server	IP address or hostname of the FTP server.	transaction-logs export ftp-server {hostname servipaddrs} login passw directory
Name	Name of the user.	
Password	Password for the user.	
Confirm Password	Confirms the password for the user.	
Directory	Name of a working directory that will contain the transaction logs.	transaction-logs export sftp-server {hostname servipaddrs} login passw directory
SFTP	Configures a secure FTP server.	

Using WMT Transaction Logging

For certain companies, streaming media is a source of revenue and needs to be tracked closely. Because these companies charge their customers to stream on-demand content and live broadcasts, they must rely on logged information to track what content a particular customer viewed, how long they viewed it, and what their viewing quality was. Consequently, the accuracy and reliability of transaction logging is very important to these companies.

The Windows Media Services 9 Series provides a more robust logging model than Windows Media Services version 4.1. ACNS 5.5 software supports Windows Media Services 9 logging.



Note

In ACNS software (release 5.1 and earlier), only the standard Windows Media Services version 4.1 and the extended Windows Media Services version 4.1 logging formats were supported.

In ACNS 5.5 software, the following logging formats are supported for WMT transaction logging:

- Standard Windows Media Services version 4.1
- Extended Windows Media Services version 4.1
- Standard Windows Media Services version 9.0
- Extended Windows Media Services version 9.0

The extended versions of the logging formats contain additional fields that are Content-Engine specific (For example, the CE-action field specifies a cache hit or miss, and the CE-bytes field specifies the number of bytes that were sent from the Content Engine.)

The Content Engine's transaction logging format for WMT streaming is consistent with that of the Windows Media Services and the World Wide Web Consortium (W3C)-compliant log format. A log line is written for every stream accessed by the client. The location of the log is not configurable. These logs can be exported using FTP. When transaction logging is enabled, daemons create a separate *working.log* file in /local1/logs/export for WMT transactions.

All client information in the transaction logs is sent to the origin server by default.

Log Formats Accepted by Windows Media Services 9

Windows Media Players connect to a Windows Media server using the following protocols:

- Windows Media Players earlier than Version 9.0 use HTTP/1.0 or the MMS protocol.
- Windows Media Player Version 9.0 uses HTTP/1.1 and RTSP.

Depending on the version of the Windows Media Player, logs are sent in different formats, such as text, binary, or Extensible Markup Language (XML). [Table 19-6](#) describes the log formats accepted by Windows Media Services 9.

Table 19-6 Windows Media Services 9 Log Formats

Protocol	Player and Distributor	Log Type
HTTP/1.0	Windows Media Player earlier than Version 9.0 Content Engine (caching and proxy server) is running Windows Media Services Version 9.0 and streaming from a WMT server that is running Windows Media Services 4.1	World Wide Web Consortium (W3C) standard space-delimited text log
MMS	Windows Media Player earlier than Version 9.0	Binary structure log
HTTP/1.1	Windows Media Player Version 9.0 Distribution server is running Windows Media Services 9.0 Content Engine (caching and proxy server) is running Windows Media Services 9.0	XML structure log
RTSP	Windows Media Player Version 9.0 Distribution server is running Windows Media Services 9.0 Content Engine (caching and proxy server) is running Windows Media Services 9.0	XML structure log

**Note**

ACNS 5.5 software supports Extensible Markup Language (XML) logging for MMS-over-HTTP and MMS-over-RTSP (RTSP over Windows Media Services 9). The posted XML log file from the Windows Media Player to the Content Engine (Windows Media server) can be parsed and saved to the normal WMT transaction logs that are stored on the Content Engine.

Configuring WMT Transaction Logging

To configure WMT transaction logging, follow these steps:

- Step 1** Navigate to the Transaction Logs Settings for Content Engine window.
- From the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Click the **Edit** icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.
 - From the Contents pane, choose **Request Processing > Transaction Logs**. The Transaction Log Settings window appears. (See [Figure 19-1](#).)
- Step 2** Under the WMT Settings heading (see [Figure 19-4](#)), to generate transaction logs for WMT streaming sessions, check the **Enable WMT Settings** check box.

Figure 19-4 Transaction Log Settings Window—WMT and Logging Settings

The screenshot shows the Cisco Application and Content Networking System interface in a Microsoft Internet Explorer browser. The main title bar reads "Cisco Application and Content Networking System - Microsoft Internet Explorer". The interface includes a top navigation bar with "System Status" and "Devices" tabs. The "Devices" tab is active, showing a list of devices and a "Transaction Log Settings for Content Engine, CONTENTENGINE" window. The "Transaction Log Settings" window has a left sidebar with a "Contents" pane listing various settings like "Device Home", "Device Activation", "Assignments", "Prepositioning", "Request Routing", "Request Processing", "URL Filter", "Websense Server", "Enable Rules", "Service Rules", "ICAP", "ICAP Services", "Transaction Logs", "Applications", "General Settings", and "Device Monitoring". The "Transaction Logs" option is selected. The main content area of the "Transaction Log Settings" window is divided into two sections: "Windows Media Settings" and "Logging Settings". In the "Windows Media Settings" section, "Enable Windows Media Settings" is checked, and "Log File Format" is set to "wms-41". In the "Logging Settings" section, "Enable" is unchecked, "Facility" is set to "Do Not Set", "Entry Type" is set to "request-auth-failures", "Enable Host Settings" is unchecked, "Hostname" is a required field (indicated by a red asterisk), "Port" is set to "514", and "Rate Limit" is set to "0" (with a range of "0-10000"). A note at the bottom states "Note: * - Required Field". The window has "Submit" and "Cancel" buttons at the bottom right.

- Step 3** From the **Log File Format** drop-down list, choose the logging format for WMT transaction logs. [Table 19-7](#) describes the drop-down list options.

Table 19-7 WMT Log File Format Options

Log Format	Description
extended	Specifies the WMT extended configurations for transaction logs. Enables username logging in the WMT transaction log.
wms-41	Sets WMT to generate transaction logs in extended Windows Media Services 4.1 format. When you use this option, the Content Engine uses the standard Windows Media Services 4.1 format to generate the transaction log and includes the following three additional fields in the transaction log: • CE_action (cache hit or cache miss) • CE-bytes (number of bytes sent from the Content Engine for a cache hit) • username (username of the WMT request when NTLM, Negotiate, Digest, and basic authentication is used)
wms-90	Sets WMT to generate transaction logs in extended Windows Media Services 9 format. When you use this option, the Content Engine uses the standard Windows Media Services 9 format to generate the transaction log and includes the following three additional fields in the transaction log: • CE_action (cache hit or cache miss) • CE-bytes (number of bytes sent from the Content Engine for a cache hit) • username (username of the WMT request when NTLM, Negotiate, Digest, and basic authentication is used)
wms-41	Sets WMT to generate transaction logs in the standard Windows Media Services 4.1 format.
wms-90	Sets WMT to generate transaction logs in the standard Windows Media Services 9 format.

To configure WMT transaction logging from the CLI, use the following global configuration command:

```
wmt transaction-logs format {extended {wms-41 | wms-90} | wms-41 | wms-90}
```

Using Real-Time Transaction Logging

You can monitor transaction logs in real-time for particular errors such as authentication errors. By sending HTTP transaction log messages to a remote syslog server, you can monitor the remote syslog server for HTTP request authentication failures in real-time. This real-time transaction log feature allows you to monitor transaction logs in real-time for particular errors such as HTTP request authentication errors. The existing transaction logging to the local file system remains unchanged.

For this purpose, you must configure the Content Engine to send transaction log messages to a remote syslog server using UDP as the transport protocol. Because UDP is an unreliable transport protocol, message transport to remote syslog host is not reliable and you must monitor the syslog messages received at the remote syslog server. You can limit the rate at which the transaction logging module is allowed to send messages to the remote syslog server. The format of the syslog message is in standard syslog message format with the transaction log message as the payload of the syslog message.

Real-time transaction logging to a remote syslog server uses the standard syslog message format with the message payload as the transaction log entry. A new syslog error identifier is defined for this type of real-time transaction log message. You can configure the Content Engine to send transaction log messages in real-time to one remote syslog host. The message format of the transaction log entry to the remote syslog host is the same as in the transaction log file and prepended with Cisco's standard syslog header information.

The following is an example of the format of the real-time syslog message sent from the transaction logging module (Content Engine) to the remote syslog host:

```
fac-pri Apr 22 20:10:46 ce-host cache: %CE-TRNSLG-6-460012: translog formatted msg
```

The fields in the message are described as follows:

- *fac-pri* denotes the facility parameter and priority for transaction log messages encoded (as in standard syslog format) as a 32-bit decimal value between 0 and 1023 (0x0000 and 0x03FF). The least significant 3 bits denote priority (0-7) and the next least significant 7 bits denote facility (0-127).

The facility parameter used by the transaction logging module when a real-time transaction log message is logged to the remote syslog host is "user". The same facility is sent to the remote syslog host unless you configure a different facility parameter for transaction logging. The priority field is always set to LOG_INFO for real-time transaction log messages.

In the above example, the default value of *fac-pri* is 14 (0x000E) where facility = user (LOG_USER (1)) and priority = LOG_INFO (6).

- The next field in the message is the date, which follows the format as shown in the above example.
- *ce-host* is the hostname or IP of the Content Engine that is sending the message.
- *cache* is the name of the process on the Content Engine that is sending the message.
- %CE-TRNSLG-6-460012 is the Cisco standard formatted syslog header on the Content Engine for a real-time transaction log message. This identifier indicates a priority level of 6, which denotes informational messages.

**Note**

The Content Engine system syslog messages report communication errors with the remote syslog host that is configured for transaction logging. These syslog messages are in the error message range: %CE-TRNSLG-6-460013 to %CE-TRNSLG-3-460016. Note that the last error message (%CE-TRNSLG-3-460016), shows level “3” (for error-level messages) instead of “6” (for information-level messages). Information-level messages are reported when messages are dropped due to rate limiting and the number of dropped messages are reported. For more information about these syslog messages, refer to the ACNS Syslog Error Book.

- *translog formatted msg* is the transaction log message as it appears in the transaction log file.

**Note**

The total length of the real-time syslog message is 1024 characters. If the actual transaction log entry exceeds this limit, it is truncated.

To include the username and domain name in the transaction log, check the **Log Windows Domain** check box or use the **transaction-logs log-windows-domain** global configuration command.

When the remote syslog server logs this message to a file, the format appears as follows:

```
Apr 22 20:10:46 ce-host cache: %CE-TRNSLG-6-460012: translog formatted msg
```

where ce-host is the host name of the Content Engine that sent the real-time transaction log message to the remote syslog server.

The configuration of host settings for transaction logs is identical to the configuration settings for syslog messages except that you need not specify the priority level of the message for real time transaction logs. All messages are associated with the priority level of 6 (LOG_INFO). You are not required to filter messages based on priority levels.

Configuring Real-Time Transaction Logging

To configure real-time transaction logging, follow these steps:

-
- Step 1** Navigate to the Transaction Logs Settings for Content Engine window.
- From the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Click the **Edit** icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.
 - From the Contents pane, choose **Request Processing > Transaction Logs**. The Transaction Logs settings window appears. (See [Figure 19-1](#).)
- Step 2** Under the Logging Settings section (see [Figure 19-4](#)), to enable real-time transaction logging, check the **Enable** check box. You can retain the logging host configuration for transaction logs even if you temporarily disable real-time transaction logging by unchecking the check box. This new logging option applies only to the cache's HTTP transaction log entries. The real-time transaction logging feature is disabled by default.

**Note**

You must check the **Transaction Log Enable** check box in the General settings section before you check the **Enable** check box. Otherwise, the second setting does not apply unless you enable transaction logging for the entire Content Engine.

- Step 3** From the Facility drop-down list, choose the appropriate transaction log facility.
- This drop-down list is set to an initial value of *Do not set*. This setting denotes that the facility sent to the syslog host will be the facility on the local host that is sending the syslog message. For instance, in the case of the transaction logging module that sends the real-time transaction log message, the facility is the “user” facility.
- Step 4** From the Entry Type drop-down list, choose the type of transactions that you want to log from the Content Engine to the remote syslog host. Choose **request-auth-failures** to send only those transactions associated with HTTP request authentication failures to the remote syslog host. Choose **all** to send all of the transaction messages to the remote syslog host. See the [“About Specifying the Transaction Log Entry Type when Logging to a Remote Syslog Host”](#) section on page 19-21.
- Step 5** To enable sending of transaction log files to a remote syslog host, check the **Enable Host Settings** check box.
- Step 6** In the Hostname field, enter a host name or an IP address of the remote syslog server to which transaction logs must be sent. No remote syslog server is specified by default.
- Step 7** In the Port field, specify the destination port on the remote syslog host to which the Content Engine should send the message. The default port number is 514. This port is a well-known port for system logging.
- Step 8** In the Rate Limit field, specify the number of messages that are allowed to be sent to the remote syslog host per second. To limit bandwidth and other resource consumption, messages to the remote syslog host can be rate limited. If this limit is exceeded, the specified remote syslog host drops the messages. There is no default rate limit (rate-limit is set to 0), and by default all syslog messages are sent to all of the configured syslog hosts. The range is 1 to 10,000 messages per second.
- Step 9** To save the settings, click **Submit**.
- A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking **Reset**. The **Reset** button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.
- If you try to leave this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box only appears if you are using the Internet Explorer browser.

About Specifying the Transaction Log Entry Type when Logging to a Remote Syslog Host

You can configure the Content Engine to send only the transactions associated with HTTP request authentication failures, or to send all of the transactions.

Typically, organizations are interested in only HTTP request authentication failures for security purposes. By monitoring these types of authentication failures in real time, it enables organizations to identify which end users failed to be authenticated through the Content Engine.

Only the authentication failure transaction that is associated with an end user who is attempting to contact the authentication server is logged. The “pending” transactions that are waiting for a response from the transaction that contacted the authentication servers are not logged. This approach provides you with the information needed to determine which user fails to authenticate with the Content Engine and minimizes the traffic to the syslog host. To track which users failed to authenticate, you must configure a transaction log format that logs the username by configuring either Extended Squid-style format or the custom log format with the custom format token %u . For more information about specifying the format of the transaction log, see the [“Understanding Transaction Log Formats”](#) section on page 19-1 and the [“Custom Format Transaction Logging”](#) section on page 19-2.

When you check the **Enable** check box to enable real-time transaction logging (or specify the **transaction-logs logging enable** global configuration command), the logging of only HTTP request authentication failures is the default. If you want to change this default and log all transactions, then you must choose **all** from the Entry Type drop-down list (or enter the **transaction-log logging entry-type all** global configuration command on the Content Engine). However, if you log all transactions, there may be a significant UDP drop rate if your syslog host cannot handle the rate of incoming traffic.



PART 6

System Administration and Configuration





Configuring Platform and System Settings

This chapter explains how to configure platform and system settings. It contains the following sections:

- [Configuring Remote Login Access Control Settings, page 20-1](#)
- [Configuring Network Services, page 20-7](#)
- [Configuring Date and Time Settings, page 20-11](#)
- [Configuring System-Wide Distribution QoS Settings Based on Channel Priority, page 20-16](#)
- [Modifying System Default Properties, page 20-18](#)
- [Configuring Faster Detection of Offline Content Engines, page 20-21](#)
- [Printing ACNS Network Data, page 20-23](#)
- [Creating Proxy Error Messages, page 20-23](#)
- [Enabling Offline Operation of Network Devices, page 20-27](#)

Configuring Remote Login Access Control Settings

This section explains how to configure connection settings for various means of remote login and access to the system software.

Enabling the Telnet Service on the Content Engine

To enable the Telnet service on the Content Engine, follow these steps:

-
- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
 - Step 2** Click the **Edit** icon next to the Content Engine for which you want to enable Telnet. The Device Home for Content Engine window appears.
 - Step 3** In the Contents pane, choose **General Settings > Login Access Control > Telnet**. The Telnet Settings for Content Engine window appears.
 - Step 4** To enable the terminal emulation protocol for remote terminal connection, check the **Telnet Enable**.

**Note**

You must use a console connection instead of a Telnet session to define device network settings on the Content Engine. However, after you have used a console connection to define the device network settings for this Content Engine, you can use a Telnet session to perform subsequent configuration tasks (for example, configure DNS caching).

- Step 5** To save the settings, click **Submit**. A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking the Reset button. The Reset button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

If you try to leave this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box only appears if you are using the Internet Explorer browser.

To enable the Telnet service on the Content Engine from the CLI, use the **telnet enable** global configuration command.

Configuring Secure Shell Settings


Secure Shell (SSH) consists of a server and a client program. Like Telnet, you can use the client program to remotely log in to a machine that is running the SSH server, but unlike Telnet, messages transported between the client and the server are encrypted. The functionality of SSH includes user authentication, message encryption, and message authentication.

The SSH management window in the Content Distribution Manager GUI allows you to specify the key length, login grace time, and maximum number of password guesses allowed when logging in.

To enable the SSH daemon on the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to enable SSH. The Device Home for Content Engine window appears.
- Step 3** In the Contents pane, choose **General Settings > Login Access Control > SSH**. The SSH Configuration window appears. (See [Figure 20-1](#).) [Table 20-1](#) describes the fields in this window and provides the corresponding CLI global configuration commands.

Figure 20-1 SSH Configuration Window



- Step 4** To enable the SSH feature, check the **Enable** check box. SSH enables login access to the Content Engine through a secure and encrypted channel.
- Step 5** In the Length of key field, specify the number of bits needed to create an SSH key. The default is 1024. When you enable SSH, be sure to generate both a private and a public host key, which client programs use to verify the server's identity. (See the [“About the Secure Shell Host Key”](#) section on page 20-4.)
- Step 6** In the Login grace time field, specify the number of seconds for which an SSH session will be active during the negotiation (authentication) phase between client and server before it times out.
- Step 7** In the Maximum number of password guesses field, specify the maximum number of incorrect password guesses allowed per connection. The default is 3. (See the [“About Secure Shell Daemon Password Guesses”](#) section on page 20-4.)
- Step 8** To allow clients to connect using SSH protocol version 1, check the Enable SSHv1 check box.
- Step 9** To allow clients to connect using SSH protocol version 2, check the Enable SSHv2 check box.



Note You can enable both SSHv1 and SSHv2, or you can enable one version and not the other. You cannot disable both versions of SSH unless you disable the SSH feature by unchecking the **Enable** check box. (See [Step 4](#).)

- Step 10** To save the settings, click **Submit**. A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking the Reset button. The Reset button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

If you try to leave this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box only appears if you are using the Internet Explorer browser.

Table 20-1 SSH Settings

GUI Parameter	Function	CLI Command
Enable	Enables the SSH feature.	sshd enable
Length of key	Number of bits needed to create an SSH key.	ssh-key-generate key-length <i>length</i>
Login grace time	Number of seconds for which an SSH session will be active during the negotiation (authentication) phase between client and server before it times out.	sshd timeout <i>seconds</i>
Maximum number of password guesses	Maximum number of incorrect password guesses allowed per connection.	sshd password-guesses <i>num</i>
Enable SSHv1	Allows clients to connect using SSH protocol version 1.	sshd version 1
Enable SSHv2	Allows clients to connect using SSH protocol version 2.	sshd version 2

About the Secure Shell Host Key

When you use an SSH client and log in to the Content Engine, the public key for the SSH daemon running on the Content Engine is recorded in the client machine `known_hosts` file in your home directory. If the Content Engine administrator subsequently regenerates the host key by specifying the number of bits in the Length of key field, you must delete the old public key entry associated with the Content Engine in the `known_hosts` file before running the SSH client program to log in to the Content Engine. When you use the SSH client program after deleting the old entry, the `known_hosts` file is updated with the new SSH public key for the Content Engine.

About Secure Shell Daemon Password Guesses

Although the value in the Maximum number of password guesses field specifies the number of allowed password guesses from the SSH server side, the actual number of password guesses for an SSH login session is determined by the combined number of allowed password guesses of the SSH server and the SSH client. Some SSH clients limit the maximum number of allowed password guesses to three (or to one in some cases), even though the SSH server allows more than this number of guesses.

When you specify n allowed password guesses, certain SSH clients interpret this number as $n + 1$. For example, when configuring the number of guesses to two for a particular device, SSH sessions from some SSH clients will allow three password guesses.

Configuring Message of the Day Settings for the Content Engine

The Message of the Day (MOTD) feature enables you to provide information bits to the users when they log in to a Content Engine that is part of your ACNS network. There are three types of messages that you can set up:

- MOTD Banner
- EXEC Process Creation Banner
- Login Banner



Note

When you run an SSH version 1 client and log in to the device, the MOTD and login banners are not displayed. You need to use SSH version 2 to display the banners when you log in to the device.

To configure the banner settings, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
 - Step 2** Click the **Edit** icon next to the Content Engine for which you want to set up the message of the day. The Device home window for the chosen Content Engine appears.
 - Step 3** In the Contents pane, choose **General Settings > Login Access Control > Message of the Day** from the Contents Pane. The MOTD Configuration window for the chosen Content Engine appears.
 - Step 4** To enable the MOTD settings, check the **Enable** check box. The Message of the Day (MOTD) banner, EXEC process creation banner, and Login banner fields become enabled.
 - Step 5** In the Message of the Day (MOTD) Banner field, enter a string that you want to display as the MOTD banner when a user attempts to login to the Content Distribution Manager.
- 

Note In the Message of the Day (MOTD) Banner, EXEC Process Creation Banner, and Login Banner fields, you can enter a maximum of 1024 characters. A new line character (or Enter) is counted as two characters, as it is interpreted as \n by the system. You cannot use special characters such as ` , % , ^ , and " in the MOTD text. If your text contains any of these special characters, ACNS software removes it from the MOTD output.
- Step 6** In the EXEC Process Creation Banner field, enter a string to be displayed as the EXEC process creation banner when a user enters into the EXEC shell of the Content Engine.
 - Step 7** In the Login Banner field, enter a string to be displayed after the MOTD banner, when a user attempts to login to the Content Engine.
 - Step 8** To save the configuration, click **Submit**.

To configure banner messages from the CLI, use the **banner** global configuration command.

```
banner {motd | login | exec} message message_text
```

```
banner enable
```

Configuring Exec Timeout Settings

To configure the length of time that an inactive Telnet session remains open on the Content Engine, follow these steps:

-
- Step 1** From the Content Distribution Manger GUI, choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
 - Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure the exec timeout. The Device Home for Content Engine window appears.
 - Step 3** In the Contents pane, choose **General Settings > Login Access Control > Exec Timeout**. The Exec Timeout Settings for Content Engine window appears.
 - Step 4** In the Exec Timeout field, specify the number of minutes after which an active session times out. The default is 15 minutes.

A Telnet session with the Content Engine can remain open and inactive for the period specified in this field. When the exec timeout period elapses, the Content Engine automatically closes the Telnet session.

- Step 5** To save the settings, click **Submit**. A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking the Reset button. The Reset button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

If you try to leave this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box only appears if you are using the Internet Explorer browser.

To configure the length of time that an inactive Telnet session remains open on the Content Engine from the CLI, use the **exec-timeout** *minutes* global configuration command.

Configuring Line Console Carrier Detection

Carrier detection should be enabled if the Content Engine, Content Router, or Content Distribution Manager is to be connected to a modem for receiving calls. If you are using a null-modem cable with no carrier detect pin, the device might appear unresponsive on the console until the carrier detect signal is asserted. To recover from a misconfiguration, the device should be rebooted and the 0x2000 bootflag should be set to ignore the Carrier Detect (CD) setting.

This feature is disabled by default in the Content Distribution Manager GUI.

To enable console line carrier detection using the Content Distribution Manager, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices** (or **Devices > Device Groups**).
 - Step 2** Click the **Edit** icon next to the device (or device group) that you want to configure.
 - Step 3** In the Contents pane, choose **General Settings > Login Access Control > Console Carrier Detect**. The Console Carrier Detect Settings window appears.
 - Step 4** To enable the window for configuration, check the **Enable console line carrier detection before writing to the console** check box.

Step 5 To save this setting, click **Submit**.

To configure line console carrier detection from the CLI, use the **line console carrier-detect** global configuration command. To disable console line carrier detection, use the **no** form of the command.

Configuring Network Services

This section explains how to configure settings for various network services.

Configuring CDP Settings

CDP (Cisco Discovery Protocol) is a device discovery protocol that runs on all Cisco-manufactured devices. With CDP, each device in a network sends periodic messages to all devices in the network. All devices listen to periodic messages that are sent by others to learn about neighboring devices and determine the status of their interfaces.


With CDP, network management applications can learn the device type and the Simple Network Management Protocol (SNMP) agent address of neighboring devices. Applications are then able to send SNMP queries within the network. Also, CiscoWorks2000 discovers the Content Engine by means of the CDP packets that are sent by the Content Engine after booting.

Content Engine-related tasks require that the Content Engine platform support CDP to be able to notify the system manager of the existence, type, and version of the Content Engine platform.

To configure CDP settings using the Content Distribution Manager GUI, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.
 - Step 3** From the Contents pane, choose **General Settings > Services > CDP**. The CDP Settings window appears. (See [Figure 20-2](#).)

Figure 20-2 CDP Settings Window



- Step 4** To enable CDP support, check the **Enable** check box.
- Step 5** In the Hold Time field, enter the time (in seconds) to specify the length of time that a receiver is to keep the CDP packets. The range is 10 to 255 seconds. The default is 180 seconds.
- Step 6** In the Packet Send Rate field, enter a value (in seconds) for the interval between CDP advertisements. The range is 5 to 254 seconds. The default is 60 seconds.
- Step 7** To save the settings, click **Submit**.

To configure CDP in the CLI, use the following global configuration command:

```
cdp {enable | holdtime seconds | timer seconds}
```

To enable CDP on a particular interface, use the following global configuration command:

```
interface interface slot/port cdp enable
```

Enabling RCP Services on the Content Engine

Remote Copy Protocol (RCP) lets you download, upload, and copy configuration files between remote hosts and a switch. Unlike TFTP, which uses User Datagram Protocol (UDP), a connectionless protocol, RCP uses TCP, which is connection oriented.

RCP is a subset of the UNIX rshell service, which allows UNIX users to execute shell commands on remote UNIX systems. It is a UNIX built-in service. This service uses TCP as the transport protocol and listens for requests on TCP port 514. RCP service can be enabled on Content Engines that use ACNS

software. Inetd (an Internet daemon) is a program that listens for connection requests or messages for certain ports and starts server programs to perform the services associated with those ports. RCP copies files between devices.

To enable RCP services on the Content Engine, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to enable RCP services. The Device Home for Content Engine window appears.
- Step 3** In the Contents pane, choose **General Settings > Services > Inetd RCP**. The Inetd RCP Settings for Content Engine window appears.
- Step 4** To enable the RCP service on the Content Engine, check the **Inetd Rcp Enable** check box.



Note The Inetd daemon listens for FTP, RCP, and TFTP services. For Inetd to listen to RCP requests, it must be explicitly enabled for RCP service.

- Step 5** To save the settings, click **Submit**. A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking the Reset button. The Reset button is visible only when you have applied default or group settings to change the current device settings but you have not yet submitted the changes.

If you try to leave this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box only appears if you are using the Internet Explorer browser.

To enable RCP services on the Content Engine from the CLI, use the **inetd enable rcp** global configuration command.

Configuring the Content Router for Interoperability with a Layer 4 Switch

When a request for content is made, the Layer 4 Cisco Content Services Switch uses owner and content rules to translate the virtual IP address of the owner to the IP address of the service where the content resides. Depending on the content rules configured, the Content Services Switch checks for services that match the content rules and determines the service that can best serve the content request.



Note

A Cisco Content Services Switch enables you to configure owners and content rules to direct requests for content to a specific destination service (for example, a server or a port on a server). By configuring services, owners, and content rules, you optimize and control how the Content Services Switch handles each request for specific content.

- A *service* is a destination location where a piece of content physically resides (such as, a local or remote server and port).
- An *owner* is generally the person or company who contracts the web hosting service to host their web content and allocate bandwidth as required.


- A *content rule* is a hierarchical rule set containing individual rules that describe which content (for example, .html files) is accessible by visitors to the web site, how the content is mirrored, on which server the content resides, and how the Content Services Switch should process requests for the content. Each rule set must have an owner.

If you have a Content Services Switch deployed in your ACNS network for load balancing, this feature of redirection of requests with the origin server name allows you to access the content using Content Router redirection without advertising a separate Content Router Fully Qualified Domain Name (FQDN).

To configure and enable Layer 4 switch interoperability, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the name of the Content Router that you want to view. The Modifying Content Router window appears.
- Step 3** In the Contents pane, choose **General Settings > Network > L4 Switch**. The Layer 4 Switch Settings for Content Router window appears. (See Figure 20-3.)

Figure 20-3 Layer 4 Switch Settings for Content Router Window



- Step 4** To enable Layer 4 switch redirection interoperability using HTTP, check the **Enable HTTP I4-switch** check box.
- Step 5** To enable Layer 4 switch redirection interoperability using RTSP, check the **Enable RTSP I4-switch** check box.

Step 6 To save the configured settings, click **Submit**.

Configuring Date and Time Settings

This section explains how to configure date and time settings for your ACNS network devices.


Configuring NTP Settings

Cisco ACNS 5.x software allows you to configure the Content Engine time and date settings using an NTP (Network Time Protocol) host on your network. NTP allows the synchronization of time and date settings for the different geographical locations of the Content Engines on your ACNS network.

To configure Content Engine or device group NTP settings, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to configure. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **General Settings > Services > Date/Time > NTP**. (See [Figure 20-4](#).)

Figure 20-4 NTP Settings Window



- Step 4** To enable NTP settings, check the **Enable** check box.
- Step 5** In the NTP Server field, enter a host name or IP address.
- Step 6** To save the settings, click **Submit**.

Configuring Device Clock and Time Zone Settings

If you have an outside source on your network that provides time services (such as a Network Time Protocol [NTP] server), you do not need to set the system clock manually. When manually setting the clock, enter the local time. The Content Engine calculates Coordinated Universal Time (UTC) based on the time zone set by the **clock timezone** global configuration command. Customized summer time settings configured from the Content Distribution Manager GUI correspond to the **clock summertime** global configuration command.



Note

Two clocks exist in the system: the software clock and the hardware clock. The software uses the software clock. The hardware clock is used only at startup to initialize the software clock. The **clock set EXEC** command sets the software clock.

To configure the time zone for proper system configuration and to set summer time (daylight savings time) on the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure the time zone. The Device Home for Content Engine window appears.
- Step 3** In the Contents pane, choose **General Settings > Services > Date/Time > Time Zone**. The Time Zone Settings for Content Engine window appears.
- Step 4** To configure a standard time zone on the Content Engine, follow these steps:
- Under the Time Zone Settings section, click the **Standard Time Zone** radio button. The default is UTC (offset = 0) with no summer time configured. When you configure a standard time zone, the system is automatically adjusted for the UTC offset, and the UTC offset need not be specified.

The standard convention for time zones uses a *Location/Area* format in which *Location* is a continent or a geographic region of the world and *Area* is a time zone region within that location. For a list of standard time zones that can be configured and their UTC offsets, see [Table 20-2](#).
 - From the drop-down list, choose a location for the time zone. The window refreshes, displaying all area time zones for the chosen location in the second drop-down list.
 - Choose an area for the time zone. The UTC offset (hours and minutes ahead or behind UTC) for the corresponding time zone is displayed next to the second drop-down list. During summer time savings, the offset may be different and will be displayed accordingly.



Note

Some of the standard time zones (mostly time zones within the United States) have daylight savings time zones configured automatically.

- Step 5** To configure a customized time zone on the Content Engine, follow these steps:
- Under the Time Zone Settings section, click the **Customized Time Zone** radio button.
 - In the Customized Time Zone field, specify the name of the time zone. The time zone entry is case-sensitive and can contain up to 40 characters including spaces. If you specify any of the standard time zone names, an error message is displayed when you click **Submit**.
 - For UTC Offset, choose the + or – sign from the first drop-down list to specify whether the configured time zone is ahead or behind UTC. Also, choose the number of hours (0–23) and minutes (0–59) offset from UTC for the customized time zone. The range for the UTC offset is from –23:59 to 23:59, and the default is 0:0.

- Step 6** To configure customized summer time, follow these steps under the Customized Summer Time Savings section.



Note Customized summer time can be specified for both standard and customized time zones.

- To configure absolute summer time, click the **Absolute Dates** radio button.
The start date and end date for summer time can be configured in two ways: absolute dates or recurring dates. Absolute date settings apply only once and must be set every year. Recurring dates apply repeatedly for many years.
 - In the Start Date and End Date fields, specify the month (January through December), day (1–31), and year (1993–2032) on which summer time must start and end in mm/dd/yyyy format. Make sure that the end date is always later than the start date.

Alternatively, click the **Calendar** icon next to the Start Date and End Date fields to display the Date Time Picker popup window. By default the current date is highlighted in yellow. In the Date Time Picker popup window, use the left or right arrow icons to choose the previous or following years, if required. Choose a month from the drop-down list. Click a day of the month. The chosen date is highlighted in blue. Click **Apply**. Alternatively, click **Set Today** to revert to the current day. The chosen date will be displayed in the Start Date and End Date fields.
 - To configure recurring summer time, click the **Recurring Dates** radio button.
 - From the Start Day drop-down list, choose a day of the week (**Monday–Sunday**) to start.
 - From the Start Week drop-down list, choose an option (**first**, **2nd**, **3rd**, or **last**) to set the starting week. For example, choose **first** to configure summer time to recur beginning the first week of the month or **last** to configure summer time to recur beginning the last week of the month.
 - From the Start Month drop-down list, choose a month (**January–December**) to start.
 - From the End Day drop-down list, choose a day of the week (**Monday–Sunday**) to end.
 - From the End Week drop-down list, choose an option (**first**, **2nd**, **3rd**, or **last**) to set the ending week. For example, choose **first** to configure summer time to end beginning the first week of the month or **last** to configure summer time to stop beginning the last week of the month.
 - From the End Month drop-down list, choose a month (**January–December**) to end.
- Step 7** From the Start Time drop-down lists, choose the hour (0–23) and minute (0–59) at which daylight saving time should start. From the End Time drop-down lists, choose the hour (0–23) and minute (0–59) at which daylight saving time should end.

Start Time and End Time fields for summer time are the times of the day when the clock is changed to reflect summer time. By default, both start and end times are set at 00:00.
- Step 8** In the Offset field, specify the minutes offset from UTC (0–1439). (See [Table 20-2](#).)

The summer time offset specifies that the number of minutes that the system clock moves forward at the specified start time and backward at the end time.

Step 9 To not specify a summer or daylight saving time for the corresponding time zone, click the **No Customized Summer Time Configured** radio button.

Step 10 To save the settings, click **Submit**. A “Click Submit to Save” message appears in red next to the Current Settings line when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured settings by clicking the Reset button. The Reset button is visible only when you have applied default or group settings to change the current device settings but have not yet submitted the changes.

If you attempt to leave this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box only appears if you are using the Internet Explorer browser.

Table 20-2 Timezone—Offset from UTC

Time Zone	Offset from UTC (in hours)
Africa/Algiers	+1
Africa/Cairo	+2
Africa/Casablanca	0
Africa/Harare	+2
Africa/Johannesburg	+2
Africa/Nairobi	+3
America/Buenos_Aires	−3
America/Caracas	−4
America/Mexico_City	−6
America/Lima	−5
America/Santiago	−4
Atlantic/Azores	−1
Atlantic/Cape_Verde	−1
Asia/Almaty	+6
Asia/Baghdad	+3
Asia/Baku	+4
Asia/Bangkok	+7
Asia/Colombo	+6
Asia/Dacca	+6
Asia/Hong_Kong	+8
Asia/Irkutsk	+8
Asia/Jerusalem	+2
Asia/Kabul	+4.30
Asia/Karachi	+5
Asia/Katmandu	+5.45
Asia/Krasnoyarsk	+7

Table 20-2 *Timezone—Offset from UTC (continued)*

Time Zone	Offset from UTC (in hours)
Asia/Magadan	+11
Asia/Muscat	+4
Asia/New Delhi	+5.30
Asia/Rangoon	+6.30
Asia/Riyadh	+3
Asia/Seoul	+9
Asia/Singapore	+8
Asia/Taipei	+8
Asia/Tehran	+3.30
Asia/Vladivostok	+10
Asia/Yekaterinburg	+5
Asia/Yakutsk	+9
Australia/Adelaide	+9.30
Australia/Brisbane	+10
Australia/Darwin	+9.30
Australia/Hobart	+10
Australia/Perth	+8
Australia/Sydney	+10
Canada/Atlantic	−4
Canada/Newfoundland	−3.30
Canada/Saskatchewan	−6
Europe/Athens	+2
Europe/Berlin	+1
Europe/Bucharest	+2
Europe/Helsinki	+2
Europe/London	0
Europe/Moscow	+3
Europe/Paris	+1
Europe/Prague	+1
Europe/Warsaw	+1
Japan	+9
Pacific/Auckland	+12
Pacific/Fiji	+12
Pacific/Guam	+10
Pacific/Kwajalein	−12

Table 20-2 *Timezone—Offset from UTC (continued)*

Time Zone	Offset from UTC (in hours)
Pacific/Samoa	–11
US/Alaska	–9
US/Central	–6
US/Eastern	–5
US/East–Indiana	–5
US/Hawaii	–10
US/Mountain	–7
US/Pacific	–8

UTC was formerly known as Greenwich mean time (GMT). The offset time (number of hours ahead or behind UTC) as displayed in the table is in effect during winter time. During summer time or daylight savings time, the offset may be different from the values in the table and is calculated and displayed accordingly by the system clock.

Configuring System-Wide Distribution QoS Settings Based on Channel Priority

Cisco ACNS software allows you to configure system-wide QoS settings based on channel priority for proper bandwidth usage. You can create a mapping between the channel priority and the QoS settings for unicast content distribution. You can choose the corresponding QoS settings for unicast replication traffic for each channel priority group (low, medium, and high).



Note

When a single URL is associated with more than one channel, the Web content is distributed only one time to all of the Content Engines subscribed to each channel. When different QoS settings are configured for different channels that contain the same content, the channel priority setting determines which QoS settings are applied to the content distribution. The channel with the higher priority dictates which QoS settings are used.

To configure system-wide QoS settings based on channel priority, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **System > Configuration**. The Config Priorities window appears.
- Step 2** In the Contents pane, click **Distribution QoS**. The Distribution QoS Settings window appears. (See [Figure 20-5](#).)

Figure 20-5 *Distribution QoS Settings Window*

Step 3 To enable QoS settings for unicast data based on channel priority, check the **Set QoS for unicast data** check box.

Step 4 To set the QoS value for a channel with low priority, choose a Differentiated Service Code Point (DSCP) value from the QoS value for channel with low priority drop-down list. Alternatively enter a decimal value in the corresponding field.



Note See the next section, “[Setting DSCP Values for QoS Packets](#).” You can override the system-wide settings for unicast data by configuring QoS settings on a per-channel basis from the Creating a New Channel window. (See the “[Creating a Channel](#)” section on page 5-11.)

Step 5 To set the QoS value for a channel with medium priority, choose a DSCP value from the QoS value for channel with medium priority drop-down list. Alternatively enter a decimal value in the corresponding field.

Step 6 To set the QoS value for a channel with high priority, choose a DSCP value from the QoS value for channel with high priority drop-down list. Alternatively enter a decimal value in the corresponding field.

Step 7 To enable QoS settings for metadata replication traffic, check the **Set QoS value for metadata replication** check box.

Step 8 To set the QoS value for metadata replication traffic, choose a DSCP value from the QoS value for metadata replication drop-down list.

Step 9 To save the settings, click **Submit**.

Setting DSCP Values for QoS Packets

Cisco ACNS software allows you to set Differentiated Services Code Point (DSCP) values for Unicast QoS packets. DSCP values define relative priority levels for the packets. You can either choose a DSCP keyword from the drop-down list or enter a value in the corresponding field. (See [Table 20-3](#).)

Table 20-3 **DSCP Values**

Keyword	Description and Value
af11	Sets packets with AF11 DSCP (001010). Note The number in parentheses denotes the DSCP value for each per-hop behavior keyword.
af12	Sets packets with AF12 DSCP (001100).
af13	Sets packets with AF13 DSCP (001110).
af21	Sets packets with AF21 DSCP (010010).
af22	Sets packets with AF22 DSCP (010100).
af23	Sets packets with AF23 DSCP (010110).
af31	Sets packets with AF31 DSCP (011010).
af32	Sets packets with AF32 DSCP (011100).
af33	Sets packets with AF33 DSCP (011110).
af41	Sets packets with AF41 DSCP (100010).
af42	Sets packets with AF42 DSCP (100100).
af43	Sets packets with AF43 DSCP (100110).
cs1	Sets packets with CS1 (precedence 1) DSCP (001000).
cs2	Sets packets with CS2 (precedence 2) DSCP (010000).
cs3	Sets packets with CS3 (precedence 3) DSCP (011000).
cs4	Sets packets with CS4 (precedence 4) DSCP (100000).
cs5	Sets packets with CS5 (precedence 5) DSCP (101000).
cs6	Sets packets with CS6 (precedence 6) DSCP (110000).
cs7	Sets packets with CS7 (precedence 7) DSCP (111000).
default	Sets packets with the default DSCP (000000).
ef	Sets packets with EF DSCP (101110).

Modifying System Default Properties

You can modify the following system properties:


- `cdm.session.timeout`—Length of a Content Distribution Manager session (in minutes).
- `DeviceGroup.overlap`—Content Engine feature overlapping (enable or disable).
- `System.CmsUnsProgramSync.Interval`—Interval by which CMS synchronizes program import UNS objects (in minutes). The default is 1440 minutes.

- `System.datafeed.pollRate`—Poll rate between the Content Engine or the Content Router and the Content Distribution Manager (in seconds).
- `System.device.recovery.key`—Device identity recovery key. This property enables a device to be replaced by another node in the ACNS network.
- `System.guiServer.fqdn`—Scheme to use (IP address or FQDN) to launch the Content Engine Cache GUI.
- `System.healthmonitor.collectRate`—Sets the collect and send rate in seconds for the CMS device health (or status) monitor. If the rate is set to 0, the health monitor is disabled.
- `System.lcm.enable`—Local and central management feature (enable or disable). This property allows settings that are configured using the local device CLI or the central Content Distribution Manager GUI to be stored as part of the ACNS network configuration data.
- `System.monitoring.collectRate`—Rate at which the Content Engine collects and sends the monitoring report to the Content Distribution Manager (in seconds). The default is 300 seconds.
- `System.monitoring.dailyConsolidationHour`—Hour at which the Content Distribution Manager consolidates hourly and daily monitoring records.
- `System.monitoring.enable`—Content Engine statistics monitoring (enable or disable).
- `System.monitoring.monthlyConsolidationFrequency`—Frequency (in days) with which the Content Distribution Manager consolidates daily monitoring records into monthly records.
- `System.monitoring.recordLimitDays`—Maximum number of days of monitoring data to maintain in the system.
- `System.repstatus.updateEnabled`—Replication status periodic calculations on a Content Engine (enable or disable).
- `System.repstatus.updateRate`—Rate of replication status periodic updates calculated on a Content Engine (in minutes).
- `System.repstatus.updateRateSec`—Rate of replication status periodic updates calculated on a Content Engine (in seconds). The default is 600 seconds. Setting this rate will override the update rate set in minutes.
- `System.repstatus.updateSyncEnabled`—Sending summary replication status with requested detailed status (enable or disable).

The Config Properties window displays information about existing system configuration properties and their current values. To modify the value of a system property, follow these steps:


-
- Step 1** From the Content Distribution Manager GUI, choose **System > Configuration**. The Config Properties window appears. (See [Figure 20-6](#).)

Figure 20-6 Config Properties Window—Page 1



Step 2 To see the second page of this window, click **Page 2**. (See [Figure 20-7](#).)

Figure 20-7 Config Properties Window—Page 2



- Step 3** Next to the system property that you want to change, click the **Edit** icon. The Modifying Config Property window appears.
- Step 4** From a drop-down list, enter a new value or choose a new parameter, depending on the system property that you want to change.
- Step 5** To save the setting, click **Submit**.

Configuring Faster Detection of Offline Content Engines

You can detect offline Content Engines more quickly if you enable the fast detection of offline Content Engines. A Content Engine is declared as offline when it has failed to contact the Content Distribution Manager for a getUpdate (get configuration poll) request for at least two polling periods. (See [“About Faster Detection of Offline Content Engines”](#) section on page 20-22 for more information about this feature.)

To configure fast detection of offline Content Engines, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **System > Configuration**. The Config Properties window appears.
- Step 2** In the Contents pane, choose **Fast CE Offline Detection**. The Configure Fast CE Offline Detection window appears.

**Note**

The fast detection offline Content Engines feature is in effect only when the Content Distribution Manager receives the first UDP heartbeat packet and a getUpdate request from a Content Engine.

- Step 3** To enable the Content Distribution Manager detect the offline status of Content Engines quickly, check the **Enable** check box.

This check box is unchecked by default so that if UDP traffic between Content Engines and the Content Distribution Manager is blocked, then all Content Engines using ACNS 5.1 software or earlier are not declared offline when they are upgraded to ACNS 5.2 software.

- Step 4** In the Heartbeat Rate (Seconds) field, specify how often Content Engines should transmit a UDP heartbeat packet to the Content Distribution Manager.

- Step 5** In the Heartbeat Fail Count field, specify the number of UDP heartbeat packets that can be dropped during transmission from Content Engines to the Content Distribution Manager before a Content Engine is declared offline.

- Step 6** In the Heartbeat UDP Port field, specify the port number using which Content Engines will send UDP heartbeat packets to the primary Content Distribution Manager.

The Maximum Offline Detection Time field displays the product of the failed heartbeat count and heartbeat rate.

Maximum Offline Detection Time = Failed heartbeat count * Heartbeat rate

If you have not enabled the fast detection of offline Content Engines feature, then the Content Distribution Manager waits for at least two polling periods to be contacted by the Content Engine for a getUpdate request before declaring the Content Engine to be offline. However, if you enable the fast detection of offline Content Engines feature, then the Content Distribution Manager waits until the value displayed in the Maximum Offline Detection Time field is exceeded.

If the Content Distribution Manager receives the Cisco Discovery Protocol (CDP) from a Content Engine, then the Content Distribution Manager displays the Content Engine as offline after a time period of $2 * (\text{heartbeat rate}) * (\text{failed heartbeat count})$.

- Step 7** To save the settings, click **Submit**.

About Faster Detection of Offline Content Engines

Communication between the Content Engine and Content Distribution Manager using User Datagram Protocol (UDP) allows faster detection of Content Engines that have gone offline. UDP heartbeat packets are sent at a specified interval from each Content Engine to the primary Content Distribution Manager in an ACNS network. The primary Content Distribution Manager tracks the last time that it received a UDP heartbeat packet from each Content Engine. If the Content Distribution Manager has not received the specified number of UDP packets, it displays the status of the nonresponsive Content Engines as offline. Because UDP heartbeats require less processing than a getUpdate request, they can be transmitted more frequently, and the Content Distribution Manager can detect offline Content Engines much faster.

You can enable or disable this feature, specify the interval between two UDP packets, and configure the failed heartbeat count. Heartbeat packet rate is defined as the interval between two UDP packets. Using the specified heartbeat packet rate and failed heartbeat count values, the Content Distribution Manager GUI displays the resulting offline detection time as a product of heartbeat rate and failed heartbeat count.

If the fast detection of offline Content Engines is enabled, the Content Distribution Manager detects Content Engines that are in network segments that do not support UDP and uses getUpdate (get configuration poll) request to detect offline Content Engines.

By default, the feature to detect offline Content Engines more quickly is not enabled. Content Engines using releases earlier than ACNS 5.2 software might not allow UDP traffic from Content Engines to the Content Distribution Manager, so when this feature is enabled, erroneous offline status reporting might occur because these Content Engines might be in network segments that do not support UDP, and UDP heartbeat packets might not be sent at the specified interval to the primary Content Distribution Manager. If you disable the fast detection of offline Content Engines, all Content Engines can be upgraded without any error in their status being displayed in the Content Distribution Manager GUI.

However, because the UDP packets are sent as clear text, they can be spoofed by a hacker, causing the Content Distribution Manager to incorrectly report the status of a Content Engine as online even when it is actually offline. This problem can be avoided by forcing the Content Distribution Manager to display a Content Engine as offline when the Content Distribution Manager does not receive the specified small number of secure getUpdate requests from the Content Engine.

In ACNS networks with heavy traffic, dropped UDP packets can cause the Content Distribution Manager to incorrectly report the status of Content Engines as offline. To avoid this problem, configure a higher value for dropped UDP heartbeat packets.

Printing ACNS Network Data

Using the features of the Content Distribution Manager, you can print any tabular data about your ACNS network, including lists of content providers, websites, and locations, or any of the resources of your ACNS network, such as Content Engines, channels, and so on.

To print data from the Content Distribution Manager, follow these steps:

-
Step 1	From the Content Distribution Manager GUI, locate the information that you wish to print. For example, if you wanted to print data about the locations defined for your ACNS network, you would choose Network > Locations .
Step 2	To print your ACNS network data by using the default printer on your operating system, click the Printer icon.
-

Creating Proxy Error Messages

The configurable proxy error messages feature allows you to download default error message files from an FTP server or an HTTP server to the Content Engine. When these default error message files are downloaded to the Content Engine, they are stored in the /ruby/errmsg directory. The naming convention for these files is ERR_error_reason. These files are statically mapped to response error codes. The Content Engine sends a user the appropriate error message based on the mapping between the response error code and the message file. These static error files can serve as templates for creating custom proxy error messages.

After you have downloaded the error message files, you can use the upload feature to upload the file to an FTP server in your network for editing and customizing. Custom error messages are stored in the Content Engine /local/local1/ directory with the file naming convention CUSTOM_ERR_error_reason.

Configuring Proxy Error Message Download Settings for the Content Engine

You can download default error message files from a list of filenames that have been previously configured. The list of proxy error message files is fixed. The protocols that can be used for downloading default error files are FTP, HTTP, and HTTPS. You cannot download two error messages with the same name and different URLs.

To download the error message files to the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices listing window appears.
- Step 2** Click the **Edit** icon next to the Content Engine to which you want to download the error message file. The Content Engine Device Home window appears with the Contents pane on the left.
- Step 3** In the Contents pane, choose **Applications > Web > Proxy Error Messages > Download**. The Proxy Error Message Download for Content Engine window appears, listing the proxy error messages and their download URLs.
- Step 4** In Aggregate Settings, the **Yes** radio button is chosen by default. This specifies that the proxy error message configurations for the Content Engine and the device groups with which the Content Engine is associated are displayed. They cannot be modified or deleted. You can only view the proxy error messages created for device groups. Alternatively, click the **No** radio button to apply the proxy error message configurations for only the Content Engine.
- Step 5** In the taskbar, click the **Create New Proxy Error Messages** icon. The Creating New Proxy Error Message for Content Engine window appears.
- Step 6** From the Proxy Error Message drop-down list, choose a proxy error message that you wish to download to the Content Engine. (See [Table 20-4](#) for a description of the proxy error messages that can be downloaded.)
- Step 7** In the Download URL field, enter the host name or IP address of the server from which the proxy error message is to be downloaded to the Content Engine.
- Step 8** To save the settings, click **Submit**.

Table 20-4 Proxy Error Messages

Proxy Error Message Name	Description
blocked-dueto-filter-error	Error response when a request is blocked because of a filter
cache-read-error	Error response when a cache file system (cfs) read fails
cache-write-error	Error response when a cfs write fails
cdn-not-found-error	Error response when an ACNS network is not found
client-access-denied-msg	Error response when client access is denied
client-connection-broken-error	Error response when a client connection is lost
cr-domain-not-found-err	Error response when a Content Router could not be found
cr-general-error	Error response when a Content Router is not operational
cr-not-in-cz-error	Error response when a Content Router is not found in a coverage zone
cr-unavailable-error	Error response when a Content Router is not available
dns-not-available-error	Error response when DNS is unavailable for resolution
error-signature	Signature that will be appended to the error messages

Table 20-4 Proxy Error Messages (continued)

Proxy Error Message Name	Description
expect-failed-error	Error response when the Expect specifier in the HTTP request header cannot be met
ftp-bad-login-error	Error response when an FTP login fails
ftp-bad-url-error	Error response when an FTP request receives a bad URL
ftp-disabled-error	Error response when FTP is disabled
ftp-failure-error	Error response when an FTP failure occurs
ftp-internal-error	Error response when an FTP interval is exceeded
ftp-not-found-error	Error response when an FTP reports file not found
ftp-put-created-msg	Error response when an FTP PUT operation is successful
ftp-put-error	Error response when an FTP PUT operation fails
ftp-put-modified-msg	Response when an FTP update is successful
ftp-unavailable-msg	Error response when an FTP file is unavailable
http-blocked-port-msg	Error response when an HTTP request comes through a blocked port
https-blocked-port-msg	Error response when an HTTPS request comes through a blocked port
icap-processing-error	Error response when an error has occurred in ICAP processing
invalid-port-error	Error response when an invalid port is accessed
looped-req-error	Error response when a looped request is unsuccessful
not-enough-resources-error	Error response when not enough resources are available for the request process
not-in-cache	Error response when an object is not found in the cache
offline-miss-error	Error response when a Content Engine that is offline finds a cache miss
outgoing-proxy-fail-error	Error response when all outgoing proxies fail
proxy-allow-domain-error	Error response when the domain is not allowed to authenticate in proxy mode
proxy-no-default-domain-error	Error response when no default domain is available to authenticate in proxy mode
proxy-unauthenticated-error	Error response when proxy authentication fails
radius-redirect-error	Response for a RADIUS redirect message
request-blocked-msg	Error response when a request is blocked
request-malformed-error	Error response when request headers are malformed
rev-dns-not-available-msg	Error response when DNS is not available
server-connection-broken-error	Error response when a server connection is lost
ssl-server-error	Error response when an SSL handshake with the requested server fails
unsupported-cr-method-error	Error response when an unsupported Content Router method is used

Table 20-4 Proxy Error Messages (continued)

Proxy Error Message Name	Description
www-allow-domain-error	Error response when domain is not allowed to authenticate
www-no-default-domain-error	Error response when no default domain is available
www-unauthenticated-error	Error response when server authentication fails

To delete proxy error messages from the Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to delete download error messages. The Content Engine Device Home window appears with the Contents pane on the left.
- Step 3** In the Contents pane, choose **Applications > Web > Proxy Error Messages > Download**. The Proxy Error Message Download for Content Engine window appears, listing the proxy error messages and their download URLs.
- Step 4** Delete the proxy error message or messages:

- To delete all download proxy error messages configured for the Content Engine, click the **Remove All Proxy Error Message Configurations in this CE/Device Group** icon in the taskbar.

The system displays a dialog box asking you to confirm your decision. Click **OK** to confirm.



Note This action does not delete the error messages configured through device groups.

- To delete a single proxy error message, from the Proxy Error Message Download for Content Engine window, click the **Edit** icon of the proxy error message that you want to delete. The Modifying Proxy Error Message window appears. In the taskbar of this window, click the **Delete Proxy Error Message** icon to delete the download settings for the corresponding error message.

The system displays a dialog box, asking you to confirm whether you want to permanently delete the error message configuration. To confirm, click **OK**.

Configuring Proxy Error Message Upload Settings for the Content Engine

Proxy error messages can be uploaded from the Content Engine to a specified FTP server from which they can be edited and customized. You can only upload those error files for which you have previously configured the download settings. All the fields in the window are disabled if no download settings have been configured.

To upload an error file from the Content Engine to the specified FTP server, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the desired Content Engine. The Contents pane appears on the left.

- Step 3** In the Contents pane, choose **Applications > Web > Proxy Error Messages > Upload**. The Proxy Error Message Upload Settings for Content Engine window appears.
- Step 4** From the Proxy Error Message drop-down list, choose the proxy error message to be uploaded to the FTP server. All proxy error messages for which download settings have been configured are displayed here.
- Step 5** In the FTP Server Address field, enter the host name or IP address of the FTP server.
- Step 6** In the FTP Server Directory field, enter the remote directory where the error message needs to be saved.
- Step 7** In the FTP Server Filename field, enter the name of the custom error file.
A custom error file with the name entered here is created for each proxy error message in the directory specified in the previous field. If you specify the same filename for more than one proxy error message, the previously uploaded error message to that file will be overwritten.
- Step 8** In the FTP Server Username field, enter the name of the user who needs to access the FTP server.
- Step 9** In the FTP Server Password field, enter the password to authenticate the user who logs in to the server.
- Step 10** In the Confirmation Password field, reenter the password entered in the previous field.
- Step 11** To save the settings, click **Submit**.
-

Enabling Offline Operation of Network Devices

Cisco ACNS software allows you to configure your network devices to support offline operation when the external network links are disrupted.

To enable offline operation, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the device for which you want to enable offline operation.
- Step 3** In the Contents pane, choose **Applications > Web > Offline Operations**. The Offline Operations Settings window appears.
- Step 4** Check the **Enable Offline Operation** check box.
Only when this check box is checked can users continue to access preloaded and cached content if external connections are disrupted. Content objects are then delivered to users directly from the cache.
- Step 5** To save the configuration, click **Submit**.
-

To enable offline operation of network devices from the CLI, use the **offline-operation enable** global configuration command.



Monitoring and Troubleshooting the ACNS Network

This chapter provides information on monitoring and troubleshooting devices and content replication in the ACNS network, and on using the system message logs, transaction logs, and SNMP.


This chapter contains the following sections:

- [Monitoring System Status, page 21-1](#)
- [Monitoring System Events Using the System Message Log, page 21-5](#)
- [Alarm Overload Detection, page 21-10](#)
- [Monitoring Device Status, page 21-12](#)
- [Monitoring Device Performance, page 21-12](#)
- [Monitoring Specified HTTP URLs, page 21-20](#)
- [Enabling Kernel Debugger, page 21-23](#)

Monitoring System Status

The ACNS 5.5 Content Distribution Manager GUI displays the system status in a system status bar that is located above the navigation tabs in every window. The system status bar presents the overall device and content health of the system. You can use this feature to monitor devices and content replication in your ACNS network. The system status bar helps you immediately identify any problems on the network, allowing you to act and respond to problems quickly. (See [Figure 21-1](#).)

Figure 21-1 **System Status Bar**




The system status reporting mechanism uses four alarm lights to identify problems that need to be resolved. Each light represents a different alarm level, as follows:

- Green—No alarms (the system is in excellent health)
- Yellow—Minor alarms

- Orange—Major alarms
- Red—Critical alarms


When you roll your mouse over an alarm light in the system status bar, a popup message provides further details about the device or channel status. (See [Figure 21-2](#).) When you click the alarm light, a troubleshooting window opens (Troubleshooting Devices or Troubleshooting Content), listing the individual devices or channels that need attention.

Figure 21-2 Status Details



When you roll your mouse over an item under the Alarm Information column in the Troubleshooting Devices or Troubleshooting Content window, a contextual popup menu appears. The popup menu provides links to all the diagnostic tools, troubleshooting tools, logs, and monitoring applications for troubleshooting and resolving the problem. [Figure 21-3](#) shows the troubleshooting tools menu for device alarms.

Figure 21-3 Troubleshooting Tools Menu



Device Alarms

Device alarms are associated with device objects and pertain to applications and services running on Content Engines, Content Routers, and Content Distribution Managers. Device alarms are defined by the reporting application or service. Device alarms can also reflect reporting problems between the device and the Content Distribution Manager. (See [Table 21-1](#).)

Table 21-1 *Device Alarms for Reporting Problems*

Alarm	Alarm Severity	Device Status	Description
Device is offline	Critical	Offline	The device has failed to communicate with the Content Distribution Manager.
Device is pending	Major	Pending	The device status cannot be determined.
Device is inactive	Minor	Inactive	The device has not yet been activated or accepted by the Content Distribution Manager.
Device has lower software version	Minor	Online	The device is not interoperable with the Content Distribution Manager because it has an earlier software version.

Troubleshooting Devices Using the System Status Bar

To troubleshoot a device from the system status bar, follow these steps:

- Step 1** In the system status bar, click the Devices alarm light or click the alarm message next to the Devices alarm light panel. The Troubleshooting Devices window pops up as a separate window.
- Step 2** In the Alarm Information column, hold your mouse over the alarm message until the Troubleshooting tools menu appears. (See [Figure 21-3](#).)
- Step 3** Choose the troubleshooting tool that you want to use, and click the link. The link takes you to the appropriate window in the Content Distribution Manager GUI. [Table 21-2](#) describes the tools available for all device alarms.

Table 21-2 *Troubleshooting Tools for Device Alarms*

Item	Navigation	Description
Get Alarm Description(s)	None	Replaces alarm counts with alarm descriptions
Telnet to Device	Opens a Telnet window	Initiates a Telnet session using the device IP address
View Device Logs	Devices > Device Monitoring > Logs	Displays system message logs filtered for this device
Edit Device	Device Home	Displays device home window for configuration
Monitor Device	Device Home	Displays device home window for monitoring
Run Show Commands	Devices > Device Monitoring > Show/Clear Commands > Show Commands	Displays device show command tool

Using the Content Distribution Manager GUI show Command Tool

To use the Content Distribution Manager GUI **show** command tool, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the name of the device for which you want to issue a **show** command.
- Step 3** In the Contents pane, choose **Device Monitoring > Show/Clear Commands** and then click **Show Commands**.
- Step 4** From the drop-down list, choose a **show** command.
- Step 5** Enter arguments for the command, if any. (Refer to the *Cisco ACNS Software Command Reference, Release 5.x* publication for more command information.)
- Step 6** To display the **show** command output, click **Submit**. A window appears, displaying the **show** command output for that device.

Content Alarms

Content alarms pertain to content replication problems and are associated with channels. Content alarms are raised by the Content Distribution Manager based on replication status reports or by the node health manager based on acquisition and distribution errors.

If the same fault is reported by replication status and by node health manager, the Content Distribution Manager GUI reports both; one appears as the true alarm and the other as an error. The Content Distribution Manager does not corollate nor attempt to consolidate the errors generated by replication status and by node health manager.

To view the content alarms, click the **Content** alarm light or click the **Channels** link next to the alarm light in the status bar. The Troubleshooting Content window pops up. (See [Figure 21-4](#).) [Table 21-3](#) lists the content alarms.

Figure 21-4 Troubleshooting Content—Content Alarms

Channel Name	State	Severity	Alarm Information
Boltzano	N/A	Minor	Replication Status is Unknown.
channel2	N/A	Minor	Replication Status is Unknown.
export_exportSample	N/A	Minor	Replication Status is Unknown.
export_exportSample0	N/A	Minor	Replication Status is Unknown.
mcastCh	N/A	Minor	Replication Status is Unknown.
TestVT	Failed	Critical	Replication Status is Failed.

Table 21-3 Content Alarms for Channel Replication Status

Alarm	Severity	Description
Replication Status is Failed	Critical	The number of Content Engines in the channel that failed to replicate the content is greater than zero.
Replication Status is Pending	Minor	The number of Content Engines in the channel with content replication status unknown is greater than zero.

Troubleshooting Content Replication Issues Using the System Status Bar

To troubleshoot content replication issues from the system status bar, follow these steps:

- Step 1** In the system status bar, click the Content alarm light or click the alarm message next to the Content alarm light panel. The Troubleshooting Content window pops up as a separate window.
- Step 2** In the Alarm Information column, hold your mouse over the alarm message until the Troubleshooting tools menu appears.
- Step 3** Choose the troubleshooting tool that you want to use, and click the link. The link takes you to the appropriate window in the Content Distribution Manager GUI. [Table 21-4](#) describes the tools available for all content alarms.

Table 21-4 Troubleshooting Tools for Content Alarms

Item	Navigation	Description
View Replication Status	Content > Channels > Replication Status	Displays second-level replication status for a channel.
Edit Channel	Content > Channels > Definition	Opens the Modifying Channel window.

Monitoring System Events Using the System Message Log

Use the ACNS system logging feature to set specific parameters for the system log file (syslog). This file contains authentication entries, privilege level settings, and administrative details. System logging is always enabled. The system log file is located on the system file system (sysfs) partition as /local1/syslog.txt.

You can use either of the following methods to configure the Content Engine to send varying levels of event messages to disk, console, or host.

- Content Distribution Manager GUI, as described in the [“Configuring System Event Logging Using the Content Distribution Manager GUI”](#) section on page 21-5
- Content Engine CLI, as described in the [“Configuring System Event Logging Using CLI Commands”](#) section on page 21-8

Configuring System Event Logging Using the Content Distribution Manager GUI

To enable system logging, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to enable system logging. The Contents pane appears on the left.
- Step 3** From the Contents pane, choose **General Settings > Notification and Tracking > System Logs**. The System Log Settings for Content Engine window appears. (See [Figure 21-5](#).)

Figure 21-5 Syslog Settings Window

System Log Settings for Content Engine, CONTENTENGINE

Current applied settings from Content Engine, CONTENTENGINE

System Log Settings

Enable: ☒

Facility:

Console Settings

Enable: ☒

Priority:

Disk Settings

Enable Disk Settings: ☒

File Name:

Priority:

Recycle: (1000000-500000000)

Host Settings

Enable: ☒

	Hostname	Priority	Port	Rate Limit (0-10000 messages per second)
1 *	1.2.1.1	warning	514	0
2		warning	514	0

- Step 4** Under the System Log Settings heading, check the **Enable** check box to enable system logging.
- Step 5** From the Facility drop-down list, choose the appropriate facility.
- Step 6** To save the settings, click **Submit**.

You can enable sending syslog files to the console, disk, or a host.

To enable syslog files to be sent to the console, follow these steps under Console Settings:

- Step 1** Check the **Enable** check box to enable sending syslog files to the console.
- Step 2** From the Priority drop-down list, choose the severity level of the message that should be sent to the specified remote syslog host. The default priority-code is “warning” (level 4). Each syslog host is capable of receiving a different level of event messages. (See [Table 21-6](#) for a list of priority levels.)
- Step 3** To save the settings, click **Submit**.

To enable syslog files to be sent to disk, follow these steps under Disk Settings:

- Step 1** Check the **Enable Disk Settings** check box to enable sending syslog files to disk.
- Step 2** In the File Name field, enter a path and a filename where the syslog files will be stored on disk.

- Step 3** From the Priority drop-down list, choose the severity level of the message that should be sent to the specified remote syslog host. The default priority-code is “warning” (level 4). Each syslog host is capable of receiving a different level of event messages. (See [Table 21-6](#) for a list of priority levels.)
- Step 4** In the Recycle field, specify the size of the syslog file (in bytes) that can be recycled when it is stored on disk. The default value of the file size is 10000000.
- Whenever the current log file size surpasses the recycle size, the log file is rotated. (The default recycle size for the log file is 10,000,000 bytes.) The log file cycles through at most five rotations, and each rotation is saved as *log_file_name.[1-5]* under the same directory as the original log.
- The rotated log file is configured in the File Name field (or by using the **logging disk filename** command).
- Step 5** To save the settings, click **Submit**.

To enable syslog files to be sent to a host, follow these steps under Host Settings:

- Step 1** Check the Enable Host Settings check box to enable sending syslog files to a host. You can configure up to four hosts to which syslog messages can be sent. (See the next section, “[About Multiple Hosts for System Logging](#).”)
- Step 2** In the Hostname field, enter a host name or IP address of the remote syslog host. Specify up to three more remote syslog hosts in the Hostname fields 2 through 4. You must specify at least one host name if you have enabled system logging to a host.
- Step 3** From the Priority drop-down list, choose the severity level of the message that should be sent to the specified remote syslog host. The default priority-code is “warning” (level 4). Each syslog host is capable of receiving a different level of event messages. (See [Table 21-6](#) for a list of priority levels.)
- Step 4** In the Port field, specify the destination port on the remote host to which the Content Engine should send the message. The default port number is 514.
- Step 5** In the Rate Limit field, specify the number of messages that are allowed to be sent to the remote syslog host per second. To limit bandwidth and other resource consumption, messages to the remote syslog host can be rate limited. If this limit is exceeded, the specified remote syslog host drops the messages. There is no default rate limit, and by default all syslog messages are sent to all of the configured syslog hosts.
- Step 6** To save the settings, click **Submit**.

About Multiple Hosts for System Logging

Each syslog host can receive different priority levels of syslog messages. Therefore, you can configure different syslog hosts with a different syslog message priority code to enable the Content Engine to send varying levels of syslog messages to the four external syslog hosts. For example, the Content Engine can be configured to send messages that have a priority code of “error” (level 3) to the remote syslog host that has an IP address of 172.31.2.160 and messages that have a priority code of “warning” (level 4) to the remote syslog host that has an IP address of 172.31.2.161.

However, if you want to achieve syslog host redundancy or failover to a different syslog host, you must configure multiple syslog hosts on the Content Engine and assign the same priority code to each configured syslog host (for example, assigning a priority code of “critical” (level 2) to syslog host 1, syslog host 2, and syslog host 3).

In addition to configuring up to four logging hosts, you can also configure the following for multiple syslog hosts:

- A port number different from the default port number, 514, on the Content Engine to send syslog messages to a logging host.
- A rate limit for the syslog messages, which limits the rate at which messages are sent to the remote syslog server (messages per second) to control the amount of bandwidth used by syslog messages.

Configuring System Event Logging Using CLI Commands

Use the **logging** command to set specific parameters for the system log file (syslog). This file contains authentication entries, privilege level settings, and administrative details. System logging is always enabled. The system log file is located in the system file system (sysfs) storage area as /local1/syslog.txt.

Logging can be configured to send messages to the console, to disk, or to an external syslog host. See [Table 21-5](#) for a list of the key system logging parameters, descriptions, and CLI commands.

Table 21-5 System Logging Parameter Settings

GUI Parameter	Function	CLI Command
Host Server	IP address or host name of the host that receives syslog messages from the device group.	logging host <i>hostname</i> or <i>ip-address</i>
Console	Enables sending syslog files to the console.	logging console enable
Disk	Enables sending syslog files to disk.	logging disk enable
Priority	Determines priority level when system logging is enabled to disk, console, or host.	logging disk priority, logging console priority, logging host priority
Syslog file	Path of the syslog file on the hard drive.	logging disk filename <i>filename</i>
Recycle	Rewrites the syslog file when it surpasses the specified recycle size.	logging recycle size

Logging can be configured to send various levels of messages to disk, console, or host using the **priority** option of the logging command. [Table 21-6](#) lists the different priority levels of detail to send to the recipient of the syslog messages for a corresponding event.

Table 21-6 System Logging Priority Levels and Descriptions

Priority Code	Condition	Description
0	Emergency	System is unusable.
1	Alert	Immediate action needed.
2	Critical	Critical condition.
3	Error	Error conditions.
4	Warning	Warning conditions.
5	Notice	Normal but significant conditions.

Table 21-6 System Logging Priority Levels and Descriptions (continued)

Priority Code	Condition	Description
6	Information	Informational messages.
7	Debug	Debugging messages.

Viewing the System Message Log

Using the system message log feature of the Content Distribution Manager GUI, you can view information about events that have occurred in your ACNS network.

To view logged information for your ACNS network, follow these steps.



Note

The Content Distribution Manager logs messages only of the severity level “critical” or higher from registered nodes.

Step 1

From the Content Distribution Manager GUI, choose **System > Logs > System Messages**. The System Message Log window appears. (See [Figure 21-6](#).)

Figure 21-6 System Message Log (Showing All Messages)

Time	Node Type	Node Name	Module	Severity	Description	Message
Sun Aug 28 05:36:58 UTC 2005	CE	iptv-dev-510.cisco.com	Server	warning	Unexpected CLI command failure on the node	no vmt multicast station mcast 239.1.1.2 40000 http://10.77.155
Sun Aug 28 05:36:57 UTC 2005	CE	iptv-dev-510.cisco.com	Server	warning	Unexpected CLI command failure on the node	no vmt multicast station test 239.1.1.1 40000 http://10.77.155
Wed Aug 17 03:12:38 UTC 2005	CE	iptv-dev-510.cisco.com	Server	info	Server started	none
Wed Aug 17 03:08:16 UTC 2005	CE	iptv-dev-510.cisco.com	Server	info	Server is shutting down	exitCode=104
Tue Aug 16 10:48:10 UTC 2005	CE	iptv-dev-510.cisco.com	Server	info	Server started	none
Tue Aug 16 10:43:50 UTC 2005	CE	iptv-dev-510.cisco.com	Server	info	Server is shutting down	exitCode=104
Tue Aug 16 10:36:34 UTC 2005	CE	iptv-dev-510.cisco.com	Server	warning	Unexpected CLI command failure on the node	no vmt multicast station mcast_sspi 229.156.29.1 2345 http://
Tue Aug 16 10:25:42 UTC 2005	CE	iptv-dev-510.cisco.com	Server	info	Server started	none
Tue Aug 16 10:25:38 UTC 2005	CE	iptv-dev-510.cisco.com	Server	warning	Critical message on the node	%CE-TBD-2-210103: Some of tables which are accessed by
Thu Aug 11 11:08:58 UTC 2005	CDM	CDM.cisco.com	Server	warning	The device is about to disconnect from the network.	Device stream-dev1 with id CeConfig_4319 came offline

Step 2 Choose one of the following types of messages to display from the System Message Log drop-down list:

- All
- CLI
- Critical
- Database

A table listing and describing each message is displayed.

Step 3 To sort the messages chronologically, click a column heading by node type, node name, module, severity, or message text. By default, messages are listed chronologically. (See [Figure 21-6.](#))



Note If no name is available for a node, the name displayed is “Unavailable.” This might occur if the node has been deleted or has been reregistered with Cisco ACNS software.

Step 4 If you have many event messages, you might need to view multiple pages to view the activity in which you are interested. Click the forward (>>) and back (<<) buttons to move between pages. Alternatively, click the link for a specific page number to jump to that page.



Note You can choose the number of rows to be displayed in the System Message Log window by choosing a number from the Rows drop-down list.

Alarm Overload Detection

Content Engines that are running ACNS software can track the rate of incoming alarms from the Node Health Manager. If the rate of incoming alarms exceeds a certain threshold that is configured as the high water mark (HWM), then the Content Engine enters an alarm overload state. This condition occurs when multiple applications raise alarms at the same time to report error conditions. When a Content Engine is in an alarm overload state, then the following events occur:

- SNMP traps for subsequent alarm raise-and-clear operations are suspended. The trap for the raise alarm-overload alarm and the clear alarm-overload alarm are sent. However, traps related to alarm operations between the raise-alarm-overload alarm and the clear-alarm-overload alarm operations are suspended.
- Alarm overload raise-and-clear notifications are not blocked. The alarm overload state is communicated to SNMP and the Configuration Management System (CMS). However, in the alarm overload state, SNMP and the CMS are not notified of individual alarms. Individual alarm information is available from the CLI only.
- The Content Engine remains in an alarm overload state until the rate of incoming alarms decreases to the point that the alarm rate is less than the low water mark (LWM).
- If the incoming alarm rate falls below the LWM, the Content Engine comes out of the alarm overload state and begins to report the alarm counts to SNMP and the CMS.

When the Content Engine is in an alarm overload state, the Node Health Manager continues to record the alarms being raised on the Content Engine and keeps track of the incoming alarm rate. Alarms that have been raised on a Content Engine can be listed by using the Content Engine CLI commands shown

in Table 21-7. These CLI commands allow you to systematically drill down to the source of an ACNS software alarm (the cause of the problem). You can use these CLI commands to identify the source of a problem without wading through numerous ACNS software logs.

Table 21-7 Viewing Content Engine Alarms

Command	Syntax	Description
show alarms		Displays a list of all currently raised ACNS software alarms (critical, major, and minor alarms) on the Content Engine.
	show alarms critical	Displays a list of only currently raised ACNS software critical alarms on the Content Engine.
	show alarms major	Displays a list of only currently raised ACNS software major alarms on the Content Engine.
	show alarms minor	Displays a list of the currently raised ACNS software minor alarms on the Content Engine.
	show alarms detail	Displays detailed information about the currently raised ACNS software alarms.
	show alarms history	Displays a history of ACNS software alarms that have been raised and cleared on the Content Engine. The CLI retains the last 100 alarm raise and clear events only.
	show alarms status	Displays the counts for the currently raised ACNS software alarms on the Content Engine. Also lists the alarm overload state and the alarm overload settings.

Configuring Alarm Overload Detection Settings for the Content Engine

To configure alarm overload detection for a Content Engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears, listing all the device types configured in the ACNS network.
- Step 2** Click the **Edit** icon next to the name of the Content Engine for which you want to configure the alarm overload state. The Device Home for Content Engine window appears.
- Step 3** In the Contents pane, choose **General Settings > Notification and Tracking > Alarm Overload Detection**. The Alarm Overload Detection Settings for Content Engine window appears.
- Step 4** Uncheck the **Enable Alarm Overload Detection** check box if you do not want to configure the Content Engine to suspend alarm raise and clear operations when multiple applications report error conditions. This check box is checked by default.
- Step 5** In the Alarm Overload Low Water Mark (Clear) field, enter the number of incoming alarms per second below which the Content Engine comes out of the alarm overload state. Low water mark is the level to which the number of alarms must drop before alarms can be restarted. The default value is 1. The low water mark value should be less than the high water mark value.
- Step 6** In the Alarm Overload High Water Mark (Raise) field, enter the number of incoming alarms per second above which the Content Engine enters the alarm overload state. The default value is 10.


Step 7 To save the settings, click **Submit**

If you try to leave this window without saving the modified settings, a warning dialog box prompts you to submit the changes. This dialog box appears only if you are using the Internet Explorer browser.

Monitoring Device Status

The Device Home window provides device and alarm status for individual devices. (See Figure 21-7.)


Figure 21-7 Monitoring Device Status



Monitoring Device Performance

It is often useful to be able to assess the performance of Content Engines across your network. You can do this using the Content Engine statistics feature in the Content Distribution Manager GUI. The System Home window displays system-wide statistics in a graphical form and enables you to view, at a glance, which Content Engines are online, as well as assess their available resources, the volume of traffic being routed to them, and their performance in serving requests. (See Figure 21-8.)

Figure 21-8 System Home



The information displayed is based on a snapshot of your ACNS network that represents the state of your Content Engines at the end of every two polling periods. You can configure the interval between polls in the Content Distribution Manager GUI (**System > Configuration > System Properties > System.datafeed.pollRate**). The default polling rate is 300 seconds (5 minutes).

The Content Distribution Manager GUI also allows you to view statistics for Content Engines, device groups, and Content Routers in tabular form based on the type of server. (See the next three sections.)

Viewing Content Engine Statistics

To view Content Engine statistics in tabular form, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Statistics**.
- Step 2** In the Contents pane, click **Content Engines**, and then choose one of the following submenu options to view the Content Engine statistics for that option:
- Cisco Streaming Engine
 - HTTP
 - RTSP
 - Windows Media
- [Table 21-8](#) explains the meaning of each Content Engine statistic presented in the GUI.
- Step 3** To print the statistics data, click the **Printer** icon.

Table 21-8 Content Engine and Device Group Statistics

Content Engine Property	Description
Cisco Streaming Engine	
Bandwidth (bits/sec)	Current bandwidth output by the server in bits per second.
Total Bytes	Total bytes output by the server since it was started.
Total Packets	Total packets output by the server since it was started.
RTSP Connections	Number of clients currently connected over RTSP.
All Connections	Number of clients currently connected since startup.
Updated	Time stamp indicating when the statistics were updated.
HTTP	
Requests/Sec	Number of requests per second.
Bytes/Sec	Number of bytes per second.
Request Latency	Average number of seconds per HTTP request. Corresponds to the output from the show statistics http performance EXEC command as the “Seconds / Request Avg.”
Hit Rate	Average number of content items per minute successfully served from the cache of the Content Engine or from all the Content Engines in the channel or virtual ACNS network during the preceding quarter hour.
Updated	Time stamp indicating when the statistics were updated.
RTSP	
Requests	Number of requests. Corresponds to the output from the show statistics rtsp proxy media-real savings command.
Bytes	Number of bytes. Corresponds to the output from the show statistics rtsp proxy media-real savings command.

Table 21-8 *Content Engine and Device Group Statistics (continued)*

Content Engine Property	Description
Hit Rate	Average number of content items per minute successfully served from the cache of the Content Engine or from all the Content Engines in the channel or virtual ACNS network during the preceding quarter hour.
Updated	Time stamp indicating when the statistics were updated.
Windows Media	
Concurrent Requests	Number of simultaneous requests the WMT proxy/server is serving at the current time.
Kbits/Sec	Number of kilobits per second.
Cache Hit Rate	Average number of content items per minute successfully served from the cache of the Content Engine or from all the Content Engines in the channel or virtual ACNS network during the preceding quarter hour.
Updated	Time stamp indicating when the statistics were updated.

Viewing Device Group Statistics

To view device group statistics for your ACNS network, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Statistics**.
- Step 2** In the Contents pane, click **Device Groups**, and then choose one of the following submenu options to view the Content Engine statistics for that option:
- Cisco Streaming Engine
 - HTTP
 - RTSP
 - Windows Media
- Statistics for device groups are the same as for Content Engines. (See [Table 21-8](#) for an explanation of the meaning of each device group statistic presented in the GUI.)
- Step 3** To print the statistics data, click the **Printer** icon.
-

Viewing Routing Statistics

To view routing statistics for Content Routers and routing Content Engines, follow these steps:

- Step 1** From the Content Distribution Manager GUI, choose **Devices > Statistics**.
- Step 2** In the Contents pane, click **Routing Statistics**, and then choose one of the following submenu options to view the Content Engine statistics for that option:
- Routing Requests
 - Routing Redirects

[Table 21-9](#) explains the meaning of each Content Router statistic presented in the GUI.

Table 21-9 Content Router Statistics

Content Router Property	Description
Routing Requests	
Total Requests	Total number of content requests received by simplified hybrid routing.
Http Requests	Number of ASX and traditional HTTP web requests received.
Rtsp Requests	Number of RTSP requests received.
Updated	Time stamp indicating when the statistics were updated.
Routing Redirects	
Total Requests	Total number of content requests received by simplified hybrid routing.
Reqs Redirected	Total number of content requests received by simplified hybrid routing that were redirected.
Reqs Not Directed	Total number of content requests received by simplified hybrid routing that were not redirected.
Updated	Time stamp indicating when the statistics were updated.

- Step 3** To print the statistics data, click the **Printer** icon.

You can also view additional content routing statistics through the Content Distribution Manager GUI using the **show** command tool. (See the [“Using the Content Distribution Manager GUI show Command Tool”](#) section on page 21-4.) The **show** commands are further documented in the *Cisco ACNS Software Command Reference, Release 5.5* publication.

Viewing Streaming Statistics

CLI commands provide the following statistical information for incoming streams and multicast streams:

- Incoming statistics
- Outgoing statistics

- Multicast statistics
- Bandwidth Savings

The following statistics are not provided:

- Unicast per-stream statistics
- Multicast per-program statistics

Incoming Statistics

Incoming stream statistics are provided in following EXEC commands:

- **show statistics rtsp server cisco-streaming-engine bytes incoming**

The output shows incoming bytes for both unicast and multicast streams.

- For unicast, the command output shows the total incoming bytes and incoming packets and the aggregate bytes and packets for both RTSP and UDP source streams.
- For multicast, the command output shows the total incoming bytes and total incoming packets.

- **show statistics rtsp server cisco-streaming-engine usage**

The output of this command shows the incoming unicast connection details, incoming unicast bandwidth, and incoming multicast bandwidth statistics.

- For unicast incoming connection statistics, the command output provides statistics for current RTSP connections, current RTP connections, total RTSP connections, and total RTP connections.
- For unicast incoming bandwidth statistics, the command output provides statistics for current and average RTP bandwidth for RTSP source streams, current and average UDP bandwidth for UDP source streams, and the aggregate current and average bandwidth for both RTSP and UDP source streams.
- For multicast incoming statistics, the command output provides statistics for current and average multicast incoming bandwidth.

Outgoing Statistics

Outgoing stream statistics are provided in following EXEC commands:

- **show statistics rtsp server cisco-streaming-engine bytes outgoing**

The output shows outgoing bytes for both unicast and multicast streams.

- For unicast, the command output shows total and aggregate outgoing bytes and packets for VOD and live streams.
- For multicast, the command output shows total outgoing bytes and packets.

- **show statistics rtsp server cisco-streaming-engine usage**

The command output shows unicast outgoing connection details, unicast outgoing bandwidth, and multicast outgoing bandwidth statistics.

- For Unicast outgoing connection statistics, the command output shows statistics for current and total RTSP connections, current and total RTP connections, current RTP-over-RTSP connections, and current and total RTSP connections for VOD and live streams.
- For unicast outgoing bandwidth statistics, the command output shows statistics for current and average RTP bandwidth.

- For multicast outgoing statistics, the command output shows statistics for current and average multicast outgoing bandwidth.

Multicast Statistics

Multicast stream statistics are provided in the following EXEC commands:

- **show statistics rtsp server cisco-streaming-engine bytes incoming**
- **show statistics rtsp server cisco-streaming-engine bytes outgoing**
- **show statistics rtsp server cisco-streaming-engine usage**

The command output provides the following statistical information:

- Total incoming bytes and total incoming packets.
- Current and average incoming bandwidth.
- Total outgoing bytes and total outgoing packets.
- Current and average outgoing bandwidth.

Bandwidth Savings Statistics

Bandwidth savings statistics for unicast live splitting are provided in the following EXEC command:

- **show stat rtsp server cisco-streaming-engine bw-savings**

The command output shows statistics for total incoming, outgoing, and saved bytes.

Revised Command Syntax

The following existing commands have been modified to resemble the **wmt** command syntax:

```
CE# show statistics rtsp server cisco-streaming-engine ?
  all           Display all statistics
  bytes         Display bytes statistics
  requests      Display unicast request statistics
  bw-savings    Display unicast savings statistics
  usage         Display usage statistics
  performance   Display server performance

CE# show statistics rtsp server cisco-streaming-engine bytes ?
  incoming      Display incoming bytes statistics
  outgoing      Display outgoing bytes statistics
  <cr>
```

Clearing the Statistics

The following commands clear all Cisco Streaming Engine statistics without restarting the streaming server:

- **clear statistics rtsp server cisco-streaming-engine**
- **clear statistics all**

Configuring Report Options

The Content Distribution manager GUI provides four different monitoring reports. Report data is available in both tabular and graphical form.

Bytes Served Report

The Bytes Served report allows you to view Content Engine usage over time and to identify usage trends. To view the Bytes Served report and configure the reporting options, follow these steps:

-
- Step 1** Choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon for the Content Engine that you want to view or configure.
 - Step 3** In the Contents pane, choose **Device Monitoring > Statistics > Bytes Served**. The Bytes Served Report for Content Engine window appears, displaying the statistical data.
 - Step 4** To change the report parameters and display characteristics, modify the report options as desired.
 - Step 5** To reset the report parameters, click **Update**.
-

Bandwidth Efficiency Gain Report

After a Content Engine has been in use for some time and has collected statistics, the bandwidth efficiency gain report can demonstrate the value of the Content Engine in terms of bandwidth savings.

To view the Bandwidth Efficiency Gain report and configure the reporting options, follow these steps:

-
- Step 1** Choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon for the Content Engine that you want to view or configure.
 - Step 3** In the Contents pane, choose **Device Monitoring > Statistics > Bandwidth Efficiency Gain**. The Bandwidth Efficiency Gain Report for Content Engine window appears, displaying the statistical data.
 - Step 4** To change the report parameters and display characteristics, modify the report options as desired.
 - Step 5** To reset the report parameters, click **Update**.
-

Streaming Sessions Report

The Streaming Sessions report lists the total number of streaming sessions in progress at the collection time. It allows you to plan for future hardware provisioning and licensing requirements based on utilization data.

To view the Streaming Sessions report and configure the reporting options, follow these steps:

-
- Step 1** Choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon for the Content Engine that you want to view or configure.
 - Step 3** In the Contents pane, choose **Device Monitoring > Statistics > Streaming Sessions**. The Streaming Sessions Report for Content Engine window appears, displaying the statistical data.

- Step 4** To change the report parameters and display characteristics, modify the report options as desired.
 - Step 5** To reset the report parameters, click **Update**.
-

CPU Utilization Report

To view the CPU Utilization report and configure the reporting options, follow these steps:

- Step 1** Choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon for the Content Engine that you want to view or configure.
 - Step 3** In the Contents pane, choose **Device Monitoring > Statistics > CPU Utilization**. The CPU Utilization Report for Content Engine window appears, displaying the statistical data.
 - Step 4** To change the report parameters and display characteristics, modify the report options as desired.
 - Step 5** To reset the report parameters, click **Update**.
-

Monitoring Specified HTTP URLs

ACNS 5.5 software supports HTTP URL monitoring using either the Content Distribution Manager GUI or the CLI. This feature allows you to specify the time interval (in seconds) for monitoring the specified URL(s), as well as the acceptable delay (in seconds) before the URL is retrieved. You can monitor up to 10 HTTP URLs.

To configure HTTP URL monitoring using the Content Distribution Manager, follow these steps:

- Step 1** From the Content Distribution manager GUI, choose **Devices > Devices** or **Devices > Device Groups**.
 - Step 2** Click the **Edit** icon next to the name of the Content Engine (or device group) that you want to configure for URL monitoring.
 - Step 3** In the Contents Pane, choose **Applications > Web > HTTP > HTTP Monitor URL**. The HTTP Monitored URLs window appears.
 - Step 4** To add a URL to the list of monitored URLs, click the **Create New HTTP Monitored URL** icon in the taskbar. The Creating New HTTP Monitor URL window appears.
 - Step 5** In the HTTP URL field, enter the URL that you want to monitor.
 - Step 6** To configure the time interval for monitoring the specified URL, enter a number from 1 to 3600 (seconds) in the Monitoring Interval field. The default setting is 60 seconds.
 - Step 7** To configure the acceptable delay before the URL is retrieved, enter a number from 1 to 3600 (seconds) in the Acceptable Delay field. The default setting is 60 seconds.
 - Step 8** To save the settings, click **Submit**.
-

To configure a Content Engine to monitor the performance of specific URLs using the CLI, use the **http monitor url url** global configuration command.

This command enables you to specify up to 10 URLs that you want the Content Engine to monitor. The Content Engine maintains statistics about the various response characteristics for each of the monitored URLs. (You can use the new **show statistics http monitor** command to view these statistics, as described later in this section.)

```
ContentEngine(config)# http monitor url ?
WORD URL for monitoring
```

The **http monitor url url** command has two command options, the **acceptable-delay** and **interval** options. As the following sample output indicates, the **acceptable-delay** option is used to specify the acceptable delay in seconds (the maximum number of seconds that the specified monitored URL should be retrieved within). The default acceptable delay is 60 seconds.

```
ContentEngine(config)# http monitor url http://www.abc.com/ ?
acceptable-delay Threshold time in seconds before which the URL should be
retrieved.(default is 60 seconds)
interval Interval in seconds for monitoring the URL.(default is 60 seconds)
<cr>
```

As the following sample command output indicates, the **acceptable-delay** option is used to specify the acceptable delay, which is the maximum number of seconds that the specified URL should be retrieved within.

```
ContentEngine(config)# http monitor url http://www.abc.com/ acceptable-delay ?
<1-3600> Acceptable delay in seconds
```



Note

If you use the **http monitor url url** command to configure the same URL with a different interval or acceptable-delay setting, the most recently configured setting takes precedence and overrides any previously configured settings for that particular URL.

As the following sample command output indicates, the **interval** option specifies the monitoring interval (that is, how frequently the Content Engine should monitor requests for a specific URL). The monitoring interval is specified in seconds. The default monitoring interval is 60 seconds.

```
ContentEngine(config)# http monitor url http://www.abc.com/ acceptable-delay 100
interval ?
<1-3600> Monitor interval in seconds
```

In the following example, the Content Engine is configured to monitor the URL named “http://www.abc.com/” using the default values (an interval of 60 seconds and an acceptable delay of 60 seconds):

```
http monitor url http://www.abccorp.com/
```

In the following example, the Content Engine is configured to monitor the URL named “http://www.abc.com/.” The Content Engine is configured to wait up to 100 seconds for the URL to be retrieved and to monitor requests for this URL every 100 seconds.

```
ContentEngine(config)# http monitor url http://www.abc.com/ acceptable-delay 100
interval 100
```

If it takes more than 100 seconds for the URL to be retrieved, the specified acceptable delay is exceeded. The Content Engine tracks the response time (minimum and maximum delay time) as well as the number of times that the acceptable delay is exceeded for a particular URL. These statistics are shown in the output from the new **show statistics http monitor EXEC** command. (An example of the output from the **show statistics http monitor EXEC** command is provided below.)

To display statistics for the monitored URLs, use the **statistics http monitor** EXEC command. The following example shows the statistics that are reported for each of the monitored URLs:

```
ContentEngine# show statistics http monitor
HTTP Monitor URL statistics
-----

Monitor URL                               = http://www.abc.com/
Total requests                             = 118
Failed requests                           = 30
Requests above acceptable delay           = 37
Minimum response time                     = 8.183 seconds
Maximum response time                     = 210.021 seconds

Monitor URL                               = http://www.abccorp.com/
Total requests                             = 275
Failed requests                           = 44
Requests above acceptable delay           = 26
Minimum response time                     = 0.071 seconds
Maximum response time                     = 164.061 seconds
```

In the command output shown above, note the following information:

- “Failed requests” are requests that did not succeed (for example, the request failed to resolve the domain name of that URL).
- “Requests above acceptable delay” are the requests that took longer than the specified acceptable delay (the maximum number of seconds specified by the acceptable-delay setting).

The output of the **show running-configuration** EXEC command includes information about the URL monitoring configuration. In the following excerpt from the **show running-configuration** command output, this particular information is highlighted in bold:

```
ContentEngine# show running-configuration
! ACNS version 5.2.3
!
!
hostname sust-7320-cel
!
http persistent-connections timeout 300
http proxy incoming 8080
http proxy outgoing preserve-407
http tcp-keepalive enable
http monitor url http://www.abc.com/ interval 100 acceptable-delay 100
http monitor url http://www.abccorp.com/
!
ftp proxy incoming 8080
!
clock timezone US/Eastern -5 0
!
.
.
.
```

Only the non-default values are displayed in the output from the **show running-configuration** command. Consequently, because the Content Engine was configured to use the default values to monitor the URL “http://www.abccorp.com,” the above sample output does not display these values for that URL.

To display a list of monitored URLs, including the interval and acceptable delay setting for each monitored URL, use the **show http monitor EXEC** command:

```
ContentEngine# show http monitor

Monitor URL: http://www.abc.com/
Monitor Interval: 100
Acceptable Delay: 100

Monitor URL: http://www.abccorp.com/
Monitor Interval: 60
Acceptable Delay: 60
```

Enabling Kernel Debugger

Cisco ACNS software allows you to enable or disable access to the kernel debugger (kdb) from the Content Distribution Manager GUI. Once enabled, kernel debugger is automatically activated when kernel problems occur, or you can manually activate it from the local console for the ACNS device by pressing the required key sequence.

To enable the kernel debugger from the Content Distribution Manager GUI, follow these steps:

-
- Step 1** From the Content Distribution Manager GUI, choose **Devices > Devices**.
 - Step 2** Click the **Edit** icon next to the name of the Content Engine that you want to debug.
 - Step 3** In the Contents Pane, choose **General Settings > Troubleshooting > Kernel Debugger**. The Kernel Debugger window appears.
 - Step 4** To enable the kernel debugger, check the **Enable** check box, and click **Submit**.
-

To enable the kernel debugger from the CLI, use the **kernal kdb** global configuration command.



Monitoring with SNMP

ACNS 5.x software supports Simple Network Management Protocol (SNMP), which is an interoperable standards-based protocol that allows external monitoring of the Content Engine through an SNMP agent.

This chapter explains how to configure SNMP traps, recipients, community strings and group associations, user security model groups, and user access permissions. It contains the following sections:

- [Supported SNMP Versions, page 22-1](#)
- [Configuring SNMP Traps, page 22-2](#)
- [Configuring SNMP Community Settings, page 22-5](#)
- [Configuring SNMP Group Settings, page 22-6](#)
- [Configuring SNMP User Settings, page 22-8](#)
- [Configuring SNMPv2 View Settings, page 22-9](#)
- [Configuring SNMP Host Settings, page 22-10](#)
- [Configuring SNMP Asset Tag Settings, page 22-11](#)
- [Supported MIBs, page 22-12](#)
- [Key SNMP CLI Commands, page 22-13](#)
- [Configuring SNMP Traps Using the CLI, page 22-5](#)

Supported SNMP Versions

ACNS 5.x software supports the following versions of SNMP:

- Version 1 (SNMPv1)—This is the initial implementation of SNMP. Refer to RFC 1157 for a full description of its functionality.
- Version 2 (SNMPv2c)—This is the second release of SNMP, described in RFC 1902. It provides additions to data types, counter size, and protocol operations.
- Version 3 (SNMPv3)—This is the most recent version of SNMP, defined in RFC 2271 through RFC 2275.

SNMPv1 and SNMPv2c do not have any security (that is, authentication or privacy) mechanisms to keep SNMP packet traffic on the wire confidential. As a result, packets on the wire can be detected and SNMP community strings compromised.

To solve the security shortcomings of SNMPv1 and SNMPv2c, SNMPv3 provides secure access to Content Engines by authenticating and encrypting packets over the network. In ACNS 5.x software, SNMPv3 features are added to the SNMP agent in addition to SNMPv1 and SNMPv2c features.

**Note**

Registering a new SNMP agent with your ACNS network, as well as modifying or removing a registered SNMP agent, causes each of your ACNS nodes to restart as they register the configuration change. Restarting your ACNS devices results in a temporary interruption in service across your ACNS network for the time it takes for each of your devices to come back online—usually a few minutes.

Each Cisco device running ACNS 5.x software contains the software necessary to communicate information about device configuration and activity using the SNMP protocol. Before you can begin logging SNMP data, you must acquire and deploy an SNMP manager application for use with the ACNS network.

Configuring SNMP Traps

To enable the Content Engine to send SNMP traps, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the Content Engine that you want to enable. The Content Engine Device Home window appears with the Contents pane on the left.
- Step 3** In the Contents pane, choose **General Settings > Notification and Tracking > SNMP > General Settings**. The SNMP General Settings for Content Engine window appears. (See [Figure 22-1](#).) [Table 22-1](#) describes the fields in this window.

Figure 22-1 SNMP General Settings Window




Table 22-1 SNMP General Settings

GUI Parameter	Function	CLI Command
Traps		
Content Engine	Enables SNMP Content Engine traps: - Disk Read—Enables disk read error trap. - Disk Write—Enables disk write error trap. - Disk Fail—Enables disk failure error trap. - Overload Bypass—Enables WCCP overload bypass error trap. - Transaction Logging—Enables transaction log write error trap.	snmp-server enable traps content-engine
SNMP	Enables SNMP-specific traps: - Authentication—Enables authentication trap. - Cold Start—Enables cold start trap.	snmp-server enable traps snmp authentication snmp-server enable traps snmp cold-start

Table 22-1 *SNMP General Settings (continued)*

GUI Parameter	Function	CLI Command
CE Alarm	Enables Content Engine alarm traps: • Raise Critical—Enables raise-critical alarm trap • Clear Critical—Enables clear-critical alarm trap • Raise Major—Enables raise-major alarm trap • Clear Major—Enables clear-major alarm trap • Raise Minor—Enables raise-minor alarm trap • Clear Minor—Enables clear-minor alarm trap	snmp-server enable traps alarm [clear-critical clear-major clear-minor raise-critical raise-major raise-minor]
Entity	Enables SNMP entity traps.	snmp-server enable traps entity
Event	Enables the Event MIB.	snmp-server enable traps event
Config	Enables CiscoConfigManEvent error traps.	snmp-server enable traps config
Miscellaneous Settings		
Contact	Specifies a string for the MIB-II object sysContact. The name identifies the contact person for this managed node.	snmp-server contact <i>line</i>
Location	Specifies a string for the MIB-II object sysLocation. The name identifies the physical location of the node.	snmp-server location <i>line</i>
MIB Persistent Event	Enables persistence for the SNMP Event MIB.	snmp-server mib persist event
Notify Inform	Enables the SNMP notify inform request.	snmp-server notify inform

Step 4 To enable SNMP traps, check the appropriate check boxes.

Step 5 To save the settings, click **Submit**.

A “Click Submit to Save” message appears in red next to the current settings when there are pending changes to be saved after you have applied default or device group settings. You can also revert to the previously configured window settings by clicking **Reset**. The Reset button is visible only when you apply default or device group settings to change the current device settings but the settings have not yet been submitted.

Configuring SNMP Traps Using the CLI

To enable the Content Engine to send SNMP traps, use the **snmp-server enable traps** global configuration command. If you do not enter the **snmp-server enable traps** command, no traps are sent. Use the **no** form of this command to disable all SNMP traps or only SNMP authentication traps.

The **snmp-server enable traps** command is used in conjunction with the **snmp-server host** command. Use the **snmp-server host** command to specify which hosts or hosts receive SNMP traps. To send traps, you must configure at least one host.

For a host to receive a trap, both the **snmp-server enable traps** command and the **snmp-server host** command for that host must be used.

In addition, SNMP must be enabled with the **snmp-server community** command.

To disable the sending of the MIB-II SNMP authentication trap, you must enter the command **no snmp-server enable traps snmp authentication**.

The following example enables the Content Engine to send all traps to the host 172.31.2.160 by using the community string *public*:

```
ContentEngine(config)# snmp-server enable traps
ContentEngine(config)# snmp-server host 172.31.2.160 public
```

The following example disables all traps:

```
Content Engine (config)# no snmp-server enable traps
```

Configuring SNMP Community Settings



Note

Community strings are listed in the order in which they have been created. The maximum number of SNMP communities that can be created is ten. By default, an SNMP agent is disabled, and a community string is not configured. When a community string is configured, it permits read-only access to all agents by default.

To enable the SNMP agent and configure a community string to permit access to the SNMP agent, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to configure an SNMP community setting. The Contents pane appears on the left.
- Step 3** In the Contents pane, choose **General Settings > Notification and Tracking > SNMP > Community**. The SNMP Community Strings for Content Engine window appears.
- Step 4** In the taskbar, click the **Create New SNMP Community String** icon. The Creating New SNMP Community String for Content Engine window appears. [Table 22-2](#) describes the fields in this window.

Table 22-2 *SNMP Community Settings*

GUI Parameter	Function
Community* ¹	Community string used as a password for authentication when you access the SNMP agent of the Content Engine. The “Community Name” field of any SNMP message sent to the Content Engine must match the community string defined here in order to be authenticated. Entering a community string enables the SNMP agent. You can enter a maximum of 256 characters in this field.
Groupname/rw*	Group to which the community string belongs. The rw option allows a read or write group to be associated with this community string. The rw option permits access to only a portion of the MIB subtree. Choose one of the following three options from the drop-down list: • None—Choose this option if you do not want to specify a group name to be associated with the community string. If you enter a group name after selecting this option, an error message is displayed. The Group Name field remains disabled if you select this option. • Rw—Choose this option if you want to allow read-write access to the group associated with a community string. The Group Name field remains disabled if you select this option. • Group—Choose this option if you want to specify a group name.
Group Name*	Name of the group to which the community string belongs. You can enter a maximum of 256 characters in this field. This field is available only if you have chosen the group option in the previous field.

1. * = Required field.

Step 5 In the appropriate fields, enter the community string, choose whether or not read-write access to the group is allowed, and enter the group name.

Step 6 To save the settings, click **Submit**.

To enable the SNMP agent and configure a community string from the CLI, use the following global configuration command:

```
snmp-server community string [group groupname | rw]
```

Configuring SNMP Group Settings



Note

Groups are listed in the order in which they have been created. The maximum number of SNMP groups that can be created is ten.

To define a user security model group, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to create an SNMP group. The Device Home window appears with the Contents pane on the left.
- Step 3** In the Contents pane, choose **General Settings > Notification and Tracking > SNMP > Group**. The SNMP Group Strings for Content Engine window appears.
- Step 4** In the taskbar, click the **Create New SNMP Group String** icon. The Creating New SNMP Group String for Content Engine window appears. [Table 22-3](#) describes the fields in this window.

Table 22-3 *SNMP Group Settings*

GUI Parameter	Function
Name	Name of the SNMP group. You can enter a maximum of 256 characters.
Sec Model	Security model for the group. Choose one of the following options from the drop-down list: - v1—Version 1 security model (SNMP Version 1 [noAuthNoPriv]). - v2c—Version 2c security model (SNMP Version 2 [noAuthNoPriv]). - v3-auth—User security level SNMP Version 3 AuthNoPriv. - v3-noauth—User security level SNMP Version 3 noAuthNoPriv. - v3-priv— User security level SNMP Version 3 AuthPriv.
Read View	Name of the view (a maximum of 64 characters) that enables you only to view the contents of the agent. By default, no view is defined. In order to provide read access to users of the group, a view must be specified.
Write View	Name of the view (a maximum of 64 characters) that enables you to enter data and configure the contents of the agent. By default, no view is defined.
Notify View	Name of the view (a maximum of 64 characters) that enables you to specify a notify, inform, or trap. By default, no view is defined.

- Step 5** In the appropriate fields, enter the SNMP group configuration name, the security model, and the names of the read, write, and notify views.
- Step 6** To save the settings, click **Submit**.

To define a user security model group from the CLI, use the following global configuration command:

```
snmp-server group name {v1 [notify name] [read name] [write name] | v2c [notify name] [read name] [write name] | v3 {auth [notify name] [read name] [write name] | noauth [notify name] [read name] [write name] | priv [notify name] [read name] [write name]}}
```

Configuring SNMP User Settings



Note

Users are listed in the order in which they have been created. The maximum number of users that can be created is ten.

To define a user who can access the SNMP engine, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to create an SNMP user. The Contents pane appears on the left.
- Step 3** In the Contents pane, choose **General Settings > Notification and Tracking > SNMP > User**. The SNMP Users for Content Engine window appears.
- Step 4** In the taskbar, click the **Create New SNMP User** icon. The Creating New SNMP User window appears. [Table 22-4](#) describes the fields in this window.

Table 22-4 *SNMP User Settings*

GUI Parameter	Function
Name	String representing the name of the user (256 characters maximum) who can access the Content Engine.
Group	Name of the group (256 characters maximum) to which the user belongs.
Remote SNMP ID	Globally unique identifier for a remote SNMP entity. To send an SNMPv3 message to the Content Engine, at least one user with a remote SnpID must be configured on the Content Engine. The SnpID must be entered in octet string format.
Authentication Algorithm	Authentication algorithm that ensures the integrity of SNMP packets during transmission. Choose one of the following three options from the drop-down list: - No-auth—Requires no security mechanism to be turned on for SNMP packets. - MD5—Provides authentication based on the hash-based Message Authentication Code Message Digest 5 (HMAC-MD5) algorithm. - SHA—Provides authentication based on the hash-based Message Authentication Code Secure Hash (HMAC-SHA) algorithm.
Authentication Password	String (256 characters maximum) that configures the user authentication (HMAC-MD5 or HMAC-SHA) password. The number of characters is adjusted to fit the display area if it exceeds the limit for display. This field is optional if the no-auth option is chosen for the authentication algorithm. Otherwise, this field must contain a value.

Table 22-4 *SNMP User Settings (continued)*

GUI Parameter	Function
Confirmation Password	Authentication password for confirmation. The reentered password must be the same as the one entered in the previous field.
Private Password	String (256 characters maximum) that configures the authentication (HMAC-MD5 or HMAC-SHA) parameters to enable the SNMP agent to receive packets from the SNMP host. The number of characters will be adjusted to fit the display area if it exceeds the limit for display.
Confirmation Password	Private password for confirmation. The reentered password must be the same as the one entered in the previous field.

Step 5 In the appropriate fields, enter the user name, the group to which the user belongs, the engine identity of the remote entity to which the user belongs, the authentication algorithm used to protect SNMP traffic from tampering, the user authentication parameters, and the authentication parameters for the packet.

Step 6 To save the settings, click **Submit**.

To define a user who can access the SNMP engine from the CLI, use the following global configuration command:

```
snmp-server user name group [auth {md5 password [priv password] | sha password [priv password]}] | remote octetstring [auth {md5 password [priv password] | sha password [priv password]}]]
```

Configuring SNMPv2 View Settings



Note

Views are listed in the order in which they have been created. The maximum number of views that can be created is ten.

To define a Version 2 SNMP (SNMPv2) MIB view, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to create an SNMPv2 view. The Contents pane appears on the left.
- Step 3** In the Contents pane, choose **General Settings > Notification and Tracking > SNMP > View**. The SNMP Views for Content Engine window appears.
- Step 4** In the taskbar, click the **Create New View** icon. The Creating New SNMP View window appears. [Table 22-5](#) describes the fields in this window.

Table 22-5 *SNMPv2 View Settings*

GUI Parameter	Function	CLI Command
Name	String representing the name of this family of view subtrees (256 characters maximum). The family name must be a valid MIB name such as ENTITY-MIB.	snmp-server view <i>viewname</i>
Family	Object identifier (256 characters maximum) that identifies a subtree of the MIB.	snmp-server view <i>viewname</i> <i>MIBfamily</i>
View Type	View option that determines the inclusion or exclusion of the MIB family from the view. Choose one of the following two options from the drop-down list: - Included—The MIB family is included in the view. - Excluded—The MIB family is excluded from the view.	snmp-server view <i>viewname</i> <i>MIBfamily</i> { excluded included }

Step 5 In the appropriate fields, enter the view name, the family name, and the view type.

Step 6 To save the settings, click **Submit**.

Configuring SNMP Host Settings



Note

Hosts are listed in the order in which they have been created. The maximum number of SNMP hosts that can be created is four.

To configure SNMP host settings, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to define an SNMP host. The Device Home window appears with the Contents pane on the left.
- Step 3** In the Contents pane, choose **General Settings > Notification and Tracking > SNMP > Host**. The SNMP Hosts for Content Engine window appears.
- Step 4** In the taskbar, click the **Create New SNMP Host** icon. The Creating New SNMP Host window appears. [Table 22-6](#) describes the fields in this table.

Table 22-6 *SNMP Host Settings*

GUI Parameter	Function	CLI Command
Trap Host	Host name or IP address of the SNMP trap host that is sent in SNMP trap messages from the Content Engine.	snmp-server host {hostname ipaddress}
Community/User	Name of the SNMP community or user (256 characters maximum) that is sent in SNMP trap messages from the Content Engine.	snmp-server host {hostname ipaddress} communitystring
Authentication	Security model to use for sending notification to the recipient of an SNMP trap operation. Choose one of the following options from the drop-down list: - No-auth—Sends notification without any security mechanism. - v2c-noauth—Sends notification using Version 2c security. - Model v3-auth—Sends notification using SNMP Version 3 AuthNoPriv. - Security Level v3-noauth—Sends notification using SNMP Version 3 NoAuthNoPriv security. - Level v3-priv—Sends notification using SNMP Version 3 AuthPriv security.	snmp-server host {hostname ipaddress} communitystring [v2c [retry num] [timeout seconds] v3 {auth [retry num] [timeout seconds] noauth [retry num] [timeout seconds] priv [retry num] [timeout seconds]}]
Retry	Number of retries (1–10) allowed for the inform request. The default is 2 tries.	
Timeout	Timeout for the inform request in seconds (1–1000). The default is 15 seconds.	

Step 5 Enter the host name or IP address of an SNMP trap host, SNMP community or user name, security model to send notification, and retry count and timeout for inform requests.

Step 6 To save the settings, click **Submit**.

Configuring SNMP Asset Tag Settings

To configure SNMP asset tag settings, follow these steps:

- Step 1** Choose **Devices > Devices**. The Devices window appears.
- Step 2** Click the **Edit** icon next to the Content Engine for which you want to define an SNMP asset tag. The Contents pane appears on the left.
- Step 3** In the Contents pane, choose **General Settings > Notification and Tracking > SNMP > Asset Tag**. The SNMP Asset Tag Settings for Content Engine window appears.

- Step 4** In the Asset Tag Name field, enter a name for the asset tag.
- Step 5** To save the settings, click **Submit**.
-

To configure SNMP asset tag settings from the CLI, use the **asset tag** global configuration command.

Supported MIBs

The SNMP agent supports the following MIBs:

- CISCO-CDP-MIB
- ENTITY-MIB
- CISCO-ENTITY-ASSET-MIB
- MIB-II
- CISCO-CONFIG-MAN-MIB
- CISCO-CONTENT-ENGINE-MIB (supports streaming media-related MIB objects)
- HOST-RESOURCES-MIB
- EVENT-MIB

The EVENT-MIB can set the threshold on any MIB variables supported by ACNS 5.x software and store the threshold permanently on disk. The HOST-RESOURCES-MIB provides statistics on system resources.

The CISCO-CONTENT-ENGINE-MIB is extended to incorporate MIB objects related to streaming. The WMT, Cisco Streaming Engine, and RealProxy MIB groups incorporate statistics about the WMT server or proxy, Cisco Streaming Engine, and Real Proxy. For each 64-bit counter MIB object, a 32-bit counter MIB object is implemented so that SNMP clients using the SNMPv1 protocol can retrieve data associated with 64-bit counter MIB objects. The MIB objects of each of these groups are read-only.

- The WMT MIB group provides statistics about WMT proxy and server performance. Twenty-eight MIB objects are implemented in this group. Six of these MIB objects are implemented as 64-bit counters.
- The Cisco Streaming Engine MIB group provides statistics about RTSP streaming engine performance. Seven MIB objects are implemented in this group. Two of these MIB objects are implemented as 64-bit counters.
- The RealProxy MIB group provides statistics about RealProxy performance. Fourteen MIB objects are implemented in this group.

Use the following link to access these MIBs:

<ftp://ftp.cisco.com/pub/mibs/v2/>



Note

If your browser is located behind a firewall or you are connecting to the Internet with a DSL modem and you are unable to access this file folder, you must change your web browser compatibility settings. In the Internet Explorer (IE) web browser, choose **Tools > Internet options > Advanced**, and check the “Use Passive FTP” checkbox.

Key SNMP CLI Commands

Use the **snmp-server group** global configuration command to select one of the three SNMP versions, SNMPv1, SNMPv2c, or SNMPv3. Use the **no** form of this command to remove the specified version. Refer to the *Cisco ACNS Software Command Reference, Release 5.5* for more information on how to use this and other SNMP commands.

The **snmp-server community string** command provides view-based access control for SNMPv1, SNMPv2c, and SNMPv3 but also continues to provide backward compatibility between different versions. The previous version of this command did not have an option to create a community string that allows SNMP messages to execute a set operation on a MIB object. An **rw** option has been introduced for this purpose. Also, the previous version of the SNMP agent did not provide selective access control to MIB objects. Access to any MIB object was denied or granted based on authentication of the SNMP community string. With the introduction of view-based access control, it is now possible to configure a community string that grants access to only part of the MIB subtree. To provide backward compatibility with the previous version of this command, a default read group or default write group (if the **rw** option is specified on the command line) is associated with the community string if no group name is specified. Both of these default groups are hidden from users and not displayed in the configuration file or in the **show snmp group** command, but they are created during initialization of the SNMP agent.



Note

The SNMP agent is disabled by default, and a community string is not configured.

The following example enables the SNMP agent and assigns the community string *comaccess* to SNMP:

```
507-1(config)# snmp-server community comaccess
```

The preceding example defines a community string *comaccess* used as a password for authentication when you access the SNMP agent of the Content Engine. Any SNMP message sent to the Content Engine must have the “Community Name” field of the message match the community string defined here to be authenticated. Entering a community string enables the SNMP agent.

The following example disables the SNMP agent and removes the previously defined community string.

```
507-1(config)# no snmp-server community
```




PART 7

Appendices





Creating Manifest Files

This appendix describes the process for creating manifest files used to acquire and distribute content with ACNS software, Release 5.0.3 and later releases. (Manifest files are also used in the ACNS 5.0 software release; however, this chapter contains some new and changed information that does not pertain to ACNS 5.0 software.) This chapter is divided into three major sections:

- [Manifest File User Guidelines, page A-1](#)

This section provides:

- A general overview and purpose of manifest files in the context of a Cisco ACNS network
- A quick start section to get you up and running immediately
- Useful sample manifest files
- A syntax validation utility
- Explanation of live content distribution

- [Manifest File Reference, page A-33](#)

This section describes:

- Detailed manifest file structure and syntax
- XML schema
- Running of automated manifest file scripts
- Manifest file time zone tables

- [Using the ACNS Manifest Generator Tool, page A-86](#)

This section provides instructions for using the ACNS software GUI-based manifest file tool to create a basic manifest file.

Manifest File User Guidelines

This first major section contains the following topics:

- [Overview, page A-2](#)
- [Quick Start, page A-2](#)
- [Working with Manifest Files, page A-6](#)
- [More Sample Manifest Files, page A-21](#)
- [Manifest Validator Utility, page A-28](#)

Overview

Cisco ACNS 5.x software manages the acquisition and distribution of pre-positioned content through an Extensible Markup Language (XML)-based reference file called the *manifest file*. The manifest file lists content that is to be used to populate Content Engines registered on a Cisco ACNS network. There should be one manifest file per channel.

The manifest file is placed on an origin server and identified by a unique URL. The location of the manifest file is specified when you enter the manifest file URL in the Modifying Channel window of the Content Distribution Manager GUI. Unlike the treatment of content by Cisco ACNS 4.x software, pre-positioned content is not stored on the Content Distribution Manager in ACNS 5.x software but is fetched from origin servers and distributed to Content Engines by a Content Engine that is a root Content Engine for the channel.

The Content Distribution Manager disseminates the manifest file URL to each of the root Content Engines on the ACNS network. The root Content Engine then parses the file and checks for any new or different information. After the root Content Engine determines what content is new, it fetches only that new content from the specified pre-positioned or live content from one or more origin servers.

The manifest file has the following features:

- Administrators and content providers can provide content on an origin server.
- Files can be imported over HTTP, HTTPS, or FTP while they are served using another streaming protocol based on a designated type of media playserver to play back the requested file.

Content acquisition and distribution can be controlled by setting prescheduled content availability dates and times. Two content acquisition methods can be configured within the manifest file. The first method specifies the acquisition of a single <item>. The second method specifies content acquisition by crawling a website or FTP server with the <crawler> feature. Either of these two methods can schedule when the acquisition is to start and how often its content is to be checked for freshness.

Quick Start

This section will help you succeed in writing manifest files that you can use to acquire content immediately. See other sections of this chapter to learn more about specifying useful attributes to customize the manifest files further and to obtain more information on the correct manifest file syntax.



Note

The username and password specified in the Content Distribution Manager GUI Creating New Channel or Modifying Channel window (**Channels > Channels > Basic Settings > Definition**) serves only to fetch the manifest file. The actual content acquisition process does not use this username and password. For fetching actual content, the username and password need to be specified in the <server> <host> tag.

Writing XML Tags

The manifest file is a text file written in XML format. An XML text file consists of a series of XML tags. The following is an example of a simple XML tag:

```
<item attr1="value1" attr2="value2" />
```

In the preceding example, “item” is the name of the XML tag, so this tag is called the “item” tag. A tag can have many attributes in the form of name=“value.” The value field must be bounded by double quotation marks. There are two attributes inside the “item” tag shown in the example. The first attribute, called “attr1,” has a value called “value1.” The second attribute, called “attr2,” has a value called “value2.”

Tags typically start with “<” (left angle bracket) and end with “/>” (forward slash and right angle bracket), but they can start with “<” and end with “>.” If a tag ends with “>,” it means its scope is not yet complete. To complete its scope, a tag called “tag-name” must end with “</tag-name>,” as shown in this example:

```
<server name="name" >
  <host name="name" />
</server>
```

On the first line of the example, the <server> tag ends with a “>,” but its scope does not end on the first line. Its scope ends on the third line with the tag </server>. Because the <host> tag is inside the <server> tag, the <host> tag is called a subtag of the <server> tag, and the <server> tag is considered the parent tag of the <host> tag.

Two tag relationships can exist between XML tags: peer and subtag. In the following example, the two “item” tags have a peer relationship:

```
<item src="url1" />
<item src="url2" />
```

The key to identifying their peer relationship is that the first tag ends with a “/>” before the second tag starts. In the following example, the <server> and <host> tags have a subtag relationship:

```
<server name="cisco">
  <host name="url" />
</server>
```

The key to identifying their subtag relationship is that the first tag ends with a “>” before the second tag starts. The <host> tag is the subtag of the <server> tag; that is, the <server> tag is the parent of the <host> tag.

Important Manifest Tags

This section lists and briefly describes the important manifest tags for you to better understand manifest files.

- **<CdnManifest> tag**

The <CdnManifest> </CdnManifest> tag set must be the highest-level tag for the ACNS 5.x software manifest file. The tag set is required and marks the beginning and end of the manifest file content. At a minimum, each <CdnManifest> tag set must contain at least one item, or content object, that is fetched and stored, as shown in this example:

```
<CdnManifest>
...
</CdnManifest>
```

- **<item> tag**

The <item> tag is used to specify a single file to be pre-positioned.

- *src* attribute

The *src* attribute is required to specify the URL. The URL can be an absolute URL as shown in the following format:

protocol://username:password@/domain-name:port/file-path/file-name

The *src* attribute can also use a relative URL, but you need to use the <server><host> tags to specify protocol, domain, username, and password.

- <crawler> tag

The <crawler> tag is used to specify a crawl job. You can use the <crawler> tag to crawl an FTP directory and its subdirectories or to crawl directories using HTTP directory indexing.

- crawl directories

Use HTTP to crawl directories to fetch files in certain directories by enabling the built-in web server directory indexing feature. If a URL points to a directory when this directory indexing feature is enabled, the web server dynamically generates an HTML page and lists all the files and subdirectories. By parsing such an HTML page, the ACNS software can identify those files it can fetch from that particular directory.

- *start-url* attribute

The *start-url* attribute is required to specify the starting URL for crawling. It can be an absolute URL or a relative URL. (See the [•<item> tag](#) discussion.)

- *depth* attribute

The crawler *depth* attribute specifies the directory depth of a web crawl. A *depth* value of 0 allows only a crawl of the starting URL page, whereas a *depth* value of 1 allows a crawl of the starting URL page and its links.

Writing a Single-Item HTTP Manifest File

The following sample shows the simplest way to write a manifest file that fetches content using the HTTP protocol:

```
<CdnManifest>
<item src="http://www.my-server.com/project-one.html" />
<item src="http://www.my-server.com/my-eng-group/project-two.html" />
<item src="http://www.my-server.com/project-three.html" />
</CdnManifest>
```

The <CdnManifest> tag set is required to specify a manifest file. Upon execution, the preceding manifest file sample instructs the ACNS software to fetch the following items using the HTTP protocol:

- http://www.my-server.com/project-one.html
- http://www.my-server.com/my-eng-group/project-two.html
- http://www.my-server.com/project-three.html

Writing a Single-Item FTP Manifest File

The following sample shows the simplest way to write a manifest file that fetches content using the FTP protocol:

```
<CdnManifest>
<item src="ftp://johnw:my-pass-word@myftp.cisco.com/file1.txt" />
<item src="ftp://johnw:my-pass-word@myftp.cisco.com/full-path/file2.txt" />
</CdnManifest>
```

Notice that after the FTP server name “myftp.cisco.com,” the first URL has only one slash (“/”) and the second URL has two slashes (“//”). The first file path, “file1.txt,” is the relative file name and is relative to the home directory of the FTP user login. The second file path, “/full-path/file2.txt,” is the absolute path.

Writing an FTP Crawler Manifest File

The following sample shows the simplest way to write a crawler manifest file that fetches content using the FTP protocol:

```
<CdnManifest>
<crawler start-url="ftp://ftp-server/folder/" depth="10" ttl = "10" />
</CdnManifest>
```

The preceding manifest file sample instructs the ACNS software to start crawling from ftp://ftp-server/folder to ten directory levels deep and to check those directories every 10 minutes for freshness.

The <crawler> tag specifies the crawl task. The *start-url* attribute specifies where the web crawler is to start crawling. The *depth* attribute of ten specifies how many levels of subdirectories the crawler should check to obtain the required content. The *ttl* attribute specifies how often the file is to be checked for freshness. The *ttl* attribute can be specified as an attribute in a single-item manifest file as well.

Writing an HTTPS Crawler Manifest File

The following manifest file sample instructs the ACNS software to start crawling from https://www.cisco.com/jobs/eng/ to a depth of five levels:

```
<CdnManifest>
<crawler start-url="https://www.cisco.com/jobs/eng/" depth="5" />
</CdnManifest>
```

As with the single-item manifest file, the <CdnManifest> tag set is required to specify a manifest file. If directory indexing is enabled for the jobs/eng directory and its subdirectories, then the crawler will go to a depth of five directory levels to retrieve the files. If the origin server is running HTTP or HTTPS services, directory indexing must be enabled for these directories. Enabling directory indexing allows ACNS 5.x software to crawl the directory and return a list of files from it.

Validating Manifest Files

It is a good practice to use the Manifest Validator utility to validate your manifest file after it is created. See the [“Manifest Validator Utility” section on page A-28](#) for more information on the Manifest Validator utility.

Migrating from ACNS 4.x Software to ACNS 5.x Software

Unlike ACNS 4.2 software, ACNS 5.x software requires one or more origin servers where source files can be stored for the pre-positioning of content. These origin servers require that remote access servers be installed to support HTTP, FTP, or HTTPS services so that ACNS network devices can fetch pre-positioned files.

Files associated with a particular channel are typically stored in the same directory or subdirectory on the origin server. If the origin server is running HTTP or HTTPS services, directory indexing must be enabled for these directories. Enabling directory indexing allows ACNS 5.x software to crawl the directory and return a list of files from it.

Once the content is uploaded to a suitable origin server and is available, you can use the following simple manifest file to specify content acquisition:

```
<CdnManifest>
<crawler start-url="http://my-server/my-path/" ttl="10" />
</CdnManifest>
```

In the preceding example, a crawl job is specified for the associated channel to check the “my-path” directory for freshness every 10 minutes. Once this setup is complete, the root Content Engine associated with this channel monitors this directory to determine if there are any new or updated files and then automatically fetches them.

Using the preceding manifest file example achieves the same objective as that in ACNS software releases earlier than 5.0, where users copied pre-positioned files into a Content Distribution Manager import folder. However, using the manifest file is more powerful than the Content Distribution Manager import feature. For example, if you store content at different locations but the content must be distributed through the same channel, you can create multiple crawl jobs in the manifest file to monitor these locations. The following sample manifest file allows you to monitor different locations:

```
<CdnManifest>
<crawler start-url="http://my-server/my-path-http/" ttl="10" />
<crawler start-url="ftp://my-server/my-path-ftp/" ttl="10" />
</CdnManifest>
```

You are monitoring a directory “my-path-http” in an HTTP server and a directory “my-path-ftp” in an FTP server.

Working with Manifest Files

This section provides manifest file samples for carrying out specific tasks. Each sample has an associated explanation of its purpose and function. The manifest file can specify a single content object, a website crawler job, or an FTP server crawler job to acquire pre-positioned content or live content that is distributed to edge Content Engines later.

Specifying a Single Content Item

The following manifest file example specifies a single content item:

```
<CdnManifest>
<item src="http://www.my-server/test.html" />
<item src="test.html" />
<server name="my-origin-server-one"> <host name="http://www.my-server-one.com/eng/" />
</server>
<server name="my-origin-server-two"> <host name="http://www.my-server-two.com/eng/" />
</server>
<item src="project-two.html" />
<item server="my-origin-server-one" src="project-one.html" />
</CdnManifest>
```

For a single item, the key is to specify the item URL in the *src* attribute. There are two ways to specify the item URL:

- Specify the *src* attribute with the absolute URL as shown using the following format:

```
proto://username:password@/domain-name:port/file-path/file-name
```

The first `<item>` tag uses the full path.

- Specify the origin server information using the `<server><host>` tags and use the *src* attribute to only specify the relative path.

In the preceding example, every `<item>` tag except the first one uses a relative path. The second `<item>` tag uses the manifest server, where `test.html` is relative to the manifest file URL. The second `<item>` tag, “project-two.html,” uses “my-origin-server-two.” The third `<item>` tag, “project-one.html,” uses “my-origin-server-one.”

Use the `<item>` tag to specify a single content item, object, or URL. The required attribute *src* is used to specify the relative path portion of the URL. If the server *name* attribute is omitted, the server *name* attribute in the last specified `<server>` tag above it is used. If there are no `<server>` tags close by in the manifest file, the server that hosts the manifest file will be used, which means that the relative URL will be relative to the manifest file URL.



Note

Before any content can be acquired, you must enter the URL that defines the location of the manifest file in the Content Distribution Manager GUI (**Services > Web > Channels > Channel Content**). When defining the manifest settings, enter the location URL of the manifest file, its Time To Live (TTL), the username, and the password required to access the manifest file (if the location is password-protected). (See the “[Configure Manifest File and Proxy Information for the Channel](#)” section on page 6-26.)

Specifying a Crawl Job

The web crawler application methodically and automatically searches acceptable websites and makes a copy of the visited pages for later processing. The web crawler starts with a list of URLs to visit and identifies every web link in the page, adding these links to the list of URLs to visit. The process ends after one or more of the following conditions are met:

- Links have been followed to a specified depth.
- The maximum number of objects has been acquired.
- The maximum content size has been acquired.

By crawling a site at regular intervals using the Time To Live (or *ttl*) attribute, these links and their associated content can be updated regularly to keep the content fresh. For more information on the *ttl* attribute, see the “[Refreshing and Removing Content](#)” section on page A-19.

Use the `<crawler>` tag to specify the website or FTP server crawler attributes. [Table A-1](#) lists the attributes, states whether these attributes are required or optional, and describes their functions.

Table A-1 Website or FTP Server Crawler Job Attributes

Attribute	Description
<i>start-url</i>	(Required) Defines the URL to start from for the specified crawl job. It can be the full path or a relative path. If it is a relative path, the <server><host> tags are required to specify the origin server information.
<i>depth</i> (0, 1,-1)	(Optional) Defines the level of depth to crawl the specified website. The depth is defined as the level of a website's URL links or FTP server's directory, where 0 is the URL or directory from which the crawler job starts. 0 = acquire only the starting URL 1, 2, 3,... = acquire the starting URL and its referred files to the depth specified -1 = infinite or no depth restriction The default is 20 if the depth is not specified. Note It is not advisable to specify a depth of -1 because it will take a long time to crawl a large website and is wasteful if all of the content on that particular website is not required.
<i>prefix</i>	(Optional) Combines the host name from the <server> value and this field to create a full prefix. Only content whose URLs match the full prefix is acquired, as shown in this example: <pre><server name="xx"> <host name="www.cisco.com" proto="https" port=433/> </server></pre> and in a <crawler> tag: <pre>prefix="marketing/eng/"</pre> The full prefix is "https://www.cisco.com:433/marketing/eng/." Only URLs that match this prefix are crawled. If a web page refers to ../marketing/ops, the marketing/ops page and its children are not acquired. If the prefix is omitted, the crawler checks the default full prefix, which is the host name portion of the URL from the server. In the previous example, the default full prefix is "https://www.cisco.com:433."
<i>accept</i>	(Optional) Uses a regular expression to define acceptable URLs to crawl, in addition to having acceptable URLs match a prefix. For example, accept="stock" means that only URLs that meet two conditions are crawled: the URL matches the prefix and also contains the regular expression string "stock."
<i>reject</i>	(Optional) Uses a regular expression to reject a URL if it matches the expression. The URL is first checked for a possible prefix match and then checked for a reject regular expression. If a URL does not match the prefix, it is immediately rejected. If a URL matches both the prefix and the reject regular expression, it is rejected by the expression.
<i>max-number</i>	(Optional) Specifies the maximum number of crawler job objects that can be acquired.

Table A-1 Website or FTP Server Crawler Job Attributes (continued)

Attribute	Description
<i>max-size-in-B</i> <i>max-size-in-KB</i> <i>max-size-in-MB</i>	(Optional) Specifies the maximum size of content that this crawler job can acquire. The size can be expressed in bytes (B), kilobytes (KB), or megabytes (MB). Note The size of the file that is acquired is going to be less than the amount of disk space required to store the file. Files, when stored, contain overhead that contributes to the amount of disk space used for the channel. This overhead is approximately 20 KB per file. File size and storage overhead need to be taken into account when you are configuring the channel disk quota. (See the “ Channel Quota ” section on page 6-25 and the “ Creating a Channel ” section on page 5-11.)
<i>externalPrefixes</i>	(Optional) Specifies additional prefixes for crawl jobs to crawl multiple protocols or multiple websites. Prefixes are separated with a bar ().
<i>externalServers</i>	(Optional) Specifies additional hosts for crawling jobs. Can be used for multiple host crawling jobs where each host has a different user account. This attribute can be used to refer to the <host> tag with the proper authentication information.

**Note**

If you specify both the *max-number* and *max-size* attributes as the criteria to use to stop a crawler job, the condition that is met first takes precedence. That is, the crawler job stops either when the maximum number of objects is acquired or when the maximum content size is reached, whichever occurs first. For example, if the crawler job has acquired the maximum number of objects specified in the manifest file but has not yet reached the maximum content size, the crawler job stops.

The following is an example of a website crawler job:

```
<server name="cisco">
  <host name="http://www.cisco.com/jobs/" />
</server>
<crawler
  server="cisco"
  start-url="eng/index.html"
  depth="10"
  prefix="eng/"
  reject=".pl"
  max-size-in-MB="200"
/>
```

This website crawler job example contains the following attributes:

- The *start-url* path is *http://www.cisco.com/jobs/eng/index.html*.
- Search to a website link *depth* of 10.
- Search URLs with the *prefix* *http://www.cisco.com/jobs/eng/*.
- Reject URLs containing *.pl* (Perl script pages).
- Crawl only until 200 megabytes in total content size is acquired.

If the *server name* attribute is omitted, the *server name* in the last specified <server> tag above it is used. If there are no <server> tags close by in the manifest file, the server that hosts the manifest file will be used, which means that the relative URL will be relative to the manifest file URL.

Writing Common Regular Expressions

A regular expression is a formula for matching strings that follow a recognizable pattern. The following special characters have special meanings in regular expressions:

. * \ ? [] ^ \$

If the regular expression string does not include any of these special characters, then only an exact match satisfies the search. For example, “stock” must match the exact substring “stock.”

For more information about writing regular expressions, refer to the following website:

<http://yenta.www.media.mit.edu/projects/Yenta/Releases/Documentation/regex-0.12/>

Scheduling Content Acquisition

Two attributes, *ttl* and *prefetch*, are used to schedule content acquisition. Use *ttl* to specify the frequency of checking the content for freshness, in minutes. For example, to check for page freshness every day, enter *ttl*="1440."

In the following example, page freshness is scheduled to be checked once a day:

```
<item
  src="index.html"
  ttl="1440"
/>
```

In the following example, page freshness is scheduled to be crawled and checked every hour to a link *depth* value of 2:

```
<crawler
  start-url="index.html"
  depth="2"
  ttl="60"
/>
```

If the content is not yet available at a particular URL, the *prefetch* attribute can be used to specify the start time for acquisition at that specified URL. For example, *prefetch*="2002-06-28 18:35:21" means that the content acquisition job can only start on June 28, 2002 at this specific time.

The following example schedules a crawl of this website every hour to a link *depth* value of 2 to start on November 9, 2001 at 8:45 a.m.

```
<crawler
  start-url="index.html"
  depth="2"
  prefetch="2001-11-09 08:45:12"
  ttl="60"
/>
```

Specifying Shared Attributes

Attributes in single *<item>* tags can be shared or have the same attribute values. Instead of writing these attributes individually for every *<item>* tag, you can extract them and place them in a higher-level tag called *<item-group>*, where these attributes can be shared from this higher-level tag. You can create an *<item-group>* tag at a level below the *<CdnManifest>* tag, and write *<item>* tags into it as subtags, moving shared attributes into the *<item-group>* tag, as shown in the following example:

```
<?xml version="1.0"?>
<CdnManifest>
```

```

<server name="cisco-cco">
  <host name="http://www.cisco.com"
    proto="http" />
</server>

<item-group
  server="cisco-cco"
  ttl="1440"
  type="prepos" >

  <item src="jobs/index.html"/>
  <item src="jobs/index1.html"/>
  <item src="jobs/index2.html"/>
  <item src="jobs/index3.html"/>
  <item src="jobs/index4.html"/>
  <item src="jobs/index5.html"/>

</item-group>

</CdnManifest>

```

You can also use the `<options>` tag to share attributes at the topmost level of the manifest file. Shared attributes in the `<options>` tag can be shared by every `<item>` tag or by the `<crawler>` tag in the manifest file. However, if a shared attribute is specified in both the `<item-group>` and the `<item>` tags or the `<options>` and `<item>` tags, attribute values in the `<item>` tags take precedence over the `<item-group>` and `<options>` tags. (For a list of shared attributes, see the [“options” section on page A-40](#).)

The following example illustrates this precedence rule. The first `<item>` tag takes the TTL value 1440 from the `<options>` tag, but the second `<item>` uses its own TTL value of 60.

```

<options
  ttl="1440" >
<item src="index.html" />
<item src="index1.html" ttl="60" />

```

If you need to specify many single `<item>` tags and if a manifest file with many content items or URLs must be created, Perl scripts are available to create such single `<item>` tags. (See the [“Manifest File Automated Scripts” section on page A-78](#) to use automated Perl scripts.)

Specifying a Crawler Filter

With a rule-based crawler filter, you can crawl an entire website and only acquire contents with certain predefined characteristics. In contrast, crawler attributes in the `<crawler>` tag do not act as filters but only define the attributes for crawling. The `<matchRule>` tag is designed to act as a rule-based filter. You can define rule-based matches for file extensions, size, content type, and time stamp. In the following example, the crawl job is instructed to crawl the whole website starting at “index.html,” but to acquire only files with the .jpg extension and those larger than 50 kilobytes.

```

<crawler
  start-url="index.html" >
  <matchRule>
    <match size-min-in-KB="50" extension="jpg" />
  </matchRule>
</crawler>

```

There can be multiple `<match>` subtags within a `<matchRule>` tag. [Table A-2](#) lists and describes the `<match>` subtag attributes.

Table A-2 *<match> Subtag Attributes*

Attribute	Description
<i>mime-type</i>	Specifies match of these MIME-types.
<i>extension</i>	Specifies match of files with these extensions.
<i>time-before</i>	Specifies match of files modified before this time (using the Greenwich mean time [GMT] time zone) in yyyy-mm-dd hh:mm:ss format. (See the “options” section on page A-40 for a description of the <i>timezone</i> attribute.)
<i>time-after</i>	Specifies match of files modified after this time (using the Greenwich mean time [GMT] time zone) in yyyy-mm-dd hh:mm:ss format.
<i>size-min-in-MB</i> <i>size-min-in-KB</i> <i>size-min-in-B</i>	(Optional) Specifies match of content size equal to or larger than this value. The size can be expressed in megabytes (MB), kilobytes (KB), or bytes (B).
<i>size-max-in-MB</i> <i>size-max-in-KB</i> <i>size-max-in-B</i>	(Optional) Specifies match of content size equal to or smaller than this value. The size can be expressed in megabytes (MB), kilobytes (KB), or bytes (B).
<i>prefix</i>	The <i>prefix</i> attribute is optional and specifies a prefix as a match rule to filter out websites during a crawl job.
<i>url-pattern</i>	The <i>url-pattern</i> attribute is optional and specifies a regular expression as a match rule to filter out certain URLs.

A *<match>* subtag can specify multiple attributes. Attributes within a *<match>* tag have a Boolean AND relationship. In the following example, to satisfy this match rule, a file must have an .mpg type file extension and its size must be larger than 50 kilobytes.

```
<match extension="mpg" size-min-in-KB="50" />
```

There is a Boolean OR relationship between the *<match>* rules themselves. A *<matchRule>* tag can have multiple *<match>* subtags, but only one of these subtags must be matched. The *<matchRule>* tag can be specified as a subtag of the *<crawler>* tag, or a subtag of the *<item-group>* tag. If there is a subtag in an *<item-group>* tag, it is shared by every *<crawler>* tag within that *<item-group>* tag.

**Note**

The *accept* or *reject* attributes can be mistakenly used in the *<crawler>* tag for a crawler filter.

For example, to crawl files with the extension .mpg, simply specifying *accept="\.mpg"* is not correct. In this case, although specifying *accept="\.mpg"* is not technically incorrect, no crawling occurs. Pages whose URLs do not match the *accept* constraint are not searched. For example, if the starting URL is index.html, this HTML file is parsed and any links not containing .mpg are rejected. If the .mpg files are located in the second or lower link levels, they are not fetched because the links connecting them have been rejected.

To properly crawl for the .mpg extension, use *<matchRule>*. Specify *<matchRule>* *<match extension=".mpg" />*. The whole site is crawled and only those files with the .mpg extension are retained.

Specifying Content Priority

A priority can be assigned to content objects to define their order of importance. The ACNS software determines the order of processing from the level of priority of the content. The higher the content priority, the sooner the acquisition of content from the origin server and the sooner the content is distributed to the Content Engines.



Note

Every content object acquired by running a crawler job has the same priority.

Three factors combine to determine content priority:

- Channel priority—Content Distribution Priority drop-down list in the in the Acquisition and Distribution Properties area of the Modifying Channels window in the Content Distribution Manager GUI
- Item index—Content order listed in the manifest file
- Item priority—Priority of the attributes specified in the <item> or <crawler> tag

To calculate content priority, use one of the following formulas:

- If there is a priority value for this content specified in the manifest file *priority* attribute, use the following formula:

$$\text{content priority} = \text{channel priority} * 10000 + \text{item priority}$$

where *item priority* can be any integer and is unrestricted



Tip

If you want a particular content object to have the highest priority, specify a very large integer value for item priority in the content priority formula.

- If an object does not have a priority value specified in the manifest file *priority* attribute, use the following formula:

$$\text{content priority} = \text{channel priority} * 10000 + 10000 - \text{item index}$$

where *item index* is the order in which content is listed in the manifest file



Note

If there is no priority specified for any items, content is processed in the order listed in the manifest file.

Generating a Playserver List

ACNS 5.x software supports playservers that play back the following pre-positioned content types on the ACNS network: HTTP, HTTPS, and RTSP (RealMedia and QuickTime Streaming Server [QTSS]). The ACNS software checks whether the requested protocol matches the list in the playserver table. If it matches, the request is delivered. If it does not match, the request is rejected.

You can generate a playserver list in the following ways:

- The manifest file, by configuring playserver attributes in an <item> tag
- The <playServerTable> tag, by configuring playserver MIME-type extension names

To create the playserver list directly through the manifest file, configure playserver attributes of the playserver list in an <item> tag. If an <item> tag does not have a playserver attribute, its playserver list is generated through the <playServerTable> tag. If the <playServerTable> tag is omitted in the manifest file, a built-in default <playServerTable> tag is used to generate the playserver list. Multiple servers are separated by commas, as shown in the following example:

```
<item src="video.mpg" playServer="real,wmt" />
```

You can also generate the playserver list that supports these streaming media types through the <playServerTable> tag. The <playServerTable> tag maps content into a playserver list based on the MIME-type extension name. If there is a <playServerTable> tag in the manifest file, use the <playServerTable> tag in the manifest file.

To generate the playserver list through the <playServerTable> tag, use MIME-type extension names to configure which playserver can play the particular pre-positioned content, as shown in the following example:

```
<playServerTable>
<playServer name="real">
  <contentType name="application/x-pn-realaudio" />
  <contentType name="application/vnd.rn-rmadriver" />
  <extension name="rm" />
  <extension name="ra" />
  <extension name="rp" />
  <extension name="rt" />
  <extension name="smi" />
</playServer>
<playServer name="wmt">
  <extension name="wmv" />
  <extension name="wma" />
  <extension name="wmx" />
  <extension name="asx" />
  <extension name="asf" />
  <extension name="avi" />
</playServer>
<playServer name="http">
  <contentType name="application/pdf" />
  <contentType name="application/postscript" />
  <extension name="pdf" />
  <extension name="ps" />
</playServer>
</playServerTable>
```

The <playServerTable> tag is used to generate a playserver list for each content type. In the preceding example, any file with a PDF file or a PostScript extension uses HTTP to play the content.

Customized Manifest Playserver Tables and the HTTP or HTTPS Playserver

In general, you do not need to specify your own playserver table or playserver in the manifest file. The ACNS 5.x software contains a default playserver table that maps appropriate file extensions or MIME-types to the proper playservers. (See the [“Default PlayServerTable Schema”](#) section on [page A-70](#) to view the default playserver table.)

When you use the default playserver table, the HTTP and HTTPS playservers are always included in the playserver list, so that ACNS software always allows pre-positioned content to be played using the HTTP or HTTPS protocol. If the default playserver table does not meet your needs, you can customize your playserver lists by defining your own playserver table or by specifying a *playServer* attribute in the manifest file.

However, there are subtle implementation differences between ACNS 5.0.1 software and ACNS 5.0.3 software (and later releases) when you use customized playserver tables or attributes. In ACNS 5.0.1 software, even if you do not specify HTTP or HTTPS in a customized playserver table, the HTTP and HTTPS playserver are always included automatically, so that in addition to your customized list of playserver, all pre-positioned content can always be played using the HTTP or HTTPS protocol. In ACNS 5.0.3 software, the HTTP playserver is included in the default playserver table; however, if you specify your own playserver table or *playServer* attribute in the <item> or <crawler> tags, you must add the HTTP or HTTPS playserver in order to play HTTP or HTTPS content or other content using the HTTP or HTTPS protocol. The HTTP and HTTPS playserver are not automatically included, as is the case in ACNS 5.0.1 software.

In the following example, ACNS 5.0.1 software allows both the HTTP or HTTPS playserver and the WMT playserver to play the content, even though only the WMT playserver is specified:

```
<item src="video.mpg" playServer="wmt" />
```

However, in ACNS 5.0.3 software and later releases, only the WMT playserver plays the content. If you want to allow the HTTP or HTTPS playserver to play this content, you must specify the HTTP or HTTPS playserver, as follows:

```
<item src="video.mpg" playServer="wmt,http" />
```

or

```
<item src="video.mpg" playServer="wmt,https" />
```

Furthermore, in the following example, the manifest produces different results depending on whether you are using ACNS 5.0.1 software or ACNS 5.0.3 software and later releases:

```
<CdnManifest>
<playServerTable>
<playServer name="wmt">
<extension name="asf" />
<extension name="wmv" />
</playServer>
<playServer name="qtss">
<extension name="mov" />
</playServer>
</playServer>
</playServerTable>
<server name="server">
<host name="http://server.com/" proto="http" />
</server>
<crawler start-url="root" depth="3" ttl="45" />
</CdnManifest>
```

The contents crawled in this sample manifest file would have the following features:

- In ACNS 5.0.1 software, files with the extension .asf and .wmv are played by both the HTTP playserver and the WMT playserver, and files with the extension .mov are played by the HTTP playserver and the QTSS (QuickTime) playserver because the HTTP playserver is always included automatically, even though it is not specified.
- In ACNS 5.0.3 software and later releases, files with the extension .asf and .wmv are played only by the WMT playserver, and files with the extension .mov are played only by the QTSS playserver. To include the HTTP playserver, you must add <playServer name="http"> and list all your extensions, as shown in the following example:

```
<playServerTable>
<playServer name="wmt">
<extension name="asf" />
```

```

<extension name="wmv" />
</playServer>
<playServer name="qtss">
<extension name="mov" />
</playServer>
<playServer name="http">
<extension name="avi" />
<extension name="mpeg" />
<extension name="mpg" />
<extension name="mp3" />
<extension name="rm" />
<extension name="ram" />
</playServer>
</playServerTable>

```

If you do not specify your own playserver table or *playServer* attribute, the default playserver table is used. The default playserver table always includes the HTTP playserver.

In the following manifest example, ACNS 5.0.1 software and ACNS 5.0.3 software and later releases behave the same. Because there is no `<playServerTable>` tag or *playServer* attribute defined, the default playserver table is used to generate a playserver for this item, which includes the HTTP playserver. The default playserver table automatically allows files to be played with the HTTP playserver as well as the associated playserver for that type of file extension.

```

<CdnManifest>
<item src="video.asf" />
</CdnManifest>

```

Specifying Attributes for Content Serving

Certain attributes in the manifest file can be specified to control the manner in which content is served by the Content Engines. These attributes can be specified in the `<item>` and `<crawler>` tags. These same attributes can also be specified in the `<item-group>` or `<options>` tags, so they can be shared by their `<item>` and `<crawler>` subtags. [Table A-3](#) lists and describes these content-serving attributes.

Table A-3 *Attributes for Content Serving*

Attribute	Description
<i>noRedirectToOrigin</i>	(Optional) Sets the redirection to the origin server to true or false. A false setting allows the ACNS network Content Engine to redirect content requests to the origin server if the content is not available at that device. A true setting does not allow the ACNS network Content Engine to redirect content requests to the origin server and generates an error. The default setting is false. This attribute is a per-content attribute, meaning that if the content has been removed, the redirection settings do not apply. But if the software merely fails to acquire the content, then the settings apply.
<i>serveStartTime</i>	(Optional) Designates a time in yyyy-mm-dd hh:mm:ss format at which the ACNS software is allowed to start serving the content. If the serving start time is omitted, content is ready to serve once it is distributed to the Content Engine.

Table A-3 **Attributes for Content Serving (continued)**

Attribute	Description
<i>serveStopTime</i>	(Optional) Designates a time in yyyy-mm-dd hh:mm:ss format at which the ACNS software temporarily stops serving the content. If the serving stop time is omitted, the ACNS software serves the content to the Content Engine until the content is removed by modifying the manifest file or renaming the channel.
<i>alternateUrl</i>	(Optional) If the content requested by the user is not ready in the ACNS network, the ACNS network redirects the request to this alternative URL, which can be configured as an error reporting page. This attribute supports only the full URL. Note This attribute applies to the content being requested. If the content is deleted, the attribute does not apply. But if the content failed to be acquired, then this attribute should apply.
<i>requireAuth</i>	(Optional) Determines whether users need to be authenticated before the specified content is played. When <i>requireAuth</i> is set to true, the Content Engine requires authentication to play back the specified content to users and communicates with the origin server to check credentials. If the requests pass the credential check, the content is played back from the Content Engine. If this attribute is omitted, a heuristic approach is used to determine the setting: if the specified content is acquired by using a username and password, <i>requireAuth</i> is set to true; otherwise, it is set to false. For FTP, if the username is anonymous, <i>requireAuth</i> is set to false. Note If <i>requireAuth</i> is true, the Origin Server field in the Content Distribution Manager GUI Web Site window for this channel needs to point to the server that can authenticate the users. When users want to play back the content, the server specified in the Origin Server field is checked for authentication.
<i>bitrate-in-bps</i> <i>bitrate-in-kbps</i> <i>bitrate-in-mbps</i>	Specifies the bit rate of the content. The ACNS network uses this value in playback download, as shown in this example: item src="http://www.cisco.com/Prod/ACNS.swf" bitrate-in-kbps="500" / means that ACNS software downloads this file at 500 kbps, if it is requested for download and playback.

Table A-3 **Attributes for Content Serving (continued)**

Attribute	Description
<i>playDuration</i>	Specifies the playtime duration, in seconds, for a video file. This value is used for these purposes: - To download the video file and play it over HTTP If the video file is played using HTTP, ACNS software uses the <i>playDuration</i> attribute to calculate the HTTP download bit rate. - To schedule a video TV-out program If the playtime duration value is omitted, the ACNS software attempts to calculate this value for an MPEG file.
<i>ignoreQueryString</i>	Playback attribute that can be used with the <options>, <item-group>, <item>, and <crawler> tags. If <i>ignoreQueryString</i> is set to true, then the ACNS software ignores any string after a question mark (?) in the request URL for playback. If this attribute is omitted, then the default value is false. For example, content with the request URL <code>url=http://web-server/foo</code> has been pre-positioned. If a user requests content with the URL <code>url=http://web-server/foo?id=xxx</code> and the <i>ignoreQueryString</i> attribute is set to false, then the ACNS software will not use pre-positioned content from the request URL <code>http://web-server/foo</code>. However, if the <i>ignoreQueryString</i> attribute is set to true, then the ACNS software treats the request URL <code>http://www-server/foo?id=xxx</code> the same as <code>http://www-server/foo</code> and returns pre-positioned content.

Specifying Metadata for Content Serving

In certain situations, you must specify the metadata for content playback. For example, if content is acquired from an FTP server but must be played back with HTTP, the HTTP playback metadata, such as MIME type and cache control, must be specified.

The <http-meta-data> subtag is used to specify HTTP metadata. Within the <http-meta-data> subtag shown in the following example, the *name=value* attributes are *content-type="video/x-asf"* and *app-data="hh and dd."* These subtags are specified so that the ACNS software passes them directly to the end user when the HTTP content is played back.

```
<http-meta-data content-type="video/x-asf" app-data="hh and dd" />
```

The <http-meta-data> tag is a subtag of the <item> or <crawler> tags. For every <item> or <crawler> tag in the <item-group> tag that is to share the metadata, configure the <http-meta-data> tag to be a subtag of the <item-group> tag. If a <crawler> tag has an <http-meta-data> subtag, each of its crawled content objects shares this metadata.

Specifying Time Values in the Manifest File

The following attributes require that you enter a time value in the format `yyyy-mm-dd hh:mm:ss` (year-month-day hour:minute:second).

- *prefetch*
- *serveStartTime*
- *serveStopTime*

- *expires*
- *time-before*
- *time-after*

In the manifest file, the time string conforms to the yyyy-mm-dd hh:mm:ss format. A time zone designation can be optionally specified at the end of a time string to indicate the particular time zone used. If a time zone designation is omitted, the GMT time zone is used. For a complete list of time zone designations and their GMT offsets, see the “[Manifest File Time Zone Tables](#)” section on page A-71. Note that automatic conversion between daylight saving time and standard time within a time zone is not supported, but a special designation for daylight saving time can be used, such as PDT for Pacific daylight saving time. In the following example, the prefetch time is September 5, 2002 at 09:09:09 Pacific daylight saving time:

```
<options timeZone="PDT" />
<item src="index.html" prefetch="2002-09-05 09:09:09 PDT" />
```

Refreshing and Removing Content

Use the *ttl* (Time To Live) and *expires* attributes of the manifest file to monitor and control the freshness of content objects.

The *ttl* attribute is expressed in minutes and specifies how frequently the software checks the freshness of the content at the origin server. If the *ttl* attribute is specified inside an `<item>` tag, it applies to that item; if it is specified inside a `<crawler>` tag, the attribute applies to the crawl job.

For example, if you give the *ttl* attribute a value of 10, the software checks the item or crawl job every 10 minutes. If the item has been updated, then the updated file is reacquired.



Note

Sometimes a crawl job can be very large, crawling over thousands of files. The recrawl speed is 5000 files per hour for small files. It is time-consuming to recheck so many files. We strongly recommend that you specify a large *ttl* value for such crawl jobs (for example, 1440 minutes [daily]). Otherwise, the software continues to crawl the site over and over again, blocking other acquisition tasks.

If you omit the *ttl* attribute in the manifest file, the Time To Live is assumed to be zero and the software does not recheck that item after it is acquired. A value of 0 (zero) for *ttl* means that the content is fetched only once and is never checked again unless you click the Fetch Manifest Now button in the Content Distribution Manager GUI or use the **acquirer start-channel EXEC** command in the root Content Engine CLI.

The Fetch Manifest Now button is located in the **Services > Web > Channels > Channel Content** window in the Content Distribution Manager GUI. When you click this button, the software checks to see if the manifest file has been updated, and the updated manifest file is downloaded and reparsed. Also, regardless of whether the manifest file has been updated, all content in the channel is rechecked and the updated content is downloaded.

If you assign a negative value to the *ttl* attribute, such as -1, that item is never to be rechecked. A negative *ttl* attribute value prevents the software from checking item freshness, even if you click the Fetch Manifest Now button or use the **acquirer start-channel** command.



Note

Configuring the update interval in the Content Distribution Manager GUI (**Services > Web > Channels > Channel Content**) sets the interval for checking updates to the manifest file itself. This setting only pertains to checking the manifest file; it does not pertain to checking the content.

For further information, see the discussion of the *ttl* attribute in the “item” section on page A-44.

The *failRetryInterval* attribute is sometimes confused with the *ttl* attribute. The fail and retry feature acts upon failed content or failed updates. If the acquisition of a single item or of some crawled content fails, the software automatically tries to refetch these failed objects after a default interval of 5 minutes. The fail and retry interval can also be specified by using the *failRetryInterval* attribute in the manifest file.

The difference between the *failRetryInterval* attribute and the *ttl* attribute is that the *ttl* attribute is for successfully acquired content and the *failRetryInterval* attribute is for content acquisition failures. The *ttl* attribute must be specified for the software to recheck the content freshness, whereas the *failRetryInterval* attribute does not need to be specified unless you want to change the retry interval.

For further information, see the discussion of the *failRetryInterval* attribute in the “item” section on page A-44.

The *expires* attribute designates the time that the content is to be removed from the ACNS network. If you do not specify a time when you set the *expires* attribute, content is stored in the ACNS network until it is explicitly removed when you modify the manifest file. The *expires* attribute uses the format yyyy-mm-dd hh:mm:ss (year-month-day hour:minute:second). In the following example, the content expires on June 12, 2003 at 2:00 p.m.

```
expires="2003-06-12 14:00:00 PST"
```

If the *expires* attribute is specified inside an <item> tag, it applies to that item; if it is specified inside a <crawler> tag, the attribute applies to the crawl job. For further information, see the discussion of the *expires* attribute in the “item” section on page A-44.

You can monitor the status of content replication and freshness by enabling and then viewing the transaction log files that reside on the Content Engines. To verify whether or not a content object or file was successfully imported to or refreshed on a particular Content Engine, take these actions:

- Enable the transaction log function on the Content Engine you want to monitor.
- View the transaction log entries for the content object or filename that resides on that Content Engine.

Specifying Live Content

The two types of live content that you can specify in a manifest file are:

- wmt-live
- real-live

Use the <item> tag and specify the *type* attribute as either wmt-live or real-live, as shown in the following example. The live stream for the wmt-live content type is

url=rtsp://www.company-web-site.org/tmp/ceo-talk. The stream for the real-live content type is url=rtsp://real-server/tmp/funny-video.

```
<CdnManifest>
<server name="wmt-server">
<host name="rtsp://www.company-web-site.org" />
</server>
<item src="/tmp/ceo-talk" type="wmt-live" >
</item>
<!--
This is a "wmt-live" streaming content type specified by the "type" attribute. The live
stream URL is
rtsp://www.company-web-site.org/tmp/ceo-talk.
-->
<server name="real-server">
<host name="real-server" proto="rtsp" />
```

```

</server>
<item src="tmp/funny-video" type="real-live" />
<!--
This is "real-live" streaming content type specified by the "type" attribute. The stream
URL is rtsp://real-server/tmp/funny-video.
-->
</CdnManifest>

```

More Sample Manifest Files

This section contains five sample manifest files. In XML, text between `<!--` and `-->` contains comments that have no effect on the execution of the file. In these five samples, narrative comments have been added immediately below certain tags or groups of tags to provide you with a better understanding of what these particular tags mean. You can copy an entire sample file, save it to a text file, and then view it with Microsoft Internet Explorer.

Additionally, cross-reference links from the first occurrence of a tag to the [“Manifest File Reference” section on page A-33](#) have been embedded in the narrative comments of each sample to provide you with a more in-depth explanation of the tag.

To download these sample files from Cisco.com, see the [“Downloading the Sample Manifest Files” section on page A-28](#).

The following paragraphs describe the topics covered by the five sample manifest files:

Sample 1

- How to use HTTP, HTTPS, and FTP protocols to acquire content
- How to specify a username and password when authentication is required

Sample 2

- How to specify attributes for acquisition:
 - *ttl*—Sets the time interval between content freshness checks
 - *prefetch*—Specifies the time when the ACNS software can start to acquire content from the origin server
- How to specify acquisition and distribution priorities
- How to specify the following playback attributes:
 - *serveStartTime*—Sets time and date to start serving this content
 - *serveStopTime*—Sets time and date to stop serving this content
 - *alternativeURL*—Provides an alternative URL to use if content has not yet been replicated at the Content Engine
 - *requireAuth*—Requires authentication credentials from users to play back the content
 - *expires*—Sets content expiration time and date
 - *playServer*—Chooses which play servers can play the specified content
 - *noRedirectToOrigin*—If set to false and content has not yet been replicated, does not redirect the incoming request to the origin server
 - *<http-meta-data>*—Adds metadata for HTTP playback

Sample 3

- A simple crawl job
 - FTP crawl of a directory
 - HTTP crawl of a directory
 - HTTP crawl of a website
- A simple crawl job using the <matchRule> tag
 - FTP crawl of a directory to fetch only MPEG files
 - HTTP crawl of a directory to fetch only files larger than 10 MB
 - HTTP crawl of a web site, to fetch only if-modified-since files

Sample 4

- Items with the <contains> tag included

Sample 5

- RealMedia and WMT streaming live content

Sample 1

Sample 1 shows how to use HTTP, HTTPS, and FTP protocols to acquire content and how to specify a username and password when authentication is required.

```
<CdnManifest>
<!--
The CdnManifest tag pair is absolutely essential for a manifest file. It has to be the
first tag and used only as a super tag for other tags.
-->

<item •src="http://my-server.xyz.com/myphotocollection/index.html" />
<item •src="http://my-server.xyz.com/myphotocollection/myname/000001.jpg" />
<!--
The preceding XML code defines single items, using the <item/> tag, to fetch from the
origin server.
-->

<item
•src="http://my-name:my-password@my-auth-server.xyz.com/mymoviecollection/index.html" />
<item
•src="http://my-name:my-password@my-auth-server.xyz.com/mymoviecollection/myname/000001.w
mv" />
<!--
This origin server requires user authentication. Notice that there is my-name:my-password
in the URL to specify the username and password to access the file.
-->

<item •src="https://my-server.xyz.com/my_secure_photocollection/index.html" />
<item •src="https://my-server.xyz.com/my_secure_photocollection/myname/000001.jpg" />
<!--
From this origin server, the content is to be acquired using HTTPS, or HTTP over SSL, so
that the protocol specified is HTTPS.
-->

<!--
The following XML code defines the origin server from which to obtain content. The content
is to be acquired using HTTPS, or HTTP over SSL, so that the protocol specified is HTTPS.
This origin server also requires user authentication. The "user" and "password" attributes
```

specify the required username and password to access content from the origin server. The "sslAuthType" attribute specifies either "weak" or "strong" SSL certification. For example, "weak" certification allows expired, self-signed certification.

```
-->
<item
  •src="https://myself:mypwd@my-auth-server.xyz.com:443/my_auth_moviecollection/index.html"
  sslAuthType="weak" />
<item
  •src="https://myself:mypwd@my-auth-server.xyz.com:443/my_auth_moviecollection/myname/0000
01.wmv" sslAuthType="weak"/>

<item •src="ftp://my-server.xyz.com//my-doc-root/myphotocollection/index.html" />
<item •src="ftp://my-server.xyz.com/my-doc-root/myphotocollection/file1.jpg" />
<!--
The preceding XML code defines two single items to obtain from the origin
server. Notice that the first item has two "/" characters (forward slashes). This means
the path is absolute or relative to the root directory. The second item only has one "/"
character (forward slash). This means that the content path is relative to the default
login directory for an anonymous user.
To understand absolute and relative paths, consider the following directory listings:
The first directory lists the contents of /my-doc-dir, and the second directory lists the
contents of anonymous-default-dir, where anonymous-default-dir is the default directory
for the "anonymous" user.
xyz# ls -lR /my-doc-root/
/my-doc-root/:
total 1
drwxrwxrwx 2 admin root 1024 Dec 28 01:46 myphotocollection
/my-doc-root/myphotocollection:
total 1
-rw-rw-rw- 1 admin root 4 Dec 28 01:46 index.html
xyz# ls -lR /anonymous-default-dir/
/anonymous-default-dir/:
total 1
drwxrwxrwx 3 admin root 1024 Dec 28 01:53 my-doc-root
/anonymous-default-dir/my-doc-root:
total 1
drwxrwxrwx 2 admin root 1024 Dec 28 01:53 myphotocollection
/anonymous-default-dir/my-doc-root/myphotocollection:
total 1
-rw-rw-rw- 1 admin root 4 Dec 28 01:53 index.html
The single item with the following absolute path
<item •src="/my-doc-root/myphotocollection/index.html" />
fetches the file /my-doc-root/myphotocollection/index.html.
The single item with the following relative path
<item •src="my-doc-root/myphotocollection/file1.jpg" />
fetches the file /anonymous-default-dir/my-doc-root/myphotocollection/file1.jpg.
You must be careful to specify exactly what you want.
-->

<item
  •src="ftp://my-name:my-password@my-auth-server.xyz.com//my-doc-root/mymoviecollection/ind
ex.html" />
<item
  •src="ftp://my-name:my-password@my-auth-server.xyz.com/my-own-moviecollection/wedding/fil
e1.wmv" />
<!--
The preceding XML code defines two single items to obtain from the origin server. Notice
that the first item specifies an absolute path, and the second one specifies a relative
path. In this case, the relative path is relative to the default login directory for the
user "my-name."
-->
</CdnManifest>
```

Sample 2

Sample 2 is a manifest file written to show how to specify attributes.

```
<CdnManifest>

<item •src="ftp://my-ftp-server/data/video.mpg"
    •ttl="60"
    •type="prepos"
    •prefetch="2003-03-20 10:00:00 PST"
    •noRedirectToOrigin="true"
    •requireAuth="true"
    playServer="http,wmt"
    •expires="2003-06-12 14:00:00 PST"
    •alternateUrl="http://my-web-server.com/video-error-page.htm"
    •priority="50000"
    •serveStartTime="2003-01-12 14:00:00 PST"
    •serveStopTime="2099-04-12 14:00:00 PST"
    •playDuration="240"
    •failRetryInterval="60"
/>

<!--comments:
src: specifies the file location, and is required.
prefetch: specifies the time when the ACNS software can start to acquire content from the
origin server.
ttl: checks whether this file is updated every 60 minutes, and is required.
noRedirectToOrigin: when false, does not redirect the request to the origin server if the
content has not yet been replicated on the Content Engine.
requireAuth: when true, requires authentication to play back this content to users. User
requests are redirected to the origin server to check credentials. If the requests pass
the credential check, the content is played back from the Content Engine.
playServer: allows the HTTP Apache server and WMT server to play back this content. That
is, the supported playback protocols for this content are HTTP and RTSP.
expires: removes content from the ACNS network when it expires on the specified date.
alternateUrl: redirects the user to this URL when the request to play back the content is
received but the content has not yet been replicated on the Content Engine.
priority: specifies the item priority. Content acquisition and distribution is processed
in the order set by the overall priority. This means that the higher the overall priority,
the earlier the content is acquired and distributed. The overall priority is calculated as
channel-priority * 10000 + item-priority. Channel priority is 250 for low, 500 for normal,
and 750 for high. Item priority is 10000 - (index of the item in the manifest file) if a
priority is not specified. For example, assume there are two items in a manifest file. The
first item does not have a "priority" attribute, but the second item does and its priority
is 20000. The item priority of the first item is 10000 - 1 = 9999 and the item priority of
the second item is 20000. In this example, the item priority for this item is 50000.
serveStartTime: specifies the time that ACNS software can start to serve this content.
serveStopTime: specifies the time that ACNS software stops serving this content.
playDuration: specifies the video play duration in seconds. It will be
used by the HTTP playserver to figure out the downloading bit rate.
failRetryInterval: specifies the interval to retry in minutes if this
acquisition fails. In this example, the system will try again in 60 minutes.
end of comments -->

<crawler •start-url="ftp://my-ftp-server//root/data/video-files/"
    •depth="3"
    •ttl="60"
    •noRedirectToOrigin="false"
    •requireAuth="true"
    playServer="http,wmt"
    •expires="2003-06-12 14:00:00 PST"
```

```

    •alternateUrl="http://my-web-server.com/video-error-page.htm"
    •priority="50000"
    •serveStartTime="2003-01-12 14:00:00 PST"
    •serveStopTime="2003-04-12 14:00:00 PST"
    •failRetryInterval="60"
  />

  <!-- comments
  start-url: specifies the directory where crawling is to start "/root/data/video-files/."
  depth: specifies a crawl level of three subdirectories.
  noRedirectToOrigin: if false and the crawled items are not replicated to the Content
  Engine, the request for that content is not redirected to the origin server.
  requireAuth: if true, authentication is required for all crawled content.
  playServer: all crawled content can be played back by an HTTP web server and WMT streaming
  server.
  expires: all crawled content is expired and deleted at the time specified.
  alternateUrl: if any of the crawled items are not replicated to the Content Engine, the
  request for that content is redirected to this URL.
  priority: all crawled items have the same item priority of 50000. Because they are in the
  same channel, they have the same overall-priority.
  serveStartTime: all crawled content can be served after the specified time.
  serveStopTime: all crawled content cannot be served after the
  specified time.
  failRetryInterval: if some of the content in this crawl job fails to be acquired, ACNS
  software will try to acquire the content again at the specified interval. In this case,
  ACNS software will retry in 60 minutes.
  end of comments -->

  <item •src="ftp://my-ftp-server/data/video2.mpg"
    playServer="http,wmt" >
    <http-meta-data content-type="video/mpeg" />
  </item>

  <!-- comments
  The <http-meta-data> tag can be used to specify any metadata for HTTP playback. For
  example, because this item is acquired using FTP, this tag must be used to specify the
  content type for this MPEG file.
  end of comments
  -->

  <crawler •start-url="ftp://my-ftp-server/data/mpeg-files-2/"
    •depth="3" >
    <http-meta-data content-type="video/mpeg" />
  </crawler>

  <!-- comments
  For this crawl job, three subdirectory levels are crawled under the "data/mpeg-files-2/"
  folder. The <http-meta-data> tag is used to specify the content type for HTTP playback for
  all crawled content.
  end of comments -->

</CdnManifest>

```

Sample 3

Sample 3 is a manifest file that shows how to use the crawl feature.

```

<CdnManifest>

<crawler •start-url="ftp://ftp-server/pub-data/video-files/" •ttl="10" •depth="1" />

```

```

<!-- The preceding XML code specifies that an FTP crawl job is to crawl the "ftp-server"
using the <crawler> </crawler> tag pair. Its starting directory is "pub-data/video-files/"
and its crawl depth is 1. The files in this folder are monitored at 10-minute intervals.
If files are updated, removed, or added, the resulting change is reflected in the ACNS
software. -->
<crawler •start-url="ftp://ftp-server/video-files/" •ttl="10" •depth="1" >
  <matchRule>
    <match extension="mpg" />
  </matchRule>
</crawler>

<!-- This crawl job is similar to the preceding crawl job, except that it includes a
<matchRule> </matchRule> tag pair to specify the kind of content that is to be acquired.
In this case, only files with the "mpg" file extension are acquired.
-->

<crawler •start-url="http://www.ftp-server.com/pub-data/video-files/" •depth="5" />

<!-- This is an HTTP directory crawl job. The HTTP server must be configured to enable the
directory indexing feature for those directories that you want to crawl. For the Apache
server, you must modify the Apache configuration file such that it looks like the
following:<crawl directories pub-data/video-files/ > Options Index </crawl directories> If
the request URL points to a directory, the web server dynamically generates an HTML page
with a list of files contained in that directory. In this crawl job, the directory
"pub-data/video-files/" and its subdirectories are crawled to a depth level of 5. -->

<crawler •start-url="http://www.ftp-server.com/pub-data/video-files/" •depth="5" >
  <matchRule>
    <match size-min-in-KB="10" extension="mpg" />
  </matchRule>
</crawler>

<!-- This crawl job is similar to the preceding crawl job, except that it includes the
<matchRule> </matchRule> tag pair. This matchRule tag fetches only files that match files
with the extension "mpg" and file sizes equal to or larger than 10 MB.
-->

<crawler •start-url="http://www.cnn.com/sport/index.htm" •prefix="sport/" •depth="3"
size-min-in-KB="1000" />

<!-- This crawl job attempts to crawl part of cnn.com. The start URL is
http://www.cnn.com/sport/index.htm. This job only crawls URLs with the prefix
http://www.cnn.com/sport/, acquiring only files from the directory "sport/." The "depth=3"
means the job is to crawl only up to 3 link levels. The max-size-in-MB means that crawling
stops if the total of crawled items reaches 1000 MB.
-->

<crawler •start-url="http://www.cnn.com/movie/index.htm" •depth="3" >
  <matchRule>
    <match extension="mpg" •time-after="2002-01-02 00:00:00" />
    <match extension="asf" •time-after="2002-07-02 00:00:00" />
  </matchRule>
</crawler>

<!-- This crawl job attempts to crawl part of www.cnn.com. Its start URL is
http://www.cnn.com/movie/index.htm. and its crawl depth is 3. This matchRule acquires only
files with the "mpg" extension created after than Jan. 2, 2002 or files with the "asf"
extension created after July 2, 2002.
-->

</CdnManifest>

```

Sample 4

Sample 4 is a manifest file written to show the purpose and use of the <contains> tag. The <contains> tag is designed to prevent serving content if the required files are not present on the Content Engine. Typically, the delivery of a presentation consists of serving multiple files. For example, if an ASF video file uses several JPG or HTML files for its presentation, but the JPG or HTML files are not present, then the ASF video is not served.

```
<CdnManifest>
<item •src="http://my-origin-server/images/intro.html" />
<item •src="http://my-origin-server/images/intro.jpg" />

<!--
These are just two regular single item acquisition jobs.
-->

<item •src="http://my-origin-server/movie/movie1.asf" >
<contains •cdn-url="images/intro.html" />
<contains •cdn-url="images/intro.jpg" />
</item>

<!-- The preceding item, movie1.asf, contains two other items, intro.html and intro.jpg.
The items are contained using the <contains/> tag. If these two contained items are not
present on the Content Engine, then the ACNS software does not serve the container file
movie1.asf.
-->

</CdnManifest>
```

Sample 5

Sample 5 is a manifest file written to show how to specify live content. ACNS 5.x software supports two types of live content: wmt-live and real-live. You need to use the *type* attribute to specify live streaming content.

```
<CdnManifest>
<item •src="rtsp://www.company-web-site.org/tmp/ceo-talk" •type="wmt-live" />

<!-- This is a "wmt-live" streaming content type specified by the "type" attribute. The
live stream URL is rtsp://www.company-web-site.org/tmp/ceo-talk.
-->

<item •src="rtsp://real-server/tmp/funny-video" •type="real-live" />

<!-- This is "real-live" streaming content type specified by the "type" attribute. The
stream URL is rtsp://real-server/tmp/funny-video.
-->

</CdnManifest>
```

Downloading the Sample Manifest Files

To download the preceding five sample manifest files from Cisco.com, follow these steps:

-
- Step 1** Go to the following URL to find the sample files:
<http://www.cisco.com/cgi-bin/tablebuild.pl/acns52>
- Step 2** When prompted, log in to Cisco.com using your designated Cisco.com username and password.
The Cisco ACNS Software download page appears, listing the available software updates for the Cisco ACNS Software product.
- Step 3** Locate the file named **ACNS-5.2.1.5-manifest-samples.zip**. This file is a ZIP archive that contains the sample manifest files.
- Step 4** Click the link for the **ACNS-5.2.1.5-manifest-samples.zip** file. The download page appears.
- Step 5** Click **Software License Agreement**.
A new browser window opens, displaying the license agreement.
- Step 6** After you have read the license agreement, close the browser window that displays the agreement, and return to the Software Download page.
- Step 7** Click the filename link labeled **Download**.
- Step 8** To temporarily store the zipped file containing the sample files, click **Save to file**, and then choose a location on your workstation.
- Step 9** Use your preferred unzip program to unpack the scripts to a location on your workstation or your network.
After you have unzipped the sample files, you are ready to begin using them to create your own manifest files for your website.
-


Manifest Validator Utility

Because correct manifest file syntax is so important to the proper deployment of pre-positioned content on your ACNS network, Cisco makes available a manifest file syntax validator. The Manifest Validator, a Java-based command-line interface that verifies the correctness of the syntax of the manifest file you have written or modified, is built into the Content Distribution Manager GUI.

The Manifest Validator utility tests each line of the manifest file to identify syntax errors where they exist and determine whether or not the manifest file is valid and ready for use in importing content into your ACNS network. The results of these syntax validation tests are logged into a text file at a location that you name.

Running the Manifest Validator Utility

The Manifest Validator utility is built into the Content Distribution Manager GUI. [Figure A-1](#) shows the Manifest Validator GUI window.

Figure A-1 Manifest Validator Content Distribution Manager GUI Window

After you create a channel, you can access the Manifest Validator using the following GUI path:

Services > Web > Channels > Edit or Create Channel Icon > Tools > Manifest Validator

**Note**

You must first create a new channel or edit an exiting channel before you can access the Manifest Validator.

Under Validate Manifest File, enter the URL of the manifest file that you want to test in the Manifest File field and click Validate. The Manifest Validator checks the syntax of your manifest file to make sure that source files are named for each content item in the manifest. It then checks the URL for each content item to verify that the content is placed correctly and then displays the output in the lower part of the GUI window. The manifest file validator does not determine the size of the item.

Valid Manifest File Example

The following text is an example of a valid manifest file:

```
<CdnManifest>
<item
  src="tmp/mao's.html"
  priority="20"
/>
<server name="my-dev'box">
<host name="http://128.107.150.26"
  proto="http" />
</server>
```



```

<item
    src="tmp/lu.html"
    priority="300"
  />
<item
    src="/tmp/first_grader.html"
  />
<server name="server0">
  <host name="http://umark-u5.cisco.com:8080/" />
</server>
<item src="a.gif"/>
<server name="server1">
  <host name="http://unicorn-web" />
</server>
<item src="Media/wmtfiles/DCA%20Disk%201/Microsoft_Logos/Logos_100k.wmv" />

</CdnManifest>

```

The final lines of the manifest file validator output indicate whether the manifest is valid or not. Wait until the following message is displayed, indicating that the validator has completed processing the manifest file:

```

Total Number of Error: 0
Total Number of Warning: 0
Manifest File is CORRECT.

```

If errors are found, the error messages reported appear before the preceding message.

Invalid Manifest File Example

The following text is an example of an invalid manifest file:

```

<CdnManifest>
<item
    src="tmp/mao's.html"
    priority="20"
  />
<server name="my-dev'box">
<host name="http://128.107.150.26"
    proto="http" />
</server>
<item
    src="tmp/lu.html"
    priority="300"
  />
<item
    src="/tmp/first_grader.html"
  />
<server name="server0">
  <host name="http://umark-u5.cisco.com:8080/" >
</server>
<item src="a.gif"/>
<server name="server1">
  <host name="http://unicorn-web" />
</server>
<item src1="Media/wmtfiles/DCA%20Disk%201/Microsoft_Logos/Logos_100k.wmv" />
</CdnManifest>

```

In the preceding example, although there are no warnings, two errors are found, and this manifest file is syntactically incorrect, as shown in the following message:

```

ERROR (/state/dump/tmp.xml.1040667979990 line: 23 col: 1 ):No character data is allowed by
content model

```

```

ERROR (/state/dump/tmp.xml.1040667979990 line: 23 col: 9 ):Expected end of tag 'host'
Manifest File: /state/dump/tmp.xml.1040667979990
Total Number of Error: 2
Total Number of Warning: 0
Manifest File is NOT CORRECT!

```

The following full-text output is an example of the invalid manifest file after the Manifest Validator checks the file:

```

Manifest validated: http://qiwzhang-lnx/nfs-obsidian/Unicorn/my-single-bad.xml
The manifest is downloaded as /state/dump/tmp.xml.1040667979990 for validation, this file
will be removed when validation is completed.
Start CdnManifest
Start item
    priority=20
    src=tmp/mao's.html
End item

Start server
    name=my-dev'box
Start host
    name=http://128.107.150.26
    proto=http
    uuencoded=false
End host

End server

Start item
    priority=300
    src=tmp/lu.html
End item

Start item
    src=/tmp/first_grader.html
End item

Start server
    name=server0
Start host
    name=http://umark-u5.cisco.com:8080/
    uuencoded=false
ERROR (/state/dump/tmp.xml.1040667979990 line: 23 col: 1 ):No character data is allowed by
content model
ERROR (/state/dump/tmp.xml.1040667979990 line: 23 col: 9 ):Expected end of tag 'host'
Manifest File: /state/dump/tmp.xml.1040667979990
Total Number of Error: 2
Total Number of Warning: 0
Manifest File is NOT CORRECT!

```

Understanding Manifest File Validator Output

The manifest file validator messages appear below the Manifest File field in the Manifest Validator window of the Content Distribution Manager GUI.

Each output file has a similar structure and syntax. It clearly identifies any errors or warning messages arising from incorrect manifest file syntax. Manifest files are determined by the validator to be either:

- **CORRECT**—Contains possible syntax irregularities but is syntactically valid and ready for deployment on your ACNS network
- **INCORRECT**—Contains syntax errors and is unsuitable for deployment on your ACNS network

Syntax Errors

The manifest file validator issues syntax errors only when it cannot identify a source file for a listed content item, either because it is not listed or because it is listed using improper syntax. Files containing syntax errors are marked *INCORRECT*.

Syntax errors are identified in the output with the *ERROR* label. In addition to the label, the line and column number containing the error are provided, as well as the manifest attribute for which the error was issued. An error appears in the following example:

```
ERROR (/state/dump/tmp.xml.1040667979990 line: 23 col: 1 ):No character data is allowed by content model
```

where:

- `/state/dump/tmp.xml.1040667979990` is the manifest file name
- `line: 23 col: 1` is the manifest file line and column number where the error occurs
- `No character data is allowed by content model` describes the type of manifest file error

Syntax Warnings

The manifest file validator issues syntax warnings for a wide variety of irregularities in the manifest file syntax. Files containing syntax warnings may be marked *CORRECT* or *INCORRECT*, depending on whether or not syntax errors have also been issued.

Syntax warnings are identified in the output with the *WARNING* label. In addition to this warning label, the line number for which the warning is issued is provided, as well as the manifest attribute, valid options, and the default value for that attribute for which the warning was issued.

Correcting Manifest File Syntax

Once you have identified syntax warnings, errors, and messages using the output from the manifest file validator, you can correct your manifest file syntax and then rerun the manifest file validator on the corrected file to verify its correctness.

To correct syntax warnings and errors in your manifest file, follow these steps:

-
- Step 1** Open your manifest file using your preferred XML editor.
 - Step 2** Referring to your manifest file validator output, use the line numbers provided by the manifest file validator to locate the syntax violations in your manifest file.

It is a good idea to review every warning and error in your manifest file. Some warnings, although they still allow the manifest file validator to find your manifest file syntax to be correct, can be the source of problems when you deploy the identified content to your ACNS network.
 - Step 3** After you have made the necessary corrections for syntax warnings and errors, click **Save**.
 - Step 4** Process the manifest file through the manifest file validator again and review the validator output for new or unresolved errors and warnings.
 - Step 5** Repeat [Step 1](#) through [Step 4](#) until every error and warning has been adequately resolved and the manifest validator indicates that your manifest file syntax is correct.
-

Manifest File Reference

This major section contains the following topics:

- [Manifest File Structure and Syntax](#)
- [XML Schema](#)
- [Manifest File Automated Scripts](#)
- [Manifest File Time Zone Tables](#)

The most efficient and least error-prone methods of creating a manifest file are:

- Modify one of the sample manifest files in this chapter to suit your particular needs, ensuring that your XML syntax is correct.
- Use the two sample Perl scripts that can be downloaded from Cisco.com as is, or customize these downloaded scripts for your own purposes.

You can start with one of the prewritten sample XML manifest files presented in this chapter. Choose a sample manifest file that is closest to matching your content acquisition and pre-positioning needs, and then modify the XML code accordingly, while ensuring that your XML syntax is correct.

Alternatively, use the sample Perl scripts that Cisco provides (see the [“Obtaining the Perl Scripts” section on page A-79](#)).

Once you have created a suitable manifest file, you can verify its correctness by running the Manifest Validator utility (see the [“Manifest Validator Utility” section on page A-28](#)) on your newly written XML code from the Content Distribution Manager GUI.

Manifest File Structure and Syntax

The Cisco ACNS 5.x software manifest file provides powerful features for representing and manipulating ACNS network data that can be easily edited using any simple text editor. [Table A-4](#) provides a summary list of the manifest file tags, their corresponding attributes and subelements, and a brief description of each tag. [Table A-5](#) shows an example of how tags are nested in a manifest file. The sections that follow provide a more detailed description of the manifest file tags, the data they contain, and their attributes.

Table A-4 *Manifest File Tag Summary*

Tag Name	Subelements	Attributes		Description
CdnManifest	<playServerTable/> <options/> <server/> <item/> <item-group/> <crawler/>	None		Marks the beginning and end of the manifest file content.
playServerTable	<playServer/>	None		(Optional) Sets default mappings for media types.
playServer	<contentType/> <extension/>	<i>name</i> ¹		Names the media server type on the Content Engine responsible for playing content types and files with extensions mapped to it using <contentType> tags.

Table A-4 Manifest File Tag Summary (continued)

Tag Name	Subelements	Attributes		Description
contentType	None	name		(Optional, but must have either <content-type> or <extension> tag.) Names the MIME-type content mapped to a playserver.
extension	None	name		(Optional, but must have either <content-type> or <extension> tag.) Names the file extension that is mapped to a playserver.
options	<schedule><repeat>	<i>alternateUrl</i> <i>bitrate-in-bps</i> <i>bitrate-in-kbps</i> <i>bitrate-in-mbps</i> <i>discardAllHttpHeaders</i> <i>discardHttpHeaders</i> <i>enableCookies</i> <i>expires</i> <i>failRetryInterval</i> <i>ignoreOriginPort</i> <i>ignoreQueryString</i>	<i>keepHttpHeaders</i> <i>noRedirectToOrigin</i> <i>notFoundURL</i> <i>playDuration</i> <i>prefetch</i> <i>priority</i> <i>requireAuth</i> <i>server</i> <i>sslAuthType</i> <i>timeZone</i> <i>ttn</i> <i>type</i>	(Optional) Defines attributes specific to the manifest file that can be shared.
server	<host/>	name		Defines <i>only</i> one host from which content is to be retrieved.
host	None	name <i>disableBasicAuth</i> <i>noProxy</i> <i>ntlmUserDomain</i> <i>password</i> <i>port</i> <i>proto</i>	<i>proxyServer</i> <i>sslAuthType</i> <i>user</i> <i>userDomainName</i> <i>uuencoded</i>	Defines a web server or live server from which content is to be retrieved and later pre-positioned. The host name can be specified as this: proto://user:password@hostname:port
proxyServer	None	serverName <i>disableBasicAuth</i> <i>ntlmUserDomain</i> <i>password</i>	<i>port</i> <i>user</i> <i>uuencoded</i>	Specifies proxy server information.

Table A-4 Manifest File Tag Summary (continued)

Tag Name	Subelements	Attributes		Description
item	<code><contains/></code> <code><http-meta-data/></code> <code><schedule><repeat></code>	src <i>alternateUrl</i> <i>authCookie</i> <i>bitrate-in-bps</i> <i>bitrate-in-kbps</i> <i>bitrate-in-mbps</i> <i>cdn-url</i> <i>disableBasicAuth</i> <i>discardAllHttpHeaders</i> <i>discardHttpHeaders</i> <i>enableCookies</i> <i>expires</i> <i>failRetryInterval</i> <i>host</i> <i>ignoreOriginPort</i> <i>ignoreQueryString</i> <i>keepHttpHeaders</i> <i>noProxy</i> <i>noRedirectToOrigin</i> <i>ntlmUserDomain</i>	<i>password</i> <i>playDuration</i> <i>playServer</i> <i>port</i> <i>prefetch</i> <i>priority</i> <i>proto</i> <i>proxyServer</i> <i>requireAuth</i> <i>server</i> <i>serveStartTime</i> <i>serveStopTime</i> <i>sslAuthType</i> <i>ttnl</i> <i>type</i> <i>user</i> <i>userDomainName</i> <i>uuencoded</i>	Identifies specific content that is to be acquired from the origin server.
crawler	<code><http-meta-data/></code> <code><matchRule/></code> <code><schedule><repeat></code>	start-url <i>accept</i> <i>alternateUrl</i> <i>authCookie</i> <i>bitrate-in-bps</i> <i>bitrate-in-kbps</i> <i>bitrate-in-mbps</i> <i>cdnPrefix</i> <i>depth</i> <i>disableBasicAuth</i> <i>discardAllHttpHeaders</i> <i>discardHttpHeaders</i> <i>enableCookies</i> <i>expires</i> <i>externalPrefixes</i> <i>externalServers</i> <i>failRetryInterval</i> <i>host</i> <i>ignoreOriginPort</i> <i>ignoreQueryString</i> <i>keepExpiredContent</i> <i>keepFolder</i> <i>keepHttpHeaders</i> <i>keepNoCacheContent</i> <i>keepQueryUrl</i> <i>max-number</i> <i>max-size-in-MB</i>	<i>noProxy</i> <i>noRedirectToOrigin</i> <i>ntlmUserDomain</i> <i>password</i> <i>playDuration</i> <i>playServer</i> <i>port</i> <i>prefetch</i> <i>prefix</i> <i>priority</i> <i>proto</i> <i>proxyServer</i> <i>reject</i> <i>reportBrokenLinks</i> <i>requireAuth</i> <i>serveStartTime</i> <i>serveStopTime</i> <i>server</i> <i>srcPrefix</i> <i>sslAuthType</i> <i>ttnl</i> <i>type</i> <i>user</i> <i>userDomainName</i> <i>uuencoded</i>	Supports crawling of a website or FTP server.

Table A-4 Manifest File Tag Summary (continued)

Tag Name	Subelements	Attributes		Description
item-group	<item/> <crawler/> <item-group/> <http-meta-data/> <schedule><repeat>	<i>alternateUrl</i> <i>bitrate-in-bps</i> <i>bitrate-in-kbps</i> <i>bitrate-in-mbps</i> <i>cdnPrefix</i> <i>cdn-url</i> <i>disableBasicAuth</i> <i>discardAllHttpHeaders</i> <i>discardHttpHeaders</i> <i>enableCookies</i> <i>expires</i> <i>failRetryInterval</i> <i>host</i> <i>ignoreOriginPort</i> <i>ignoreQueryString</i> <i>keepHttpHeaders</i> <i>noProxy</i> <i>noRedirectToOrigin</i> <i>ntlmUserDomain</i>	<i>password</i> <i>playDuration</i> <i>playServerport</i> <i>prefetch</i> <i>priority</i> <i>proto</i> <i>proxyServer</i> <i>requireAuth</i> <i>serveStartTime</i> <i>serveStopTime</i> <i>server</i> <i>src</i> <i>srcPrefix</i> <i>sslAuthType</i> <i>ttl</i> <i>type</i> <i>user</i> <i>userDomainName</i> <i>uuencoded</i>	Places shared attributes under one tag so that they can be shared by every <item> and <crawler> tag within that group.
matchRule	<match>	None		(Optional) Defines additional filter rules for crawler jobs.
match	None	<i>extension</i> <i>mime-type</i> <i>prefix</i> <i>size-min-in-B</i> <i>size-min-in-KB</i> <i>size-min-in-MB</i>	<i>size-max-in-B</i> <i>size-max-in-KB</i> <i>size-max-in-MB</i> <i>time-after</i> <i>time-before</i> <i>url-pattern</i>	(Optional) Specifies the acquisition criteria of content objects before they can be acquired by the ACNS network.
contains	None	<i>cdn-url</i>		(Optional) Identifies content objects that are embedded within the content item currently being described.
http-meta-data	None	name=value		(Optional) Sends HTTP response headers to end user HTTP requests to specify content type for FTP acquired content.

1. Attributes that are required for a tag are shown in **boldface** font.

Table A-5 Manifest File Nested Tag Relationships

<CdnManifest>				
	<playServerTable> <playServer>			
		<contentType/> <extension/>		
			</playServerTable> </playServer>	
	<options>			
		Manifest file shared attributes	</options>	
	<server>			
		<host/>		
			</server>	
	<item>			
		<contains/> <http-meta-data/>		
			</item>	
	<crawler>			
		<http-meta-data/> <matchRule/>		
			</crawler>	
	<item-group>			
		<contains/> <http-meta-data/>		
			</item-group>	
				</CdnManifest>

CdnManifest

The <CdnManifest> </CdnManifest> tag set is required and marks the beginning and end of the manifest file content. At a minimum, each <CdnManifest> tag set must contain at least one item, or content object, that is fetched and stored.

Attributes

None

Subelements

The <CdnManifest> tag set can contain the following subelements:

- playServerTable

The <CdnManifest> tag set can only contain one playServerTable subelement.

- options

The <CdnManifest> tag set can only contain one options subelement.

- server
- item
- item-group
- crawler

Example

```
<CdnManifest>
  <server name="origin-server">
    <host name="www.name.com" proto="http" port="80" />
  </server>
  <item cdn-url= "logo.jpg" server="originserver" src= "images/img.jpg" type="prepos"
    playServer="http" ttl="300"/>
</CdnManifest>
```

playServerTable

The <playServerTable> </playServerTable> tag set is optional and provides a means for you to set default mappings for a variety of media types. Mappings can be set for both MIME-type content (the preferred mapping) and file extensions. Playserver tables allow you to override default mappings on the Content Engine for content types from a particular origin server. Playservers can be any one of the following five streaming servers: WMT, RealMedia, HTTP, QTSS, or TV-out. If no <playServerTable> tag is configured in the manifest file, a default <playServerTable> tag is used.

Using the manifest file, you can map groups of single items as well as individual content objects to an installed playserver. The following are content item and manifest file playserver mappings:

- Content item URL

Playserver mappings appear immediately after the origin server name in place of the default <cdn-media> tag.

- Manifest file as an attribute of the <item> or <item-group> tag

Playserver mappings placed at this location are identified using the *playserver* attribute and only apply to the named item or group of items.

- Manifest file as a playserver table

Mappings are grouped within the <playServerTable> and <playServer> tags and are applied to content served from the origin server as directed by the manifest file.

- System-level

Playserver mappings are configured during ACNS software startup.

The <playServerTable> tags are enclosed within the <CdnManifest> tags and name at least one of four playservers, such as RealServer, to which certain MIME-types and file extensions are mapped.

Attributes

None

Subelements

The <playServerTable> element must contain at least one <playServer> tag.

playServer

The `<playServer>` `</playServer>` tag set is required for the `<playServerTable>` tag and names the media server type on the Content Engine that is responsible for playing the content types and files with extensions mapped to it using the `<contentType>` tags. The `<playServer>` tag is enclosed within `<playServerTable>` tags.

**Note**

Do not confuse the `<playServer>` tag with the `playserver` setting in an `<item>` or `<item-group>` tag. An `<item>` or `<item-group>` tag specifies a server type to be used for an individual content object or group of related content objects. Although both `playserver` settings accomplish the same task, `<item>` tag-level `playserver` settings take precedence over the content type and file extension mappings specified by the `<playServer>` tags in the `<playServerTable>` tag.

Attributes

The `<playServer>` tag name is required. Each `<playServer>` tag names the type of server to which content is mapped using the `name` attribute. In ACNS 5.x software, Content Engines support five types of `playserver`s:

- `real`: RealMedia RealServer
- `http`: HTTP web server
- `https`: HTTPS web server
- `qtss`: Apple QuickTime Streaming Server
- `wmt`: Microsoft Windows Media Technologies
- `tvout`: Excludes other streaming services from playing the pre-positioned files

Subelements

At least one of the following subelements must be present in a `<playServer>` tag set.

- `<contentType/>`
- `<extension/>`

contentType

The `<contentType/>` tag is optional but either a `<contentType/>` or an `<extension/>` subelement must be present in a `<playServer>` tag set. The `<contentType/>` tag names MIME-type content that is to be mapped to a `playserver`. The `<contentType/>` tag must be enclosed within a `<playServer>` tag set. When both `<contentType/>` and `<extension/>` tags are present in a `<PlayServerTable>` tag for a particular media type, the `<contentType/>` mapping takes precedence.

Attributes

Each `<contentType/>` tag names a media content type that is to be mapped to the `playserver` using the `name` attribute. The `name` attribute is required.

Subelements

None

extension

The `<extension/>` tag is optional but either a `<contentType/>` or an `<extension/>` subelement must be present in a `<playServer>` tag set. The `<extension/>` tag names the file extension that is being mapped to a playserver.

The `<extension/>` tag follows the `<playServer>` tag. When both `<contentType/>` and `<extension/>` tags are present in the `<playServer>` tag for a particular media type, the `<contentType/>` mapping takes precedence.

Attributes

The *name* attribute is required and provides the file extension for a mapped content type. When files with the named extension are requested, the mapped playserver is used to serve them.

Subelements

None

Example

```
<CdnManifest>
<playServerTable>
  <playServer name="real">
    <contentType name="application/x-pn-realaudio" />
    <contentType name="application/vnd.rn-rmadriver" />
    <extension name="rm" />
    <extension name="ra" />
    <extension name="rp" />
    <extension name="rt" />
    <extension name="smi" />
  </playServer>
  <playServer name="wmt">
    <extension name="asx" />
    <extension name="asf" />
    <extension name="avi" />
  </playServer>
  <playServer name="http">
    <contentType name="application/pdf" />
    <contentType name="application/postscript" />
    <extension name="pdf" />
    <extension name="ps" />
  </playServer>
</playServerTable>
<server name="test.origin.com/">
  <host name="http://tst.orgn.com" proto="http" />
</server>
<item
  src="pic1.mpg"
/>
</CdnManifest>
```

options

The `<options/>` tag is optional and used to define attributes specific to the manifest file. Shared attributes can be inherited by `<item>` and `<crawler>` tags in the manifest file. For example, *timeZone* is an attribute specific to the manifest file that is used to set the time zone for all time-related values. Attributes such as *ttl* and *alternativeUrl* can exist as `<options/>` tags, and their values can be shared by all `<item>` and `<crawler>` tags within the manifest file.

The `<options/>` tag set is enclosed within the `<CdnManifest>` tag set and specifies at least one global setting. No more than one `<options>` tag is allowed per manifest file.

If parameters are defined within the manifest file `<options/>`, `<item-group>`, or `<item>` tags, the order of precedence from lowest to highest is `<options/>`, `<item-group>`, and `<item>`.

Attributes

The `timeZone` attribute specifies the time zone for time values of attributes such as `expire` and `prefetch`.

The `notFoundURL` attribute is optional. This attribute is similar to the `alternateUrl` attribute. If content requested by the user is not ready in the ACNS network, the ACNS software redirects the request to this alternative URL, which can be configured as an error reporting page. The `notFoundURL` attribute supports only the full URL.

The following list of attributes can be shared by `<item>` and `<crawler>` tags. See the [“item” section on page A-44](#) for descriptions of the following attributes:

- `alternateUrl`
- `bitrate-in-bps/bitrate-in-kbps/bitrate-in-mbps`
- `discardAllHttpHeaders`
- `discardHttpHeaders`
- `enableCookies`
- `expires`
- `failRetryInterval`
- `ignoreOriginPort`
- `ignoreQueryString`
- `keepHttpHeaders`
- `noRedirectToOrigin`
- `playDuration`
- `prefetch`
- `priority`
- `requireAuth`
- `server`
- `sslAuthType`
- `ttl`
- `type`

Subelements

`<schedule><repeat>`

(See the [“item” section on page A-44](#) for descriptions of these subelements.)

Example

```
<CdnManifest>
<options
  noRedirectToOrigin= "true"
  timeZone="PST" />

</CdnManifest>
```

server

The <server> and <host> tag fields configure the origin content source server. The <host> tag field inside the <server> tag field configures the content source host. Having multiple <host> tag fields in one <server> tag field is not supported in ACNS 5.x software.

Each <item> or <item-group> tag can have a *server* attribute that refers to this <server> tag field. The <server> </server> tag set is required and defines only one host from which content is to be retrieved. The <server> tags are contained within <CdnManifest> tags and contain one <host> tag that identifies the host from which content is retrieved.

Attributes

name

The *name* attribute is required and can be any name as long as it matches the *server* attribute values in the <item> or <crawler> tags.

Subelements

<host/>

The <server> tag set can only contain one <host/> subelement.

host

The <host/> tag is required and defines a web server or live server from which content is to be retrieved and later pre-positioned. Only one host can be defined within a single <server> tag set. The <host/> tag must be enclosed within <server> tags.

Attributes

- *disableBasicAuth*

The *disableBasicAuth* attribute is optional; if specified, basic authentication is disabled.

- *name*

The *name* attribute is required and identifies the domain name or IP address of the host, unless the *proto* attribute field is empty. If the *proto* attribute field is empty, the *name* attribute must be a fully qualified URL, including scheme and domain name or IP address. It can also include subdirectories, such as `http://www.abc.com/media`.

The *name* attribute can also contain the UNC path to an SMB server. For example, `\\SMBserver\directory\`.

- *noProxy*

The *noProxy* attribute is optional. If set to true, no proxy is used for the origin server. The default is false.

- *ntlmUserDomain*

The *ntlmUserDomain* attribute is optional and specifies the user domain name for NTLM authentication.

- *password*

The *password* attribute is optional and identifies the password for the user account that is required to access the host server.

- *port*

The *port* attribute is optional and identifies the TCP port through which traffic to and from the host passes. The port used depends on the protocol used. The default port for HTTP is 80. The *port* attribute is only required for a nonstandard port assignment. The port attribute can also be specified in the *name* attribute, such as `name="http://www.cisco.com:8080/."`

- *proto*

The *proto* attribute is optional and identifies the communication protocol that is used to fetch content from the host. Supported protocols are HTTP, HTTPS, MMS-over-HTTP, or FTP. The default *proto* attribute is HTTP. The *proto* attribute can be empty if the *name* attribute is a fully qualified domain name (FQDN).

- *proxyServer*

The *proxyServer* attribute is optional and specifies which proxy server to use if there are multiple `<proxyServer>` tags in the manifest file. If no proxy server is specified, the server in the closest `<proxyServer>` tag is used.

- *sslAuthType*

The *sslAuthType* attribute is optional and has two possible values for the type of SSL certificate verification:

- *strong*—Strong authentication. If any errors occur during certificate verification by the acquirer module, content from that site will not be acquired. The default *sslAuthType* attribute setting is strong.
- *weak*—Weak authentication. If certain errors occur during certificate verification by the acquirer module, content from that site will continue to be acquired. These errors are as follows:

Unable to decode issuer's public key

Certificate has expired

Self-signed certificate

Self-signed certificate in certificate chain

Unable to get local issuer certificate

Subject issuer mismatch

Authority and issuer serial number mismatch

The root Content Engine is not marked as trusted

Unable to verify the first certificate

Certificate is not yet valid

Certificate has invalid purpose

- *user*

The *user* attribute is optional and identifies the secure login used for host access.

- *userDomainName*

See the “[item](#)” section on page A-44 for a description of this attribute.

- *uuencoded*

The *uuencoded* attribute is optional. If set to true, the password is not encoded. The *uuencoded* attribute default setting is false.

Subelements

None

proxyServer

The <proxyServer> tag specifies proxy server information. This tag is supported only in ACNS 5.1 and later software releases. The <proxyServer> tag must be located at the top level of the manifest file, directly under the <CdnManifest> tag; it cannot be used as a subtag of any other tags, as shown in this example:

```
<CdnManifest>
<proxyServer>
...
</CdnManifest>
</proxyServer>
```

Attributes

- *disableBasicAuth*

The *disableBasicAuth* attribute is optional; if specified, basic authentication is disabled.

- *ntlmUserDomain*

The *ntlmUserDomain* attribute is optional and specifies the user domain name for NTLM authentication.

- *password*

The *password* attribute is optional and identifies the password for the user account that is required to access the proxy server.

- *port*

The *port* attribute is optional and specifies the proxy port.

- *serverName*

The *serverName* attribute is required and identifies the domain name or IP address of the proxy server.

- *user*

The *user* attribute is optional and identifies the secure login used for proxy authentication.

- *uuencoded*

The *uuencoded* attribute is optional and designates whether the password is to be encoded.

Subelements

None

item

The <item> </item> tag set identifies the specific content that is to be acquired. The <item> tag names a single piece of content or a content object on the origin server, such as a graphic, MPEG video, or RealAudio sound file. Content items can be listed individually or grouped using the <item-group> tag.

The `<item>` tag must be enclosed within the `<CdnManifest>` tag set and can also be enclosed within `<item-group>` tags.

Attributes

- *src*

The *src* attribute is required and identifies the URL from which to fetch the content. The URL can be a full URL or a relative URL. A full URL has the following format:

proto://username:password@/domain-name:port/file-path/file-name

Protocols supported in item *src* are HTTP, HTTPS, MMS-over-HTTP, FTP, SMB, and file. For SMB, the URL must be written in UNC format (\\SMBserver\directory\file).

If a relative path is used, the `<server>` and `<host>` tags are required to specify origin server information, as shown in this example:

```
<item src="http://user:password@www.cisco.com/HR/index.html" />
<server name="ftp-server" >
  <host name="ftp://ftp-server" user="johw" password="www" />
</host>
<item src="data/video.asf" />
```



Note A URL containing a question mark (?) is not supported. A manifest parsing error will occur if you specify a URL that contains a question mark.



Note A URL containing a pound sign (#) will be modified. All characters that follow a pound sign will be discarded, including the pound sign itself.

- *host*

Specify the host name if the source URL of the *src* attribute is a relative URL.

- *server*

The *server* attribute is optional and refers to the server name in the `<server>` tag. If the *server* attribute is omitted, the server listed in the closest `<server>` tag is used. If there is no `<server>` tag close to this item, the manifest server is used.

- *cdn-url*

The *cdn-url* attribute is optional and is used when content needs to be acquired from one URL (the content acquisition URL) and published using another URL (the publishing URL). The *cdn-url* attribute is the relative ACNS network URL that end users use to access this content. If no *cdn-url* attribute is specified, then the *src* attribute is used as the relative ACNS network URL.

In the following sample manifest file, the content item being acquired contains the file path /RemAdmin/InternalReview/firstpage.htm. By specifying a new file path (RemAdmin/Production/firstpage.htm) using the *cdn-url* attribute, the publishing URL disguises the fact that the content originated from an internal review.

```
<CdnManifest>
<server name="ultra-server">
  <host name="http://ultra-server" />
</server>
<item src="RemAdmin/InternalReview/firstpage.htm"
cdn-url="RemAdmin/Production/firstpage.htm" />
</CdnManifest>
```


In the preceding example, *src* is the content acquisition URL and *cdn-url* is the publishing URL.



Note The content item file path (RemAdmin/InternalReview/firstpage.htm) is controlled by the manifest file. The *cdn-url* attribute associates a file path with the content item in the manifest file. The manifest file allows the file path for the *cdn-url* attribute to be specified independently of the file path from which the content items are to be acquired from the origin server (*src* attribute), allowing the publishing URL to differ from the content acquisition URL. (See the [“Generate the Publishing URL” section on page 6-30.](#))

If the content requires playback authentication or is live content, the origin server from which the content is acquired has to be contacted. Therefore, two URLs must exist for the same content item, and the URL specified in the *cdn-url* attribute must exist on the origin server at all times.

For example, if the content item “RemAdmin/Production/firstpage.htm” in the preceding example requires playback authentication, this content must exist on the “ultra-server” origin server. Otherwise, pre-positioned content playback will fail.

In general, you should not use the *cdn-url*, *cdnPrefix*, or *srcPrefix* attributes if playback authentication is required or if the content is live.

If you use FTP to acquire content and the content type is not specified in the manifest file and the *cdn-url* attribute is specified to alter your publishing URL, the *cdn-url* attribute must have the correct file path extension. Otherwise, the incorrect content type will be generated and you cannot play the content.

The following example correctly shows the publishing URL with the same file path extension (.jpg) as the origin server URL.

```
<item src="ftp://ftp-server.abc.com/pictures/pic.jpg" cdn-url="pic.jpg" />
```

The following example is incorrectly written, because it does not specify the file path extension (.jpg) in the *cdn-url* attribute.

```
<item src="ftp://ftp-server.abc.com/pictures/pic.jpg" cdn-url="pic" />
```

- *type*

The *type* attribute is optional and defines whether content is to be pre-positioned or live on the ACNS network. The three type attributes are *prepos*, *wmt-live*, and *real-live*. The *wmt-live* and *real-live* attributes are used to deliver live content. If this field is left blank, the default type is *prepos*.

- *noRedirectToOrigin*

The *noRedirectToOrigin* attribute is optional and sets the redirection to the origin server to true or false. A false setting allows the ACNS Content Engine or other edge device to redirect content requests to the origin server if the content is available at that device. A true setting does not allow the ACNS Content Engine or edge device to redirect content requests to the origin server, and it generates an error. The default *noRedirectToOrigin* setting is false. This attribute is a per-content attribute, meaning that if the content has been removed, the redirection settings do not apply. But if the software merely fails to acquire the content, then the settings apply.

(For the effect of the *noRedirectToOrigin* attribute on pre-positioned content freshness, see the [“Configuring the Freshness of Pre-Positioned Content” section on page A-64.](#))

- *playServer*

The *playServer* attribute is optional and names the server used to play back the content. Valid playervers are real (RealServer), wmt (Windows Media Technologies), qtss (QuickTime Streaming Server), http and https (web servers), and tvout. The value in this field is either one playserver or multiple playervers separated by commas. If a value for this attribute is not specified, the

<PlayServerTable> tag in the manifest file is used to generate the playserver list for this content. If the manifest file does not have the <PlayServerTable> tag specified, it uses the default <PlayServerTable> tag.

- *prefetch*

The *prefetch* attribute is optional and specifies a time (in yyyy-mm-dd hh:mm:ss [year-month-day hour:minute:second] format) for the first content acquisition or re-check after the manifest file is parsed. The time zone for the time can be specified in the <options> tag. Note that the autoconversion between daylight saving time and standard time within a time zone is not supported, but a special designation for daylight saving time can be used, such as PDT for Pacific daylight saving time. In the following example, the prefetch time is September 5, 2002 at 09:09:09 Pacific daylight saving time.

```
<options timeZone="PDT" />
<item src="index.html" prefetch="2002-09-05 09:09:09 PDT" />
```

This attribute is used when you want to specify a future time for the acquirer to begin fetching content from the origin server. When a future time is specified, the acquirer will not acquire content before this time; however, it will check content freshness during its scheduled *ttl* interval. If a *prefetch* time is omitted, the content is acquired immediately.

This attribute behaves as follows:

- ACNS 5.1—After the manifest file is parsed, if the *prefetch* attribute specifies a future time, regardless whether items or crawl tasks have changed or new ones have been added, the acquirer checks and fetches the content or re-crawls the crawl jobs at the time specified by the *prefetch* attribute.
- ACNS 5.2—After the manifest file is parsed, if any items or crawl tasks have changed or new ones have been added and if the *prefetch* attribute specifies a future time, the acquirer checks and fetches the content or re-crawls the crawl jobs at the time specified by the *prefetch* attribute.

- *expires*

The *expires* attribute is optional and designates a time in yyyy-mm-dd hh:mm:ss format when the content is to be removed from the ACNS network. Additionally, you can specify the GMT time zone. (See the [“Specifying Time Values in the Manifest File” section on page A-18.](#)) If a time value is omitted, content is stored until it is removed when you modify the relevant manifest file code. (For the effect of the HTTP header *expires* attribute on pre-positioned content freshness, see the [“Configuring the Freshness of Pre-Positioned Content” section on page A-64.](#))

- *ttl*

The *ttl* attribute is optional and designates a time interval, in minutes, for revalidation of the content. If a time value is omitted, the content is fetched only once and its freshness is never checked again. Usually the *ttl* attribute is a positive value; however, you can also assign a negative value to the *ttl* attribute. The following list describes *ttl* attribute value ranges.

<i>ttl</i> Attribute Value	Action
<i>ttl</i> > 0	Content is rechecked every <i>ttl</i> minute. Content is also rechecked if the manifest is reparsed and the content specification in the manifest file has changed or if you click the Refetch button.

<i>ttl</i> Attribute Value	Action
<i>ttl</i> = 0	Content is fetched only once and never checked again. Content is only rechecked if the manifest is reparsed and the content specification in the manifest file has changed or if you click the Refetch button.
<i>ttl</i> < 0	Content is fetched only once and never checked again. Content will <i>not</i> be rechecked if the manifest is reparsed or if you click the Refetch button.

- *serveStartTime*

The *serveStartTime* attribute is optional and designates a time in yyyy-mm-dd hh:mm:ss format when the ACNS software is allowed to start serving the content. If the time to serve is omitted, content is ready to serve once it is distributed to the Content Engine or other edge device.

- *serveStopTime*

The *serveStopTime* attribute is optional and designates a time in yyyy-mm-dd hh:mm:ss format when the ACNS software temporarily stops serving the content. If the time to stop serving is omitted, the ACNS software serves the content until it is removed when you modify the relevant manifest file code. For the effect of the *serveStopTime* attribute on pre-positioned content freshness, see the “[Configuring the Freshness of Pre-Positioned Content](#)” section on page A-64.

- *alternateUrl*

The *alternateUrl* attribute is optional. If content requested by the user is not ready in the ACNS network, the ACNS software redirects the request to this alternative URL, which can be configured as an error reporting page. The *alternateUrl* attribute supports only the full URL.

This attribute applies to the content being requested. If the content is deleted, the attribute does not apply. But if the content failed to be acquired, then this attribute should apply.

- *priority*

The *priority* attribute is optional and can be any integer value to specify the content processing priority. If a priority value is omitted, its index order within the manifest file is used to set the priority.

- *requireAuth*

The *requireAuth* attribute is optional and determines whether users need to be authenticated before the specified content is played. When true, the Content Engine requires authentication to play back the specified content to users and communicates with the origin server to check credentials. If the requests pass the credential check, the content is played back from the Content Engine. If this attribute is omitted, a heuristic approach is used to determine the value: if the specified content is acquired by using a username and password, *requireAuth* is set to true; otherwise, it is set to false. For FTP, if the username is anonymous, *requireAuth* is set to false.



Note

If *requireAuth* is true, the Origin Server field in the Website window for this channel needs to point to the server that can authenticate users. When users want to play back the content, the server specified in the Origin Server field of the Website window is checked for authentication.

- *playDuration*

The *playDuration* attribute specifies the playtime duration, in seconds, for a video file. This value is used to:

- Download the video file and play it over HTTP.

If the video file is played using HTTP, the ACNS software uses the *playDuration* attribute to calculate the HTTP download bit rate.

- Schedule a video TV-out program.

If the time duration value is omitted, the ACNS software attempts to calculate this value for an MPEG file.

- *failRetryInterval*

This value specifies the retry interval, in minutes, when content acquisition fails. For example, *failRetryInterval*="10" means the ACNS software retries content acquisition every 10 minutes after acquisition has failed. If the retry universal value is not specified, the default value is 5 minutes. (The minimum *failRetryInterval* value accepted.) If a value of less than 5 minutes is specified, that value is converted to 5 minutes.

The behavior differs between failed content acquisition of a single item and failed content acquisition of a crawl item.

- For single item failure:

```
if ( ttl != 0, ttl < retryInterval)
```

The item is rechecked in accordance with the *ttl* attribute. Otherwise, the item is rechecked at the interval specified in the *failRetryInterval* attribute.

- For crawl item failure:

If some items are not acquired (excluding 300 and 400 series status error codes), only failed items are rechecked as specified in the *failRetryInterval* attribute.

When the *ttl* attribute interval occurs, all pages are recrawled.

```
if ( ttl != 0 and ttl < retryInterval )
always re-crawl
```

For example, if *ttl* = 10, and *failRetryInterval* = 4, the following will occur:

Number of Minutes	Action
0	Crawl
4	Recheck failed
8	Recheck failed
10	Recrawl
14	Recheck
18	Recheck
20	Recrawl

- *bitrate-in-bps/bitrate-in-kbps/bitrate-in-mbps*

This value specifies the bit rate of the content. ACNS software uses this value in playback download. In the following example, ACNS software downloads the file at 500 kbps if it is requested for download and play:

```
<item src="http://www.cisco.com/Prod/ACNS.swf" bitrate-in-kbps="500" />
```

- *ignoreQueryString*

This is a playback attribute that can be used with the <options>, <item-group>, <item>, and <crawler> tags. If the value is set to true, then ACNS software ignores any string after a question mark (?) in the request URL for playback. If this attribute is omitted, then the default value is false.

For example, content with the request URL `url=http://web-server/foo` has been pre-positioned. If a user requests content with the URL `url=http://web-server/foo?id=xxx` and the *ignoreQueryString* attribute value is false, then ACNS software does not use the pre-positioned content from the request URL `http://web-server/foo`.

However, if the *ignoreQueryString* attribute is set to true, then the ACNS software treats the request URL `http://www-server/foo?id=xxx` the same as `http://www-server/foo` and returns with pre-positioned content.

- *ignoreOriginPort*

This attribute allows playback of pre-positioned content from a port other than the standard port. If the *ignoreOriginPort* attribute is set to true, content can be played back without regard to the port specified in the request URL. The default for this attribute is false.

This attribute is not intended to be used for content that is routed using a Content Router. It is intended to work only for transparent WCCP routing or explicit proxy routing. A typical usage scenario for the *ignoreOriginPort* attribute might be as follows:

- The origin web server is not using port 80; it is using a non-standard port number in the URL.
- Users are using WCCP or explicit proxy routing, where the original URL containing the non-standard port number is used for playback from the Content Engine.

Pre-positioned content cannot be played back using a non-standard port; pre-positioned content is served only on ports that are standard for the protocol. If the incoming URL contains a port number other than the protocol's standard port, you must set the *ignoreOriginPort* attribute to true for playback to succeed.

- *userDomainName*

This attribute is used in two instances: (1) for NTLM authentication and (2) for the SMB file import feature. If the origin server is using NTLM authentication, you must use this attribute to specify the user domain name for NTLM authentication. If a shared folder is protected and the user account is part of a domain, you must use this attribute to specify the domain name of the configured shared folder.



Note

Both *userDomainName* and *ntlmUserDomain* cannot co-exist in the manifest file; only one attribute can be used at a time.

- *discardAllHttpHeaders*

This attribute can be used for HTTP header filtering. When this attribute is set to "1" or true, no HTTP headers obtained from origin servers are stored in the CDNFS file system. When this attribute is set to "0" or false, only nontransit HTTP headers are stored, while transit headers are filtered out. The default value is false. Transit headers are automatically filtered out unless otherwise specified in the *keepHttpHeaders* attribute.

Transit HTTP headers that are automatically filtered out are as follows:

- accept-ranges	- proxy-connection
- connection	- set-cookie
- content-length	- transfer-encoding
- date	- upgrade
- keep-alive	- via
- proxy-authenticate	- www-authenticate

- *discardHttpHeaders*

This attribute can be used to filter specific non-transit headers that would normally be stored by default in the CDNFS file system. Nontransit headers include “server,” “last-modified,” and “content-type.” To discard specific headers, list the headers separated by a bar (|), as shown in the following example:

```
<item discardHttpHeaders="Server|Last-modified">
```

- *keepHttpHeaders*

This attribute is used to store only specific HTTP headers within the CDNFS file system. To store specific headers, list the headers separated by a bar (|), as shown in the following example:

```
<item keepHttpHeaders="Connection|Date">
```

When you use the *keepHttpHeaders* attribute, only headers that are returned from the server with a successful GET request can be stored. Headers such as, www-authenticate or proxy-authenticate, are not returned with a successful GET request and cannot be stored.



Note The three HTTP header filtering attributes (*discardAllHttpHeaders*, *discardHttpHeaders*, and *keepHttpHeaders*) are mutually exclusive. You can specify only one attribute at a time within a single tag.

For example, the following configuration is not valid and will produce an error:

```
<item src="abc.html" discardHttpHeaders="Server" keepHttpHeaders="Date"/>
```

You can specify one attribute in the options or item-group tag and one attribute in the item or crawler tag; however, the attribute that is specified in the item or crawler tag will override the attribute that is inherited from the options or item-group tag. For example, the following configuration is valid:

```
<options discardAllHttpHeaders="true" />
<item src="one.html" /> <!--This uses the attribute inherited from options-->
<crawler start-url="two.html" discardHttpHeaders="Server"/> <!--This uses one given
with crawler -->
<item src="Three.html" keepHttpHeaders="Date"/> <!--This uses one given with item -->
```

- *enableCookies*

This attribute enables cookie support for the item. When this attribute is set to true, the root Content Engine, after sending a request for an item to the origin server, parses the server response for cookie name value pairs. If the server response contains a cookie that is valid and has not expired, the root Content Engine stores the cookie in main memory.

The root Content Engine then returns the valid cookie to the server the next time the root Content Engine sends a request for the item.

A cookie is rejected if it contains any of the following rejection criteria, as found in RFC 2965:

- The value for the Path attribute is not a prefix of the request URI.

For example, if the request is www.abc.com/aaa/bbb/ccc.html and the Path attribute of the cookie returned is /aaa/ccc, then it is not valid because /aaa/ccc is not a prefix of /aaa/bbb/ccc [URL].

- The value for the Domain attribute contains no embedded dots or does not start with a dot.
- The value for the request host is not a domain-match of the Domain attribute.
- The request host is a FQDN (not an IP address) and has the form HD, where D is the value of the Domain attribute and H is a string that contains one or more dots.
- The Path attribute is not a prefix match of the request URL.



Note The root Content Engine does not use persistent memory to store cookies. If the Content Engine is restarted, all cookie information is lost.

The *enableCookies* attribute can be used with the <item>, <crawler>, <itemgroup>, and <options> tags.

- *authCookie*

This attribute enables the processing and sending of authentication cookies for the item. To enable this feature, the *authCookie* attribute must be set to true for the particular item that passes the user credentials and for which the server sends back the authentication cookies.

For example:

```
<item src=http://abc.com/auth.cgi?id=10000 authCookie="true"/>
```

The *authCookie* attribute can be used with the <item> and <crawler> tags.

The following attributes described under the <host> tag attributes can also be specified by the <item> tag.

- *disableBasicAuth*
- *noProxy*
- *ntlmUserDomain*
- *password*
- *port*
- *proto*
- *proxyServer*
- *sslAuthType*
- *user*
- *uuencoded*

Subelements

- <contains/>
- <http-meta-data/>
- <schedule/> <repeat/>

The <schedule/> <repeat/> subelement and its attributes specify a time for a recrawl or an item refetch to begin. You can have multiple <repeat> subelements under the <schedule> subelement. The attributes *time*, *start*, and *end* specify the day of the month or day of the week and the duration of the specified repeat. The *time* attribute is required, whereas *start* and *end* are optional attributes.



Note The schedule element takes precedence over the *ttl* attribute.

The *time* attribute uses either of the following formats:

time="dom:hh:mm" or

time="dow:hh:mm"

where dom is the day of the month (0–30), dow is the day of the week (Sun, Mon, Tue, Wed, Thu, Fri, Sat, or *), hh is the clock hour (0–23 or *), and mm is the minute (0–59).

For example,

```
<schedule>
  <repeat time="*:*:0" /> <!-- repeat every hour on the hour -->
  <repeat time="*:13:0" /> <!-- repeat at 1300 every day -->
  <repeat time="Sun:2:30" /> <!-- repeat on Sundays at 2:30 -->
  <repeat time="4:2:30" /> <!-- repeat at 2:30 on the fourth day of the month -->
  <repeat time="Mon:*:30" /> <!-- On Monday, repeat every hour at 30 minutes past
    the hour -->
</schedule>
```

The *start* and *end* attributes use the following format:

start="yyyy-mm-dd hh:mm:ss"

end="yyyy-mm-dd hh:mm:ss"

For example,

```
<CdnManifest>
<item>
  <schedule>
    <repeat time="Sun:02:30" />
    <repeat time="*:*:34" start="2003-09-11 11:11:11 PST" end="2004-09-11 11:11:21
PST"/>
    <repeat time="21:02:35" start="2003-09-11 11:11:11 PST" end="2004-09-11
11:11:21 PST"/>
    <repeat time="21:02:35" end="2004-09-11 11:11:21 PST"/>
  </schedule>
</item>
.
.
.
```

Example

```
<item
  src="index.html"
  server="cisco.com"
  ttl="3000"
  alternateUrl="http://www.cisco.com/cdn-error.html"
/>
```


crawler

The `<crawler>` `</crawler>` tag set supports crawling a website or an FTP server.

Attributes

- *start-url*

The *start-url* attribute is required. It defines the URL at which to start the process of crawling the website or FTP server. It is identical to the *src* attribute used in the `<item>` tag. (See the “[•src](#)” section on page A-45.)

- *host*

The *host* attribute specifies the host name if the starting URL specified in the *start-url* attribute is a relative URL.

- *depth*

The *depth* attribute is optional and defines the link depth to which a website is to be crawled or directory depth to which an FTP server is to be crawled. If the depth is not specified, the default is 20. The following are the general depth values:

0 = acquire only the starting URL

1 = acquire the starting URL and its referred files

–1 = infinite or no depth restriction

Depth is defined as the level of a website or the directory level of an FTP server, where 0 is the starting URL.

- *prefix*

The *prefix* attribute is optional and combines the host name from the `<server>` tag with the value of the *prefix* attribute to create a full prefix. Only content whose URLs match the full prefix is acquired, as shown in this example:

```
<server name="xx"> <host name="www.cisco.com" proto="https" port=433 /> </server>
```

and in a `<crawler>` tag:

```
prefix="marketing/eng/"
```

The full prefix is “https://www.cisco.com:433/marketing/eng/.” Only URLs that match this prefix are crawled.

If a prefix is omitted, the crawler checks the default full prefix, which is the host name portion of the URL from the server. In the previous example, the default full prefix is “https://www.cisco.com:433.”

- *accept*

The *accept* attribute is optional and uses a regular expression to define acceptable URLs to crawl in addition to matching the prefix. For example, `accept="stock"` means that only URLs that meet two conditions are searched: the URL matches the prefix and contains the string “stock.” (See the “[Writing Common Regular Expressions](#)” section on page A-10 for more information on using regular expressions.)

Note the following two key differences between the *accept* attribute and the *prefix* attribute:

- The *prefix* attribute uses an exact string match, while the *accept* attribute uses a regular expression.
- The *prefix* attribute applies to a URL including all of its links or subdirectories. However, the *accept* attribute allows the URL and its links and subdirectories to be evaluated separately.

- *reject*

The *reject* attribute is optional and uses a regular expression to reject a URL if it matches the reject regular expression. The reject regular expression is checked after checking for a prefix URL match. If a URL does not match the prefix, it is immediately rejected. If a URL matches the prefix and the reject parameters, it is rejected by the particular reject constraint. (See the [“Writing Common Regular Expressions”](#) section on page A-10 for more information on using regular expressions.)

Note the following two key differences between the *reject* attribute and the *prefix* attribute:

- The *prefix* attribute uses an exact string match, while the *reject* attribute uses a regular expression.
- The *prefix* attribute applies to a URL including all of its links or subdirectories. However, the *reject* attribute allows the URL and its links and subdirectories to be evaluated separately.

- *max-number*

The *max-number* attribute is optional and specifies the maximum number of crawler job objects that can be acquired.

- *maxTotalSizeInB/KB/MB*

The *maxTotalSizeInB/KB/MB* attribute is optional and specifies the maximum total content size in bytes, kilobytes or megabytes that this crawler job can acquire. The size attribute can be expressed in megabytes (MB), kilobytes (KB), or bytes (B).

This attribute replaces the *max-size-in-B/KB/MB* attribute, which continues to be supported for backward compatibility only.

- *srcPrefix*

The *srcPrefix* attribute is optional and must be used in conjunction with the *cdnPrefix* attribute to form a relative ACNS network URL. If a *srcPrefix* attribute is not specified, or if the prefix of the relative source URL does not match the *srcPrefix* attribute, then the relative ACNS network URL is the *cdnPrefix* value combined with the relative source URL. For example, if these content objects have the same source URL prefix “acme/pubs/docs/online/Design/” and you want to replace this prefix with a simple “online/,” then specify *srcPrefix*=“acme/pubs/docs/online/Design/” and *cdnPrefix*=“online/.”

- *cdnPrefix*

The *cdnPrefix* attribute is optional and must be used in conjunction with the *srcPrefix* attribute.

- *requireAuth*

The *requireAuth* attribute is optional and determines whether users need to be authenticated before the specified content is played. When true, the Content Engine requires authentication to play back the specified content to users and communicates with the origin server to check credentials. If the requests pass the credential check, the content is played back from the Content Engine. If this attribute is omitted, a process of discovery approach is used to determine its value: if the specified content is acquired by using a username and password, *requireAuth* is set to true; otherwise, it is set to false. For FTP, if the username is anonymous, *requireAuth* is set to false.



Note

If *requireAuth* is set to true, the Origin Server field in the Website window for this channel needs to point to the server that can authenticate the users. When users want to play back the content, the server specified in the Origin Server field of the Website window is checked for authentication.

- *externalPrefixes*

The *externalPrefixes* attribute is optional and specifies additional prefixes for crawl jobs to crawl multiple protocols or multiple websites. Prefixes are separated with a bar (|).

- *externalServers*

The *externalServers* attribute is optional and can be used for multiple host crawling jobs where each host has a different user account. This attribute can be used to refer to the <host> tag with the proper authentication information.

- *keepExpiredContent*

This attribute can be used to acquire content during an HTTP or HTTPS crawl that is expired. When this attribute is set to true, expired content will be fetched. When this attribute is set to false, expired content is discarded. If this attribute is not specified, the default is false.

- *keepFolder*

This attribute is used to fetch folders (a folder is indicated when the request URL ends with a forward slash "/"). If this attribute is set to false, folder URLs are not acquired.

- *keepNoCacheContent*

This attribute can be used to acquire content during an HTTP or HTTPS crawl that would normally not be cached. When this attribute is set to true, the acquirer will fetch the content even though the content contains an HTTP cache control header indicating that the content is not to be cached. If this attribute is not specified, the default is false.

- *keepQueryUrl*

This attribute can be used to fetch URLs that contain "?" in the URL string. If this attribute is set to true, URLs with "?" will be fetched during HTML parsing for a crawl job if the URL meets the other crawling criteria set forth in the manifest file.



Note

This attribute is supported in ACNS 5.3 and later releases. You cannot use URLs containing a "?" in a manifest file for a channel unless all Content Engines in the channel are using a version of ACNS software that supports this attribute.

This attribute is useful when you want to acquire content from a database, for example, where multiple files are differentiated in the portion of the URL string after the "?". When this attribute is not set, the portion of the URL after the "?" is discarded. If multiple URLs are found where the portion of the URL string in front of the "?" is the same, these URLs appear as duplicates, and only the last "duplicate" URL found is fetched.

- *reportBrokenLinks*

This attribute is used to report links on an HTML web page that cannot be fetched. If this attribute is set to true, all broken links encountered during a website crawl will be reported as errors. This attribute only applies to a website crawl, not to an index crawl. The default is false: broken links are not reported as errors.

The following attributes described under the <host> tag attributes can also be specified by the <crawler> tag:

- *disableBasicAuth*
- *noProxy*
- *ntlmUserDomain*
- *password*

- *port*
- *proto*
- *proxyServer*
- *sslAuthType*
- *user*
- *uuencoded*

The following attributes described under the `<item>` tag attributes can also be specified by the `<crawler>` tag.

- *alternateUrl*
- *authCookie*
- *bitrate-in-bps/bitrate-in-kbps/bitrate-in-mbps*
- *discardAllHttpHeaders*
- *discardHttpHeaders*
- *enableCookies*
- *expires*
- *failRetryInterval*
- *ignoreOriginPort*
- *ignoreQueryString*
- *keepHttpHeaders*
- *noRedirectToOrigin*
- *playDuration*
- *playServer*
- *prefetch*
- *priority*
- *serveStartTime*
- *serveStopTime*
- *server*
- *ttl*
- *type*
- *userDomainName*

Subelements

- `<http-meta-data>`
- `<matchRule></matchRule>`
- `<schedule><repeat>`

(See the “[item](#)” section on page A-44 for descriptions of the `<schedule><repeat>` subelements.)

Example

```
<server name="cisco">
  <host name="http://www.cisco.com/jobs/" />
</server>
<crawler
  server="cisco"
  start-url="eng/index.html"
  depth="10"
  prefix="eng/"
  reject=".pl"
  max-size-in-MB="200"
/>
```

item-group

The `<item-group></item-group>` tag set is used to place shared attributes under one tag so that they can be shared by every `<item>` and `<crawler>` tag within that group. When attributes are shared, it means that attributes can be defined at either the `<item-group>` tag level for group-wide control or on a per `<item>` or per `<crawler>` tag basis. For example, if every `<item>` tag is using the same *server* and *ttl* attributes, you can create an `<item-group>` tag on top of these `<item>` tags and place the *server* and *ttl* attributes in the `<item-group>` tag.

Using shared attributes makes any manifest file with many `<item>` tags more efficient by consolidating the `<item>` tags with shared attributes. If the same attribute value exists in both the `<item-group>` and `<item>` tags, the value in the `<item>` tag takes precedence over that value in the `<item-group>` tag.

The `<item-group>` tag must be enclosed within the `<CdnManifest>` tag set and contain one or more `<item>` or `<crawler>` tags.

Attributes

If an attribute value is present only at the `<item-group>` tag level, then it is inherited by its inner element in the `<item>` tag. If an attribute value is present in a crawler job, its attributes, whether inherited or owned, are propagated to the content fetched by the crawler job.

The following attributes can be shared across many `<item>` and `<crawler>` tags and are candidates for the `<item-group>` level tag. See the “[item](#)” section on page A-44 for detailed descriptions of the following attributes:

- *alternateUrl*
- *bitrate-in-bps/bitrate-in-kbps/bitrate-in-mbps*
- *cdn-url*
- *discardAllHttpHeaders*
- *discardHttpHeaders*
- *enableCookies*
- *expires*

- *failRetryInterval*
- *host*
- *ignoreOriginPort*
- *ignoreQueryString*
- *keepHttpHeaders*
- *noRedirectToOrigin*
- *playDuration*
- *playServer*
- *prefetch*
- *requireAuth*
- *serveStartTime*
- *serveStopTime*
- *server*
- *src*
- *priority*
- *ttl*
- *type*
- *userDomainName*

The following attributes described under the [<host>](#) tag attributes can also be specified by the [<item-group>](#) tag.

- *disableBasicAuth*
- *noProxy*
- *ntlmUserDomain*
- *password*
- *port*
- *proto*
- *proxyServer*
- *sslAuthType*
- *user*
- *uuencoded*

Additionally, the following two attributes can be placed within the [<item-group>](#) tag. See the [“crawler” section on page A-54](#) for a detailed description of the following two attributes:

- *srcPrefix*
- *cdnPrefix*

These two attributes convert the prefix of the *src-url* (content acquisition URL) to the *cdn-url* (publishing URL) for multiple content objects. These content objects are either implicitly specified by multiple [<item>](#) tags or acquired through a crawler job.

These two attributes can also be specified in the <crawler> tag. If you explicitly specify the *srcPrefix* attribute and *cdnPrefix* attribute for an individual <crawler> job, the <crawler> tag-level specification takes precedence over the <item-group> tag-level settings. If you do not specify these attributes for an individual <crawler> job, the <item-group> tag-level specification is inherited by the <crawler> job.

The *srcPrefix* and *cdnPrefix* attributes generate the relative ACNS network URL using the following rules:

- If the *cdn-url* attribute is present in the <item> tag, the relative ACNS network URL contains both the *cdnPrefix* attribute plus the *cdn-url* attribute. For example, if *cdnPrefix*="eng/spec" and *cdn-url*="e/f.html," the relative path in the URL is "eng/spec/e/f.html."
- If the *srcPrefix* attribute is not present in the <item> tag, the relative ACNS network URL is the *cdnPrefix* attribute plus the relative source URL.
- If the prefix of the relative source URL does not match the *srcPrefix* attribute, the relative ACNS network URL is the *cdnPrefix* attribute plus the source relative URL.
- To generate a relative ACNS network URL, remove the matched prefix from the relative source URL and replace it with the *cdnPrefix* attribute.

The relative ACNS network URL of <item> in the following example is "acme/default.htm."

```
<item-group cdnPrefix="acme/" >
  <item src="design/index.html" cdn-url="default.html" />
</item-group>
```

In the following example, content objects with the *srcPrefix* attribute, such as "design/plan/," have the relative ACNS network URL as "acme/" plus relative source URLs stripped of "design/plan/." Other content objects whose prefix attribute does not match "design/plan/" have "acme/" plus their original relative source URL.

```
<crawler
  start-url="design/plan/index.html"
  depth="-1"
  srcPrefix="design/plan/"
  cdnPrefix="acme/" />
```

Subelements

- <http-meta-data/>
- <crawler></crawler>
- <item-group>
- <item></item>
- <schedule><repeat>

(See the "item" section on page A-44 for descriptions of the <schedule><repeat> subelements.)

Example

```
<!--grouped content items-->
<item-group server="origin-web-server" type="prepos" ttl="300" cdnPrefix="unicorn/" >
  <item cdn-url="newHQpresentation.rm" src="newHQpresentation.rm" />
  <item cdn-url="animatedlogo.mpg" src="animlogo.mpg" />
  <item cdn-url="companytheme.mp3" src="cotheme.mp3" />
  <item cdn-url="newHQlayout.avi" src="newHQ.mov" />
</item-group>
```

matchRule

The `<matchRule>` `</matchRule>` tag set is optional and defines additional filter rules for crawler jobs. It affects only `<crawler>` tasks and is not used by single `<item>` tags. The crawler parameters defined in the `<crawler>``</crawler>` tag set determine primarily the scope of a crawl search. If a content object does not meet the criteria specified by the crawler parameter, neither it nor its children are searched.

The `<matchRule>` tag, however, determines only whether or not the content objects should be acquired regardless of the scope of the search. If a web page matches the crawler parameters without the `<matchRule>` feature, its children are searched even though its content objects are not acquired.

In the following crawler job example that uses the `<matchRule>` tag, the entire website is searched, but only files with the .jpg file extension larger than 50 kilobytes are acquired.

```
<crawler start-url="index.html" depth="-1" >
  <matchRule>
    <match size-min-in-KB="50" extensiioin="jpg" />
  </matchRule>
</crawler>
```

The `<matchRule>` element can be nested within an `<item-group>` tag to define group-wide filter rules for `<crawler>` tags contained in the group. It can also be a subelement of a particular `<crawler>` job. The `<crawler>` tag-level setting overrides the `<item-group>` tag-level setting when both tags are present.

If you define criteria locally for individual `<crawler>` jobs, any existing group-level criterion is entirely discarded for that `<crawler>` job. That is, if your `<item-group>` tag match rule is set to A and your `<crawler>` tag specifies another match rule set to B, only B is to be used for the `<crawler>` tag rather than a combination of A and B. You can define at most one `<matchRule>` tag per `<item-group>` tag and at most one `<matchRule>` tag per `<crawler>` tag.

Attributes

None

Subelements

At least one `<match>` tag

match

The `<match>` `</match>` tag is optional and specifies the acquisition criteria of content objects before they can be acquired by ACNS software. Every attribute within a single `<match>` tag is ANDed (to form a logical conjunction) with the other attributes.

You can specify multiple `<match>` tags within the `<matchRule>` tag. The `<match>` tags are ORed (to form a logical inclusion) with other `<match>` tags. You must specify at least one `<match>` tag per `<matchRule>` tag.

Attributes

- *mime-type*

The *mime-type* attribute specifies MIME types.

- *extension*

The *extension* attribute specifies file extensions.

- *time-before*

The *time-before* attribute can provide both an absolute time (modified before yyyy-mm-dd hh:mm:ss) or a relative time (modified within ddd:hh:ss), relative to the present time, to download content. Time parameters should be expressed in GMT time zones. (For GMT offsets, see the “Manifest File Time Zone Tables” section on page A-71.)

- *time-after*

The *time-after* attribute can provide both an absolute time (modified after yyyy-mm-dd hh:mm:ss) or a relative time (modified within ddd:hh:ss), relative to the present time, to download content. Time parameters should be expressed in GMT time zones. (For GMT offsets, see the “Manifest File Time Zone Tables” section on page A-71.)

**Note**

Relative time is calculated based on current time. We recommend that you synchronize the the server clock and the Content Engine clock so that relative time calculations will be accurate.

- *minFileSizeInB/KB/MB*

The *minFileSizeInB/KB/MB* attribute specifies that the acquired content size must be larger than this number of bytes, kilobytes, or megabytes. The size attribute can be expressed in bytes (B), kilobytes (KB), or megabytes (MB).

The *minFileSizeInB/KB/MB* attribute replaces the *size-min-in-B/KB/MB* attribute, which continues to be supported for backward compatibility only.

- *maxFileSizeInB/KB/MB*

The *maxFileSizeInB/KB/MB* attribute specifies that the acquired content size must be smaller than this number of bytes, kilobytes, or megabytes. This attribute can be expressed in bytes (B), kilobytes (KB), or megabytes (MB).

The *maxFileSizeInB/KB/MB* attribute replaces the *size-max-in-B/KB/MB* attribute, which continues to be supported for backward compatibility only.

- *prefix*

The *prefix* attribute is optional and specifies a prefix as a match rule to filter out websites during a crawl job.

- *url-pattern*

The *url-pattern* attribute is optional and specifies a regular expression as a match rule to filter out certain URLs.

Subelements

None

Examples

```
<!-- crawling item group -->
<item-group server="origin-server" type="prepos">
  <matchRule>
    <match time-before="2000-05-05 12:0:0"/>
  </matchRule>
  <crawler start-url="eng/index.html" depth="-1"/>
  <crawler start-url="hr/index.html" depth="3">
    <matchRule>
      <match size-min-in-KB="1" extension="xxx"/>
    </matchRule>
  </crawler>
```

```
</item-group>
```

To download content that was created or modified within the last 90 days, use the relative time format, as shown in the following example:

```
<match time-after="90:00:00"/>
```

To download content that was not modified within the last 2 weeks, use the relative time format, as shown in the following example:

```
<match time-before="14:00:00"/>
```

To download content that has been modified after January 30, 2003, 10:30PM, use the absolute time format, as shown in the following example:

```
<match time-after="2003-01-30 10:30:00"/>
```

contains

The `<contains/>` tag is optional and identifies content objects that are embedded within the content item currently being described. For example, the components of a SMIL (Synchronized Multimedia Integration Language) file request for an item using `<contains/>` links are only accepted after ACNS software determines that dependent content objects are present in the Content Engine.

The `<contains />` tag must be enclosed within the `<item>` `</item>` tag.

The `<contains />` tag is used to include embedded files for some video files, such as .asf or .rp. The ACNS software does not serve this item unless every contained item is present.

Attributes

The `cdn-url` attribute is required and is the relative ACNS network URL of one of the embedded contents.

Subelements

None

Example

```
<item src="house/img08.jpj" cdn-url="img08.jpg" />
<item src="house/img09.jpj" cdn-url="img09.jpg" />
<item cdn-url="house.rp"src="house/house.rp">
  <contains cdn-url="img08.jpg"/>
  <contains cdn-url="img09.jpg"/>
</item>
```

http-meta-data

The `<http-meta-data>` tag is optional and used for HTTP playback of content. If a content object is requested through HTTP, these attributes are sent to the end users as HTTP response headers. This type of response header is useful when you specify the content type for FTP acquired content.

The element can also be nested within an `<item>` or a `<crawler>` tag. The `<item>` or `<crawler>` tag-level settings override the `<item-group>` tag-level settings.

Attributes

The *name=value* attribute can be both standard HTTP header metadata and customized application metadata. If nested within the <item-group> tag, then the attribute applies to the content contained within the group.

This attribute can also be nested within an <item> or a <crawler> tag. The <item> or <crawler> tag-level settings override the <item-group> tag-level settings.

Subelements

None

Configuring the Freshness of Pre-Positioned Content

Four different manifest file configurations are possible to configure and manage the freshness of your pre-positioned content using the *serveStopTime* and *noRedirectToOrigin* attributes. The following configurations are possible:

- Both the *serveStopTime* and *noRedirectToOrigin* attributes are included in the manifest file, making the condition *noRedirectToOrigin* true. The conditions for this first case are shown in [Table A-6](#).
- Only the *serveStopTime* attribute is included in the manifest file but the *noRedirectToOrigin* attribute is not, making the condition *noRedirectToOrigin* false. The conditions for this second case are shown in [Table A-7](#).
- Neither the *serveStopTime* nor the *noRedirectToOrigin* attribute is included in the manifest file, making the condition *noRedirectToOrigin* false. The conditions for this third case are shown in [Table A-8](#).
- Only the *noRedirectToOrigin* attribute is included in the manifest file but the *serveStopTime* attribute is not, making the condition *noRedirectToOrigin* true. The conditions for this fourth case are shown in [Table A-9](#).

Depending on whether or not the *serveStopTime* and *noRedirectToOrigin* attributes are included and the timing combinations of the *serveStopTime* value and the HTTP header expiration, the conditions and corresponding results are listed in [Table A-6](#) through [Table A-9](#). In the following tables, *now* is defined as the time the end user request arrives. These tables use the end user request arrival time to make content delivery decisions.

Table A-6 Both *serveStopTime* and *noRedirectToOrigin* Attributes Included, *noRedirectToOrigin=true*

Condition	Result
Now is past the <i>serveStopTime</i> value	Content is not served, and an error message appears.
Now is before the <i>serveStopTime</i> value but is past the HTTP <i>expires</i> header	Content is served from the cdnfs, but the content can be stale.
Now is before the <i>serveStopTime</i> value and is before the HTTP <i>expires</i> header	Content is served from the cdnfs.
Now is before the <i>serveStopTime</i> value and no HTTP <i>expires</i> header exists	Content is served from the cdnfs.

Table A-7 *Only serveStopTime Attribute Included, noRedirectToOrigin=false*

Condition	Result
Now is past the <i>serveStopTime</i> value	Content is served by proxy from the origin server.
Now is before the <i>serveStopTime</i> value but is past the HTTP <i>expires</i> header	Content is served from the cdnfs, but content can be stale.
Now is before the <i>serveStopTime</i> value and is before the HTTP <i>expires</i> header	Content is served from the cdnfs.
Now is before the <i>serveStopTime</i> value and no HTTP <i>expires</i> header exists	Content is served from the cdnfs.

Table A-8 *Neither serveStopTime nor noRedirectToOrigin Attributes Included, noRedirectToOrigin=false*

Condition	Result
Now is past the HTTP <i>expires</i> header	Content is served by proxy from the origin server.
Now is before the HTTP <i>expires</i> header	Content is served from the cdnfs.
No HTTP <i>expires</i> header exists	Content is served from the cdnfs.

Table A-9 *Only the noRedirectToOrigin Attribute Included, noRedirectToOrigin=true*

Condition	Result
Now is past the HTTP <i>expires</i> header	Content is served from the cdnfs, but content can be stale.
Now is before the HTTP <i>expires</i> header	Content is served from the cdnfs.
No HTTP <i>expires</i> header exists	Content is served from the cdnfs.

XML Schema

In the case of the manifest file, an XML schema defines the custom markup language of the manifest file and the appearance of a given set of XML documents. The XML schema specifies which tags or elements you can use in your documents, the attributes those tags can contain, and their arrangement.

Manifest XML Schema

An XSD is a library that provides an application programming interface (API) for manipulating the components of an XML schema. For more information on an XSD, go to http://www.w3schools.com/schema/schema_intro.asp.

The following XML code is the manifest XML schema (CdnManifest.xsd).

```
<?xml version="1.0"?>

<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema">

<xs:include schemaLocation="PlayServerTable.xsd"/>
```

```

<xs:element name="CdnManifest">
  <xs:complexType>
    <xs:sequence>
      <xs:element ref="playServerTable" minOccurs="0" maxOccurs="1"/>
      <xs:element ref="options" minOccurs="0" maxOccurs="1"/>
      <xs:choice minOccurs="0" maxOccurs="unbounded">
        <xs:element ref="proxyServer" maxOccurs="unbounded"/>
        <xs:element ref="server" maxOccurs="unbounded"/>
        <xs:element ref="item-group" maxOccurs="unbounded"/>
        <xs:element ref="item" maxOccurs="unbounded"/>
        <xs:element ref="crawler" maxOccurs="unbounded"/>
      </xs:choice>
      <xs:choice minOccurs="0" maxOccurs="unbounded">
        <xs:element ref="item-group" maxOccurs="unbounded"/>
        <xs:element ref="item" maxOccurs="unbounded"/>
        <xs:element ref="crawler" maxOccurs="unbounded"/>
      </xs:choice>
    </xs:sequence>
  </xs:complexType>
</xs:element>

<xs:element name="options">
  <xs:complexType>
    <xs:attribute name="timeZone" type="xs:string" use="optional"/>
    <xs:attribute name="notFoundUrl" type="xs:string" use="optional"/>
    <xs:attribute name="alternateUrl" type="xs:string" use="optional"/>
    <xs:attribute name="noRedirectToOrigin" type="xs:boolean" use="optional"
default="false" />
    <xs:attribute name="requireAuth" type="xs:boolean" use="optional"/>
    <xs:attribute name="ignoreQueryString" type="xs:boolean" use="optional"/>
    <xs:attribute name="playDuration" type="xs:unsignedInt" use="optional"/>
    <xs:attribute name="bitrate-in-bps" type="xs:unsignedInt" use="optional"/>
    <xs:attribute name="bitrate-in-kbps" type="xs:unsignedInt" use="optional"/>
    <xs:attribute name="bitrate-in-mbps" type="xs:unsignedInt" use="optional"/>
    <xs:attribute name="ttl" type="xs:unsignedInt" use="optional"/>
    <xs:attribute name="failRetryInterval" type="xs:unsignedInt" use="optional"/>
    <xs:attribute name="prefetch" type="xs:string" use="optional"/>
    <xs:attribute name="type" use="optional" default="prepos">
      <xs:simpleType>
        <xs:restriction base="xs:string">
          <xs:enumeration value="prepos"/>
          <xs:enumeration value="wmt-live"/>
          <xs:enumeration value="real-live"/>
        </xs:restriction>
      </xs:simpleType>
    </xs:attribute>
    <xs:attribute name="manifest-id" type="xs:string" use="optional"/>
    <xs:attribute name="clearlog" type="xs:boolean" use="optional" default="false"/>
    <xs:attribute name="rd" type="xs:string" use="optional"/>
    <xs:attribute name="prepos-tag" type="xs:string" use="optional"/>
    <xs:attribute name="live-tag" type="xs:string" use="optional"/>
  </xs:complexType>
</xs:element>

<xs:element name="server">
  <xs:complexType>
    <xs:sequence>
      <xs:element ref="host" minOccurs="1" maxOccurs="1"/>
    </xs:sequence>
    <xs:attribute name="name" type="xs:string" use="required"/>
  </xs:complexType>
</xs:element>

<xs:attributeGroup name="hostAttr">

```

```

<xs:attribute name="proxyServer" type="xs:string" use="optional"/>
<xs:attribute name="noProxy" type="xs:boolean" use="optional"/>
<xs:attribute name="proto" use="optional">
  <xs:simpleType>
    <xs:restriction base="xs:string">
      <xs:enumeration value="http"/>
      <xs:enumeration value="https"/>
      <xs:enumeration value="ftp"/>
      <xs:enumeration value="rtsp"/>
      <xs:enumeration value="file"/>
    </xs:restriction>
  </xs:simpleType>
</xs:attribute>
<xs:attribute name="port" type="xs:unsignedShort" use="optional"/>
<xs:attribute name="user" type="xs:string" use="optional"/>
<xs:attribute name="password" type="xs:string" use="optional"/>
<xs:attribute name="ntlmUserDomain" type="xs:string" use="optional"/>
<xs:attribute name="disableBasicAuth" type="xs:boolean" use="optional"/>
<xs:attribute name="uuencoded" type="xs:boolean" use="optional"/>
<xs:attribute name="sslAuthType" use="optional">
  <xs:simpleType>
    <xs:restriction base="xs:string">
      <xs:enumeration value="weak"/>
      <xs:enumeration value="strong"/>
    </xs:restriction>
  </xs:simpleType>
</xs:attribute>
</xs:attributeGroup>

<xs:element name="host">
  <xs:complexType>
    <xs:attribute name="name" type="xs:string" use="required"/>
    <xs:attributeGroup ref="hostAttr"/>
  </xs:complexType>
</xs:element>

<xs:element name="proxyServer">
  <xs:complexType>
    <xs:attribute name="serverName" type="xs:string" use="required"/>
    <xs:attribute name="host" type="xs:string" use="optional"/>
    <xs:attribute name="port" type="xs:unsignedShort" use="optional"/>
    <xs:attribute name="user" type="xs:string" use="optional"/>
    <xs:attribute name="password" type="xs:string" use="optional"/>
    <xs:attribute name="ntlmUserDomain" type="xs:string" use="optional"/>
    <xs:attribute name="uuencoded" type="xs:string" use="optional" default="false"/>
    <xs:attribute name="disableBasicAuth" type="xs:boolean" use="optional"
default="false"/>
  </xs:complexType>
</xs:element>

<xs:attributeGroup name="contentAttr">
  <xs:attribute name="server" type="xs:string" use="optional"/>
  <xs:attribute name="playServer" type="xs:string" use="optional"/>
  <xs:attribute name="type" use="optional">
    <xs:simpleType>
      <xs:restriction base="xs:string">
        <xs:enumeration value="prepos"/>
        <xs:enumeration value="wmt-live"/>
        <xs:enumeration value="real-live"/>
      </xs:restriction>
    </xs:simpleType>
  </xs:attribute>
  <xs:attribute name="playDuration" type="xs:unsignedInt" use="optional"/>
  <xs:attribute name="bitrate-in-bps" type="xs:unsignedInt" use="optional"/>

```

```

<xs:attribute name="bitrate-in-kbps" type="xs:unsignedInt" use="optional"/>
<xs:attribute name="bitrate-in-mbps" type="xs:unsignedInt" use="optional"/>
<xs:attribute name="noRedirectToOrigin" type="xs:boolean" use="optional"/>
<xs:attribute name="requireAuth" type="xs:boolean" use="optional"/>
<xs:attribute name="ignoreQueryString" type="xs:boolean" use="optional"/>
<xs:attribute name="alternateUrl" type="xs:string" use="optional"/>
<xs:attribute name="ttl" type="xs:int" use="optional"/>
<xs:attribute name="failRetryInterval" type="xs:unsignedInt" use="optional"/>
<xs:attribute name="priority" type="xs:unsignedInt" use="optional"/>
<xs:attribute name="prefetch" type="xs:string" use="optional"/>
<xs:attribute name="expires" type="xs:string" use="optional"/>
<xs:attribute name="serve" type="xs:string" use="optional"/>
<xs:attribute name="serveStartTime" type="xs:string" use="optional"/>
<xs:attribute name="serveStopTime" type="xs:string" use="optional"/>
</xs:attributeGroup>

<xs:attributeGroup name = "prefixAttr">
  <xs:attribute name="cdnPrefix" type="xs:string" use="optional"/>
  <xs:attribute name="srcPrefix" type="xs:string" use="optional"/>
</xs:attributeGroup>

<xs:element name="item-group">
  <xs:complexType>
    <xs:sequence>
      <xs:element ref="matchRule" minOccurs="0" maxOccurs="1"/>
      <xs:element ref="http-meta-data" minOccurs="0" maxOccurs="1"/>
      <xs:choice minOccurs="1" maxOccurs="unbounded">
        <xs:element ref="item-group" maxOccurs="unbounded"/>
        <xs:element ref="item" maxOccurs="unbounded"/>
        <xs:element ref="crawler" maxOccurs="unbounded"/>
      </xs:choice>
    </xs:sequence>
    <xs:attributeGroup ref="contentAttr"/>
    <xs:attributeGroup ref="prefixAttr"/>
    <xs:attribute name="host" type="xs:string" use="optional"/>
    <xs:attributeGroup ref="hostAttr"/>
  </xs:complexType>
</xs:element>

<xs:element name="item">
  <xs:complexType>
    <xs:sequence>
      <xs:element ref="contains" minOccurs="0" maxOccurs="unbounded"/>
      <xs:element ref="http-meta-data" minOccurs="0" maxOccurs="1"/>
    </xs:sequence>
    <xs:attribute name="src" type="xs:string" use="required"/>
    <xs:attribute name="cdn-url" type="xs:string" use="optional"/>
    <xs:attributeGroup ref="contentAttr"/>
    <xs:attribute name="host" type="xs:string" use="optional"/>
    <xs:attributeGroup ref="hostAttr"/>
  </xs:complexType>
</xs:element>

<xs:element name="crawler">
  <xs:complexType>
    <xs:all>
      <xs:element ref="matchRule" minOccurs="0" maxOccurs="1"/>
      <xs:element ref="http-meta-data" minOccurs="0" maxOccurs="1"/>
    </xs:all>
    <xs:attribute name="start-url" type="xs:string" use="required"/>
    <xs:attribute name="externalPrefixes" type="xs:string" use="optional"/>
    <xs:attribute name="externalServers" type="xs:string" use="optional"/>
    <xs:attribute name="depth" type="xs:short" use="optional"/>
    <xs:attribute name="prefix" type="xs:string" use="optional"/>
    <xs:attribute name="accept" type="xs:string" use="optional"/>
  </xs:complexType>
</xs:element>

```

```

<xs:attribute name="reject" type="xs:string" use="optional"/>
<xs:attribute name="max-number" type="xs:unsignedInt" use="optional"/>

<xs:attribute name="max-size-in-B" type="xs:unsignedInt" use="optional"/>
<xs:attribute name="max-size-in-KB" type="xs:unsignedInt" use="optional"/>
<xs:attribute name="max-size-in-MB" type="xs:unsignedInt" use="optional"/>
<xs:attributeGroup ref="contentAttr"/>
<xs:attributeGroup ref="prefixAttr"/>
<xs:attribute name="host" type="xs:string" use="optional"/>
<xs:attributeGroup ref="hostAttr"/>
</xs:complexType>
</xs:element>

<xs:element name="contains">
  <xs:complexType>
    <xs:attribute name="cdn-url" type="xs:string" use="required"/>
  </xs:complexType>
</xs:element>

<xs:element name="matchRule">
  <xs:complexType>
    <xs:sequence>
      <xs:element ref="match" minOccurs="1" maxOccurs="unbounded"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>
<xs:element name="match">
  <xs:complexType>
    <xs:attribute name="mime-type" type="xs:string" use="optional"/>
    <xs:attribute name="time-before" type="xs:string" use="optional"/>
    <xs:attribute name="time-after" type="xs:string" use="optional"/>
    <xs:attribute name="size-min-in-B" type="xs:int" use="optional"/>
    <xs:attribute name="size-max-in-B" type="xs:int" use="optional"/>
    <xs:attribute name="size-min-in-KB" type="xs:int" use="optional"/>
    <xs:attribute name="size-max-in-KB" type="xs:int" use="optional"/>
    <xs:attribute name="size-min-in-MB" type="xs:int" use="optional"/>
    <xs:attribute name="size-max-in-MB" type="xs:int" use="optional"/>
    <xs:attribute name="extension" type="xs:string" use="optional"/>
    <xs:attribute name="prefix" type="xs:string" use="optional"/>
    <xs:attribute name="url-pattern" type="xs:string" use="optional"/>
  </xs:complexType>
</xs:element>

<xs:element name="http-meta-data">
  <xs:complexType>
    <xs:anyAttribute processContents="skip" />
  </xs:complexType>
</xs:element>

</xs:schema>

```

PlayServerTable XML Schema

The following XML code defines the PlayServerTable schema (playServerTable.xsd) for the CdnManifest.xsd.

```

<?xml version="1.0"?>

<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema">

  <xs:element name="playServerTable">

```



```

    <xs:complexType>
      <xs:sequence>
        <xs:element ref="playServer" minOccurs="1" maxOccurs="unbounded"/>
      </xs:sequence>
    </xs:complexType>
  </xs:element>

  <xs:element name="playServer">
    <xs:complexType>
      <xs:choice minOccurs="1" maxOccurs="unbounded">
        <xs:element ref="contentType"/>
        <xs:element ref="extension"/>
      </xs:choice>
      <xs:attribute name="name" use="required">
        <xs:simpleType>
          <xs:restriction base="xs:string">
            <xs:enumeration value="real"/>
            <xs:enumeration value="wmt"/>
            <xs:enumeration value="http"/>
            <xs:enumeration value="qtss"/>
          </xs:restriction>
        </xs:simpleType>
      </xs:attribute>
    </xs:complexType>
  </xs:element>

  <xs:element name="contentType">
    <xs:complexType>
      <xs:attribute name="name" type="xs:string" use="required"/>
    </xs:complexType>
  </xs:element>

  <xs:element name="extension">
    <xs:complexType>
      <xs:attribute name="name" type="xs:string" use="required"/>
    </xs:complexType>
  </xs:element>
</xs:schema>

```

Default PlayServerTable Schema

The following XML code defines the default PlayServerTable schema (PlayServerTable.xsd).

```

<?xml version="1.0"?>

<playServerTable xmlns:xsi = "http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation = "PlayServerTable.xsd">

  <!-- playServer http and https can always play all prepositon
  contents unless users use customized <playServerTable>
  or "playServer" attribute in the manifest file
  -->

  <playServer name="real">
    <!-- MIME type staken from
    http://service.real.com/help/library/guides/server8/htmfiles/custmizg.htm
    -->
    <contentType name="audio/x-pn-realaudio" />
    <contentType name="audio/x-pn-realaudio-plugin" />
    <contentType name="application/x-pn-realmedia" />
    <contentType name="application/smil" />
    <contentType name="application/vnd.rn-rmadriver" />
    <extension name="rm" />
  </playServer>

```

```

    <extension name="rms" />
    <extension name="ra" />
    <extension name="rp" />
    <extension name="rt" />
    <extension name="smi" />
</playServer>

<playServer name="qtss">
  <contentType name="video/quicktime" />
  <extension name="mov" />
  <extension name="qt" />
  <extension name="mp4" />

  <!-- extension avi could go here, but is also supported by wmt -->
</playServer>

<playServer name="wmt">
  <!-- MIME types taken from
  http://msdn.microsoft.com/workshop/imedia/windowsmedia/server/mime.asp
  -->
  <contentType name="video/x-ms-asf" />
  <contentType name="audio/x-ms-wma" />
  <contentType name="video/x-ms-wmv" />
  <contentType name="video/x-ms-wm" />
  <contentType name="application/x-ms-wmz" />
  <contentType name="application/x-ms-wmd" />

  <extension name="wma" /> <!-- audio content -->
  <extension name="wmv" /> <!-- audio/video content -->
  <extension name="asf" /> <!-- audio/video content (legacy) -->
  <extension name="wm" /> <!-- reserved for future use -->

  <!-- extension avi could go here, but is also supported by qtss -->
</playServer>

</playServerTable>

```

Manifest File Time Zone Tables

To convert to local time, you must know the time difference between Greenwich mean time (GMT) and local time for both standard time and summer time (daylight saving time). [Table A-10](#) through [Table A-25](#) list the time zones supported by the manifest file. The format for writing the time zone is:

<zonenumber>:[+|-:]hh:mm per line

where <zonenumber> is the name of the time zone or standard time zone abbreviation (see [Table A-10](#)) without spaces before or after the colon (":"), and "[+|-:]hh:mm" is the GMT offset in hours and minutes. The GMT offset default is "+."

Table A-10 Standard Time Zones and GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
ACT:+:09:30	Etc/GMT+7:-:07:00	HST:-:10:00
ADT:-:03:00	Etc/GMT+8:-:08:00	IET:-:05:00
AET:+:10:00	Etc/GMT+9:-:09:00	IST:+:05:30
AGT:-:03:00	Etc/GMT-0:00:00	JST:+:09:00

Table A-10 **Standard Time Zones and GMT Offsets (continued)**

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
ART:+:02:00	Etc/GMT-10:+:10:00	MDT:-:06:00
AST:-:09:00	Etc/GMT-11:+:11:00	MET:+:01:00
BET:-:03:00	Etc/GMT-12:+:12:00	MIT:-:11:00
BST:+:06:00	Etc/GMT-13:+:13:00	MST7MDT:-:07:00
CAT:+:02:00	Etc/GMT-14:+:14:00	MST:-:07:00
CDT:-:05:00	Etc/GMT-1:+:01:00	NET:+:04:00
CET:+:01:00	Etc/GMT-2:+:02:00	NST:+:12:00
CNT:-:03:30	Etc/GMT-3:+:03:00	NZ-CHAT:+:12:45
CST6CDT:-:06:00	Etc/GMT-4:+:04:00	NZ:+:12:00
CST:-:06:00	Etc/GMT-5:+:05:00	Navajo:-:07:00
CTT:+:08:00	Etc/GMT-6:+:06:00	PDT:-:07:00
EAT:+:03:00	Etc/GMT-7:+:07:00	PLT:+:05:00
ECT:+:01:00	Etc/GMT-8:+:08:00	PNT:-:07:00
EDT:-:04:00	Etc/GMT-9:+:09:00	PRC:+:08:00
EET:+:02:00	Etc/GMT0:00:00	PRT:-:04:00
EST5EDT:-:05:00	Etc/GMT:00:00	PST8PDT:-:08:00
EST:-:05:00	Etc/Greenwich:00:00	PST:-:08:00
Etc/GMT+0:00:00	Etc/UCT:00:00	ROK:+:09:00
Etc/GMT+10:-:10:00	Etc/UTC:00:00	SST:+:11:00
Etc/GMT+11:-:11:00	Etc/Universal:00:00	UCT:00:00
Etc/GMT+12:-:12:00	Etc/Zulu:00:00	UTC:00:00
Etc/GMT+1:-:01:00	GB-Eire:00:00	Universal:00:00
Etc/GMT+2:-:02:00	GB:00:00	VST:+:07:00
Etc/GMT+3:-:03:00	GMT0:00:00	W-SU:+:03:00
Etc/GMT+4:-:04:00	GMT:00:00	WET:00:00
Etc/GMT+5:-:05:00	Greenwich:00:00	Zulu:00:00
Etc/GMT+6:-:06:00	HDT:-:09:00	

Table A-11 **Africa GMT Offsets**

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Africa/Abidjan:00:00	Africa/Djibouti:+:03:00	Africa/Maputo:+:02:00
Africa/Accra:00:00	Africa/Douala:+:01:00	Africa/Maseru:+:02:00
Africa/Addis_Ababa:+:03:00	Africa/El_Aaiun:00:00	Africa/Mbabane:+:02:00
Africa/Algiers:+:01:00	Africa/Freetown:00:00	Africa/Mogadishu:+:03:00
Africa/Asmera:+:03:00	Africa/Gaborone:+:02:00	Africa/Monrovia:00:00

Table A-11 Africa GMT Offsets (continued)

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Africa/Bamako:00:00	Africa/Harare:+:02:00	Africa/Nairobi:+:03:00
Africa/Bangui:+:01:00	Africa/Johannesburg:+:02:00	Africa/Ndjamena:+:01:00
Africa/Banjul:00:00	Africa/Kampala:+:03:00	Africa/Niamey:+:01:00
Africa/Bissau:00:00	Africa/Khartoum:+:03:00	Africa/Nouakchott:00:00
Africa/Blantyre:+:02:00	Africa/Kigali:+:02:00	Africa/Ouagadougou:00:00
Africa/Brazzaville:+:01:00	Africa/Kinshasa:+:01:00	Africa/Porto-Novo:+:01:00
Africa/Bujumbura:+:02:00	Africa/Lagos:+:01:00	Africa/Sao_Tome:00:00
Africa/Cairo:+:02:00	Africa/Libreville:+:01:00	Africa/Timbuktu:00:00
Africa/Casablanca:00:00	Africa/Lome:00:00	Africa/Tripoli:+:02:00
Africa/Ceuta:+:01:00	Africa/Luanda:+:01:00	Africa/Tunis:+:01:00
Africa/Conakry:00:00	Africa/Lubumbashi:+:02:00	Africa/Windhoek:+:01:00
Africa/Dakar:00:00	Africa/Lusaka:+:02:00	—
Africa/Dar_es_Salaam:+:03:00	Africa/Malabo:+:01:00	—

Table A-12 America GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
America/Adak:-:10:00	America/Grenada:-:04:00	America/Noronha:-:02:00
America/Anchorage:-:09:00	America/Guadeloupe:-:04:00	America/North_Dak/Ctr:-:06:00
America/Anguilla:-:04:00	America/Guatemala:-:06:00	America/Panama:-:05:00
America/Antigua:-:04:00	America/Guayaquil:-:05:00	America/Pangnirtung:-:05:00
America/Araguaina:-:03:00	America/Guyana:-:04:00	America/Paramaribo:-:03:00
America/Aruba:-:04:00	America/Halifax:-:04:00	America/Phoenix:-:07:00
America/Asuncion:-:04:00	America/Havana:-:05:00	America/Port-au-Prince:-:05:00
America/Atka:-:10:00	America/Hermosillo:-:07:00	America/Port_of_Spain:-:04:00
America/Barbados:-:04:00	America/Ind/Indian:-:05:00	America/Porto_Acre:-:05:00
America/Belem:-:03:00	America/Ind/Knox:-:05:00	America/Porto_Velho:-:04:00
America/Belize:-:06:00	America/Ind/Marengo:-:05:00	America/Puerto_Rico:-:04:00
America/Boa_Vista:-:04:00	America/Ind/Vevay:-:05:00	America/Rainy_River:-:06:00
America/Bogota:-:05:00	America/Indianapolis:-:05:00	America/Rankin_Inlet:-:06:00
America/Bogota:-:05:00	America/Inuvik:-:07:00	America/Recife:-:03:00
America/Buenos_Aires:-:03:00	America/Iqaluit:-:05:00	America/Regina:-:06:00
America/Cambridge_Bay:-:07:00	America/Jamaica:-:05:00	America/Rio_Branco:-:05:00
America/Cancun:-:06:00	America/Jujuy:-:03:00	America/Rosario:-:03:00
America/Caracas:-:04:00	America/Juneau:-:09:00	America/Santiago:-:04:00
America/Catamarca:-:03:00	America/Ken/Louisville:-:05:00	America/Santo_Domingo:-:04:00

Table A-12 *America GMT Offsets (continued)*

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
America/Cayenne:-:03:00	America/Ken/Monticello:-:05:0	America/Sao_Paulo:-:03:00
America/Cayman:-:05:00	America/Knox_IN:-:05:00	America/Scoresbysund:-:01:00
America/Chicago:-:06:00	America/La_Paz:-:04:00	America/Shiprock:-:07:00
America/Chihuahua:-:07:00	America/Lima:-:05:00	America/St_Johns:-:03:30
America/Cordoba:-:03:00	America/Los_Angeles:-:08:00	America/St_Lucia:-:04:00
America/Costa_Rica:-:06:00	America/Louisville:-:05:00	America/St_Thomas:-:04:00
America/Cuiaba:-:04:00	America/Maceio:-:03:00	America/St_Vincent:-:04:00
America/Curacao:-:04:00	America/Managua:-:06:00	America/Swift_Current:-:06:00
America/Danmarkshavn:00:00	America/Manaus:-:04:00	America/Tegucigalpa:-:06:00
America/Dawson:-:08:00	America/Martinique:-:04:00	America/Thule:-:04:00
America/Dawson_Creek:-:07:00	America/Mazatlan:-:07:00	America/Thunder_Bay:-:05:00
America/Denver:-:07:00	America/Mendoza:-:03:00	America/Tijuana:-:08:00
America/Detroit:-:05:00	America/Menominee:-:06:00	America/Tortola:-:04:00
America/Dominica:-:04:00	America/Merida:-:06:00	America/Vancouver:-:08:00
America/Edmonton:-:07:00	America/Mexico_City:-:06:00	America/St_Lucia:-:04:00
America/Eirunepe:-:05:00	America/Miquelon:-:03:00	America/Virgin:-:04:00
America/El_Salvador:-:06:00	America/Monterrey:-:06:00	America/Whitehorse:-:08:00
America/Ensenada:-:08:00	America/Montevideo:-:03:00	America/Winnipeg:-:06:00
America/Fort_Wayne:-:05:00	America/Montreal:-:05:00	America/Yakutat:-:09:00
America/Fortaleza:-:03:00	America/Montserrat:-:04:00	America/Yellowknife:-:07:00
America/Glace_Bay:-:04:00	America/Nassau:-:05:00	America/Virgin:-:04:00
America/Godthab:-:03:00	America/New_York:-:05:00	America/Whitehorse:-:08:00
America/Goose_Bay:-:04:00	America/Nipigon:-:05:00	America/Winnipeg:-:06:00
America/Grand_Turk:-:05:00	America/Nome:-:09:00	America/Tortola:-:04:00

Table A-13 *Antarctica/Arctic GMT Offsets*

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Antarctica/Casey:+:08:00	Antarctica/McMurdo:+:12:00	Antarctica/Vostok:+:06:00
Antarctica/Davis:+:07:00	Antarctica/Palmer:-:04:00	Arctic/Longyearbyen:+:01:00
Antarctica/DtDUrville:+:10:00	Antarctica/South_Pole:+:12:00	—
Antarctica/Mawson:+:06:00	Antarctica/Syowa:+:03:00	—

Table A-14 Asia GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Asia/Aden:+:03:00	Asia/Hong_Kong:+:08:00	Asia/Riyadh87:+:03:07
Asia/Almaty:+:06:00	Asia/Hovd:+:07:00	Asia/Riyadh88:+:03:07
Asia/Amman:+:02:00	Asia/Irkutsk:+:08:00	Asia/Riyadh89:+:03:07
Asia/Anadyr:+:12:00	Asia/Istanbul:+:02:00	Asia/Riyadh:+:03:00
Asia/Aqtau:+:04:00	Asia/Jakarta:+:07:00	Asia/Saigon:+:07:00
Asia/Aqtobe:+:05:00	Asia/Jayapura:+:09:00	Asia/Sakhalin:+:10:00
Asia/Ashgabat:+:05:00	Asia/Jerusalem:+:02:00	Asia/Samarkand:+:05:00
Asia/Ashkhabad:+:05:00	Asia/Kabul:+:04:30	Asia/Seoul:+:09:00
Asia/Baghdad:+:03:00	Asia/Kamchatka:+:12:00	Asia/Shanghai:+:08:00
Asia/Bahrain:+:03:00	Asia/Karachi:+:05:00	Asia/Singapore:+:08:00
Asia/Baku:+:04:00	Asia/Kashgar:+:08:00	Asia/Taipei:+:08:00
Asia/Bangkok:+:07:00	Asia/Katmandu:+:05:45	Asia/Tashkent:+:05:00
Asia/Beirut:+:02:00	Asia/Krasnoyarsk:+:07:00	Asia/Tbilisi:+:04:00
Asia/Bishkek:+:05:00	Asia/Kuala_Lumpur:+:08:00	Asia/Tehran:+:03:30
Asia/Brunei:+:08:00	Asia/Kuching:+:08:00	Asia/Tel_Aviv:+:02:00
Asia/Calcutta:+:05:30	Asia/Kuwait:+:03:00	Asia/Thimbu:+:06:00
Asia/Choibalsan:+:09:00	Asia/Macao:+:08:00	Asia/Thimphu:+:06:00
Asia/Chongqing:+:08:00	Asia/Magadan:+:11:00	Asia/Tokyo:+:09:00
Asia/Chungking:+:08:00	Asia/Manila:+:08:00	Asia/Ujung_Pandang:+:08:00
Asia/Colombo:+:06:00	Asia/Muscat:+:04:00	Asia/Ulaanbaatar:+:08:00
Asia/Dacca:+:06:00	Asia/Nicosia:+:02:00	Asia/Ulan_Bator:+:08:00
Asia/Damascus:+:02:00	Asia/Novosibirsk:+:06:00	Asia/Urumqi:+:08:00
Asia/Dhaka:+:06:00	Asia/Omsk:+:06:00	Asia/Vientiane:+:07:00
Asia/Dili:+:09:00	Asia/Phnom_Penh:+:07:00	Asia/Vladivostok:+:10:00
Asia/Dubai:+:04:00	Asia/Pontianak:+:07:00	Asia/Yakutsk:+:09:00
Asia/Dushanbe:+:05:00	Asia/Pyongyang:+:09:00	Asia/Yekaterinburg:+:05:00
Asia/Gaza:+:02:00	Asia/Qatar:+:03:00	Asia/Yerevan:+:04:00
Asia/Harbin:+:08:00	Asia/Rangoon:+:06:30	—

Table A-15 Atlantic GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Atlantic/Azores:-:01:00	Atlantic/Faeroe:00:00	Atlantic/South_Georgia:-:02:00
Atlantic/Bermuda:-:04:00	Atlantic/Jan_Mayen:+:01:00	Atlantic/St_Helena:00:00
Atlantic/Canary:00:00	Atlantic/Madeira:00:00	Atlantic/Stanley:-:04:00
Atlantic/Cape_Verde:-:01:00	Atlantic/Reykjavik:00:00	—

Table A-16 Australia GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Australia/ACT:+:10:00	Australia/LHI:+:10:30	Australia/Queensland:+:10:00
Australia/Adelaide:+:09:30	Australia/Lindeman:+:10:00	Australia/South:+:09:30
Australia/Brisbane:+:10:00	Australia/Lord_Howe:+:10:30	Australia/Sydney:+:10:00
Australia/Broken_Hill:+:09:30	Australia/Melbourne:+:10:00	Australia/Tasmania:+:10:00
Australia/Canberra:+:10:00	Australia/NSW:+:10:00	Australia/Victoria:+:10:00
Australia/Darwin:+:09:30	Australia/North:+:09:30	Australia/West:+:08:00
Australia/Hobart:+:10:00	Australia/Perth:+:08:00	Australia/Yancowinna:+:09:30

Table A-17 Brazil GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Brazil/Acre:-:05:00	Brazil/East:-:03:00	Brazil/West:-:04:00
Brazil/DeNoronha:-:02:00	—	—

Table A-18 Canada/Chile/Cuba GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Canada/Atlantic:-:04:00	Canada/Mountain:-:07:00	Canada/Yukon:-:08:00
Canada/Central:-:06:00	Canada/Newfoundland:-:03:30	Chile/Continental:-:04:00
Canada/East-Ssktchwan:-:06:00	Canada/Pacific:-:08:00	Chile/EasterIsland:-:06:00
Canada/Eastern:-:05:00	Canada/Saskatchewan:-:06:00	Cuba:-:05:00

Table A-19 Egypt/Eire/Europe GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Egypt:+:02:00	Europe/Kiev:+:02:00	Europe/Simferopol:+:02:00
Eire:00:00	Europe/Lisbon:00:00	Europe/Skopje:+:01:00
Europe/Amsterdam:+:01:00	Europe/Ljubljana:+:01:00	Europe/Sofia:+:02:00
Europe/Andorra:+:01:00	Europe/London:00:00	Europe/Stockholm:+:01:00
Europe/Athens:+:02:00	Europe/Luxembourg:+:01:00	Europe/Tallinn:+:02:00
Europe/Belfast:00:00	Europe/Madrid:+:01:00	Europe/Tirane:+:01:00
Europe/Belgrade:+:01:00	Europe/Malta:+:01:00	Europe/Tiraspol:+:02:00
Europe/Berlin:+:01:00	Europe/Minsk:+:02:00	Europe/Uzhgorod:+:02:00
Europe/Bratislava:+:01:00	Europe/Monaco:+:01:00	Europe/Vaduz:+:01:00
Europe/Brussels:+:01:00	Europe/Moscow:+:03:00	Europe/Vatican:+:01:00
Europe/Bucharest:+:02:00	Europe/Nicosia:+:02:00	Europe/Vienna:+:01:00
Europe/Budapest:+:01:00	Europe/Oslo:+:01:00	Europe/Vilnius:+:02:00

Table A-19 Egypt/Eire/Europe GMT Offsets (continued)

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Europe/Chisinau: +:02:00	Europe/Paris: +:01:00	Europe/Warsaw: +:01:00
Europe/Copenhagen: +:01:00	Europe/Prague: +:01:00	Europe/Zagreb: +:01:00
Europe/Dublin: 00:00	Europe/Riga: +:02:00	Europe/Zaporozhye: +:02:00
Europe/Gibraltar: +:01:00	Europe/Rome: +:01:00	Europe/Zurich: +:01:00
Europe/Helsinki: +:02:00	Europe/Samara: +:04:00	Europe/Simferopol: +:02:00
Europe/Istanbul: +:02:00	Europe/San_Marino: +:01:00	Europe/Skopje: +:01:00
Europe/Kaliningrad: +:02:00	Europe/Sarajevo: +:01:00	Europe/Sofia: +:02:00

Table A-20 Hong Kong/Iceland/India/Iran/Israel GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Hongkong: +:08:00	Indian/Cocos: +:06:30	Indian/Mauritius: +:04:00
Iceland: 00:00	Indian/Comoro: +:03:00	Indian/Mayotte: +:03:00
Indian/Antananarivo: +:03:00	Indian/Kerguelen: +:05:00	Indian/Reunion: +:04:00
Indian/Chagos: +:06:00	Indian/Mahe: +:04:00	Iran: +:03:30
Indian/Christmas: +:07:00	Indian/Maldives: +:05:00	Israel: +:02:00

Table A-21 Jamaica/Japan/Kwajalein/Libya GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Jamaica: -:05:00	Kwajalein: +:12:00	Libya: +:02:00
Japan: +:09:00	—	—

Table A-22 Mexico/Mideast GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Mexico/BajaNorte: -:08:00	Mexico/General: -:06:00	Mideast/Riyadh88: +:03:07
Mexico/BajaSur: -:07:00	Mideast/Riyadh87: +:03:07	Mideast/Riyadh89: +:03:07

Table A-23 Pacific/Poland/Portugal GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Pacific/Apia: -:11:00	Pacific/Johnston: -:10:00	Pacific/Ponape: +:11:00
Pacific/Auckland: +:12:00	Pacific/Kiritimati: +:14:00	Pacific/Port_Moresby: +:10:00
Pacific/Chatham: +:12:45	Pacific/Kosrae: +:11:00	Pacific/Rarotonga: -:10:00
Pacific/Easter: -:06:00	Pacific/Kwajalein: +:12:00	Pacific/Saipan: +:10:00
Pacific/Efate: +:11:00	Pacific/Majuro: +:12:00	Pacific/Samoa: -:11:00

Table A-23 Pacific/Poland/Portugal GMT Offsets (continued)

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Pacific/Enderbury:+:13:00	Pacific/Marquesas:-:09:30	Pacific/Tahiti:-:10:00
Pacific/Fakaofa:-:10:00	Pacific/Midway:-:11:00	Pacific/Tarawa:+:12:00
Pacific/Fiji:+:12:00	Pacific/Nauru:+:12:00	Pacific/Tongatapu:+:13:00
Pacific/Funafuti:+:12:00	Pacific/Niue:-:11:00	Pacific/Truk:+:10:00
Pacific/Galapagos:-:06:00	Pacific/Norfolk:+:11:30	Pacific/Wake:+:12:00
Pacific/Gambier:-:09:00	Pacific/Noumea:+:11:00	Pacific/Wallis:+:12:00
Pacific/Guadalcanal:+:11:00	Pacific/Pago_Pago:-:11:00	Pacific/Yap:+:10:00
Pacific/Guam:+:10:00	Pacific/Palau:+:09:00	Poland:+:01:00
Pacific/Honolulu:-:10:00	Pacific/Pitcairn:-:08:00	Portugal:00:00

Table A-24 Singapore/System V/Turkey GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
Singapore:+:08:00	SystemV/EST5:-:05:00	SystemV/PST8PDT:-:08:00
SystemV/AST4:-:04:00	SystemV/EST5EDT:-:05:00	SystemV/YST9:-:09:00
SystemV/AST4ADT:-:04:00	SystemV/MST7:-:07:00	SystemV/YST9YDT:-:09:00
SystemV/CST6:-:06:00	SystemV/MST7MDT:-:07:00	Turkey:+:02:00
SystemV/CST6CDT:-:06:00	SystemV/PST8:-:08:00	—

Table A-25 U.S. GMT Offsets

Time Zone: GMT Offset	Time Zone: GMT Offset	Time Zone: GMT Offset
US/Alaska:-:09:00	US/Eastern:-:05:00	US/Pacific-New:-:08:00
US/Aleutian:-:10:00	US/Hawaii:-:10:00	US/Pacific:-:08:00
US/Arizona:-:07:00	US/Indiana-Starke:-:05:00	US/Samoa:-:11:00
US/Central:-:06:00	US/Michigan:-:05:00	—
US/East-Indiana:-:05:00	US/Mountain:-:07:00	—

Manifest File Automated Scripts

This section contains information about automated Perl scripts that you can use to automate the creation of manifest files for your ACNS network. The most efficient method of creating a manifest file is to customize the automated Spider Perl script in combination with the Manifest Perl script, both of which are available on Cisco.com. These two Perl scripts can serve as the basis for your own automation scripts that you modify to suit your needs.

The two Perl automated scripts are called *spider.pl* and *manifest.pl*. These scripts can be used as-is. If you are proficient in using Perl, you can modify the *spider.pl* and *manifest.pl*. However, if you modify these scripts, Cisco will not support them. Both the *spider.pl* and *manifest.pl* scripts contain a "--file"

argument that is to be used in conjunction with a user-created rules file, such as .cfg. So that they can be reused, it is recommended that customers use this method to include the various arguments that they require, as opposed to running them from the command line.

First, use the Spider script, and then use the output of the Spider script as input to the Manifest script. The Spider script searches the content of the selected origin servers and outputs a database file containing a list of URLs of all content. The Manifest script uses this database file to build the manifest file with the correct syntax based on rules you stipulate from the command line or rules file. This produces an XML-based manifest file containing the URLs of only those content objects that you want made available to your users.

The following two sample automated Perl scripts are available on Cisco.com:

- Spider Perl script

The Spider script crawls over the content of a selected origin server and outputs a database file containing a list of URLs.

- Manifest Perl script

The Manifest script reads the database file output by the Spider script and uses rules that you establish to produce an XML-formatted manifest file containing the URLs of only those filtered content objects that you want to make available to users.

Installing Perl on Your Workstation

You must have Perl installed on your workstation before working with or running the Spider or Manifest scripts. It is useful to also have a Perl interpreter available. Perl is open source software and can be downloaded for free from a variety of locations on the Internet. Refer to the Comprehensive Perl Archive Network (CPAN) at these URLs:

<http://www.cpan.org>

or

<http://www.perl.com>

Obtaining the Perl Scripts

The Spider and Manifest scripts can be obtained from Cisco.com by using the same procedure that is used to obtain updated versions of Cisco ACNS software.

To obtain the Spider and Manifest scripts from Cisco.com, follow these steps:

-
Step 1	Go to the following URL to find the Spider and Manifest Perl scripts: http://www.cisco.com/cgi-bin/tablebuild.pl/acns52
Step 2	When prompted, log in to Cisco.com by using your designated Cisco.com username and password. The Cisco ACNS Software download page appears, listing the available software updates for the Cisco ACNS Software product.
Step 3	Locate the file named ACNS-5.2.1.5-manifest-tools.zip . This file is a ZIP archive containing both the Manifest and Spider Perl scripts.
Step 4	Click the link for the ACNS-5.2.1.5-manifest-tools.zip file. The download page appears.
Step 5	Click Software License Agreement . A new browser window opens, displaying the license agreement.

- Step 6** After you have read the license agreement, close the browser window displaying the agreement, and return to the Software Download page.
- Step 7** Click the filename link labeled **Download**.
- Step 8** Click **Save to file**, and then choose a location on your workstation to temporarily store the zipped file containing the scripts.
- Step 9** Use your preferred unzip program to unpack the scripts to a location on your workstation or your network.

After you have unzipped the scripts, you are ready to begin using them to build manifest files for your website. See the [“Listing Website Content Using the Spider Script”](#) section on page A-80 and the [“Selecting Live and Pre-Positioned Content Using the Manifest Script”](#) section on page A-82 for instructions on running the scripts.

Listing Website Content Using the Spider Script

In the simplest scenario, the Spider script is pointed to the address of an origin server and given the name of a database (.db) file into which it places any valid URLs it discovers on that site. For example, if you wanted to analyze the contents of www.cisco.com for content that might be pre-positioned after the manifest file is created, you would issue the following command:

```
perl spider.pl --start=www.cisco.com --db=ciscocontent.db
```

Limiting or Broadening the Scope of the Spider Script

Using the Spider script on the whole of www.cisco.com might take hours and produce much more information than you need. The Spider script contains a variety of tools that enable you to limit as well as broaden the scope of a spider's action.



Note

When running the Spider script on large websites, you must plan for the long period of time and the large amount of memory that is required for the Spider script to create a database.

For example, to limit the Spider script's search of www.cisco.com to just that part of the server containing product-related support information, you could enter the following command:

```
perl spider.pl --start=www.cisco.com/public/support/ --db=ciscocontent.db
```

To ask the Spider script to follow links from www.cisco.com to the Cisco networking professionals forum, you could enter the following Spider script command:

```
perl spider.pl --start=www.cisco.com --accept=business.cisco.com --db=ciscocontent.db
```

Spider Script Syntax Guidelines

The Spider script accepts the following syntax, as described in [Table A-26](#).

```
perl spider.pl {--start=origin_server_url [ --accept=accept_url | --depth=number | --file=filename |
--limit=number | --prefix=url_prefix | --map=url_prefix | --reject=disallowed_url |
--db=database_name.db]}
```

Table A-26 Spider Script Keywords

Keyword	Description	Command-Line Syntax
--start (required)	Names the location (URL) of the origin server that is to be analyzed.	--start=www.cisco.com
--accept (optional)	Names a location other than that specified using the start keyword that is to be accepted when it is found in URLs.	--accept=forums.cisco.com Note --accept is a more general command that can include regular expressions. For example, you can use "jobs.*tech" to accept any URLs with the string "jobs" followed by "tech."
--depth (optional)	Causes the Spider script to stop after following links to a specified number of levels deep on the origin server.	--depth=6
--file (optional)	Causes the Spider script to read its commands from a specified file, in this case the rules file, one line at a time.	--file=cisco-rules.cfg
--limit (optional)	Causes the Spider script to stop after retrieving a specified number of pages from the origin server. Specifying 0 sets no limit for the number of pages retrieved.	--limit=1000
--prefix (optional)	Specifies a URL prefix that is matched by the Spider script. The --prefix keyword is a convenient option that accepts a fixed string that it must match from the beginning, such as "http://www.cisco.com/jobs."	--prefix=http://www.cisco.com/partners/CDN/
--map (optional)	Causes the Spider script to substitute the second URL prefix (appearing after the second =) for the first in any URLs from the origin server. Or, for example, causes the Spider script to substitute the first prefix for the second when rerunning the Spider script on an origin server if links were modified to go to the Content Engine.	--map=http://www.cisco.com/public/support/tac=/support
--reject (optional)	Names a location that is rejected when it is found in URLs. Note The order in which the --accept and --reject keywords are given to the Spider script determines precedence. The first match takes precedence.	--reject=/cgi-bin
--db (optional)	Names the database file in which content URLs from the origin server and any accepted locations are to be placed.	--db=ciscocontent.db

Customizing the Spider Script

Because the Spider script anticipates certain platforms and scenarios that might not correspond to your own website configuration, Cisco provides you with the Perl source code for the Spider script, which you can modify to suit your own needs.

Selecting Live and Pre-Positioned Content Using the Manifest Script

Whereas the Spider script is used to gather a list of potential content from an origin server, the Manifest Script sifts through this information gathered by the Spider script and decides which content to actually import to the ACNS network for placement on a Content Engine.

Pre-Positioned Versus Live Content

The Manifest script distinguishes between content that needs to be pre-positioned and live, streamed content that, by definition, cannot be pre-positioned.

The result of using the **live** command is nearly the same as that of using the **prepos** command. Both commands expect you to specify what you intend to deliver as live content or to deliver as pre-positioned content with `--prepos=match()` or `--prepos=type()`. The only difference between these two commands is the tags contained in the *.xml* file that is created by *manifest.pl*. If you use the **prepos** command, then the *.xml* file that is created contains the tag `<item-group type="prepos">`. If you use the **live** command, then the *.xml* file contains the tag `<item-group type="wmt-live">` or `<item-group type="real-live">`, depending on whether the streaming data is RealMedia or WMT.

By using the **prepos** command, you identify and pre-position content that meets criteria that you specify. For example, to pre-position image files from Cisco.com that are larger than one megabyte, you would enter the following command:

```
perl manifest.pl --prepos='type(image/*) and size > 1000k' --db=ciscocontent.db
--xml=cisco.xml
```

By using the **live** command, you identify the URLs of live content. Unlike pre-positioned content, live content cannot be identified by information stored in the header, so you must devise a method of locating live content based solely on information contained in the URL of that content. For example, you can identify streamed content with the following command:

```
perl manifest.pl --live='match(http://*)'
```

Manifest Script Syntax Guidelines

The Manifest script accepts the following syntax, as described in [Table A-27](#).

```
perl manifest.pl [--file=filename | {--live='keyword_comparison' | --prepos='keyword_comparison'} |
--set='attribute=value : keyword_comparison' | --playservertable=filename |
--map=origin_server_url_prefix=cdn_prefix | --db=database_name.db |
--xml=manifest_file_name.xml]
```



Note

At least one of the `--prepos` or `--live` keywords is needed for the manifest file to be created from the Spider database. If you do not use at least one or both of these keywords, the manifest file created will be minimal and will not contain any content URLs.

The `--prepos` keyword can be used with either `--type()` or `--match()` (see [Table A-28](#)), which perform different functions. The `--match()` keyword qualifier is a text match and acts on the name of the URL. For example, to call *jpeg* files named *a.jaypeg*, use the `--match(*.jayepeg)` keyword qualifier. Another example would be to use the `--match()` keyword qualifier to find *news* in the name of the URL. The `--match()` keyword qualifier can also be used as shown in the following example:

```
perl manifest.pl --db=name.db --prepos='match(*.jpg)' --xml=xmlname.xml
```

The `--type()` keyword qualifier is used for comparing the content named in the database file to the content-type header returned by the web server. It informs the client of the object MIME type. For example, if you name your *jpeg* files **.jayscale*, the web server returns “Content-Type: image/jpeg,” which is then placed in the database.

Other examples of the `--type()` keyword qualifier include the following:

```
--prepos=type(text/html)
--prepos=type(text/plain)
--prepos=type(application/pdf)
--prepos=type(image/gif)
--prepos=type(image/jpeg)
--prepos=type(video/mpeg)
```

The following are two examples of using the `--type()` keyword qualifier in the full command line:

```
perl manifest.pl --db=name.db --prepos="type(image/jpeg)" -xml=xmlname.xml
perl manifest.pl --db=name2.db --prepos="type(application/pdf)" xml=xmlname2.xml
```



As a rule of thumb, you must use quotes only in the command line. You do not need to use quotes within a rules file for the `--file` keyword.

If the `--prepos` keyword is used in the full command line, then quotes are needed as follows:

- Windows 2000
Use double quotes instead of single quotes as shown in the preceding example.
- Linux
Use single quotes instead of double quotes.

If the `--prepos` argument is used within a rules file with the `--file` argument, then you can modify the file because quotes are not required, as shown in the following *rules.cfg* file example:

```
--start=www.cisco.com
--accept=forums.cisco.com
--reject=/cgi-bin
--limit=0
--db=ciscocontent.db
--prepos=match(image/gif) and size > 1000k
--xml=ciscomanifest.xml
```

If the quotes are not removed from within the rules file, the following message is given: “Bareword found where operator expected.”

Table A-27 **Manifest Script Keywords**

Keyword	Description	Command-Line Syntax
<code>--file</code>	Causes the Manifest script to read its commands from a specified file, one line at a time.	<code>--file=ciscocontent.cfg</code>
<code>--live</code>	Marks content URLs in the database file that match the terms of the keyword comparison as live content (<code>live=type</code>) in the manifest file.	

Table A-27 Manifest Script Keywords (continued)

Keyword	Description	Command-Line Syntax
--prepos	Marks content URLs in the database file that match the terms of the keyword comparison as pre-positioned content (prepos=type) in the manifest file. Note The <i>type</i> command matches the <i>Content-Type</i> : field in the Spider database file.	--prepos='type(image/jpeg) and size > 1000k'
--set	Sets the specified attribute to the value provided for all content items with URLs in the database file that match the keyword comparison.	--set='ttl=10000:match(* /urgent/*)'
--playservertable	Adds the playserver table in the specified file to the manifest file. Playserver tables map MIME-type content and filename extensions to specific server types to use (for example, "real" or "wmt") for the content in a specific Content Engine. For the manifest file to validate properly, move the entire playserver table to the beginning of the manifest file, as shown in the following example: <pre><CdnManifest> <playServerTable> ... </playServerTable> <server> ... </server> ... </CdnManifest></pre>	--playservertable=info.txt
--map	Causes the Manifest script to substitute the second URL prefix (appearing after the second =) for the first in any URLs from the origin server. The second URL prefix must have a full path name. The --map keyword is used to change the names in the .xml file that uses the <i>manifest.pl</i> . When you run the <i>manifest.pl</i> , you should see <item> tags that have the <i>cdn-url</i> attribute set to the requested name.	--map=http://www.cisco.com/public/support/tac=/support
--db	Names the database file in which content URLs from the origin server and any accepted locations are located. This file provides the data that the Manifest script analyzes.	--db=ciscocontent.db
--xml	Names the manifest file that is generated by the Manifest script.	--xml=ciscomanifest.xml

Table A-28 *Manifest Script Keyword Qualifiers*

Qualifier	Description	Command-Line Syntax
match(comparison)	Locates text in content URLs that are identical to a value that is provided.	--prepos='match(http://forums.cisco.com/*)'
size(comparison)	Identifies content named in the database file according to the specified file size parameter, which can be specified in kilobytes, megabytes, or gigabytes (k, KB, m, MB, g, GB).	--prepos=match(*.gif) and size > 1000k
modtime(comparison)	Identifies content named in the database file according to the time since the content was last modified (in hours). Note Do not use spaces before the word “hours,” whether used within a rules file or not. Note If using the time() keyword within a rules file, then do not use quotes. In the syntax example, the “modtime” is compared to “now” – <value>, where “now” is the current time (in seconds since 1970). The <value> is a unit of time in hours. For example, if “modtime” is the current time minus 2 hours, it would be expressed in this way: <pre>modtime < (now - 2hours)</pre>	The following example shows a .cfg file with the time comparison using the “>” character: <pre>--start=http://website.com --db=time1.db --limit=10000 --xml=time1.xml # this works for bitmaps modified within 2 hours --prepos=type(image/bmp) and modtime > (now - 2hours)</pre> The following example shows a .cfg file with the time comparison using the “<” character: <pre>--start=http://website.com --db=time2.db --limit=10000 # this works for bitmaps that have NOT been modified within last 2 hours --prepos=type(image/bmp) and modtime < (now - 2hours) --xml=time2.xml</pre>
type(comparison)	Identifies content named in the database file according to its MIME type (text, application, image, and so on).	--prepos='type(image/gif)'

Customizing the Manifest Script

Because the Manifest script anticipates certain platforms and scenarios that might not correspond to your own website configuration, Cisco provides you with the Perl source code for the Manifest script, which you can modify to suit your own needs.

Creating a Rules File for the Spider and Manifest Scripts

When using the Spider and Manifest scripts on a large web server, the parameters and rules you set for your scripts may be numerous and complex. When this is the case, it is more practical to create a separate file containing a list of your customized rules. Then you can simply point to the applicable rule rather than having to enter a long series of commands every time you want the rule applied.

Using a rules file facilitates rerunning of the Spider and Manifest scripts and ensures that the scripts are receiving identical commands each time the scripts are run. In addition, the same commands file can be read by both the Spider and the Manifest scripts without generating output errors. The Spider script simply ignores commands for the Manifest script, and vice versa.

To create a rules file for the Spider and Manifest scripts, follow these steps:

Step 1 Open your text editor.

Step 2 Enter your commands one at a time, each on its own line.

Each line of your rules file is sent to the scripts as a single argument. The following example shows a rules file for the Cisco website:

```
--start=www.cisco.com
--accept=forums.cisco.com
--reject=/cgi-bin
--limit=0
--db=ciscocontent.db
--prepos=match(image/gif) and size > 1000k
--xml=ciscomanifest.xml
```

Step 3 Save your file in a location relative to the Spider and Manifest scripts.

Step 4 Use the **--file** command to run each script using your rules file, as shown in this example:

```
perl spider.pl --file=cisco-rules.cfg
perl manifest.pl --file=cisco-rules.cfg
```

Using the ACNS Manifest Generator Tool

The ACNS Manifest Generator tool has been developed for ACNS 5.1 software and later releases. The purpose of this tool is to provide a user-friendly interface that any novice manifest file author can use to create a basic manifest file.

The tool automatically validates all user input and generates a manifest file in XML format free of syntax errors. You can save your working copy of the manifest to any desired location and reopen the file to make changes or additions.

The tool supports a limited set of XML tags; however, once the manifest file is generated, it can be modified and additional tags can be added to the XML file manually by using any text or XML editor.

Creating a manifest using the ACNS Manifest Generator tool includes the following tasks:

- Setting the host (origin server)
- Setting global settings for content items and crawl tasks
- Adding, deleting, or editing content items and crawl tasks
- Saving the manifest attributes in the ACNS Manifest Generator
- Generating the XML manifest file

Supported Platforms

The following platforms are supported for the ACNS Manifest Generator:

- Windows 98
- Windows 2000
- Windows XP

Installing the ACNS Manifest Generator


To install the ACNS Manifest Generator, follow these steps:

-
Step 1	Insert the installer CD-ROM into your CD-ROM drive.
Step 2	From the Start menu, choose Run... and enter D:\Setup.exe , replacing “D:” with your respective CD-ROM drive.
Step 3	To install the program, click OK .
-

Understanding the ACNS Manifest Generator User Interface

The ACNS Manifest Generator includes standard Windows-style GUI navigation menus, toolbars, and icons. The main window displays configured content items and crawl tasks in a split window. Lists of items and their properties can be sorted by column by clicking the column heading. (See [Figure A-2](#).)

Figure A-2 Manifest Generator Tool Main Window



1	Menu bar	4	Sortable column headings
2	Toolbar	5	List view for content items
3	Dockable global settings dialog bar	6	List view for crawl tasks

Creating a Basic Manifest File Using the ACNS Manifest Generator

To create a basic manifest file, follow these steps:

- Step 1

From the ACNS Manifest Generator, choose **File > New** to begin creating a new manifest file. The Global Manifest Settings window appears. (See [Figure A-3](#).)

Figure A-3 Global Manifest Settings

Global Manifest Settings

Host

Domain Name of Origin Server: *

Protocol Used to Fetch Content: Content Type:

Time Between Update Intervals: minutes
Range [-1 to 8000]

Authentication

User Name:

Password:

Content Serving Time

Start Time: ... yyyy-mm-dd hh:mm:ss

Stop Time: ... yyyy-mm-dd hh:mm:ss

Note

1. Values entered for attributes 'Time Between Update Intervals', 'Content Type', 'Start Time' and 'Stop Time' fields will be propagated to Content Item and Crawl Task Dialogs

2. * Required

OK Cancel

Step 2 Configure global manifest file settings:

- a. In the field provided, enter the domain name of the origin server.
- b. From the Protocol Used to Fetch Content drop-down list, choose a protocol. The default is HTTP. The choice of protocol is dependent on the content type, as follows:
 - Pre-positioned—HTTP, HTTPS, FTP
 - WMT (Live)—RTSP or HTTP
 - Real (Live)—RTSP
- c. From the Content Type drop-down list, choose a content type. The choices are **Pre-positioned**, **WMT (Live)**, or **Real (Live)**. The default is **Pre-positioned**.
- d. In the Time Between Update Intervals field, enter a value in minutes.
- e. For authentication, enter a username in the User Name field, and then enter a password in the Password field.
- f. Using the format yyyy-mm-dd hh:mm:ss, enter a start time and a stop time, or you can use the date-time picker dialog box. To configure a value using the date-time picker dialog box, click the “...” button next to the text fields. If you click the down arrow in the Select Date field, you can choose a date from the calendar that pops up.
- g. To save the settings, click **OK**. The global settings appear in the Manifest Generator tool ManifestTool window Global Settings bar and are also propagated to the Content Item and Crawl Task dialog boxes.

An item group is created with these settings and all items or crawl tasks that you add are automatically associated with this item group.

Table A-29 describes the global manifest settings fields and their corresponding XML elements and attributes.


Table A-29 Global Manifest Settings

Field Name	Description
Domain Name of Origin Server* ¹	Defines a web server or live server from which content is to be retrieved and later pre-positioned. Associated with the full URL of the <i>src</i> and <i>start-url</i> attributes of the <item> and <crawler> elements, respectively.
Protocol Used to Fetch Content	Defines the protocol used to retrieve the content from the origin server. Associated with the full URL of the <i>src</i> and <i>start-url</i> attributes of the <item> and <crawler> elements, respectively. Possible values are HTTP, HTTPS, or FTP. The default is HTTP.
Content Type	Defines the type of content to be retrieved from the origin server. Associated with the <i>type</i> attribute of the <item-group> element. Possible values are Pre-positioned, WMT (Live), or Real (Live). The default is Pre-positioned.
Time Between Update Intervals	Designates a time interval, in minutes, for revalidation of the content. This value determines how often the content at the origin server is checked for freshness. Associated with the <i>ttl</i> attribute of the <item-group> element. The valid range is –1 to 8000 minutes.
User Name	Defines a user for authentication purposes. Associated with the full URL of the <i>src</i> and <i>start-url</i> attributes of the <item> and <crawler> elements, respectively.
Password	Sets a password for authentication purposes. Associated with the full URL of the <i>src</i> and <i>start-url</i> attributes of the <item> and <crawler> elements, respectively.
Start Time	Designates a time in yyyy-mm-dd hh:mm:ss at which the Content Engine is allowed to start serving the content. If the serving start time is omitted, content is ready to serve once it is distributed to the Content Engine. Associated with the <i>serveStartTime</i> attribute of the <item-group> element.
Stop Time	Designates a time in yyyy-mm-dd hh:mm:ss format at which the ACNS network temporarily stops serving the content. If the serving stop time is omitted, the ACNS network serves the content to the Content Engine until it is removed by modifying the manifest file or renaming the channel. Associated with the <i>serveStopTime</i> attribute of the <item-group> element.
OK	Saves the changes to the Global Settings bar.
Cancel	Ignores any unsaved changes in the Global Settings bar.

1. An asterisk (*) denotes a required field.

Step 3 Add a manifest item:**a.** Choose **Edit > Add Content Item** or **Add Crawl Task**.

- If you choose **Add Content Item**, the Content Item Properties dialog box appears. (See [Figure A-4](#).)


Figure A-4 Content Item Properties Dialog


The **Content Item** dialog box contains the following fields:

- Content Item**
 - Source URL of Content Item: *
 - Acquisition Priority of this Item: 0 (Range [0 to 2147483646])
 - Content Type: Pre-positioned (dropdown)
 - Time Between Update Intervals: 10 minutes (Range [-1 to 8000])
- Content Serving Time**
 - Start Time: 2003-07-28 19:05:23 (format: yyyy-mm-dd hh:mm:ss)
 - Stop Time: 2003-07-28 19:05:26 (format: yyyy-mm-dd hh:mm:ss)

Buttons at the bottom: Ok/Add New, OK, Cancel.

- If you choose **Add Crawl Task**, the Crawl Task dialog box appears. (See [Figure A-5](#).)

Figure A-5 Crawl Task Dialog


The **Crawl Task** dialog box has two tabs: **Specification** and **Match Rules**. The **Specification** tab contains the following fields:

- Crawl Task**
 - Start URL of Crawl Task : *
 - Acquisition Priority of this Crawl: 0 (Range [0 to 2147483646])
 - Content Type: Pre-positioned (dropdown)
 - Time Between Update Intervals: 10 minutes (Range [-1 to 8000])
- Content Serving Time**
 - Start Time: 2003-07-28 19:05:23 (format: yyyy-mm-dd hh:mm:ss)
 - Stop Time: 2003-07-28 19:05:26 (format: yyyy-mm-dd hh:mm:ss)
- Crawl Limits**
 - Crawl Website up to Level: 20 (Range [-1 to 2147483646])
 - Acquire URLs Matching Prefix: http://www.cisco.com/
 - Acquire Objects up to Size: 0 (Range [0 to 2147483646])

Buttons at the bottom: OK, Cancel.

Step 4 To add a single content item, enter the following information in the Content Item Properties dialog box:**a.** In the Source URL of Content Item field, enter the source URL.**Tip**

Right-click to copy and paste URLs from a web browser to the Source URL field.

- b. In the Acquisition Priority of this Item field, specify the priority of acquisition.
- c. From the Content Type drop-down list, choose a content type, if you have not already set this attribute in the Global Manifest Settings dialog box. (See [Figure A-3](#).)
- d. In the Time Between Update Intervals field, designate a time interval in minutes for revalidation of content, if you have not already set this attribute in the Global Manifest Settings dialog box.
- e. In the Global Manifest Settings dialog box, enter the content serving start and stop times, if you have not already set this attribute.
- f. To add the item to the list of content items in the main window, click **OK** or **Ok/Add New**. (See [Figure A-2](#).)

[Table A-30](#) describes the fields in the Content Item Properties dialog box.

Table A-30 Content Item Properties Dialog Box Fields

Field Name	Description
Source URL of Content Item* ¹	Identifies the specific content that is to be acquired from the host server (also referred to as the origin server). Associated with the <i>src</i> attribute of the <item> element. The order in which the items are entered determines their order in the manifest.
Acquisition Priority of this Item	Specifies the order of importance and hence the processing order of the content item. The range is 0 to 2147483636. Associated with the <i>priority</i> attribute of the <item> element.
Content Type	Defines whether content is to be pre-positioned or live on the ACNS network. Possible values are Pre-positioned, WMT (Live), or Real (Live). The default is Pre-positioned. Note If you want this property to be set globally, configure it in the Global Manifest Settings dialog box. If you configure this property in the Content Item Properties dialog box, the type you choose here takes precedence over the global setting. Associated with the <i>type</i> attribute of the <item> element.
Time Between Update Intervals	Designates a time interval, in minutes, for revalidation of the content. Associated with the <i>ttl</i> attribute of the <item> element. The valid range is -1 to 8000 minutes. Note If you want this property to be set globally, configure it in the Global Manifest Settings dialog box. If you configure this property in the Content Item Properties dialog box, the value you enter here takes precedence over the global setting.
Start Time	Designates a time in yyyy-mm-dd hh:mm:ss format at which the Content Engine is allowed to start serving the content. If the serving start time is omitted, content is ready to serve once it is distributed to the Content Engine. Associated with the <i>serveStartTime</i> attribute of the <item> element.

Table A-30 Content Item Properties Dialog Box Fields (continued)

Field Name	Description
Stop Time	Designates a time in yyyy-mm-dd hh:mm:ss format at which the ACNS network temporarily stops serving the content. If the serving stop time is omitted, the ACNS network serves the content to the Content Engine until it is removed by modifying the manifest file or renaming the channel. Associated with the <i>serveStopTime</i> attribute of the <item> element.
Ok/Add New or OK	Adds this item configuration to the list of content items after validations.
Cancel	Ignores any unsaved changes in the Content Item Properties dialog box and closes the dialog box window.

1. * = required field. All other fields are optional.

- Step 5** To add a crawl task to the manifest file, enter the following information in the Crawl Task dialog box:
- Choose the **Specification** tab, and enter the crawl task configuration information:
 - In the Start URL of Crawl Task field, enter the starting URL. This field is required.
 - In the Acquisition Priority of this Crawl field, specify the priority of acquisition.
 - In the Time Between Update Intervals field, designate a time interval in minutes for revalidation of content, if you have not already set this attribute in the Global Manifest Settings dialog box.
 - Define the content serving times by entering a start time and a stop time, if you have not already set this attribute in the Global Manifest Settings dialog box.
 - Define the crawl limits by entering depth, prefix, and maximum size. These fields are optional.
 - In the Crawl Website up to Level field, enter a depth value.
 - Enter a prefix that you want appended to the host name.
The specified prefix is appended to the host name for the purpose of specifying a full host name to be matched. Only URLs that match the full *hostname/prefix* are crawled.
 - Set a maximum object size that will be required, and choose a unit of measure (bytes, kilobytes, or megabytes).

[Table A-31](#) provides Crawl Task dialog box field information.

Table A-31 Crawler Task Properties

Field Name	Description
Start URL of Crawl Task	Defines the URL to start from for the specified crawl job. Associated with the <i>start-url</i> attribute of the <crawler> element. Note The order in which crawler items are entered determines their order in the manifest.
Acquisition Priority of this Crawl	Specifies the order of importance and hence the processing order of the crawl job. The range is 0 to 2147483636. Associated with the <i>priority</i> attribute of the <crawler> element.

Table A-31 Crawler Task Properties (continued)


Field Name	Description
Content Type	The only option is Pre-positioned. Live content types are not supported in crawl jobs. Associated with the <i>type</i> attribute of the <crawler> element.
Time Between Update Intervals	Designates a time interval, in minutes, for revalidation of the content. Associated with the <i>ttl</i> attribute of the <crawler> element. The valid range is -1 to 8000 minutes. Note If you want this property to be set globally, configure it in the Global Manifest Settings dialog box. If you configure this property in the Crawl Task dialog box, the value you enter here takes precedence over the global setting.
Start Time	Designates a time in yyyy-mm-dd hh:mm:ss format at which the Content Engine is allowed to start serving the content. If the serving start time is omitted, content is ready to serve once it is distributed to the Content Engine. Associated with the <i>serveStartTime</i> attribute of the <crawler> element.
Stop Time	Designates a time in yyyy-mm-dd hh:mm:ss format at which the ACNS network temporarily stops serving the content. If the serving stop time is omitted, the ACNS network serves the content to the Content Engine until the content is removed by modifying the manifest file or renaming the channel. Associated with the <i>serveStopTime</i> attribute of the <crawler> element.
Crawl Website up to Level	Defines the link depth to which a website is to be crawled or directory depth to which an FTP server is to be crawled. The default depth is 20 levels. The range is -1 to 2147483636. If depth = -1, there is no depth constraint. If depth = 0, acquire only the start URL. If depth = 1, acquire the start URL and all its referred contents. Associated with the <i>depth</i> attribute of the <crawler> element.
Acquire URLs Matching Prefix	Appends the specified prefix to the host name. Only URLs that match the full <i>hostname/prefix</i> are crawled. If omitted, the host name of the server is used as the full URL prefix. Associated with the <i>prefix</i> attribute of the <crawler> element.
Acquire Objects up to Size	Specifies the maximum size of content to be acquired, in bytes, kilobytes, or megabytes. The range is 0 to 2147483636. The unit is selected from the drop-down list. Associated with the <i>maxTotalSizeInB/KB/MB</i> attribute of the <crawler> element.

- d. Click the **Match Rules** tab. The list of Match Rule entries appears. (See [Figure A-6](#).)

Items appearing in this window can be edited or deleted. By selecting one row at a time, you activate the **Edit** button. Click the **Edit** button or double-click the entry, and the Match Rules dialog box pops up, showing the values of the selected entry.


To delete an entry, select at least one row to activate the **Delete** button, and then click **Delete**.

Figure A-6 Match Rules Tab—List View



- e. Click **Add** to configure match rules for the crawl task. The Match Rule dialog box appears. (See [Figure A-7](#).)

Figure A-7 Match Rule Dialog Box



[Table A-32](#) describes the field information you need to configure match rules.

Table A-32 Match Rules Field Descriptions

Field Name	Description
MIME Type	Specifies the MIME type to be acquired. The object is acquired only if its MIME type matches the MIME type configured here. Associated with the <i>mime-type</i> attribute of the <match> element.
Extension	Specifies a file extension. The object is acquired only if its extension matches the extension configured here. Associated with the <i>extension</i> attribute of the <match> element.
Time After	Sets a date and time limit after which modified content is to be acquired. The object is acquired only if its Last Modified Date and Time is after the value specified here. The format of this field is yyyy-mm-dd hh:mm:ss. Associated with the <i>time-after</i> attribute of the <match> element.
Time Before	Sets a date and time limit before which modified content is to be acquired. The object is acquired only if its Last Modified Date and Time is before the value specified here. The format of this field is yyyy-mm-dd hh:mm:ss. Associated with the <i>time-before</i> attribute of the <match> element.
Minimum Size	Specifies that the acquired content size must be larger than or equal to this number, in bytes, kilobytes, or megabytes. The range for this field is 0 to 2147483636. Associated with the <i>minFileSizeInB/KB/MB</i> attribute of the <match> element.
Maximum Size	Specifies that the acquired content size must be smaller than or equal to this number, in bytes, kilobytes, or megabytes. The range for this field is 0 to 2147483636. Associated with the <i>maxFileSizeInB/KB/MB</i> attribute of the <match> element.

- f. Click **OK** in the Match Rules dialog box to save the configuration, or click **Ok/Add New** to save the configuration and add another set of match rules. An entry is added to the main window listing crawl tasks.
- g. When you have finished adding match rules, click **OK** in the Crawl Task dialog box to save your crawl task configurations and close the dialog box.

Step 6 Choose **File > Save** or **Save As** to save your working manifest file. The file is saved with a proprietary extension (.amg) that can only be read by the ACNS Manifest Generator tool. The default directory for saving the working manifest is the AMG Files directory.

You must save your work in the ACNS Manifest Generator if you want to edit the file at a later time using the tool. You cannot open a generated XML manifest file in the ACNS Manifest Generator for editing; however, you can edit the generated file using any plain text editor.

Step 7 Generate the manifest file.

- a. To generate the manifest in XML format with an *.xml extension, choose **File > Generate XML**. You must associate an XML viewer to view the file.

- b. To view the generated manifest file, choose **File > Preview XML**. This option launches your default XML viewer and displays the generated XML file. The Windows default viewer is Internet Explorer.

When you generate the manifest file, the Save As dialog box appears. You can save the generated XML file to any desired location. The default local directory for saving the generated XML file is the XML Files directory.

- Step 8** Save the generated XML manifest file to a unique location on an origin server that can be accessed by the Content Distribution Manager.

The Content Distribution Manager is configured to access the manifest file when you enter the manifest file URL in the Modifying Channel window of the Content Distribution Manager GUI. (See the [“Creating and Modifying Channels” section on page 5-10.](#))

Elements and Tags Not Supported in the ACNS Manifest Generator

The following elements and tags are considered advanced settings and cannot be defined using the ACNS Manifest Generator:

- Options element
 - *timeZone*—string; optional; default is GMT/UTC
 - *notFoundURL*—string; optional
 - *alternateURL*— string; optional
 - *noRedirectToOrigin*—boolean; default is false
 - *requireAuth*—boolean; optional
 - *playDuration*—numeric; optional
 - *bitrate-in-bps*—numeric; optional
 - *bitrate-in-kbps*— numeric; optional
 - *bitrate-in-mbps*—numeric; optional
 - *ttl*—numeric; optional
 - *failRetryInterval*—numeric; optional
 - *prefetch*—string; optional
 - *type*—select from prepos, wmt-live or real-live; the default is prepos
- Server or host element:
 - *<server> name*—string; required; set to be the same as the *<host> name* attribute
 - *proxyServer*—string; optional
 - *uuencoded*—boolean, optional; the default is false
 - *sslAuthType*—select from weak or strong; optional; the default is strong
- Item element:
 - *contains* element
 - *http-meta-data* element
 - *cdn-url*—string; optional
 - *server*—string; optional

- *proxyServer*—string; optional
- *playServer*—string; optional
- *playDuration*—numeric; optional
- *bitrate-in-bps*—numeric; optional
- *bitrate-in-kbps*—numeric; optional
- *bitrate-in-mbps*—numeric; optional
- *noRedirectToOrigin*—boolean; optional
- *requireAuth*—boolean; optional
- *alternateUrl*—string; optional
- *failRetryInterval*—numeric; optional
- *prefetch*—string in the format yyyy-mm-dd hh:mm:ss; optional
- *expires*—string in the format yyyy-mm-dd hh:mm:ss; optional
- *serveStartTime* and *serveStopTime*—string in the format yyyy-mm-dd hh:mm:ss; optional
- Crawler element:
 - *http-meta-data* element
 - *accept*—string; optional
 - *reject*—string; optional
 - *max-number*—numeric; optional
 - *server*—string; optional
 - *proxyServer*—string; optional
 - *playServer*—string; optional
 - *playDuration*—numeric; optional
 - *bitrate-in-bps*—numeric; optional
 - *bitrate-in-kbps*—numeric; optional
 - *bitrate-in-mbps*—numeric; optional
 - *noRedirectToOrigin*—boolean; optional
 - *requireAuth*—boolean; optional
 - *alternateUrl*—string; optional
 - *failRetryInterval*—numeric; optional
 - *prefetch*—string in the format yyyy-mm-dd hh:mm:ss; optional
 - *expires*—string in the format yyyy-mm-dd hh:mm:ss; optional
 - *serve*—string in the format yyyy-mm-dd hh:mm:ss; optional
 - *srcPrefix* and *cdnPrefix* elements
- Item-group element:
 - *matchRule* element
 - *http-meta-data* element
 - *item* element
 - *server*—string; optional

- *proxyServer*—string; optional
- *playServer*—string; optional
- *playDuration*—numeric; optional
- *bitrate-in-bps*—numeric; optional
- *bitrate-in-kbps*—numeric; optional
- *bitrate-in-mbps*—numeric; optional
- *noRedirectToOrigin*—boolean; optional
- *requireAuth*—boolean; optional
- *alternateUrl*—string; optional
- *failRetryInterval*—numeric; optional
- *priority*—numeric; optional
- *prefetch*—string in the format yyyy-mm-dd hh:mm:ss; optional
- *expires*—string in the format yyyy-mm-dd hh:mm:ss; optional
- *serve*—string in the format yyyy-mm-dd hh:mm:ss; optional
- *srcPrefix* and *cdnPrefix* elements
- *proxyServer* element—whole element is not supported
- *playServer* element—whole element is not supported



IP Multicast Addressing

Multicast delivery is enabled by setting up a multicast address on the Content Engine in the form of a multicast cloud configuration to which different devices, configured to receive content from the same channel, can subscribe. The delivering device sends content to the multicast address set up at the Content Engine, from which it becomes available to all subscribed receiving devices.

This appendix provides guidelines for IP multicast addressing and contains the following sections:

- [Unusable IP Multicast Addresses, page B-1](#)
- [Insecure Services, page B-4](#)
- [Copying Files Between Servers and Clients, page B-4](#)
- [Limited Scope Addresses, page B-4](#)
- [Source-Specific Multicast Addresses, page B-5](#)
- [GLOP Addresses, page B-5](#)

Use this information when configuring multicast address settings for multicast clouds. (See the [“Configuring Multicast Cloud Properties”](#) section on page 5-34.)

Unusable IP Multicast Addresses

The Internet Assigned Numbers Authority (IANA) controls the assignment of IP multicast addresses. IANA has assigned the IPv4 Class D address space to be used for IP multicast. Therefore, all IP multicast group addresses fall in the range from 224.0.0.0 through 239.255.255.255. However, some combinations of source and group addresses should not be routed for multicasting purposes.

[Table B-1](#) lists the unusable multicast address ranges and the reasons they should not be used.

Table B-1 **Unusable Multicast Address Assignments**

Address Range	Reason
224.0.1.2/32	Known insecure service address. See the “Insecure Services” section on page B-4.
224.0.1.3/32	Reserved for the discovery of resources within the administrative domain. See the “Limited Scope Addresses” section on page B-4.
224.0.1.22/32	Known insecure service address.
224.0.1.35/32	Reserved for the discovery of resources within the administrative domain.
224.0.1.39/32	Reserved for the discovery of resources within the administrative domain.

Table B-1 *Unusable Multicast Address Assignments (continued)*

Address Range	Reason
224.0.1.40/32	Reserved for the discovery of resources within the administrative domain.
224.0.2.2./32	Known insecure service address.
224.77.0.0/16	Used to copy files between servers and clients in a local network. See the “Copying Files Between Servers and Clients” section on page B-4.
224.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches. See the “Layer 2 Multicast Addresses” section on page B-5.
225.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
225.1.2.3/32	Used to copy files between servers and clients in a local network.
225.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
226.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
226.77.0.0/16	Used to copy files between servers and clients in a local network.
226.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
227.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
227.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
228.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
228.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
229.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
229.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
230.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
230.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
231.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
231.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
232.0.0.0/24	Source-specific multicast address. See the “Source-Specific Multicast Addresses” section on page B-5.
232.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.

Table B-1 *Unusable Multicast Address Assignments (continued)*

Address Range	Reason
233.0.0.0/8	GLOP address. See the “GLOP Addresses” section on page B-5.
233.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
233.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
234.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
234.42.42.42/32	Used to copy files between servers and clients in a local network.
234.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
234.142.142.42/31	Used to copy files between servers and clients in a local network.
234.142.142.44/30	Used to duplicate files between clients and servers in a local network.
234.142.142.48/28	Used to copy files between servers and clients in a local network.
234.142.142.64/26	Used to copy files between servers and clients in a local network.
234.142.142.128/29	Used to copy files between servers and clients in a local network.
234.142.142.136/30	Used to copy files between servers and clients in a local network.
234.142.142.140/31	Used to copy files between servers and clients in a local network.
234.142.142.142/32	Used to copy files between servers and clients in a local network.
235.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
235.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
236.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
236.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
236.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
236.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
237.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
237.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
238.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
238.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.

Table B-1 **Unusable Multicast Address Assignments (continued)**

Address Range	Reason
239.0.0.0/8	Administratively scoped address that should not be passed between administrative domains. See the “ Limited Scope Addresses ” section on page B-4 .
239.0.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.
239.128.0.0/24	Local address that maps to an Ethernet multicast address range and may overwhelm the mapping table of LAN switches.

Some of these addresses have been reserved for use by multicast applications through IANA. For example, IP address 224.0.1.1 has been reserved for the Network Time Protocol (NTP).

IP addresses reserved for IP multicast are defined in RFC 1112, *Host Extensions for IP Multicasting*. More information about reserved IP multicast addresses can be found at the following location: <http://www.iana.org/assignments/multicast-addresses>.

**Note**

You can find all RFCs and Internet Engineering Task Force (IETF) drafts on the IETF website (<http://www.ietf.org>).

Insecure Services

Applications that use multicast addresses in the 224.0.1.2/32, 224.0.1.22/23, and 224.0.2.2/32 ranges have been demonstrated to be vulnerable to exploitation, which has led to serious security problems.

Copying Files Between Servers and Clients

Some applications are used to copy files between servers and clients and to otherwise maintain groups of personal computers. These applications are intended to be used on a local subnet or within an administrative domain, but the default addresses used by the software are not within the administrative addresses used by the administratively scoped addresses listed in [Table B-1](#).

Limited Scope Addresses

Limited scope addresses are also called administratively scoped addresses. These addresses are described in RFC 2365, *Administratively Scoped IP Multicast*, to be limited to a local group or organization. Companies, universities, or other organizations can use limited scope addresses to have local multicast applications that will not be forwarded outside their domain. Routers typically are configured with filters to prevent multicast traffic in this address range from flowing outside an autonomous system (AS) or any user-defined domain. Within an autonomous system or domain, the limited scope address range can be further subdivided so that local multicast boundaries can be defined. This subdivision is called address scoping and allows for address reuse between these smaller domains.

Source-Specific Multicast Addresses

Addresses in the 232.0.0.0/24 range are reserved for source-specific multicast (SSM). SSM is an extension of the Protocol-Independent Multicast (PIM) protocol that allows for an efficient data delivery mechanism in one-to-many communications.

GLOP Addresses

RFC 2770, *GLOP Addressing in 233/8*, proposes that the 233.0.0.0/8 address range be reserved for statically defined addresses by organizations that already have an AS number reserved. This practice is called GLOP addressing. The AS number of the domain is embedded into the second and third octets of the 233.0.0.0/8 address range. For example, the AS number 62010 is written in hexadecimal format as F23A. Separating the two octets F2 and 3A results in 242 and 58 in decimal format. These values result in a subnet of 233.242.58.0/24 that would be globally reserved for AS 62010 to use.

Layer 2 Multicast Addresses

Historically, network interface cards (NICs) on a LAN segment could receive only packets destined for their burned-in MAC address or the broadcast MAC address. In IP multicast, several hosts need to be able to receive a single data stream with a common destination MAC address. Some means had to be devised so that multiple hosts could receive the same packet and still be able to differentiate between several multicast groups.

One method to accomplish this is to map IP multicast Class D addresses directly to a MAC address. Today, using this method, NICs can receive packets destined to many different MAC addresses—their own unicast, broadcast, and a range of multicast addresses.




Note

The MAC address table on an ACNS 5.x Content Engine supports up to 50,000 entries.

The IEEE LAN specifications made provisions for the transmission of broadcast and multicast packets. In the 802.3 standard, bit 0 of the first octet is used to indicate a broadcast or multicast frame. [Figure B-1](#) shows the location of the broadcast or multicast bit in an Ethernet frame.

Figure B-1 IEEE 802.3 MAC Address Format



This bit indicates that the frame is destined for a group of hosts or all hosts on the network (in the case of the broadcast address 0xFFFF.FFFF.FFFF).


IP multicast makes use of this capability by sending IP packets to a group of hosts on a LAN segment.

Ethernet MAC Address Mapping

IANA owns a block of Ethernet MAC addresses that start with 01:00:5E in hexadecimal format. Half of this block is allocated for multicast addresses. The range from 0100.5e00.0000 through 0100.5e7f.ffff is the available range of Ethernet MAC addresses for IP multicast.


This allocation allows for 23 bits in the Ethernet address to correspond to the IP multicast group address. The mapping places the lower 23 bits of the IP multicast group address into these available 23 bits in the Ethernet address. (See [Figure B-2](#).)

Figure B-2 IP Multicast to Ethernet or FDDI MAC Address Mapping



Because the upper five bits of the IP multicast address are dropped in this mapping, the resulting address is not unique. In fact, 32 different multicast group IDs map to the same Ethernet address. (See [Figure B-3](#).) Network administrators should consider this fact when assigning IP multicast addresses. For example, 224.1.1.1 and 225.1.1.1 map to the same multicast MAC address on a Layer 2 switch. If one user subscribed to Group A (as designated by 224.1.1.1) and the other users subscribed to Group B (as designated by 225.1.1.1), they would both receive both A and B streams. This situation limits the effectiveness of this multicast deployment.

Figure B-3 MAC Address Ambiguities





Mapping ACNS Software CLI Commands to the Content Distribution Manager GUI

This appendix lists the CLI commands that are supported and unsupported in the ACNS 5.5 Content Distribution Manager GUI. It contains the following sections:

- [CLI Commands Supported in the Content Distribution Manager GUI, page C-1](#)
- [CLI Commands Not Supported in the Content Distribution Manager GUI, page C-27](#)

CLI Commands Supported in the Content Distribution Manager GUI

[Table C-1](#) maps ACNS software CLI commands to the corresponding Content Distribution Manager GUI procedure. Software CLI commands are listed in alphabetical order in the first column. The second column indicates the GUI navigation path to the device group configuration window for each command. The third column provides a link to the GUI configuration procedure found elsewhere in this guide. In most cases, the procedure provides an illustration of the configuration window and a table at the end of the procedure that maps each field in the GUI window to its corresponding CLI command.



Note

Most of the procedures in this guide document the configuration steps for individual devices, whereas this table provides the GUI navigation path to the device group configuration window, unless otherwise indicated. The procedures for configuring device groups and individual devices are the same in most cases.

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map

Command	GUI Navigation Path	Procedure Reference
aaa accounting {commands {0 15} default {start-stop stop-only wait-start} tacacs exec default {start-stop stop-only wait-start} tacacs system default {start-stop stop-only} tacacs}	Devices > Device Groups > General Settings > Authentication > AAA Accounting	Configuring AAA Accounting, page 12-27
access-lists 300 {deny groupname {any [position number] groupname [position number]}} {permit groupname {any [position number] groupname [position number]}}	Devices > Devices ¹ > General Settings > Authentication > Access List > Configure Access Control List	Configuring an Access Control List for Group Authorization, page 18-1
access-lists enable	Devices > Devices > General Settings > Authentication > Access List > Enable Access Control List	Enabling the Access Control List, page 18-2
acquirer proxy authentication outgoing ipaddress port_num username [password password] [ntlm domain [basic-auth-disable]]	Devices > Device Groups > Applications > Web > HTTP > HTTP Connections	Configuring Authentication for an HTTP Proxy Using the Content Distribution Manager GUI, page 6-44
acquirer proxy authentication transparent username [password password] [ntlm domain [basic-auth-disable]]	Devices > Device Groups > Prepositioning > Acquirer WCCP Proxy Authentication	Configuring Authentication for a WCCP Proxy Using the Content Distribution Manager GUI, page 6-46
asset tag name	Devices > Device Groups > General Settings > Notification and Tracking > SNMP > Asset Tag	Configuring SNMP Asset Tag Settings, page 22-11
authentication {configuration login} {local radius tacacs} enable primary [secondary tertiary]	Devices > Device Groups > General Settings > Authentication > Login Authentication	Configuring Devices for Login Authentication and Authorization, page 12-7
authentication fail-over server-unreachable		
bandwidth {real-proxy {incoming kbits outgoing kbits} real-server kbits wmt {incoming kbits outgoing kbits} cisco-streaming-engine kbits http kbits} {start-time weekday time end-time weekday time}	Devices > Device Groups > Applications > Bandwidth Schedules	Configuring Scheduled Bandwidth Settings, page 9-42
bandwidth {real-proxy {incoming kbits outgoing kbits} real-server kbits wmt {incoming kbits outgoing kbits} cisco-streaming-engine kbits http kbits} {default max-bandwidth}	Devices > Device Groups > Applications > Default and Maximum Bandwidth	Configuring Content Services Default and Maximum Bandwidth Settings, page 9-41
banner {motd login exec} message message_text banner enable	Devices > Device Groups > General Settings > Login Access Control > Message of the Day	Configuring Message of the Day Settings for the Content Engine, page 20-5

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
bitrate http default <i>bitrate</i>	Devices > Device Groups > Applications > Streaming > Windows Media > General Settings	Configuring Windows Media Proxy Settings for the Content Engine , page 9-18
bitrate wmt incoming <i>bitrate</i>		
bitrate wmt outgoing <i>bitrate</i>		
bypass auth-traffic enable	Devices > Device Groups > Request Routing > WCCP > Bypass	Configuring WCCP Bypass Settings , page 4-23
bypass load { enable in-interval <i>seconds</i> out-interval <i>seconds</i> time-interval <i>minutes</i> }		
bypass static { <i>clientip</i> { <i>serverip</i> any-server } any-client <i>serverip</i> }	Devices > Device Groups > Request Routing > WCCP > Bypass List	Creating WCCP Bypass List Entries , page 4-25
bypass timer <i>minutes</i>	Devices > Device Groups > Request Routing > WCCP > Bypass	Configuring WCCP Bypass Settings , page 4-23
bypass gateway <i>ipaddress</i>		
error-handling { reset-connection send-cache-error transparent }		
cdp enable	Devices > Device Groups > General Settings > Services > CDP	Configuring CDP Settings , page 20-7
cdp holdtime <i>seconds</i>		
cdp timer <i>seconds</i>		
clock timezone <i>timezone</i> <i>offset</i>	Devices > Device Groups > General Settings > Services > Date/Time > Time Zone	Configuring Device Clock and Time Zone Settings , page 20-12
clock summertime <i>timezone</i> { date recurring }		
cms database maintenance full enable	Devices > Device Groups > General Settings > Services > Database Maintenance	See the <i>Cisco ACNS Software Update and Maintenance Guide</i> , Release 5.x
cms database maintenance full schedule { every-day weekday at time }		
cms database maintenance regular enable		
cms database maintenance regular schedule { every-day weekday at time }		
Content Router global configuration command: contentrouting leastloaded	Devices > Devices > Device Activation	Enabling Load-Based Routing on the Content Router , page 4-55
contentrouting servicemonitor { numberofsamples { all <i>number</i> cpu <i>number</i> disk <i>number</i> wmt <i>number</i> } sampleperiod { all <i>seconds</i> cpu <i>seconds</i> disk <i>seconds</i> wmt <i>seconds</i> } threshold { cpu <i>percentage</i> wmt <i>percentage</i> } type { all cpu disk wmt } }	Devices > Device Groups > General Settings > Notification and Tracking > Service Monitor	Configuring Content Engine Threshold Limits for Load-Based Routing , page 4-56
dns min-ttl <i>seconds</i>	Devices > Devices ¹ > Applications > DNS > DNS Caching Server	Configuring DNS Caching Server Settings for the Content Engine , page 8-71
dns max-ttl <i>seconds</i>		
dns user-expired enable		

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
dns max-cache-memory <i>mbytes</i>	Devices > Device Groups > General Settings > Services > DNS	Configuring the DNS Server for HTTP Proxy Caching, page 8-78
	Devices > Devices ¹ > Applications > DNS > DNS Caching Server	Configuring DNS Caching Server Settings for the Content Engine, page 8-71
dns enable	Devices > Device Groups > General Settings > Services > DNS	Configuring the DNS Server for HTTP Proxy Caching, page 8-78
	Devices > Devices ¹ > Applications > DNS > DNS Caching Server	Configuring DNS Caching Server Settings for the Content Engine, page 8-71
dns listen { all <i>ipaddress</i> } port <i>port_num</i> hostname <i>hostname</i>	Devices > Devices ¹ > Applications > DNS > DNS Server Binding	Configuring DNS Server Bindings for the Content Engine, page 8-75
dns use-original-server { only after-configured before-configured }	Devices > Devices ¹ > Applications > DNS > DNS Caching Server	Configuring DNS Caching Server Settings for the Content Engine, page 8-71
dns pin { forward <i>hostname ipaddress</i> reverse <i>hostname ipaddress</i> both <i>hostname ipaddress</i> cname <i>records</i> }	Devices > Device Groups > Request Routing > DNS > A-Record Mapping	Configuring DNS Address Record Mappings for the Content Engine, page 8-76
	Devices > Device Groups > Request Routing > DNS > CNAME Record Mapping	Configuring DNS Canonical Name Record Mappings for the Content Engine, page 8-77
dns retry-period <i>seconds</i>	Devices > Devices ¹ > Applications > DNS > DNS Caching Server	Configuring DNS Caching Server Settings for the Content Engine, page 8-71
dns retry-timeout <i>seconds</i>		
dns serial-lookup	Devices > Device Groups > General Settings > Services > DNS	Configuring the DNS Server for HTTP Proxy Caching, page 8-78
	Devices > Devices ¹ > Applications > DNS > DNS Caching Server	Configuring DNS Caching Server Settings for the Content Engine, page 8-71
exec-timeout <i>minutes</i>	Devices > Device Groups > General Settings > Login Access Control > Exec Timeout	Configuring Exec Timeout Settings, page 20-6
ftp-native access-list { in out } { <i>std_acl_number</i> <i>acl_name</i> }	Devices > Devices > General Settings > Network > IP ACL Feature Usage	Associating an IP ACL with an Application, page 17-12

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
EXEC mode ftp-native custom-message {download {acl-denied <i>url</i> welcome <i>url</i> } reset {acl-denied all welcome} upload <i>host directory filename</i> {welcome acl-denied}}	Devices > Devices > Applications > Web > FTP > FTP Messages > Upload Devices > Device Groups > Applications > Web > FTP > FTP Messages > Download	Configuring Native FTP Custom Messages, page 8-62
ftp-native object max-size <i>kilobytes</i>	Devices > Device Groups > Applications > Web > FTP > Native FTP Cache Freshness	Configuring Native FTP Cache Freshness, page 8-61
ftp-native proxy active-mode enable	Devices > Device Groups > Applications > Web > FTP > Native FTP Connections	Configuring Native FTP Connections, page 8-58
ftp-native proxy authentication enable		
ftp-native proxy incoming <i>portnum</i>		
ftp-over-http age-multiplier directory-listing <i>dl_time file fo_time</i> ftp-over-http max-ttl {days directory-listing <i>dlmax_days file fmax_days</i> hours directory-listing <i>dlmax_hours file fmax_hours</i> minutes directory-listing <i>dlmax_min file fmax_min</i> seconds directory-listing <i>dlmax_sec file fmax_sec</i> }	Devices > Device Groups > Applications > Web > FTP > FTP Over HTTP Cache Freshness	Configuring FTP-Over-HTTP Cache Freshness, page 8-55
ftp-over-http min-ttl <i>min_minutes</i>		
ftp-over-http object max-size <i>size</i>		
ftp-over-http proxy incoming <i>ports</i>		
ftp-over-http proxy outgoing origin-server	Devices > Device Groups > Applications > Web > FTP > FTP Over HTTP Connections	Configuring FTP-Over-HTTP Connection Settings, page 8-52
ftp-over-http proxy outgoing connection-timeout		
ftp-over-http proxy outgoing monitor		
ftp-over-http proxy outgoing host { <i>hostname</i> <i>ipaddress</i> } <i>port</i>		
ftp-over-http proxy anonymous-pswd <i>passwd</i>		
ftp-over-http proxy active-mode enable		
ftp-over-http reval-each-request {all none directory-listing}	Devices > Device Groups > Applications > Web > FTP > FTP Over HTTP Cache Freshness	Configuring FTP-Over-HTTP Cache Freshness, page 8-55
gui-server enable	Devices > Device Groups > General Settings > Login Access Control > GUI Server	Configuring the Content Engine GUI for Secure or Nonsecure Access, page 16-59
gui-server secure enable		
gui-server port <i>port_num</i>		
gui-server secure port <i>port_num</i>		

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
hostname <i>name</i>	Devices > Devices > Device Activation	Modifying Content Engine Properties, page 13-9
http add-method <i>method</i>	Devices > Device Groups > Applications > Web > HTTP > HTTP Method	Adding or Modifying Additional HTTP Request Methods for the Content Engine, page 8-23
http age-multiplier <i>text num binary num</i>	Devices > Device Groups > Applications > Web > HTTP > HTTP Cache Freshness	Configuring HTTP Cache Freshness Settings, page 8-10
http always-resolve-host	Devices > Device Groups > Applications > Web > HTTP > HTTP Caching	Configuring HTTP Cache Settings, page 8-6
http anonymizer enable		
http append via-header	Devices > Device Groups > Applications > Web > HTTP > Authenticated HTTP Caching	Configuring Authenticated HTTP Cache Settings, page 8-13
http append host-header		
http append x-forwarded-for-header		
http append x-forwarded-for-header multiple-ip-address		
http append proxy-auth-header { <i>ipaddress</i> <i>hostname</i> }		
http append www-auth-header { <i>hostname</i> <i>ipaddress</i> }		
http authenticate-strip-ntlm	Devices > Device Groups > Applications > Web > HTTP > Authenticated HTTP Caching	Configuring Authenticated HTTP Cache Settings, page 8-13
http authentication cache timeout <i>minutes</i>		
http authentication cache max-entries <i>entries</i>		
http authentication header { <i>401</i> <i>407</i> }		
http authentication realm <i>line</i>		
http cache-authenticated basic		
http cache-authenticated ntlm		
http cache-authenticated all	Devices > Device Groups > Applications > Web > HTTP > HTTP Caching	Configuring HTTP Cache Settings, page 8-6
http cache-cookies		
http cache-on-abort enable	Devices > Device Groups > Applications > Web > HTTP > Advanced HTTP Caching	Configuring Advanced HTTP Cache Settings, page 8-19
http cache-on-abort max-threshold <i>max_thresh</i>		
http cache-on-abort min-threshold <i>min_thresh</i>		
http cache-on-abort percent <i>percentthresh</i>		
http cache-vary-user-agent sub-string <i>string</i> cache-string <i>string</i>	Devices > Device Groups > Applications > Web > HTTP > HTTP Cache User Agent	Configuring HTTP Cache Vary User Agent, page 8-25

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
http cache-vary-user-agent enable	Devices > Device Groups > Applications > Web > HTTP > HTTP Caching	Configuring HTTP Cache Settings, page 8-6
http client-no-cache-request {ignore revalidate}		
http cluster max-delay <i>seconds</i>	Devices > Device Groups > Applications > Web > HTTP > Advanced HTTP Caching	Configuring Advanced HTTP Cache Settings, page 8-19
http cluster misses <i>num</i>		
http cluster http-port <i>num</i>		
http cluster heal-port <i>num</i>		
http destination-port {allow deny}	Devices > Device Groups > Web > HTTP > HTTP Destination Port Restrictions	Configuring HTTP Destination Port Restrictions, page 8-26
http dns-cache size <i>max-num</i>	Devices > Device Groups > Request Routing > DNS Caching > DNS Caching Server	Configuring DNS Caching Server Settings for the Content Engine, page 8-71
http dns-cache serial-lookup		
http fast-response enable	Devices > Device Groups > Applications > Web > HTTP > HTTP Caching	Configuring HTTP Cache Settings, page 8-6
http l4-switch {enable spoof-client-ip enable}		
http max-ttl {days text <i>textdays</i> binary <i>bindays</i> hours text <i>texthours</i> binary <i>binhours</i> minutes text <i>textminutes</i> binary <i>binminutes</i> seconds text <i>textseconds</i> binary <i>binseconds</i> }	Devices > Device Groups > Applications > Web > HTTP > HTTP Cache Freshness	Configuring HTTP Cache Freshness Settings, page 8-10
http min-ttl <i>minutes</i>		
http monitor url {url {interval <i>seconds</i> acceptable-delay <i>seconds</i> }	Devices > Device Groups > Applications > Web > HTTP > HTTP Monitor URL	Monitoring Specified HTTP URLs, page 21-20
http object max-size <i>kbytes</i>	Devices > Device Groups > Applications > Web > HTTP > HTTP Caching	Configuring HTTP Cache Settings, page 8-6
http object url-validation enable		
http persistent-connections [all server-only client-only]	Devices > Device Groups > Applications > Web > HTTP > Advanced HTTP Caching	Configuring Advanced HTTP Cache Settings, page 8-19
http persistent-connections timeout <i>seconds</i>		
http proxy incoming <i>ports</i>	Devices > Device Groups > Applications > Web > HTTP > HTTP Connections	Configuring HTTP Cache Settings, page 8-6
http proxy outgoing {connection-timeout <i>microsecs</i> host {hostname ipaddress} <i>port</i> [primary] monitor <i>seconds</i> origin-server preserve-407 }		
http request-header host unmodified	Devices > Device Groups > Applications > Web > HTTP > HTTP Caching	Configuring HTTP Cache Settings, page 8-6

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
http reval-each-request all	Devices > Device Groups > Applications > Web > HTTP > HTTP Cache Freshness	Configuring HTTP Cache Freshness Settings, page 8-10
http reval-each-request none		
http reval-each-request text		
http serve-ims text <i>percentage</i> binary <i>percentage</i>		
http smart-range enable	Devices > Device Groups > Applications > Web > HTTP > HTTP Caching	Configuring HTTP Cache Settings, page 8-6
http smart-range max-start <i>offset</i> max-interval <i>interval</i>		
http strict-request-content-length-checking enable		
http tcp-keepalive enable	Devices > Device Groups > Applications > Web > HTTP > Advanced HTTP Caching	Configuring Advanced HTTP Cache Settings, page 8-19
https proxy incoming <i>ports</i>	Devices > Device Groups > Applications > Web > HTTPS > HTTPS Proxy	Configuring HTTPS Proxy Settings, page 8-37
https proxy outgoing host {<i>ipaddress</i> <i>hostname</i>} <i>port</i>		
https destination-port {allow deny} {<i>range</i> all}		
https tcp-rw-timeout <i>seconds</i>		
https server <i>name</i>	Devices > Device Groups > Applications > Web > HTTPS > HTTPS Servers	Configuring HTTPS Servers, page 8-42
https server <i>name</i> {cert <i>cert_name</i> certgroup {chain <i>chainname</i> serverauth <i>chainname</i>} enable host {<i>hostname</i> <i>ipaddress</i>} key <i>keyname</i> password <i>password</i> protocol-version {ssl2-only ssl23-tls1 ssl3-only tls1-only} serverauth {enable ignore {cert-not-yet-valid domain-name expired-date invalid-ca}} session-cache {size <i>size</i> timeout <i>timeout</i>}	Devices > Device Groups > Applications > Web > HTTPS > Certificates	Configuring HTTPS Certificates, page 8-27
	Devices > Device Groups > Applications > Web > HTTPS > Certificate Groups	Configuring HTTPS Certificate Groups, page 8-31
	Devices > Device Groups > Applications > Web > HTTPS > Keys	Configuring HTTPS Keys, page 8-34
icap append-x-headers {x-client-ip x-server-ip}	Devices > Device Groups > Request Processing > ICAP	Configuring ICAP Settings, page 16-4
icap apply {all rules-template}		
icap bypass streaming-media		
icap logging {format standard enable}		
icap rescan-cache IStag-change		

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
icap service <i>service_id</i>	Devices > Device Groups > Request Processing > ICAP Services	Configuring ICAP Services, page 16-7
ICAP service configuration mode (config-icap-service)# icap service <i>service_id</i> [enable vector-point { reqmod-precache reqmod-postcache respmod-precache } error-handling { bypass return-error } load-balancing { round-robin client-ip-hash server-ip-hash weighted-load }]		
ICAP service configuration mode (config-icap-service)# icap service server <i>server</i> <i>server_url</i> { max-connections <i>num</i> weight <i>num</i> }	Devices > Device Groups > Request Processing > ICAP Services	Configuring an ICAP Server for ICAP Service, page 16-10
icp client add-remote-server { <i>hostname</i> <i>ipaddress</i> } { parent sibling } icp-port <i>icpport</i> http-port <i>httpport</i> [restrict <i>domainnames</i>]	Devices > Device Groups > Applications > Web > HTTP > ICP > Remote Servers	ICP Remote Server Settings, page 8-84
icp client enable	Devices > Device Groups > Applications > Web > HTTP > ICP > Client	ICP Client Settings, page 8-80
icp client exclude <i>domainnames</i>		
icp client max-fail <i>retries</i>		
icp client max-wait <i>timeout</i>		
icp client modify-remote-server { <i>hostname</i> <i>ipaddress</i> } { http-port <i>port</i> icp-port <i>port</i> parent restrict <i>domainnames</i> sibling }	Devices > Device Groups > Applications > Web > HTTP > ICP > ICP Remote Servers	ICP Remote Server Settings, page 8-84
icp server enable	Devices > Device Groups > Applications > Web > HTTP > ICP > Server	ICP Server Settings, page 8-83
icp server http-port <i>httpport</i>		
icp server port <i>icpport</i>		
icp server remote-client { <i>hostname</i> <i>ipaddress</i> } { fetch no-fetch }	Devices > Device Groups > Applications > Web > HTTP > ICP > Remote Clients	ICP Remote Client Settings, page 8-82
inetd enable ftp	Devices > Device Groups > Applications > Web > FTP > Inetd FTP	Enabling Inetd FTP Service, page 8-61
inetd enable tftp	Devices > Device Groups > Applications > Web > TFTP > TFTP General Settings	Configuring TFTP General Settings, page 8-65
inetd enable rcp	Devices > Device Groups > General Settings > Network > Inetd RCP	Enabling RCP Services on the Content Engine, page 20-8

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
interface {FastEthernet GigabitEthernet} <i>slot/port</i> cdp enable	Devices > Devices ¹ > General Settings > Network > Network Interfaces	Modifying Fast Ethernet or Gigabit Ethernet Interface Settings, page 14-7
interface {FastEthernet GigabitEthernet} <i>slot/port</i> shutdown		
interface {FastEthernet GigabitEthernet} <i>slot/port</i> autosense		
interface {FastEthernet GigabitEthernet} <i>slot/port</i> bandwidth {10 100 1000}		
interface {FastEthernet GigabitEthernet} <i>slot/port</i> [full-duplex half-duplex]		
interface {FastEthernet GigabitEthernet} <i>slot/port</i> mtu size		
interface {FastEthernet GigabitEthernet} <i>slot/port</i> ip address <i>ipaddress netmask</i> [secondary]		
interface {FastEthernet GigabitEthernet} <i>slot/port</i> standby groupnum ip <i>ipaddress netmask</i>	Devices > Devices ¹ > General Settings > Network > Network Interfaces	Configuring a Standby Interface Using the Content Distribution Manager GUI, page 14-3
interface {FastEthernet GigabitEthernet} <i>slot/port</i> standby groupnum errors num		
interface {FastEthernet GigabitEthernet} <i>slot/port</i> standby groupnum priority priority	Devices > Devices ¹ > General Settings > Network > Network Interfaces	Configuring the Interface Priority Setting, page 14-4
interface {FastEthernet GigabitEthernet} <i>slot/port</i> ip address dhcp {client-id <i>id</i> hostname <i>name</i> hostname <i>name</i> client-id <i>id</i> }		Configuring Interfaces for DHCP, page 14-13
interface FibreChannel <i>slot/port</i> [mode {autosense direct-attached switched} speed {1 2 autosense}]		Configuring the Fibre Channel Interface, page 14-9
interface PortChannel {1 2} [ip address <i>ipaddress netmask</i> shutdown]		Configuring EtherChannel, page 14-11
interface {FastEthernet GigabitEthernet} <i>slot/port</i> ip access-group {accesslistnumber accesslistname} {in out}	Devices > Devices ¹ > General Settings > Network > Network Interfaces	Applying an IP ACL to an Interface, page 17-13
ip access-list {standard extended} {acl_name acl_num} Includes other configuration commands for standard and extended IP ACLs.	Devices > Devices ¹ > General Settings > Network > IP ACL	Creating or Modifying an IP ACL, page 17-2

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
Extended ACL configuration mode (config-ext-nacl)# [insert line_num] {deny permit} {gre ip proto_num} {source_ip [wildcard] host source_ip any} {dest_ip [wildcard] host dest_ip any} [insert line_num] {deny permit} tcp {source_ip [wildcard] host source_ip any} [operator port [port]] {dest_ip [wildcard] host dest_ip any} [operator port [port]] [established] [insert line_num] {deny permit} udp {source_ip [wildcard] host source_ip any} [operator port [port]] {dest_ip [wildcard] host dest_ip any} [operator port [port]] [insert line_num] {deny permit} icmp {source_ip [wildcard] host source_ip any} {dest_ip [wildcard] host dest_ip any} [icmp_type [code]] icmp_msg delete line_num list [start_line_num [end_line_num]] move oldline_num newline_num exit	Devices > Devices ¹ > General Settings > Network > IP ACL	Adding Conditions to an IP ACL, page 17-3
Standard ACL configuration mode (config-std-nacl)# [insert line_num] {deny permit} {source_ip [wildcard] host source_ip any} delete line_num list [start_line_num [end_line_num]] move oldline_num newline_num exit	Devices > Devices ¹ > General Settings > Network > IP ACL	Adding Conditions to an IP ACL, page 17-3
ip default-gateway ipaddress	Devices > Devices ¹ > General Settings > Network > Network Interfaces	Configuring a Standby Interface Using the Content Distribution Manager GUI, page 14-3 Modifying Fast Ethernet or Gigabit Ethernet Interface Settings, page 14-7
ip domain-name name1 name2 name3	Devices > Device Groups > General Settings > Network > DNS	Configuring the DNS Server for HTTP Proxy Caching, page 8-78
ip name-server ipaddress		
ip dscp {client {cache-hit {match-server set-dscp dscp-packets set-tos tos-packets} cache-miss {match-server set-dscp dscp-packets set-tos tos-packets}} server {match-client set-dscp dscp-packets set-tos tos-packets}}	Devices > Device Groups > General Settings > Network > IP General Settings	Configuring IP Differentiated Services, page 10-9
ip path-mtu-discovery enable		

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
ip route <i>dest_addrs netmask gateway_addrs</i>	Devices > Devices ¹ > General Settings > Network > IP Routes	Configuring Static IP Routes, page 14-15
kernel kdb	Devices > Device Groups > General Settings > Troubleshooting > Kernel Debugger	Enabling Kernel Debugger, page 21-23
ldap server administrative-dn <i>name</i>	Devices > Device Groups > General Settings > Authentication > LDAP Server	Configuring LDAP Server Settings, page 15-17
ldap server administrative-passwd <i>passwd</i>		
ldap server allow-mode		
ldap server base <i>baseword</i>		
ldap server enable	Devices > Device Groups > General Settings > Authentication > Authentication Scheme	Setting the Authentication Scheme for Request Authentication, page 15-36
ldap server filter <i>filterword</i>	Devices > Device Groups > General Settings > Authentication > LDAP Server	Configuring LDAP Server Settings, page 15-17
ldap server group active-directory enable		
ldap server group { organizationUnit enable custom <i>name enable</i> static { group-attribute <i>name</i> member-attribute { member uniquemember custom-member <i>name</i> } enable nested { enable level <i>number</i> }}}	Devices > Device Groups > General Settings > Authentication > LDAP Server	Configuring LDAP Server Settings, page 15-17
ldap server host { <i>hostipaddress</i> <i>hostname</i> } [primary secondary]		
ldap server password-expiry { enable redirect-url <i>url</i> }		
ldap server policy-redirect { enable redirect-url <i>url</i> attribute <i>name</i> version-number <i>number</i> }		
ldap server policy-redirect append-request-url		
ldap server port <i>port_num</i>		
ldap server retransmit <i>retries</i>		
ldap server timeout <i>seconds</i>		
ldap server userid-attribute <i>useridword</i>		
ldap server version <i>ver_num</i>		
line console carrier-detect	Devices > Device Groups > General Settings > Login Access Control > Console Carrier Detect	Configuring Line Console Carrier Detection, page 20-6

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
logging host { <i>ipaddress</i> <i>hostname</i> }	Devices > Device Groups > General Settings > Notification and Tracking > System Logs	Configuring System Event Logging Using the Content Distribution Manager GUI, page 21-5
logging host priority { alert critical debug emergency error information notice warning }		
logging disk enable		
logging disk filename <i>filename</i>		
logging disk recycle <i>size</i>		
logging disk priority { alert critical debug emergency error information notice warning }		
logging console enable		
logging console priority { alert critical debug emergency error information notice warning }		
logging facility { auth daemon kernel local0 local1 local2 local3 local4 local5 local6 local7 mail news syslog user uucp }		
multicast license-key <i>key</i>	Devices > Device Groups > Prepositioning > Multicast Distribution License	Enabling Content Engines for Multicasting, page 5-33
multicast accept-license-agreement multicast enable		
multicast evaluate		
network-filesystem server { samba cifs } { enable authentication-mode { public ce-user cifs-user password-server } max-connections <i>num</i> }	Devices > Device Groups > General Settings > File Sharing > CIFS Server	Configuring CIFS Server Settings, page 13-17
network-filesystem server { samba cifs } { share-web-site <i>site_name</i> [share-auth-content protect-auth-content]}	Devices > Device Groups > General Settings > File Sharing > CIFS Server Access Control	Configuring CIFS Server Website Access Control Settings, page 13-19
network-filesystem client { cifs { <i>hostname</i> <i>ipaddress</i> } directory { cdnfs mediafs } reserved-disk-space <i>space</i> username <i>name</i> password <i>password</i> [domain <i>domain</i>] nfs { <i>hostname</i> <i>ipaddress</i> } directory { cdnfs mediafs } reserved-disk-space <i>space</i> }	Devices > Devices ¹ > General Settings > Storage > NAS	See “Attaching and Detaching NAS Shares Using the Content Distribution Manager GUI” in the <i>Cisco ACNS Software Update and Maintenance Guide, Release 5.x</i> .
ntlm server ad-group-search { enable enum-user gc-server groupname-attribute ldap-referral ldap-search-port membership-attribute user-objectclass username-attribute }	Devices > Device Groups > General Settings > Authentication > NTLM Server	Configuring NTLM Server Settings, page 15-23

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
ntlm server ad-group-search mem-cache {enable size <i>kilobytes</i> max-ttl <i>minutes</i> }	Devices > Device Groups > General Settings > Authentication > NTLM Server	Configuring the NTLM Server LDAP Memory Cache Settings, page 15-28
ntlm allow-domain {enable domain <i>domainname</i> }	Devices > Device Groups > General Settings > Authentication > NTLM Allow Domain	Configuring NTLM Allowed Domains for HTTP Request Authentication, page 15-31
ntlm server enable	Devices > Device Groups > General Settings > Authentication > Authentication Scheme	Setting the Authentication Scheme for Request Authentication, page 15-36
ntlm server host {hostname ipaddress}	Devices > Device Groups > General Settings > Authentication > NTLM Server	Configuring NTLM Server Settings, page 15-23
ntlm server domain name		
ntp server {hostname ipaddress}	Devices > Device Groups > General Settings > Network > NTP	Configuring NTP Settings, page 20-11
offline-operation enable	Devices > Device Groups > Applications > Web > Offline Operations	Enabling Offline Operation of Network Devices, page 20-27
port-channel load-balance {dst-ip dst-mac round-robin}	Devices > Devices ¹ > General Settings > Network > Network Interfaces	Configuring EtherChannel, page 14-11
proxy-auto-config enable	Devices > Device Groups > Applications > Web > Client Proxy Autoconfiguration	Configuring Client Proxy Autoconfiguration Settings, page 4-39
proxy-protocols outgoing-proxy exclude {enable list word}	Devices > Device Groups > Applications > Web > Outgoing Proxy Exclusions	Configuring HTTP and HTTPS Outgoing Proxy Exclusion Settings, page 8-48
proxy-protocols transparent {default-server original-proxy reset}		
qos camiant-cdn-am-service config-file filename or URL	Devices > Device Groups > Request Processing > PCMM > QoS Policy Service	Setting QoS Policy Server Settings for Content Engines, page 16-62
qos camiant-cdn-am-service enable		
radius-server key key	Devices > Device Groups > General Settings > Authentication > RADIUS Server	Configuring RADIUS Server Settings, page 15-32
radius-server timeout seconds		

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
radius-server enable	Devices > Device Groups > General Settings > Authentication > Authentication Scheme	Setting the Authentication Scheme for Request Authentication, page 15-36
radius-server host {hostname ipaddress} {auth-port port}	Devices > Device Groups > General Settings > Authentication > RADIUS Server	Configuring RADIUS Server Settings, page 15-32
radius-server retransmit retries		
radius-server redirect enable		
radius-server redirect message reply location url		
rtsp advanced max-initial-setup-delay seconds	Devices > Device Groups > Applications > Streaming > RTSP Gateway	Configuring the RTSP Gateway, page 9-8
rtsp advanced max-request-rate number		
rtsp proxy media-real enable	Devices > Device Groups > Applications > Streaming > Real Networks > Real Proxy License	Enabling RealProxy, page 9-10
rtsp proxy media-real evaluate		
rtsp proxy media-real license-key key		
rtsp proxy media-real accept-license-agreement		
rtsp server real-subscriber license-key key	Devices > Device Groups > Applications > Streaming > Real Networks > Real Subscriber License	Enabling RealSubscriber, page 9-12
rtsp server real-subscriber accept-license-agreement		
rtsp server real-subscriber evaluate		
rtsp server real-subscriber enable		
rtsp l4-switch enable	Devices > Device Groups > Applications > Streaming > RTSP Gateway	Configuring the RTSP Gateway, page 9-8
rtsp port incoming port_num		
rtsp ip-address ipaddress		
rtsp server cisco-streaming-engine enable	Devices > Device Groups > Applications > Streaming > Cisco Streaming Engine	Enabling the Cisco Streaming Engine, page 9-14
rtsp server cisco-streaming-engine broadcast port-list list_num port_num	Services > Video > Programs > Live Streaming (indirect support)	Configuring Live Stream Settings for a Cisco Streaming Engine Live Program, page 7-20
rtsp server cisco-streaming-engine broadcast id id source {source-rtsp rtsp_url source-udp url srce_ip_addr rcv_ip_addr rcv_port_list_num} {track-count 1-4} destination {destination-pull destination-udp ttl dest_ip_addr dest_list_num}		
rule enable	Devices > Device Groups > Request Processing > Enable Rules	Enabling Rule Settings, page 16-14

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
rule action rule action allow pattern-list <i>list_num</i> [protocol {all http https rtsp}] rule action append-username-header pattern-list <i>list_num</i> [protocol {all http https}] rule action block pattern-list <i>list_num</i> [protocol {all http https rtsp}] rule action cache-non-cacheable ttl {day <i>days</i> pattern-list <i>list_num</i> [protocol {all http https}] hours <i>hours</i> pattern-list <i>list_num</i> [protocol {all http https}] minutes <i>minutes</i> pattern-list <i>list_num</i> [protocol {all http https}] seconds <i>seconds</i> pattern-list <i>list_num</i> [protocol {all http https}]}	Devices > Device Groups > Request Processing > Service Rules	Configuring Service Rules, page 16-13
rule action cache-only pattern-list <i>list_num</i> [protocol {all http https}] rule action dscp client cache-hit {match-server pattern-list <i>list_num</i> [protocol {all http https}] set-dscp <i>dscpvalue</i> set-tos <i>tosvalue</i> } rule action dscp client cache-miss {match-server pattern-list <i>list_num</i> [protocol {all http https}] set-dscp <i>dscpvalue</i> set-tos <i>tosvalue</i> } rule action dscp server {match-client pattern-list <i>list_num</i> [protocol {all http https}] set-dscp <i>dscpvalue</i> set-tos <i>tosvalue</i> }	Devices > Device Groups > Request Processing > Service Rules	Configuring Service Rules, page 16-13
rule action freshness-factor <i>exp_time</i> pattern-list <i>list_num</i> [protocol {all http https}] rule action generate-url-signature [include-client-src-ip] key-id-owner <i>owner_num</i> key-id-number <i>id_num</i> pattern-list <i>list_num</i> [protocol {all http}] rule action insert-no-cache pattern-list <i>list_num</i> [protocol {all http https}] rule action no-auth pattern-list <i>list_num</i> [protocol {all http https rtsp}] rule action no-cache pattern-list <i>list_num</i> [protocol {all http https}] rule action no-persistent-connection {all client server} pattern-list <i>list_num</i> [protocol {all http https}]	Devices > Device Groups > Request Processing > Service Rules	Configuring Service Rules, page 16-13

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
rule action no-proxy pattern-list <i>list_num</i> [protocol {all http https}] rule action redirect url pattern-list <i>list_num</i> [protocol {all http https rtsp}] rule action redirect-url-for-cdn pattern-list <i>list_num</i> [protocol {all http https rtsp}] rule action refresh pattern-list <i>list_num</i> [protocol {all http https}] rule action reset pattern-list <i>list_num</i> [protocol {all http https rtsp}] rule action rewrite pattern-list <i>list_num</i> [protocol {all http https rtsp}]	Devices > Device Groups > Request Processing > Service Rules	Configuring Service Rules, page 16-13
rule action use-dns-server { <i>hostname</i> <i>ipaddress</i> } pattern-list <i>list_num</i> [protocol {all http https}] rule action use-icap-service <i>service_name</i> pattern-list <i>list_num</i> [protocol {all http https}] rule action use-proxy { <i>hostname</i> <i>ipaddress</i> } <i>port</i> pattern-list <i>list_num</i> [protocol {all http https}] rule action use-proxy { <i>hostname</i> <i>ipaddress</i> } <i>port</i> [failover] pattern-list <i>list_num</i> [protocol {all http https}] rule action use-server { <i>hostname</i> <i>ipaddress</i> } <i>port</i> pattern-list <i>list_num</i> [protocol {all http https}] rule action use-xforward-clt-ip pattern-list <i>list_num</i> [protocol {all http https}] rule action validate-url-signature error-redirect-url <i>url</i> pattern-list <i>list_num</i> [protocol {all http rtsp}]	Devices > Device Groups > Request Processing > Service Rules	Configuring Service Rules, page 16-13

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
rule pattern-list <i>list_num</i> domain <i>dn_regex</i> rule pattern-list <i>list_num</i> dst-ip <i>d_ipaddress</i> <i>d_subnet</i> rule pattern-list <i>list_num</i> dst-port <i>port</i> rule pattern-list <i>list_num</i> groupname <i>name</i> rule pattern-list <i>list_num</i> groupname-regex <i>group_name_regex</i> rule pattern-list <i>list_num</i> group-type { and or } rule pattern-list <i>list_num</i> header-field { referer <i>ref_regex</i> request-line <i>req_regex</i> user-agent <i>ua_regex</i> } rule pattern-list <i>list_num</i> header-field x-forwarded-for <i>ipaddress</i> rule pattern-list <i>list_num</i> header-field-sub { referer <i>ref_regex</i> <i>ref_sub</i> request-line <i>req_regex</i> <i>req_sub</i> user-agent <i>ua_regex</i> <i>ua_sub</i> }	Devices > Device Groups > Request Processing > Service Rules	Configuring Service Rules, page 16-13
rule pattern-list <i>list_num</i> icap-attribute <i>icap_attribute</i> <i>icap_value</i> rule pattern-list <i>list_num</i> mime-type <i>mt_regex</i> rule pattern-list <i>list_num</i> src-ip <i>s_ipaddress</i> <i>s_subnet</i> rule pattern-list <i>list_num</i> url-regex <i>url_regex</i> rule pattern-list <i>list_num</i> url-regsub <i>url_regex</i> <i>url_sub</i> rule pattern-list <i>list_num</i> username <i>user_name</i>	Devices > Device Groups > Request Processing > Service Rules	Configuring Service Rules, page 16-13
snmp-server contact <i>line</i> snmp-server enable traps [config content-engine [disk-fail disk-read disk-write overload-bypass transaction-log] entity event snmp [authentication cold-start]]	Devices > Device Groups > General Settings > Notification and Tracking > SNMP > General Settings	Configuring SNMP Traps, page 22-2
snmp-server host { <i>hostname</i> <i>ipaddress</i> } <i>communitystring</i> [v2c [retry <i>num</i>] [timeout <i>seconds</i>]] [v3 { auth [retry <i>num</i>] [timeout <i>seconds</i>] noauth [retry <i>num</i>] [timeout <i>seconds</i>] priv [retry <i>num</i>] [timeout <i>seconds</i>]}]}	Devices > Device Groups > General Settings > Notification and Tracking > SNMP > Host	Configuring SNMP Host Settings, page 22-10
snmp-server location <i>line</i>	Devices > Device Groups > General Settings > Notification and Tracking > SNMP > General Settings	Configuring SNMP Traps, page 22-2

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
snmp-server user <i>name group</i> [auth { md5 <i>password</i> [priv <i>password</i>] sha <i>password</i> [priv <i>password</i>]} remote <i>octetstring</i> [auth { md5 <i>password</i> [priv <i>password</i>] sha <i>password</i> [priv <i>password</i>]}]]]	Devices > Device Groups > General Settings > Notification and Tracking > SNMP > User	Configuring SNMP User Settings, page 22-8
snmp-server group <i>name</i> { v1 [notify <i>name</i>] [read <i>name</i>] [write <i>name</i>] v2c [notify <i>name</i>] [read <i>name</i>] [write <i>name</i>] v3 [auth [notify <i>name</i>] [read <i>name</i>] [write <i>name</i>] noauth [notify <i>name</i>] [read <i>name</i>] [write <i>name</i>] priv [notify <i>name</i>] [read <i>name</i>] [write <i>name</i>]}]}	Devices > Device Groups > General Settings > Notification and Tracking > SNMP > Group	Configuring SNMP Group Settings, page 22-6
snmp-server community <i>string</i> [group <i>groupname</i> rw]	Devices > Device Groups > General Settings > Notification and Tracking > SNMP > Community	Configuring SNMP Community Settings, page 22-5
snmp-server view <i>viewname MIBfamily</i> { excluded included }	Devices > Device Groups > General Settings > Notification and Tracking > SNMP > View	Configuring SNMPv2 View Settings, page 22-9
snmp-server notify inform snmp-server mib persist event	Devices > Device Groups > General Settings > Notification and Tracking > SNMP > General Settings	Configuring SNMP Traps, page 22-2
snmp-server access-list { <i>std_acl_number</i> <i>acl_name</i> }	Devices > Devices ¹ > General Settings > Network > IP ACL Feature Usage	Associating an IP ACL with an Application, page 17-12
sshd enable sshd timeout <i>seconds</i> sshd password-guesses <i>num</i> ssh-key-generate key-length <i>length</i>	Devices > Device Groups > General Settings > Network > SSH	Configuring Secure Shell Settings, page 20-2
tacacs enable tacacs password ascii tacacs timeout tacacs retransmit tacacs key tacacs host { <i>ipaddress</i> <i>hostname</i> } [primary]	Devices > Device Groups > General Settings > Authentication > TACACS+ Server	Configuring TACACS+ Server Settings, page 15-34

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
tcp ecn enable	Devices > Device Groups > General Settings > Network > TCP	Configuring TCP Settings for Increased Cache Performance, page 10-1
tcp type-of-service enable		
tcp cwnd-base <i>segments</i>		
tcp increase-xmit-timer-value <i>value</i>		
tcp init-ss-threshold <i>value</i>		
tcp keepalive-probe-cnt <i>count</i>		
tcp keepalive-probe-interval <i>seconds</i>		
tcp keepalive-timeout <i>seconds</i>		
tcp client-satellite		
tcp client-mss <i>max_seg_size</i>		
tcp client-receive-buffer <i>kbytes</i>		
tcp client-rw-timeout <i>seconds</i>		
tcp client-send-buffer <i>kbytes</i>		
tcp server-satellite		
tcp server-mss <i>max_seg_size</i>		
tcp server-receive-buffer <i>kbytes</i>		
tcp server-rw-timeout <i>seconds</i>		
tcp server-send-buffer <i>kbytes</i>		
telnet enable	Devices > Device Groups > General Settings > Network > Telnet	Enabling the Telnet Service on the Content Engine, page 20-1
tftp max-connections <i>number</i>	Devices > Device Groups > Applications > Web > TFTP > TFTP General Settings	Configuring TFTP General Settings, page 8-65
tftp-server dir <i>directory</i>	Devices > Device Groups > Applications > Web > TFTP > TFTP Directory	Configuring TFTP Directory Settings, page 8-69
tftp-server access-list { <i>std_acl_number</i> <i>acl_name</i> }	Devices > Devices ¹ > General Settings > Network > IP ACL Feature Usage	Associating an IP ACL with an Application, page 17-12
tftp-server gw proto ftp server { <i>hostname</i> <i>ipaddress</i> } pri <i>priority</i> [name <i>name</i> passwd <i>password</i>] [path <i>directory</i>]	Devices > Device Groups > Applications > Web > TFTP > TFTP Proxy	Configuring TFTP Proxy Server Settings, page 8-66
tftp-server gw proto http server { <i>hostname</i> <i>ipaddress</i> } [port <i>port_num</i>] pri <i>priority</i> [name <i>name</i> passwd <i>password</i>] [path <i>directory</i>]		

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
transaction-logs archive max-file-size <i>kbytes</i>	Devices > Device Groups > General Settings > Notification and Tracking > Transaction Logs	Enabling Transaction Logging with the Content Distribution Manager GUI, page 19-10
transaction-logs archive interval { <i>seconds</i> every-week [on <i>weekdays</i> at <i>hour:minute</i>] every-day {at <i>hour:minute</i> every hours } every-hour {at <i>minute</i> every minutes }}		
transaction-logs enable		
transaction-logs export enable		
transaction-logs export ftp-server { <i>hostname</i> <i>ipaddress</i> } <i>login password directory</i>		
transaction-logs export sftp-server { <i>hostname</i> <i>ipaddress</i> } <i>login password directory</i>		
transaction-logs export compress	Devices > Device Groups > General Settings > Notification and Tracking > Transaction Logs	Enabling Transaction Logging with the Content Distribution Manager GUI, page 19-10
transaction-logs export interval { <i>minutes</i> every-week [on <i>weekdays</i> at <i>hour:minute</i>] every-day {at <i>hour:minute</i> every hours } every-hour {at <i>minute</i> every minutes }}		
transaction-logs file-marker		
transaction-logs sanitize		
transaction-logs log-windows-domain		
transaction-logs format { squid extended-squid apache custom <i>string</i> }		
tvout signal { ntsc pal }	Devices > Device Groups > Applications > Set Top Box > TV-out	Enabling the TV-Out Feature, page 7-31
tvout enable		
url-filter http { bad-sites-deny good-sites-allow } { enable file <i>filename</i> }	Devices > Device Groups > Request Processing > URL Filter > HTTP > URL Filter Settings for HTTP Protocol	Configuring URL Filter Settings Using the Content Distribution Manager GUI, page 16-26
url-filter http websense allowmode enable		Configuring URL Filtering with Websense Enterprise Software, page 16-39
url-filter http websense enable		
url-filter http websense server { <i>ipaddress</i> <i>hostname</i> } port <i>1-65535</i> timeout <i>seconds</i> connections <i>1-250</i>		
url-filter http N2H2 allowmode enable		Configuring URL Filtering with the N2H2 Server, page 16-33
url-filter http N2H2 enable		
url-filter http N2H2 server { <i>ipaddress</i> <i>hostname</i> } port <i>1-65535</i> timeout		
url-filter http smartfilter enable		Configuring URL Filtering with SmartFilter Software, page 16-55

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
url-filter http custom-message <i>dirname</i>	Devices > Device Groups > Request Processing > URL Filter	Configuring URL Filter Settings Using the Content Distribution Manager GUI, page 16-26 Configuring Custom Blocking Messages Using the CLI, page 16-31
url-filter wmt { bad-sites-deny good-sites-allow } { enable file <i>filename</i> }	Devices > Device Groups > Request Processing > URL Filter > WMT > URL Filter Settings for WMT Protocol	Configuring URL Filter Settings Using the Content Distribution Manager GUI, page 16-26
url-filter rtsp { bad-sites-deny good-sites-allow } { enable file <i>filename</i> }	Devices > Device Groups > Request Processing > URL Filter > RTSP > URL Filter Settings for RTSP Protocol	
url-signature key-id-owner <i>number</i> key-id-number <i>number</i> key <i>encryption_key</i>	Devices > Device Groups > Request Processing > PCMM > URL Signature Key	Creating New URL Signature Keys for Content Engines, page 16-63
username <i>name</i> { cifs-password password privilege }	Devices > Device Groups > General Settings > Login Access Control > Users > Usernames	Creating User Accounts for Centrally Managed Devices, page 12-14
wccp access-list { <i>std_acl_number</i> <i>ext_acl_number</i> <i>acl_name</i> }	Devices > Devices ¹ > General Settings > Network > IP ACL Feature Usage	Associating an IP ACL with an Application, page 17-12
wccp custom-web-cache mask { [dst-ip-mask <i>hex_num</i>] [dst-port-mask <i>port_hex_num</i>] [src-ip-mask <i>hex_num</i>] [src-port-mask <i>port_hex_num</i>] }	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Masks for the Content Engine, page 4-18
wccp custom-web-cache router-list-num <i>num</i> port <i>port</i> [assign-method-strict] [hash-destination-ip] [hash-destination-port] [hash-source-ip] [hash-source-port] [I2-redirect] [mask-assign] [password <i>key</i>] [weight <i>percentage</i>]	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Settings for the Content Engine, page 4-10
wccp dns mask { [dst-ip-mask <i>hex_num</i>] [dst-port-mask <i>port_hex_num</i>] [src-ip-mask <i>hex_num</i>] [src-port-mask <i>port_hex_num</i>] }	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Masks for the Content Engine, page 4-18
wccp dns router-list-num <i>num</i> [hash-destination-ip] [hash-destination-port] [hash-source-ip] [hash-source-port] [I2-redirect] [mask-assign] [password <i>key</i>] [weight <i>percentage</i>]	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Settings for the Content Engine, page 4-10
wccp flow-redirect enable	Devices > Device Groups > Request Routing > WCCP > General Settings	Configuring WCCP General Settings for the Content Engine, page 4-7

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
wccp ftp router-list-num <i>num</i> [assign-method-strict] [I2-redirect] [mask-assign] [password <i>key</i>] [weight <i>percentage</i>]	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Settings for the Content Engine, page 4-10
wccp ftp mask {[dst-ip-mask <i>hex_num</i>] [src-ip-mask <i>hex_num</i>]}	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Masks for the Content Engine, page 4-18
wccp home-router <i>ipaddress</i> wccp https-cache accept-all	Devices > Device Groups > Request Routing > WCCP > General Settings	Configuring WCCP General Settings for the Content Engine, page 4-7
wccp https-cache mask {[dst-ip-mask <i>hex_num</i>] [dst-port-mask <i>port_hex_num</i>] [src-ip-mask <i>hex_num</i>] [src-port-mask <i>port_hex_num</i>]}	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Masks for the Content Engine, page 4-18
wccp https-cache router-list-num <i>num</i> port <i>port</i> [assign-method-strict] [hash-destination-ip] [hash-destination-port] [hash-source-ip] [hash-source-port] [I2-redirect] [mask-assign] [password <i>key</i>] [weight <i>percentage</i>]	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Settings for the Content Engine, page 4-10
wccp port-list <i>list_num</i> <i>port_num</i>	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Port Lists for the Content Engine, page 4-17
wccp rtsp mask {[dst-ip-mask <i>hex_num</i>] [dst-port-mask <i>port_hex_num</i>] [src-ip-mask <i>hex_num</i>] [src-port-mask <i>port_hex_num</i>]}	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Masks for the Content Engine, page 4-18
wccp rtsp router-list-num <i>num</i> [hash-destination-ip] [hash-destination-port] [hash-source-ip] [hash-source-port] [I2-redirect] [mask-assign] [password <i>key</i>] [weight <i>percentage</i>]	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Settings for the Content Engine, page 4-10
wccp reverse-proxy mask [dst-ip-mask <i>hex_num</i>] [dst-port-mask <i>port_hex_num</i>] [src-ip-mask <i>hex_num</i>] [src-port-mask <i>port_hex_num</i>]	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Masks for the Content Engine, page 4-18
wccp reverse-proxy router-list-num <i>num</i> [assign-method-strict] [I2-redirect] [mask-assign] [password <i>key</i>] [weight <i>percentage</i>]	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Settings for the Content Engine, page 4-10
wccp router-list <i>list_num</i> <i>ipaddress</i>	Devices > Device Groups > Request Routing > WCCP > Services	Creating WCCP Router Lists, page 4-14
wccp service-number <i>service_num</i> router-list-num <i>num</i> port-list-num <i>port</i> application { <i>cache</i> <i>streaming</i> } [hash-destination-ip] [hash-destination-port] [hash-source-ip] [hash-source-port] [I2-redirect] [mask-assign] [match-source-port] [password <i>key</i>] [weight <i>percentage</i>]	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Settings for the Content Engine, page 4-10

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
wccp service-number <i>service_num</i> mask {[dst-ip-mask <i>hex_num</i>] [dst-port-mask <i>port_hex_num</i>] [src-ip-mask <i>hex_num</i>] [src-port-mask <i>port_hex_num</i>]}	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Masks for the Content Engine, page 4-18
wccp shutdown max-wait <i>seconds</i>	Devices > Device Groups > Request Routing > WCCP > General Settings	Configuring WCCP General Settings for the Content Engine, page 4-7
wccp slow-start enable		
wccp spoof-client-ip enable		
wccp version {1 2}		
wccp web-cache mask {[dst-ip-mask <i>hex_num</i>] [dst-port-mask <i>port_hex_num</i>] [src-ip-mask <i>hex_num</i>] [src-port-mask <i>port_hex_num</i>]}	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Masks for the Content Engine, page 4-18
wccp web-cache router-list-num <i>num</i> [assign-method-strict] [I2-redirect] [mask-assign] [password <i>key</i>] [weight <i>percentage</i>]	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Settings for the Content Engine, page 4-10
wccp wmt mask {[dst-ip-mask <i>hex_num</i>] [dst-port-mask <i>port_hex_num</i>] [src-ip-mask <i>hex_num</i>] [src-port-mask <i>port_hex_num</i>]}	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Masks for the Content Engine, page 4-18
wccp wmt router-list-num <i>num</i> [assign-method-strict] [hash-destination-ip] [hash-destination-port] [hash-source-ip] [hash-source-port] [I2-redirect] [mask-assign] [password <i>key</i>] [weight <i>percentage</i>]	Devices > Device Groups > Request Routing > WCCP > Services	Configuring WCCP Service Settings for the Content Engine, page 4-10
wccp wmt-rtspu router-list-num <i>num</i> [assign-method-strict] [hash-destination-ip] [hash-destination-port] [hash-source-ip] [hash-source-port] [I2-redirect] [mask-assign] [password <i>key</i>] [weight <i>percentage</i>]		
websense-server enable	Devices > Device Groups > Request Processing > URL Filter > HTTP > URL Filter Settings for HTTP Protocol	Configuring URL Filter Settings Using the Content Distribution Manager GUI, page 16-26
websense-server service { edir-agent { activate edir-server administrative-passwd <i>password</i> } eim activate logon-agent activate network-agent activate policy { local activate remote host <i>hostname</i> or <i>IP address</i> port <i>portnum</i> } radius-agent activate user activate }	Devices > Device Groups > Request Processing > Websense Server	Installing Websense Server Components for the Content Engine, page 16-50
wmt accelerate live-split enable	Devices > Device Groups > Applications > Streaming > Windows Media > General Settings	Configuring Windows Media Proxy Settings for the Content Engine, page 9-18
wmt accelerate proxy-cache enable		
wmt accelerate vod enable		
wmt accept-license-agreement	Devices > Device Groups > Applications > Streaming > Windows Media > License	Enabling Windows Media Services, page 9-15

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
wmt advanced client idle-timeout <i>seconds</i>	Devices > Device Groups > Applications > Streaming > Windows Media > General Settings	Configuring Windows Media Proxy Settings for the Content Engine, page 9-18
wmt advanced client max-packet-size <i>bytes</i>		
wmt advanced server log-forwarding enable		
wmt bandwidth incoming bypass-list { <i>ipaddress</i> <i>hostname</i> } [<i>ipaddress</i> <i>hostname</i>]	Devices > Devices ¹ > Applications > Streaming > Windows Media > Bypass List	Configuring the Windows Media Incoming Bandwidth Bypass List, page 9-47
wmt broadcast { <i>alias-name name source url</i> }	Devices > Device Groups > Applications > Streaming > Windows Media > Broadcast Alias	Configuring Multicast-In Unicast-Out, page 9-34 Configuring Unicast-In Unicast-Out, page 9-36
wmt cache enable	Devices > Device Groups > Applications > Streaming > Windows Media > License	Enabling Windows Media Services, page 9-15
wmt cache max-obj-size <i>mbytes</i>	Devices > Device Groups > Applications > Streaming > Windows Media > General Settings	Configuring Windows Media Proxy Settings for the Content Engine, page 9-18
wmt disallowed-client-protocols {[<i>http</i> <i>rtsp</i> <i>rtspu</i>]}		
wmt enable	Devices > Device Groups > Applications > Streaming > Windows Media > License	Enabling Windows Media Services, page 9-15
wmt evaluate		
wmt extended transaction-log enable	Devices > Device Groups > General Settings > Notification and Tracking > Transaction Logs	Using WMT Transaction Logging, page 19-15
wmt fast-cache enable	Devices > Device Groups > Applications > Streaming > Windows Media > General Settings	Configuring Windows Media Proxy Settings for the Content Engine, page 9-18
wmt fast-cache max-delivery-rate <i>num</i>		
wmt fast-start enable		
wmt fast-start max-bandwidth <i>kbits</i>		
wmt http allow extension <i>file_extensions</i>		
wmt l4-switch enable		
wmt license-key <i>key</i>	Devices > Device Groups > Applications > Streaming > Windows Media > License	Enabling Windows Media Services, page 9-15
wmt live-url-stripping enable	Devices > Device Groups > Applications > Streaming > Windows Media > General Settings	Configuring Windows Media Proxy Settings for the Content Engine, page 9-18

Table C-1 ACNS Software CLI to Content Distribution Manager GUI Map (continued)

Command	GUI Navigation Path	Procedure Reference
wmt multicast { station-configuration <i>name</i> [{ <i>dest_addr</i> <i>dest_port</i> <i>media_source</i>] [log { local webserver <i>webserver_url</i> }] [nsc-reference { <i>ref_URL</i> }] [play-forever] [unicast-url <i>url</i>] failover { alternate-source <i>alt_src_url</i> retry-count <i>number</i> retry-interval <i>number</i> } schedule-start { <i>minute hour day month</i> now }] time-to-live <i>t1</i> }	Devices > Device Groups > Applications > Streaming > Windows Media > Multicast Stations	Configuring Unicast-In Multicast-Out, page 9-28 Configuring Multicast-In Multicast-Out, page 9-31
wmt max-concurrent-sessions <i>num</i>	Devices > Device Groups > Applications > Streaming > Windows Media > General Settings	Configuring Windows Media Proxy Settings for the Content Engine, page 9-18
wmt multicast schedule-start <i>name minute hour day month</i>	Devices > Device Groups > Applications > Streaming > Windows Media > Multicast Station Schedules	Configuring Windows Media Multicast Station Schedules, page 9-38
wmt multicast time-to-live <i>t1</i>	Devices > Device Groups > Applications > Streaming > Windows Media > General Settings	Configuring Windows Media Proxy Settings for the Content Engine, page 9-18
wmt proxy outgoing http host { <i>hostname</i> <i>ipaddress</i> } <i>port</i>		
wmt proxy outgoing rtsp host { <i>hostname</i> <i>ipaddress</i> } <i>port</i>		
wmt transaction-logs format { extended { wms-41 wms-90 } wms-41 wms-90 }	Devices > Device Groups > General Settings > Notification and Tracking > Transaction Logs	Using WMT Transaction Logging, page 19-15

1. This command can only be configured on individual devices; the Content Distribution Manager GUI does not support this command for device group configuration.

CLI Commands Not Supported in the Content Distribution Manager GUI

The following CLI commands are not supported in the Content Distribution Manager GUI:

- **auto-register enable** {FastEthernet *number number* | GigabitEthernet *number number*}
- **bandwidth advanced config-file** *string*
- **cdm ip** *string*
- **cdm** {role {primary | standby} | ui port *number*}
- **cms enable**
- **cms rpc timeout** {connection *number* transfer *number* | incoming-wait *number*}
- **content-routing-api enable**
- **device mode**
- **exception** {debug | coredump}
- **multicast back-version-compatibility** acns-5-0
- **multicast sender-delay** *number*
- **multicast max-concurrent-jobs** *number* **minimal-target-rate** *number*
- **multicast priority-weight** *number*
- **multicast fixed-carousel enable**
- **pre-load**
- **rtsp server cisco-streaming-engine broadcast port-list** *number number*
- **rtsp server cisco-streaming-engine broadcast id** *string* **source** {source-rtsp *string* | source-udp *string ipaddress ipaddress number*} {track-count *number*} **destination** {destination-pull | destination-udp *number ipaddress number*}



Symbols

* (wildcard) character [48](#)

.asx files [3](#)

.nsc files [23](#)

.pac files [39](#)

<CdnManifest> tag [3, 37](#)

<contains> tag [63](#)

<contentType> tag [39](#)

<crawler> tag [4, 54](#)

<extension> tag [40](#)

<host> tag [42](#)

<http-meta-data> tag [63](#)

<item> tag [3, 44](#)

<item-group> tag [58](#)

<match> subtag

attributes [12](#)

time-before attribute

relative time [62](#)

<matchRule> tag [61](#)

<options> tag [40](#)

<playServer> tag [38, 39](#)

<playServerTable> tag [38](#)

<proxyServer> tag [44](#)

<server> tag [42](#)

Numerics

3GPP

support in CSE [14](#)

407 messages [8](#)

A

AAA accounting

commands, list of [28](#)

configuring [27](#)

access control list. *See* ACL

ACL

configuring for group authorization [1](#)

enabling [2](#)

for group-based authorization [12](#)

modifying [3](#)

repositioning settings in [5](#)

viewing settings [5](#)

See also IP ACL

ACNS Manifest Generator tool

using [86 to 99](#)

ACNS network

administrator responsibilities [6](#)

content type mapping [39](#)

exporting data [23](#)

overview [2](#)

printing data [23](#)

topology considerations [15](#)

acquirer

commands, list of [15](#)

error codes, list of [61](#)

acquiring

files

from SMB servers [2](#)

NTLM authenticated content [42](#)

pre-positioned content

using externally hosted manifest file [23](#)

using the Content Distribution Manager GUI [9 to 23](#)

- acquisition
 - bandwidth settings
 - graphical display of [56](#)
 - choosing method for channel [6](#)
 - retry and refresh mechanisms [52](#)
 - scheduling
 - in manifest files [10](#)
 - show commands, list of [48](#)
 - verifying results of [48](#)
- acquisition and distribution
 - architecture overview [11](#)
- acquisition-related commands [15](#)
- activating
 - ACNS network devices [22](#)
- active-active interfaces [1](#)
- active-standby interfaces [1](#)
- adding
 - Content Engines to channels [16](#)
 - device groups to channels [17](#)
 - media and overlay images to playlists [33](#)
 - patterns to existing pattern lists [24](#)
- address record mappings
 - configuring [76](#)
- admin accounts
 - configuring passwords [13](#)
- advanced HTTP cache settings
 - configuring [19](#)
- alarms
 - content replication [4](#)
 - device reporting [2](#)
 - overload detection
 - about [10](#)
 - configuring [11](#)
- allow
 - Spider script [81](#)
- alternative domains
 - defining [59](#)
- Apache CLF transaction logging format [2](#)
- approving
 - devices
 - in Content Distribution Manager GUI [22](#)
- assigning
 - devices to more than one device group [30](#)
 - playservers [33](#)
- associating
 - action with existing pattern list [24](#)
- attributes
 - for content serving [16](#)
 - shared
 - specifying [10](#)
- audit trail logs
 - viewing [30](#)
- authenticated HTTP cache settings
 - configuring [13](#)
- authentication
 - configuring
 - device login [7](#)
 - proxy [43 to 47](#)
 - users [9](#)
 - HTTP requests [6](#)
 - native FTP requests [15](#)
 - pass-through mode [2](#)
 - proxy mode [2](#)
 - Windows Media requests [6](#)
- authentication bypass
 - configuring [22](#)
- authentication cache
 - for access control [7](#)
 - size adjustments [11](#)
- authentication mechanisms
 - basic [3](#)
 - NTLM [3](#)
- authentication scheme
 - configuring [36](#)
- authentication servers
 - configuring [7, 9, 17 to 34](#)
- authentication traffic bypass [22](#)
- authorization

- content authorization [12](#)
- autoregistration
 - DHCP server requirements [8](#)

B

- backup root Content Engines
 - assigning [6](#)
- bandwidth
 - incoming Windows Media bypass lists [47](#)
 - scheduling for streaming acquisition [59](#)
- bandwidth efficiency gain reports [19](#)
- bandwidth settings
 - acquisition and distribution
 - defaults [54](#)
 - graphical display [56](#)
 - scheduled [57](#)
 - content services
 - defaults [41](#)
 - graphical display of [45](#)
 - scheduled [42](#)
- banners
 - configuring messages [5](#)
- baseline groups
 - working with [31](#)
- basic authentication
 - about [3](#)
- blocking messages
 - in URL filtering [31](#)
- bytes served
 - reports [19](#)
- HTTP [6](#)
- caching
 - web traffic [3](#)
- canonical name record mappings. *See* CNAME record mappings
- carousel passes
 - configuring [27](#)
- CdnManifest.xsd [65](#)
- CDP
 - configuring [7](#)
- channel assignments
 - viewing [14](#)
- channel distribution [3, 10](#)
- channel quota [25](#)
- channel routing algorithm [20](#)
- channels
 - acquisition and distribution properties
 - specifying [24](#)
 - adding and removing
 - Content Engines [16](#)
 - device groups [17](#)
 - assigning and removing multicast clouds [43](#)
 - configuring properties of [12](#)
 - content replication options [32](#)
 - creating [11](#)
 - definition [11, 12](#)
 - deleting [19](#)
 - for set-top-box export
 - selecting [41](#)
 - live [13](#)
 - manifest file information
 - configuring [26](#)
 - modifying [15](#)
 - proxy information
 - configuring [26](#)
 - updating content [53](#)
 - viewing
 - all channels [19](#)
 - Content Engines assigned to a channel [19](#)

C

- cache freshness settings
 - FTP [55](#)
 - HTTP [10](#)
- cache settings

- content replication status [5](#)
- system-wide replication status [1](#)
- CIFS
 - configuring
 - server settings [17](#)
 - website access control [19](#)
- CIFS share names
 - customizing [21](#)
- Cisco Discovery Protocol. *See* CDP
- Cisco Streaming Engine
 - enabling [14](#)
 - programming [25](#)
 - root failover support [5](#)
- CiscoWorks2000 [5](#)
- CLF
 - log file format [1](#)
- CLI commands
 - centrally managed settings [36](#)
 - mapped to Content Distribution Manager GUI [1 to 27](#)
 - not supported in Content Distribution Manager GUI [27](#)
- client IP spoofing [9](#)
- clock
 - setting [12](#)
- CNAME record mappings
 - configuring [77](#)
- command-line interface
 - accessing [20](#)
 - to configure network settings manually [15](#)
- Common Internet File System. *See* CIFS
- Common Log Format. *See* CLF
- conditions
 - modifying or deleting from IP ACLs [11](#)
- configuring
 - ACNS network [6, 8](#)
 - Content Engine in transparent mode [7](#)
 - devices
 - locally versus centrally [35](#)
 - HTTP request authentication [17](#)
 - warm standby Content Distribution Managers [2](#)
- congestion windows
 - about [7](#)
- content
 - adding items to channel [11](#)
 - classifications [5](#)
 - creating and managing in ACNS network [6](#)
 - distributing [3, 10](#)
 - filtering [9](#)
 - forcing replication [53](#)
 - mapping [38](#)
 - pre-positioning [9](#)
 - refreshing and verifying [19](#)
 - selecting with Manifest script [82](#)
 - specifying priority [13](#)
 - updating [53](#)
- content acquisition
 - scheduling [10](#)
- content acquisition and distribution
 - overview [11](#)
- content alarms
 - description [4](#)
- content attributes
 - listing [44](#)
- content authentication
 - overview of [1](#)
- content authorization
 - overview of [1](#)
- Content Delivery Network
 - description [3](#)
- Content Distribution Manager GUI
 - entering username and password [13, 2](#)
 - exiting [12](#)
 - feature tabs [2](#)
 - icons [8](#)
 - launching [13, 1, 2](#)
 - navigating [1](#)
 - window elements [3](#)
- Content Distribution Managers
 - configuring [16](#)

- deleting [5](#)
- exiting user interface [12](#)
- logging on to [1](#)
- primary role [2, 17](#)
- registering device with [20](#)
- switching roles [3](#)
- viewing configuration [2](#)
- warm standby
 - configuring [2](#)
 - deleting [5](#)
- Content Engine
 - functional description [3](#)
- Content Engine GUI
 - configuring secure access [59](#)
- Content Engines
 - activating [22, 24](#)
 - adding
 - to channels [16](#)
 - to device groups [29, 33](#)
 - to multiple device groups [30](#)
 - assigning to playlists [37](#)
 - configuring
 - as PAC file servers [38](#)
 - as routing Content Engines [51](#)
 - for web cache services (scenarios) [27, 30](#)
 - deleting [12](#)
 - enabling for multicasting [33](#)
 - modifying configuration properties [9](#)
 - primary role [2](#)
 - rebooting [12](#)
 - receivers [42](#)
 - removing
 - from channels [16](#)
 - restarting [12](#)
 - root [12, 16](#)
 - statistics
 - viewing [14](#)
 - viewing
 - channel assignments [14](#)
 - content replication status [14](#)
 - content providers
 - creating [6](#)
 - deleting [7](#)
 - description [11](#)
 - modifying [6](#)
 - content replication
 - overview [14](#)
 - content replication status
 - aggregate channels
 - viewing [1](#)
 - channel
 - viewing [5](#)
 - Content Engine
 - viewing [14](#)
 - Content Router
 - load-based routing [55](#)
 - request overflow routing [58](#)
 - Content Router routing
 - configuring DNS server [42](#)
 - configuring request routing [41](#)
 - overview [8](#)
 - Content Routers
 - activating [22](#)
 - configuring DNS servers for [42](#)
 - deleting [24](#)
 - Layer 4 switch interoperability
 - configuring [9](#)
 - modifying [22](#)
 - primary role [2](#)
 - services
 - configuring [25](#)
 - statistics [16](#)
 - troubleshooting configurations [59](#)
 - content routing
 - request overflow [58](#)
 - content services
 - overview [8](#)
 - content serving

- specifying
 - attributes [16](#)
 - metadata [18](#)
- conventions [27](#)
- cookies [3](#)
- Coordinated Universal Time. *See* UTC
- coverage zone files
 - registering [43](#)
- coverage zones
 - applying [44](#)
 - choosing [43](#)
 - Content Engine behind NAT firewall [48](#)
 - default [44](#)
 - file elements [43](#)
 - multiple Content Engines behind NAT firewall [49](#)
 - NAT firewall with multiple IP addresses [51](#)
 - no NAT firewall [47](#)
 - user-defined [44](#)
- CPAN [79](#)
- CPU utilization reports [20](#)
- crawler filter
 - specifying [11](#)
- crawling websites [80](#)
- crawl job
 - specifying [7](#)
- crawl task
 - adding to channel [13](#)
 - configuring acquisition rules [15](#)
- custom blocking messages
 - in URL filtering [31](#)
- customized blocking messages
 - configuring [31](#)
- custom transaction log format
 - description [2](#)
 - format token values [4](#)
 - log format string values [3](#)

D

- db
 - Manifest script [84](#)
 - Spider script [81](#)
- default PlayServerTable schema [70](#)
- deleting
 - Content Engines from network [12](#)
 - content providers [7](#)
 - Content Routers from network [24](#)
 - locations [8](#)
 - warm standby Content Distribution Managers [5](#)
- denying
 - requests for HTTP URLs [32](#)
- deployment scenarios [4](#)
- depth
 - Spider script [81](#)
- device alarms
 - description [2](#)
- device clock
 - setting [12](#)
- device configuration
 - local versus central [35](#)
- device groups
 - adding and removing
 - Content Engines [29, 33](#)
 - from channels [17](#)
 - assigning to playlists [38](#)
 - creating [26](#)
 - deleting [30](#)
 - description [10](#)
 - enabling overlap [30](#)
 - removing Content Engines [29, 33](#)
 - statistics [15](#)
 - viewing [30](#)
- device mode
 - changing [36](#)
 - configuring [14](#)
- devices

- monitoring performance statistics for [12](#)
- monitoring status for [12](#)
- reports
 - bandwidth efficiency gain [19](#)
 - bytes served [19](#)
 - CPU utilization [20](#)
 - streaming sessions [19](#)
- repurposing [36](#)
- DHCP
 - configuring interfaces for [13](#)
 - for autoregistration [8](#)
 - interface-level [10](#)
- DHCP server
 - requirements for autoregistration [8](#)
- differentiated services code point. *See* DSCP
- direct proxy routing [7](#)
- disks
 - configuring sysfs space [19](#)
- distribution
 - troubleshooting [47](#)
- distribution of content [3, 10](#)
- distribution priority [24](#)
- distribution QoS
 - configuring [16](#)
- distribution-related commands [16](#)
- DNS caching
 - address record mappings [76](#)
 - CNAME record mappings [77](#)
 - commands, list of [73, 80](#)
 - resource records [74](#)
 - server bindings [75](#)
 - server settings [71](#)
- DNS server
 - configuring Content Engine settings for [78](#)
 - configuring for content routing [42](#)
- documentation
 - conventions [27](#)
 - obtaining [28](#)
 - related [28](#)

- domains
 - adding entities [24](#)
 - assigning to user accounts [26](#)
 - configuring [23](#)
 - modifying and deleting [25](#)
 - viewing [26](#)
- downloading
 - sample manifest files [28](#)
- DSCP
 - setting values for [9, 18](#)
- dynamic proxy autoconfiguration
 - overview [8, 32](#)
- dynamic WCCP redirection [6](#)

E

- eDirectory Agent
 - description [50](#)
- eDirectory Agent settings
 - specifying [52](#)
- EIM
 - through content authentication [1](#)
 - through group-based authorization [12](#)
- EMBED tag [3](#)
- entities
 - adding to domains [24](#)
 - defined [23](#)
- error codes [59](#)
- error messages
 - configuring [23](#)
- escape characters
 - in URLs [32](#)
- EtherChannel
 - configuring [11](#)
- Ethernet MAC address mapping [6](#)
- Exec timeout
 - configuring [6](#)
- exiting
 - Content Distribution Manager GUI [12](#)

- expires attribute
 - for removing content [3, 20](#)
- expires directive [5](#)
- explicit congestion notification
 - about [7](#)
- exporting
 - ACNS network data [23](#)
 - log files [9](#)
- export programs
 - about [8](#)
- Extended Squid transaction log format [2](#)
- external FTP server
 - exporting transaction logs to [9](#)
 - permanent error from [9](#)

F

- failover
 - for live programs [5](#)
- failRetryInterval attribute [20](#)
- Fast Cache
 - about [25](#)
 - configuring on Content Engines [26](#)
- Fast Ethernet interfaces
 - modifying [7](#)
- Fast Start
 - about [24](#)
 - configuring on Content Engines [24](#)
- feature tabs [2](#)
- FEC [29](#)
- fetching
 - manifest files [53](#)
- Fibre Channel interface
 - configuring [9](#)
 - modifying [10](#)
- Fibre Channel storage array
 - using [20](#)
- file
 - Manifest script [83](#)
 - Spider script [81](#)
- file attribute
 - marking for import [44](#)
- file extension
 - mapping [38](#)
 - tag [40](#)
- file resending
 - triggering [29](#)
- file sharing
 - configuring [17](#)
- filtering
 - for crawl job [11](#)
 - of content [9](#)
- firewalls
 - distributing content through [23](#)
- forced content replication [53](#)
- forward error correction. *See* FEC
- forwarders
 - setting lookup level [22](#)
 - setting selection probability [21](#)
- freshness of pre-positioned content
 - configuring [64](#)
- FTP
 - using to fetch pre-positioned content [5](#)
- FTP manifest file
 - writing crawler [5](#)
 - writing single-item [4](#)
- FTP native caching
 - cache freshness [61](#)
 - client side [60](#)
 - connections [58](#)
 - custom messages [62](#)
 - overview of [59](#)
- FTP-over-HTTP
 - cache freshness [55](#)
 - caching [51](#)
 - commands, list of [53](#)
 - connections [52](#)
- FTP servers

- crawler job attributes [8](#)
- exporting transaction logs to [9](#)
- functional description
 - Content Engine [3](#)

G

- generating a playserver list [33, 13](#)
- getting started [1](#)
- Gigabit Ethernet interfaces
 - modifying [7](#)
- GLOP addresses [5](#)
- GMT offsets [71 to 78](#)
 - See also* UTC offsets
- group-based authorization
 - ACLs [12](#)
 - rules [12](#)
- GUI icons [8](#)

H

- hardware clock [12](#)
- hinted files
 - playing [3](#)
- HTTP
 - authenticated cache
 - commands, list of [17](#)
 - settings [13](#)
 - cache freshness
 - commands, list of [12](#)
 - settings [10](#)
 - caching
 - advanced settings [19](#)
 - commands, list of [8, 22](#)
 - settings [6](#)
 - connections [2](#)
 - outgoing proxy exclusion [48](#)
 - permitting specific URLs [33](#)

- pre-positioned content [4](#)
- redirection [4](#)
- status codes of [2](#)
- http append x-forwarded-for-header command [10](#)
- HTTP destination port
 - configuring restrictions [26](#)
- HTTP header filtering
 - manifest attributes for [50](#)
- HTTP manifest file
 - writing single-item [4](#)
- HTTP playservers
 - customizing [14](#)
- HTTP proxy
 - commands, list of [4](#)
 - configuring authentication [44](#)
- HTTP proxy authentication
 - configuring [45](#)
- HTTP proxy servers
 - configuring [37](#)
- HTTP request authentication
 - configuring [17](#)
 - transparent mode [9](#)
 - using NTLM [31](#)
- HTTPS
 - configuring servers [42](#)
 - filtering servers [44](#)
 - keys
 - filtering [35](#)
 - outgoing proxy exclusion [48](#)
 - playservers
 - customizing [14](#)
 - pre-positioned content and [4](#)
 - transparent caching
 - SSL [47](#)
- HTTPS manifest file
 - writing crawler [5](#)
- HTTPS proxy
 - commands, list of [40](#)
 - configuring [37](#)

IANA 1

ICAP

- approved vendors 3

- configuration

- viewing 13

- configuring

- Content Engine settings 1

- server 10

- services 7

- maximum file size supported 4

- overview of 1

- processing

- and Java applets 12

- statistics 13

ICAP services

- description of 2

icons 8

ICP

- client settings 80

- remote client settings 82

- remote server settings 84

- server settings 83

initial configurations

- for devices

- overview 2

- using autoregistration 8

- using the CLI 14

- using the setup utility 10

- for networks

- using the Quick Start tool 12

- summary of tasks 6

- using the quick setup method 8

interface-level DHCP

- description 12

- note 9

interfaces

- manually configuring for DHCP 13

Internet Assigned Numbers Authority. *See* IANA

Internet Content Adaption Protocol. *See* ICAP

IP/TV Program Manager

- role in ACNS network 3

IP/TV programs

- configuring multicast IP addresses for 3

- hinted and nonhinted 3

IP access control lists. *See* IP ACL

IP ACL

- adding conditions to 3

- applying to interface 13

- associating with application 12

- conditions

- modifying or deleting 11

- configuration constraints 2

- creating new 3

- deleting 14

- overview

IP addresses

- configuring NAT addresses 23

- multiple

- configuring on single interface 6

- static 12

IP differentiated services

- configuring 9

IP multicasting

- and insecure services 4

- copying files between server and client 4

- Ethernet MAC address mapping 6

- GLOP addresses 5

- Layer 2 multicast address considerations 5

- limited scope addresses 4

- Protocol Independent Multicast. *See* PIM

- source-specific multicast addresses 5

- unusable addresses 1

IP routes

- configuring 15

item priority 25

J

JavaScript [3](#)

K

keepExpiredContent attribute [56](#)

keepNoCacheContent attribute [56](#)

keepQueryString attribute [56](#)

kernal debugger

enabling [23](#)

L

launching the Content Distribution Manager GUI

procedure [1](#)

username and password [13, 2](#)

Layer 2

multicast address [5](#)

LDAP server

group settings [21](#)

password expiration [19](#)

policy redirection [20](#)

LDAP server settings

configuring [17](#)

level

definition [12](#)

license keys

multicast distribution [33](#)

limit

Spider script [81](#)

limited scope addresses [4](#)

limiting scope of spider search [80](#)

line console carrier detection

configuring [6](#)

live channels

adding Content Engines to [15](#)

adding device groups to [15](#)

creating new [14](#)

priming [18, 16](#)

selecting for programs [13](#)

live command in Manifest script [82, 83](#)

live content

description [6](#)

selecting with Manifest script [82](#)

specifying [20](#)

versus pre-positioned [82](#)

live events

about [3](#)

configuration requirements [4](#)

programming [13](#)

live programs

configuring streaming settings [18](#)

root Content Engine failover for [5](#)

scheduling [23](#)

live splitting

commands related to [40](#)

live stream interruptions [4](#)

live stream settings

for CSE programs [21](#)

for Windows Media programs [18](#)

live stream splitting

multicast-in multicast-out [31](#)

multicast-in unicast-out [34](#)

unicast-in multicast-out [28](#)

unicast-in unicast-out [36](#)

weighted load balancing [38](#)

location leaders

setting preferences for [21](#)

locations

creating default [23](#)

definition [12](#)

deleting [8](#)

modifying [3, 5](#)

working with [5 to 8](#)

location tree

definition [12](#)

viewing [8](#)

- log files
 - exporting [9](#)
 - restarting export of [9](#)
- log formats
 - WMT transaction logging [16](#)
- logging
 - commands, list of [8](#)
 - configuring [8](#)
 - message priority levels [8](#)
 - multicast transmissions [28](#)
 - system events [9](#)
- login access
 - controlling [1](#)
- login authentication
 - about [2](#)
 - configuring [7, 12](#)
 - default configuration for [3](#)
 - failover [4](#)
 - local access database [5](#)
- log on
 - following timeout [11](#)
 - user interface [1](#)

M

- MAC address mapping [6](#)
- management
 - central vs. local [35](#)
- manifest files
 - <crawler> feature [2](#)
 - automated scripts [78](#)
 - channel information
 - configuring [26](#)
 - definition [13](#)
 - dissemination of [2](#)
 - examples [21 to 27](#)
 - fetching [53](#)
 - modifying sample manifest file [33](#)
 - obtaining sample PERL scripts [33](#)
 - refreshing and verifying content [19](#)
 - regular expressions [10](#)
 - schema [65](#)
 - single <item> [2](#)
 - specifying
 - crawl job [7](#)
 - live content [20](#)
 - structure and syntax
 - correcting [32](#)
 - summary [33](#)
 - tags (listing) [33](#)
 - time values [18](#)
 - time zones [71](#)
 - updating [53](#)
- manifest file validator
 - accessing [28](#)
 - output errors [31](#)
 - output examples [29](#)
- Manifest Generator. *See* ACNS Manifest Generator tool
- Manifest script
 - about [79](#)
 - customizing [85](#)
 - rules file [85](#)
 - selecting content [82](#)
 - syntax guidelines [82](#)
- manifest tags
 - summary (table) [33](#)
- Manifest Validator utility. *See* manifest file validator
- manifest XML schema [65](#)
- manual setup and registration [12](#)
- map
 - Manifest script [84](#)
 - Spider script [81](#)
- match Manifest script [85](#)
- media attribute
 - marking for import [44](#)
- media files
 - adding to playlist [33](#)
- message logs

- viewing [9](#)
- message of the day
 - configuring [5](#)
- metadata
 - for content serving [18](#)
- metadata poll interval multiplier
 - configuring [22](#)
- MIBs
 - supported by SNMP agent [12](#)
- Microsoft Active Directory
 - using [13](#)
- migrating
 - from ACNS 4.x software to ACNS 5.x software [5](#)
- MIME content types [7, 38, 39](#)
- modifying
 - content providers [6](#)
 - Content Router properties [22](#)
 - locations [3, 5](#)
- monitoring
 - system status [1](#)
- monitoring URLs [20](#)
- MPEG file specifications
 - for TV-out enabled devices [32](#)
- multicast addresses
 - for programs [46](#)
 - unusable assignments [34, 1](#)
- multicast clouds
 - adding and removing receiver Content Engines [40](#)
 - assigning and removing from channels [43](#)
 - configuring [34](#)
 - deleting [40](#)
 - description [14, 26](#)
 - modifying [39](#)
 - parameters [37](#)
 - receiver Content Engines [42](#)
 - root Content Engines [16](#)
 - viewing [40](#)
- multicast configuration
 - verifying [45](#)

- multicast distribution license key [33](#)
- multicast expert mode
 - using [30](#)
- multicasting
 - carousel passes [27](#)
 - commands, list of [46](#)
 - definition [14, 25](#)
 - enabling Content Engines for [33](#)
 - overview [14](#)
 - preparation for [25](#)
 - sender delay interval [29](#)
- multicast IP addresses
 - configuring for IP/TV programs [3](#)
- multicast logging [28](#)
- multicast sender delay
 - configuring intervals [29](#)
- multicast transmissions
 - logging information [28](#)
- multiple device groups
 - assigning Content Engines to [30](#)
- multiple IP addresses
 - coverage zones and [51](#)
 - on single interfaces
 - configuring [6](#)
- multiple network interfaces
 - configuring [1](#)

N

- N2H2
 - configuration restrictions [39](#)
 - filter settings [35, 36](#)
 - supported features [34](#)
- NACK interval multiplier [28](#)
- name server. *See* NS
- NAT firewalls
 - configuring IP addresses [23](#)
 - content routing and [47 to 51](#)
- native FTP caching

- overview of [59](#)
- proxy authentication [59](#)
- restrictions [58](#)
- native FTP requests
 - authentication [15](#)
- navigating
 - in the Content Distribution Manager GUI [1](#)
- network settings
 - configuring manually [15](#)
- Network Time Protocol. *See* NTP
- no-cache directive [4](#)
- nonhinted files
 - playing [3](#)
- nonsecure access
 - configuring [59](#)
- noProxy attribute [39](#)
- NS
 - record for content routing [42](#)
- NSC files
 - for multicasting Windows Media server-side playlists [23](#)
- NTLM
 - authentication [7](#)
 - commands, list of [27, 30](#)
 - for proxied authentication [7](#)
 - HTTP request authentication [31](#)
 - server settings
 - configuring [23](#)
- NTLM authenticated content
 - acquiring [42](#)
- NTLM authentication
 - about [3](#)
 - version 2 support [4](#)
- NTP
 - and system clock [12](#)
 - configuring [11](#)

O

- OBJECT ID tag [3](#)
- offline detection
 - about [22](#)
 - configuring [21](#)
- offline operations
 - enabling [27](#)
- on-demand content
 - description [5](#)
- one-to-one channel mapping (figure) [11](#)
- online help [11](#)
- origin server
 - replicating content from [53](#)
 - setting up [3](#)
- overflow
 - content routing [58](#)
- overlay images
 - adding to playlist [34](#)
 - assigning to media [36](#)
 - editing [35](#)
 - parameters (table) [35](#)
- overload bypass [22](#)
- overload detection
 - alarms [10](#)

P

- PAC file servers
 - configuring Content Engines as [38](#)
- PAC file templates
 - description [33](#)
 - example [34](#)
 - registering [35](#)
- packet flow
 - diagram [4](#)
- pass-through authentication [2](#)
- passwords
 - CLI user

- changing [20](#)
 - username [13, 2](#)
- pattern list
 - adding patterns to [24](#)
 - associating action with [24](#)
 - configuring [23](#)
 - verifying action [25](#)
- Perl scripts
 - for manifest file [78](#)
- Perl software
 - installing [79](#)
- permitting
 - HTTP URLs [33](#)
- PGM
 - configuring using multicast expert mode [30](#)
- PIM [5](#)
- playback guide [39](#)
- playbacks
 - schedule settings [36](#)
- playlists
 - adding
 - media [33](#)
 - overlay images [33](#)
 - assigning
 - Content Engines [37](#)
 - device groups [38](#)
 - definition [8](#)
 - MPEG file specifications [32](#)
 - setting playback schedules [36](#)
 - viewing playback guide [39](#)
- playserver lists
 - generating [33, 13](#)
- playservers
 - assigning [33](#)
 - customized in manifest file [14](#)
 - mappings
 - accepted locations [38](#)
 - settings
 - versus <playServer> tag [39](#)
- playservetable
 - Manifest script [84](#)
- PlayServerTable.xsd [69, 70](#)
- PlayServerTable schema [69](#)
- playtimes
 - definition [8](#)
- policies
 - and Rules Template [13](#)
- policy servers
 - configuring for third-party [60](#)
- polling intervals
 - configuring for content updates [22](#)
- port channel interfaces
 - configuring [11](#)
 - load balancing [13](#)
- port lists
 - configuring for WCCP [17](#)
- prefix
 - Spider script [80, 81](#)
- preloaded content
 - description [6](#)
- prepos
 - Manifest script [84](#)
- pre-positioned content
 - acquiring [23](#)
 - description [6](#)
 - freshness of [64](#)
 - overview [9](#)
 - vs. live [82](#)
- primary Content Distribution Manager
 - configuring [17](#)
- printing
 - ACNS network data [23](#)
- priority of content
 - specifying [13](#)
- privilege profiles
 - about [15](#)
- program multicast addresses
 - configuring [46](#)

- managing [46](#)
- programs
 - about [1](#)
 - Cisco Streaming Engine rebroadcast
 - creating [25](#)
 - configuring
 - Cisco Streaming Engine live [20](#)
 - Windows Media live [18](#)
 - Windows Media rebroadcast streaming settings [29](#)
 - export
 - creating [40](#)
 - selecting media files and overlay images [33](#)
 - modifying API-defined [10](#)
 - previewing [11](#)
 - scheduling live [23](#)
 - TV-out [8](#)
 - creating [31](#)
 - scheduling [36](#)
 - selecting media files and overlay images [33](#)
 - viewing [8](#)
 - Windows Media rebroadcast
 - about [7](#)
 - creating [25](#)
- promptless authentication
 - configuring for IE browsers [11](#)
- proxied authentication mode [2](#)
- proxy authentication
 - configuring [43 to 47](#)
 - using CLI [45](#)
- proxy autoconfiguration files [39](#)
- proxy configuration [32](#)
 - for channels [26](#)
- proxy error messages
 - configuring
 - download settings [24](#)
 - upload settings [26](#)
- proxy protocols parameters [49](#)
- proxy servers
 - authentication for [43](#)

- configuring
 - for manifest file acquisition [40](#)
 - using CLI [37](#)
 - using manifest [38](#)
- publishing URLs
 - generating [30](#)
 - multicast [23](#)
 - unicast [22](#)

Q

- QoS
 - configuring [16](#)
- quality of service. *See* QoS
- Quick Crawl utility
 - using [11](#)
- quick setup method [8](#)
- Quick Start tool
 - using for network configurations [12](#)

R

- RADIUS Agent
 - description [49](#)
- RADIUS Agent settings
 - specifying using the CLI [52](#)
- RADIUS server
 - commands, list of [9, 34](#)
 - configuring [7, 32](#)
- RCP services
 - enabling on Content Engines [8](#)
- RealProxy
 - enabling [10](#)
- RealSubscriber
 - enabling [12](#)
- real-time transaction logging
 - configuring [20](#)
- rebooting

- Content Engines [12](#)
- reboot option
 - Content Engines [12](#)
- rebroadcast programs
 - assigning Content Engines [28](#)
 - assigning device groups [28](#)
 - scheduling [30](#)
 - selecting media files [25](#)
- receiver Content Engines [42](#)
- registering
 - device with Content Distribution Manager [20](#)
- regular expressions
 - writing [10](#)
- reject
 - Spider script [80, 81](#)
- remote login
 - controlling access [1](#)
- removing
 - Content Engines from channels [16](#)
 - device groups from channels [17](#)
- replication [1](#)
 - overview [14](#)
- replication status
 - channel
 - viewing [5](#)
 - commands, list of [17](#)
 - Content Engine
 - viewing [14](#)
 - viewing details [9](#)
- reportBrokenLinks attribute [56](#)
- repurposing ACNS network devices [36](#)
- request interception
 - overview [7](#)
- request overflow routing
 - defining alternative domains [59](#)
- request routing
 - using a Content Router [41](#)
- resource records [74](#)
- restarting export of log files [9](#)
- retransmit time multiplier
 - about [7](#)
- retry and refresh mechanism
 - for content acquisition [52](#)
- roles
 - assigning to user accounts [22](#)
 - creating and managing [20](#)
 - modifying and deleting [21](#)
 - viewing settings [22](#)
- root Content Engine
 - failover [6](#)
- root Content Engines
 - definition [12](#)
 - designating [16](#)
 - failover [14](#)
 - failover for live programs [5, 6](#)
- root location
 - deleting Content Engines from [12](#)
- router
 - commands, list of [26](#)
 - configuring
 - for web cache service [28, 30](#)
 - enabling WCCP support [26](#)
- router lists
 - configuring for WCCP [14](#)
- router redirection
 - overview [7](#)
- routing Content Engine
 - configuring request routing using [41](#)
- routing Content Engines [51](#)
- routing statistics [16](#)
- RTSP gateway
 - commands, list of [9](#)
 - configuring [8](#)
 - overview [7](#)
- rule-based filtering
 - for crawl job [11](#)
- rules
 - for group-based authorization [12](#)

- See Rules Template
- rules file for Spider and Manifest scripts [85, 86](#)
- rules settings
 - enabling [14](#)
- Rules Template
 - actions
 - order of execution [58](#)
 - configuring
 - using CLI [23](#)
 - using Content Distribution Manager GUI [14](#)

S

- sanitized transaction logs [8](#)
- scheduled Windows Media rebroadcasts
 - about [7](#)
- secure access
 - configuring [59](#)
- secure FTP [5](#)
- secure shell
 - configuring [2](#)
 - daemon password guesses [4](#)
 - host keys [4](#)
- sender delay interval
 - configuring [29](#)
- server bindings
 - configuring [75](#)
- service masks
 - configuring for WCCP [18](#)
- service rules
 - configuring [14](#)
 - See Rules Template
 - See also Rules Template
- set Manifest script [84](#)
- setting schedules [30](#)
- setup utility
 - for configuring devices [10](#)
- SFTP servers
 - exporting transaction logs to [10](#)
- shared attributes
 - specifying [10](#)
- show command utility
 - for troubleshooting [4](#)
- Simple Network Management Protocol. *See* SNMP
- simplified hybrid routing
 - configuring DNS server [42](#)
- single content item
 - specifying [6](#)
- size Manifest script [85](#)
- slow start [9](#)
- SmartFilter software
 - and rule action no-auth command [56](#)
 - URL filtering [55](#)
- SNMP
 - configuring
 - community settings [5](#)
 - group settings [6](#)
 - host settings [10](#)
 - server settings [2](#)
 - traps [5](#)
 - user settings [8](#)
 - view settings [9](#)
 - key CLI commands [13](#)
 - manager
 - creating [2](#)
 - supported MIBs [12](#)
 - traps [5](#)
- software clock [12](#)
- source-specific multicast address [5](#)
- Spider script
 - about [79](#)
 - broadening scope of crawling [80](#)
 - customizing [81](#)
 - keywords [81](#)
 - limiting crawl scope [80](#)
 - rules file [85](#)
 - syntax guidelines [80, 82](#)
- Squid transaction log format [2](#)

- SSL
 - caching HTTPS [47](#)
- standby Content Distribution Manager
 - configuring [17](#)
- standby groups
 - of interfaces [2](#)
- standby interfaces
 - configuring [3](#)
 - priority settings [4](#)
- start Spider script [81](#)
- static IP addresses [12](#)
- static IP routes
 - configuring [15](#)
- statistics
 - Content Engines [14](#)
 - Content Router [16](#)
 - device groups [15](#)
 - ICAP services [13](#)
 - streaming [16](#)
 - Websense [53](#)
- streaming servers
 - licensing and enabling [6](#)
 - supported in ACNS software [1](#)
- streaming sessions reports [19](#)
- streaming statistics
 - viewing [16](#)
- subtag
 - attributes [11](#)
- syntax
 - Spider script
 - guidelines [80](#)
- system event logging
 - configuring
 - using CLI [8](#)
 - using GUI [5](#)
 - message priority levels [8](#)
- system events log
 - viewing [9](#)
- system message log

- using [5](#)
 - viewing [9](#)
- system properties
 - modifying [18](#)
- system status
 - monitoring [1](#)
- system status bar
 - troubleshooting
 - content replication [5](#)
 - devices [3](#)

T

- TACACS+ server
 - commands [11, 36](#)
 - configuring [9, 34](#)
- tags
 - in manifest files [33](#)
- TCP
 - congestion windows [7](#)
 - explicit congestion notification [7](#)
 - parameter settings [1](#)
 - retransmit timer [7](#)
 - slow start [8](#)
- TCP-over-satellite extensions
 - about [8](#)
- Telnet services
 - commands [2](#)
 - enabling on Content Engines [1](#)
- text conventions [27](#)
- text editor creating rules file [86](#)
- TFTP
 - configuring
 - general settings [65](#)
 - proxy server [66](#)
 - directories
 - configuring [69](#)
- TFTP gateway
 - about [65](#)

Third Generation Partnership Project. *See* 3GPP [14](#)

time

manifest file [18](#)

Manifest script [85](#)

timeout

logging on after [11](#)

time values

specifying [18](#)

time zones

in manifest files [71](#)

parameter settings for [12](#)

token strings

in custom transaction logging format [2](#)

ToS

setting values for [9](#)

transaction logging

and NTLM authentication [7](#)

enabling with Content Distribution Manager GUI [10](#)

formats [2](#)

W3C customizable [6](#)

real-time

about using [19](#)

configuring [20](#)

sanitized [8](#)

WMT

about using [15](#)

configuring [17](#)

log formats [16](#)

transparent caching

and outgoing proxy exclusion [48](#)

and SSL [47](#)

transparent interception

using WCCP [2](#)

transparent mode

hierarchical caching in [9](#)

traps

SNMP [5](#)

Trivial File Transfer Protocol. *See* TFTP

troubleshooting

acquisition [48](#)

Content Router configurations [59](#)

distribution hierarchy [47](#)

using show command utility [4](#)

ttl attribute [20](#)

TV-out devices

supported MPEG file specifications [32](#)

TV-out programs

about [8](#)

commands, list of [32](#)

type Manifest script [85](#)

Type of Service. *See* ToS [9](#)

U

unicast content forwarding [14](#)

unicasting [14, 20](#)

unified name space errors. *See* UNS errors

UNS errors [59](#)

unusable multicast addresses [34, 1](#)

update interval [26](#)

uploading

proxy error messages [26](#)

URL filtering

configuration overview [25](#)

configuring using Content Distribution Manager GUI [26](#)

custom blocking messages [31](#)

overview [9](#)

with N2H2 server [33](#)

with SmartFilter software [55](#)

with URL lists [32](#)

with Websense [43](#)

URLs

monitoring [20](#)

user accounts

about [15](#)

audit trail logs

viewing [30](#)

- CLI users
 - changing passwords [20](#)
- creating
 - using Content Distribution Manager GUI [17](#)
- domains
 - adding entities [24](#)
 - assigning to [26](#)
 - configuring [23](#)
 - modifying and deleting [25](#)
 - viewing [26](#)
- managing [16 to 20](#)
- modifying and deleting [19](#)
- roles
 - assigning to [22](#)
 - creating [20](#)
 - modifying and deleting [21](#)
 - viewing [22](#)
- viewing [20](#)
- user authentication
 - overview [9](#)
- user authentication. *See* login authentication
- user interface
 - Content Distribution Manager [3 to 12](#)
- username password [13, 2](#)
- UTC [12](#)
- UTC offsets [14](#)
 - See also* GMT offsets

V

- VBScript [3](#)
- verifying
 - action performed on pattern list [25](#)
 - content acquisition [48](#)

W

- W3C customizable logging format [6](#)

- warm standby
 - configuring Content Distribution Managers [2](#)
 - switching to primary [3](#)
- WCCP [3](#)
 - clean shutdown
 - description [9](#)
 - client IP spoofing
 - description [9](#)
 - configuring
 - general settings for the Content Engine [7](#)
 - port lists [17](#)
 - router lists [14](#)
 - service masks [18](#)
 - service settings for the Content Engine [10](#)
 - custom web cache service [7](#)
 - dynamic redirection services [6](#)
 - enabling on router [26](#)
 - HTTPS traffic service [6](#)
 - load balancing
 - description [13](#)
 - packet flow diagram [4](#)
 - service groups [5](#)
 - slow start
 - description [9](#)
 - traffic bypass
 - configuring [23](#)
 - viewing lists [25](#)
 - transparent caching options [22](#)
- WCCP bypass
 - configuring [23](#)
 - viewing lists [25](#)
- WCCP routing
 - description [7](#)
- WCCP support
 - enabling on the router [26](#)
- WCCP transparent interception [2](#)
- Web Cache Communication Protocol. *See* WCCP
- web cache service
 - clients and cache on different subnets [29](#)

- clients and cache on the same subnet [26](#)
- configuration examples
 - for different subnets [31](#)
 - for same subnet [28](#)
- web crawler [7](#)
- webmasters
 - responsibilities [6](#)
- Websense server
 - components
 - about [48](#)
 - installing [50, 52](#)
 - configuration restrictions [55](#)
 - configuring failover [42, 43](#)
 - enabling [43](#)
 - integrated server
 - port configuration [42](#)
 - usage [41](#)
- websites
 - crawler job attributes [8](#)
 - crawling
 - broadening scope [80](#)
 - limiting scope [80](#)
 - creating [8](#)
 - definition [13](#)
 - deleting [10](#)
 - description [11](#)
 - modifying [10](#)
- web traffic
 - caching [3](#)
- weighted load balancing
 - configuring for live stream splitting [38](#)
- wildcard character [48](#)
- Windows Media
 - commands, list of [21](#)
 - enabling [15](#)
 - general settings (figure) [20](#)
 - incoming bandwidth bypass list [47](#)
 - live stream splitting [28 to 36](#)
 - multicasting [34](#)

- multicast station schedules [38](#)
- proxy settings [18](#)
- request authentication [6](#)
- Windows Media rebroadcasts [7](#)
 - configuring streaming settings [29](#)
 - programming [25](#)
- Windows Media server-side playlists
 - features supported [6](#)
 - using .nsc files [23](#)
- WMT
 - Fast Cache
 - configuring on Content Engines [26](#)
 - Fast Start
 - configuring on Content Engines [24](#)
- WMT transaction logging
 - about [15](#)
 - configuring [17](#)
 - log formats [16](#)

X

- XML
 - Manifest script [84](#)
- XML schema
 - default PlayServerTable [70](#)
 - manifest [65](#)
 - PlayServerTable [69](#)
- XSD [65](#)