

Novell Volera Excelerator

2.3

www.novell.com

ADMINISTRATION GUIDE



Novell®

Legal Notices

Novell, Inc. makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc. makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

This product may require export authorization from the U.S. Department of Commerce prior to exporting from the U.S. or Canada.

Copyright © 1997-2003 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

U.S. Patent Nos. 5,870,739; 5,873,079; 5,884,304; 6,330,605. U.S. and Foreign Patents Pending.

Novell, Inc.
1800 South Novell Place
Provo, Utah 84606
U.S.A.

www.novell.com

Volera Excelerator 2.3 Administration Guide
[May 2003](#)

Online Documentation: To access the online documentation for this and other Volera products, and to get updates, see www.novell.com.

Novell Trademarks

Volera and Novell are trademarks of Novell, Inc. in the United States and other countries.

Third-Party Trademarks

All third-party trademarks are the property of their respective owners.

Contents

Part I Learning What the Excelerator Appliance Can Do

1	How Caching Accelerates Web Content Delivery	15
	Web Browsing Basics	15
	What the Excelerator Appliance Does	15
2	Individual Browser Acceleration	17
3	Network-Wide Browser Acceleration	19
4	Web Server Acceleration	21

Part II Integrating Excelerator 2.3 into Your Network

5	Managing the Appliance	25
	The Browser-Based Management Tool.	25
	The Command Line.	25
6	Preparing the Network	27
	Basic Network Configuration Setup.	27
	Configuring the Client Workstation	27
	Configuring the Excelerator Appliance	28
7	Troubleshooting the Initial Appliance Setup	31
	Appliance Problems	31
	My appliance isn't working	31
	I can't ping the appliance from my client	31
	All the numbers are correct and the appliance still won't ping	31
	Browser Problems	31
	My browser can't find the application	31
	Nothing ever comes up on my browser	32
	None of the changes I made in the browser application are taking effect	32

Part III Accelerating HTML Content

8	Accelerating Individual Browsers	35
	Overview of Forward Proxy.	35
	Key Functionality	35
	How Forward Proxy Works	35
	Benefits of Forward Proxy	36
	Forward Proxy Setup	36
9	Accelerating All the Browsers on a Network	39
	Overview of Transparent Proxy.	39
	Transparent Proxy with an L4 Switch	39
	Transparent Proxy with a WCCP-Capable Router	40
	Transparent Proxy as a Default Gateway (Router)	41
	Transparent Proxy as an Inline Router (Network Gateway)	42

Transparent/L4 Proxy Setup	43
Transparent/WCCP Proxy Setup	44
Transparent/Default Gateway Setup	47
Transparent/Inline Router Setup	49
10 Accelerating Web Servers	51
Overview of Web Server Acceleration	51
How Origin Web Server Acceleration Works	51
Benefits of Origin Web Server Acceleration	52
Web Server Accelerator Setup	52
Working with DNS	53
Configuration Considerations When Using Appliance Multihoming Features	54
Part IV Accelerating Streaming Media	
11 Streaming Media Overview	59
Streaming Media Protocols	61
QuickTime	61
Windows* Media Player*.	63
Real	64
12 Preparing to Cache Streaming Media Data	67
Identifying Your Streaming Media Types	67
Assessing Your Network Bandwidth Capacity.	67
Upstream Versus Downstream Bandwidth	67
Assessing Bandwidth for Forward and/or Transparent Streaming Proxy Services	68
Assessing Bandwidth for a Streaming Server Accelerator.	70
Configuring the Appliance to Match Your Requirements	72
Managing Streaming Bandwidth (Admission Control)	72
Managing Disk Space and Streaming Objects.	74
Managing Streaming Object Cache Freshness	74
Managing the Caching of Streaming Objects	75
Managing the Total Sessions Allowed	77
Getting QuickTime Streaming Content Through Firewalls	77
The Appliance and HTTP Tunneling.	77
Passing Streaming Content Through the Appliance Without Caching It.	78
Setting Up Your Appliance to Work with Firewalls.	78
Logging Streaming Media Transactions.	80
13 Accelerating Streaming Media to Individual Media Players	81
Overview of Forward Streaming Proxy	81
Key Functionality.	81
How Forward Streaming Proxy Works.	82
Benefits of Forward Streaming Proxy	82
Forward Streaming Proxy Setup	82
14 Accelerating Streaming Media to All Media Players on the Network	87
Overview of Transparent Streaming Proxy	87
Transparent/L7 Streaming Proxy	87
Key Functionality.	87
How Transparent Streaming Proxy Works with an L7 Switch	88
Benefits of Transparent Streaming Proxy with an L7 Switch	88
Setting Up Transparent Streaming Proxy with an L7 Switch	88
Transparent/L4 RTSP Streaming Proxy.	90
Key Functionality.	90
How Transparent RTSP Streaming Proxy Works with an L4 Switch	90
Benefits of Transparent RTSP Streaming Proxy with an L4 Switch	90

Setting Up Transparent RTSP Streaming Proxy with an L4 Switch.	91
Transparent/L4 HTTP Tunneled Streaming Proxy	92
Key Functionality	92
How Transparent HTTP Tunneled Streaming Proxy Works with an L4 Switch	92
Benefits of Transparent HTTP Tunneled Streaming Proxy with an L4 Switch	93
Setting Up Transparent HTTP Tunneled Streaming Proxy with an L4 Switch	93
15 Accelerating Streaming Media Servers	95
Overview of Streaming Media Server Acceleration.	95
How Origin Streaming Server Acceleration Works	95
Streaming Accelerator Services Are Created in Matching Pairs	96
Benefits of Origin Streaming Server Acceleration.	96
Streaming Server Accelerator Setup	97
Deploying Multiple Streaming Accelerators on the Appliance	98
16 Configuring an Upstream Proxy for the Appliance	101
17 Configuring QuickTime Media Players to Use Proxy Services	103
Setting Streaming Proxy Options.	103
Setting Streaming Transport Options.	103
Part V Hierarchies, Clusters, and Multihoming	
18 Hierarchical Caching	107
Overview	107
CERN Hierarchies	107
ICP Hierarchies	108
CERN Hierarchy Setup	109
ICP Hierarchy Setup	111
19 Clustering	113
Overview	113
Web Server Accelerator Groups	113
Appliance Clusters	114
Cluster Setup	115
About Capacity and Weight	118
20 Appliance Groups and Multihomed Configurations	121
Appliance Web Server Accelerator Group Setup.	121
Standard Multihoming for Multiple Web Sites	122
Example: Accelerating Multiple Web Servers on a Single IP Address	122
Example: Accelerating a Multihomed Web Server	123
Multihoming and Path-Based Support	123
Path-Based Multihoming Examples	124
Example One	124
Example Two	124
Domain-Based Acceleration	125
Understanding Domain-Based Acceleration	125
Creating a Domain-Based Accelerator	127
CLI Summary	127
Part VI Managing and Leveraging Excelerator's Advanced Features	
21 Installing and Upgrading Licenses	133
Obtaining Product Licenses	133
Listing Currently Installed Licenses.	133
Viewing License Information	133
Removing Licenses.	133

Installing Licenses Using a Floppy Disk	134
Installing Licenses Using FTP	134
22 Authentication Services	135
Matching Authentication Profiles to Your Requirements	135
Understanding How Profiles Work	135
A Summary of Authentication Method Pros and Cons	137
Combining Authentication Profiles	138
Understanding How Authentication Cookies Are Used	140
Setting Up Authentication Services (Overview)	140
Using Mutual (Certificate-Based) Authentication	141
How Mutual Authentication Works	141
Platforms Requirements	141
Preparing Your Network for Mutual Authentication	142
Setting Up Mutual Authentication	142
Using LDAP Authentication	143
How LDAP Authentication Works	143
Platforms Supported	144
Preparing Your Network for LDAP Authentication	144
Setting Up LDAP Authentication	145
Enabling NDS Single Sign-On for an LDAP Authentication Profile	147
Using RADIUS Authentication	148
How RADIUS Authentication Works	148
Platforms Supported	148
Preparing Your Network for RADIUS Authentication	149
Setting Up RADIUS Authentication	149
Using NDS (eDirectory) Authentication	150
How NDS (eDirectory) Authentication Works	150
Platforms Supported	150
Preparing Your Network for NDS (eDirectory) Authentication	151
Setting Up NDS (eDirectory) Authentication	151
Enabling NDS Single Sign-On for an NDS Authentication Profile	152
Using NDS (eDirectory) Single Sign-On Functionality	152
How NDS (eDirectory) Single Sign-On Works	152
Platforms Supported	153
Preparing Your Network for NDS (eDirectory) Authentication	153
Setting Up and Enabling NDS (eDirectory) Single Sign-On	154
Using Basic Authentication	155
How Basic Authentication Works	155
Platforms Supported	156
Preparing for Basic Authentication	156
Setting Up Basic Authentication	156
Using NTLM Authentication	157
How NTLM Authentication Works	157
Platforms Supported	158
Preparing Your Network for NTLM Authentication	158
Setting Up NTLM Authentication	159
NTLM Authentication Multiple Domain Support	159
23 Access Control	161
Overview	161
Process	161
Determining Your ACL Strategy	161
Wildcarding	162
Authentication	162
Examples	162

Creating an ACL	162
Implementing the ACL	163
Other Guidelines	163
24 Managing Appliance Certificates	165
Naming Certificates	165
Creating Certificates Using the Appliance CA	166
Obtaining a Certificate from an External CA	166
Requesting the CSR	166
Sending the CSR	167
Storing the Certificate	168
Viewing (Exporting) a Certificate's CA	168
Modifying a Certificate	169
Importing a Trusted Root to a Cache Device	169
Deleting a Certificate	170
Backing Up a Certificate	170
Restoring a Certificate	171
25 Transforming Content for Internet Delivery	173
Identifying the Issues	173
Understanding URL Overrides	173
Secure Excelerator: The All-in-One Solution	173
URL Overrides Transform Object Reference URLs in Cached HTML	174
Automatic Vs. Manual URL Overrides	174
Using URL Overrides	176
DNS Name Overrides	177
Port Overrides	179
Pathrule Overrides	181
Scheme Overrides	183
Reviewing URL Overrides	184
26 Cache Freshness	185
Overview	185
Managing Cache Freshness	185
How the Excelerator Appliance Checks for Object Freshness	185
How an Excelerator Appliance Keeps the Oldest Cached Objects Fresh	186
How Excelerator Handles the Freshest Objects in Cache	186
Fine-Tuning Cache Freshness on Your Appliance	187
Using Custom Cache Control Headers	187
An Overview of How Headers Work	187
Implementing Custom Cache Control Headers	187
An Implementation Example	188
27 Managing Appliance Security Features	191
The Console Lock Feature	191
Managing HTTP CONNECT Method Support	191
How the CONNECT Method Works	191
An Unverified CONNECT Connection Is a Security Risk	191
How Excelerator Protects Your Network	192
Configuring Excelerator to Meet Your CONNECT Method Requirements	192
28 Automatic Configuration Mechanisms	195
About Appliance Configuration Files	195
System-Generated Configuration Files	196
Using Customized Configuration Files to Change the System Configuration	197
Managing Configuration Files	197
Using the Browser-Based Management Tool	197

Using Telnet or the Command Line	198
Using FTP	198
Backing Up the Appliance Configuration	198
Verifying Appliance Configurations	199
Viewing the IMPORT.LOG File via Netscape	199
Viewing the IMPORT.LOG via FTP	199
Creating Appliance Configuration Shortcuts	199
Restoring Factory Settings	199
Restoring the Appliance to the Clone Image	200
Reimaging and Restoring the Appliance System	201
29 Content Filtering	203
Overview of Filtering.	203
Understanding When Filtering Actually Starts.	204
No Filtering During the Initial Download	204
Filtering During Appliance Startup	204
If You Remove a Filtering Service	204
The Filter Processing Sequence	204
Configuring a Filtering Service	205
Configuring RAM Usage on a Low-Memory Cache Device	205
Changing the Default Download Schedule for a Filtering Service.	206
The Bypass List	206
The Override List	206
Creating an Override List	206
Critical Information about Wildcards in the Override List	206
Critical Information about Filtering in CERN and ICP Hierarchies.	207
Bypassing an N2H2 Filtering Service	207
Critical Information Regarding this Feature	207
Overview.	209
Preparing to Install the Filter Bypass Software	210
Installing the Filter Bypass Software.	210
Configuring the Main Perl Script on the IIS Server	211
Configuring the IIS Server To Run the Perl Scripts	212
Changing Each Excelerator's Default Blocked Content Page	212
Setting Up Administrative Access to the List of Bypass Users	212
Adding Authorized Filter Bypass Users	213
Starting the Filtering Bypass Service	213
Running in a Hierarchy with Filtering on the Parent	213
Starting the Filtering Bypass Service Automatically	214
Tips for Non-Default Installations	214
Notes About Product Security	214
30 DNS Name Resolution	215
How the Appliance Resolves DNS Names	215
How the Appliance Formulates Subsequent DNS Queries	215
The DNSINFO.CFG File	215
An Example	215
Modifying the DNS Lookup Sequence.	217
Managing the HOSTS File.	217
Resolving DNS Names in Hierarchies.	218
Purging the Appliance's DNS Cache	218
31 Controlling Referred Access to Content	219
How Referer Header Validation Works	219
Understanding the Referer Header Configuration File	219
Configuration File Specifications.	219

Access Is Either Totally Limited or Unrestricted	221
Two Key Points About Target_URL Entries.	221
About URL Line Entries in the Configuration File: Conventions, Wildcards, Etc.	221
Including All Subdirectories and Objects in a Single Entry	223
Allowing a Web Server to Link to Its Own Objects	223
How Requests Are Processed	224
A Hypothetical Example	224
InfoAndStreams Referer Header Configuration File	225
The Net Result of the Sample Configuration File	227
Setting Up a Referer Header Control	227
Customizing Referer Header Error Messages	227
Creating the Referer Configuration Files	228
Loading the Referer Configuration File	229
Activating the Referer Control	229
Managing Referer Controls	229
De-activating a Referer Control	229
Refreshing the Referer Header Configuration File	230
Getting Information for a Referer Header Control.	230
32 Dynamic Bypass	231
Why Dynamic Bypass Is Needed.	231
What Excelerator Does	231
33 Appliance Error Messages	233
What You Need to Know about Appliance Error Messages	233
Checking the Language Directories on Your Appliance	233
Customizing the Appliance Error Template and Message Files	234
Creating a New Language	234
Customizing the Error Message Text of an Existing Language.	234
Customizing the Error Message Format of an Existing Language	235
34 Logging	237
Using Appliance Logging Services	237
Overview of Appliance Logging.	237
What the Appliance Can Log	237
The Costs of Logging	238
System Constraints	238
Planning Your Logging Strategy	240
Planning Step 1: Determining Your Logging Requirements	240
Planning Step 2: Selecting a Log File Format and Optimizing the Log Entry Size	240
Planning Step 3: Calculating Log Rollover Requirements	241
Configuring Logging Options	243
Configuration Step 1: Opening the Appropriate Log Options Dialog Box.	243
Configuration Step 2: Selecting a Log Format	243
Configuration Step 3: Specifying Rollover Options	243
Configuration Step 4: Specifying Handling of Older Files.	244
Configuration Step 5: Monitoring and Refining Your Logging Strategy	244
About the FTP Log Push Feature	245
Using FTP Push to Automatically Download and Delete Log Files.	245
Manually Downloading and Deleting Log Files	246
When to Download and Delete Log Files	246
Getting Log Filenames	246
Downloading Log Files	249
Deleting Downloaded Log Files.	249
About Extended Log Field Headers	249
Extended Logging Enhancements in Excelerator 2.3	252

Standard Header Field Logging	252
Pseudo Header Field Logging	257
35 FTP Services	261
Tips for Using Excelerator FTP Services	261
Firewalls Usually Require Passive FTP	261
Directory and File Names Cannot Contain Spaces	262
Appliance Routing and Transparent Proxy Limitations on FTP	262
Switching from Anonymous FTP to Username/Password FTP	262
Setting Up Appliance FTP Services	262
Functionality Limitations of the Appliance's Mini FTP Server	262
Starting an FTP Session with the Appliance	263
Changing the FTP Working Directory	263
Managing Configuration Files with FTP	264
Downloading a Configuration File to Your Workstation Using FTP	264
Moving a Configuration File to the Appliance from a Workstation	264
Moving a Configuration File to the Appliance and Executing It	264
Customizing the Appliance Splash Screen with FTP	264
Using Non-Anonymous FTP	265
Client Acceleration (Forward Proxy)	265
Reverse Proxy	265
Accelerating FTP Requests	265
FTP Forward Proxy	265
FTP Accelerator	266
FTP Forward Proxy Setup	266
FTP Accelerator Setup	266
36 Object Pinning	269
The Pin List	269
Configuring Pin Lists	269
URL Mask	270
Pin Type	270
Pin Links	271
Pin Images	271
Refresh Frequency/Time.	271
How URL Masks Are Processed	271
About Wildcards in Pin Lists	274
Pin List Examples	274
Excelerator Records IP Addresses When Resolving URL Masks.	276
37 Router Capabilities	277
Using Appliance Routing	277
38 Shutting Down and Restarting	279
Restarting from the Browser-Based Management Tool	279
Shutting Down and Restarting from Telnet or the Command Line	279
39 SOCKS Client Services	281
Using the SOCKS Client Service	281
40 Time Synchronization	283
Synchronizing Time	283
Using the Browser-Based Management Tool	283
Using the Command Line	284
NTP Date/Time Synchronization Is Not Immediate	284
41 Web Proxy Auto-Discovery (WPAD)	285

Customizing Web Proxy Auto-Discovery	285
How the Appliance Handles WPAD Requests	285
Creating a Default WPAD.DAT Configuration File	286
Customizing System-Created WPAD DAT Files	286
Setting Up Forward Proxy with WPAD	287

Part VII Browser-Based Tool Help

42 Using the Browser-Based Management Tool	293
Prerequisites for Running the Management Tool.	293
Starting the Management Tool	293
The Apply and Cancel Buttons	294
The Help Button	294
Encryption	294
43 Using the Home Panel	295
Introduction Tab	295
Health Status Tab	295
Certificate Maintenance Tab	297
Add Ons Tab	298
Licensing Tab.	298
44 Using the System Panel	301
Timezone Tab	301
Date/Time Tab	302
Actions Tab	303
Change Password Dialog Box	305
Purge Cache Dialog Box	306
SNMP Tab	306
Import/Export Tab	308
Upgrade Tab	310
Alerts Tab.	312
Admin ACL Tab.	314
IP QoS Tab	315
IP Access Control Tab	316
Source IP or Subnet Dialog Box	318
Destination IP or Subnet Dialog Box	319
45 Using the Network Panel	321
IP Addresses Tab.	321
TCP Options Dialog Box	322
Adapter Options Dialog Box	323
DNS Tab	324
Advanced DNS Options Dialog Box	325
Gateway/Firewall Tab	326
Additional Gateways Dialog Box	328
Routes Dialog Box	330
46 Using the Cache Panel	333
Client Accelerator Tab	333
Log Options Dialog Box.	335
FTP Log Push Configuration Dialog Box	336
Add Authentication Profiles Dialog Box	338
Custom Cache Control Header Dialog Box	340
Access Control Options Dialog Box.	341
Advanced Options (Tuning) Dialog Box.	341
Transparent Handling Tab	342

WCCP Version 1.0 Options Dialog Box	345
WCCP Version 2.0 Options Dialog Box	345
Router Options Dialog Box	347
Web Server Accelerator Tab	348
Web Server Accelerator Dialog Box	349
URL Override Dialog Box	353
Path Rule Options Dialog Box	353
FTP Tab	354
FTP Accelerator Dialog Box	356
Streaming Tab	357
Streaming Management Configuration Dialog Box	358
Forward Streaming Services (RTSP)	361
Transparent Streaming Service Dialog Box (RTSP).	362
Reverse Streaming Service Dialog Box (RTSP).	363
Streaming Media Log Options Dialog Box (RTSP)	364
Upstream Proxy Dialog Box (RTSP).	365
Forward Streaming Services (MMS).	366
Transparent Streaming Service Dialog Box (MMS)	367
Reverse Streaming Service Dialog Box (MMS)	368
Streaming Media Log Options Dialog Box (MMS).	369
Upstream Proxy Dialog Box (MMS)	370
Cluster Tab	371
Insert Client Accelerator Service (Forward Proxy) Dialog Box.	373
Insert Transparent Client Accelerator Service (Transparent Proxy) Dialog Box.	374
Insert Web Server Accelerator Service Dialog Box	375
Authentication Tab	377
Authentication Dialog Box	378
Mutual Authentication Options Dialog Box.	379
LDAP Options Dialog Box	381
Import Trusted Root Dialog Box (LDAP Authentication).	384
RADIUS Options Dialog Box.	385
NDS Options Dialog Box.	386
Basic Authentication Options Dialog.	386
NTLM Authentication Options Dialog Box	387
Download Tab	388
Scheduled Download Dialog Box	388
Filtering Tab	390
Insert Filter Service Dialog Box	393
Modify Filter Service Dialog Box.	393
Websense Options Dialog Box	395
Management Tab	396
The Pin List	397
Enable Dynamic Bypass.	398
Caching Based On URL Content	399
The Reset Button	399
Tuning Tab	399
Cache Freshness Dialog Box	401
Mini Web Tab	402

47 Using the Hierarchy Panel 405

ICP/CERN Configuration Tab	405
ICP Parent Dialog Box.	408
ICP Peer Dialog Box.	408
CERN Parent Dialog Box	409
ICP Access Control Dialog Box	410
Bypass Tab (Hierarchy)	411

48	Using the Monitoring Panel	413
	Summary Tab	413
	Services Tab	414
	Performance Tab	415
	Cache Tab	416
	ICP Tab	418
	ICP Client Statistics	419
	ICP Server Statistics	420
	CERN Statistics	420
	Cluster Tab	420
	FTP Tab	421
	Streaming Tab	423
	Streaming RTSP Statistics	423
	HTTP Tunneling Statistics (Pass Through Only)	424
	Cache Logs Tab	425
	Top Ten Sites Tab	426
	WCCP Tab	427
Part VIII	Reference Guides	
A	Command Line Reference	431
	Troubleshooting the Command Line	432
	Commands entered return an error	432
	I made several changes that don't appear when I use the GET command to display them	432
B	Connecting Through Telnet	433
	Starting a Telnet Session	433
	Setting Up an Appliance Using Telnet	434
	Additional Information	435
	Disabling Telnet Access	435
	Establishing a Null-Modem Connection	435
	Additional Information	438
	Troubleshooting Telnet	439
	Telnet never starts	439
	Telnet starts after a long time	439
	Commands at the bottom of the screen look strange or don't make sense.	439
C	Upgrading the Appliance	441
	Upgrading	441
	Critical Information	441
	Upgrading Version 1.0 Purges the Cache	441
	Making Sure You Update the Clone Image Before and After Upgrading	441
	Disabling Forward Proxy Authentication	441
	Re-Imaging the Appliance from the Vendor CD.	441
	Preserving Configuration Settings	441
	Upgrading through a Firewall	442
	Downloading and Installing the Upgrade	442
	Uninstalling the Most Recent Upgrade	443

Learning What the Excelerator Appliance Can Do

If you are new to caching and Web content acceleration, you can use the chapters in this section to:

- ♦ Become familiar with basic caching terminology and concepts
- ♦ Begin thinking about how the Excelerator appliance fits with your content delivery strategy

The following table summarizes the tasks you can accomplish using the chapters in this section.

To	See
Learn the basics of Web content caching	Chapter 1, "How Caching Accelerates Web Content Delivery," on page 15
Learn about providing Web content caching services to individual browsers	Chapter 2, "Individual Browser Acceleration," on page 17
Learn about accelerating Web content to all the browsers on a network	Chapter 3, "Network-Wide Browser Acceleration," on page 19
Learn about accelerating content delivery from Web servers to the Web	Chapter 4, "Web Server Acceleration," on page 21

1

How Caching Accelerates Web Content Delivery


This section describes:

- ♦ **Web Browsing Basics**
- ♦ **What the Excelerator Appliance Does**

Web Browsing Basics

The explanations of appliance caching features contained in this manual build on the illustration in **Figure 1** of basic Web browsing:

Figure 1




- 1 The user enters a URL in the browser, creating a DNS request.
- 2 DNS returns the numeric IP address of the origin Web server.
- 3 The browser requests objects from the origin Web server.
- 4 The origin Web server accepts the request into its queue, processes it in turn, and returns the requested objects to the browser.
- 5 The process repeats each time a browser makes the same request.

The Excelerator appliance eliminates the redundant network traffic and server processing time associated with Step 5 in **Figure 1**.

What the Excelerator Appliance Does

All appliance cache services include the basic functionality in **Figure 2**.


Figure 2



- 1 The appliance receives a browser request.
- 2 The appliance gets objects not in cache from the origin Web server.
- 3 The appliance caches objects and sends copies to the browser.

After a request has been cached, processing subsequent requests for the same objects is simpler and much faster.

Figure 3



- 1 The appliance receives a browser request and sends copies of the objects back to the browser.

Basic cache services fit in three categories:

- ♦ Accelerating browsers individually
- ♦ Accelerating all the browsers on a network
- ♦ Accelerating Web servers

Although these services are described separately, they can generally be combined and used simultaneously on a single appliance.

2

Individual Browser Acceleration

The most basic method for accelerating content delivery to browsers is to set up the appliance as an *HTTP forward proxy server*.


After this has been done, users who want to get accelerated content can configure their browsers to use the appliance's forward proxy IP address and port number as their forward proxy server.

Configuration steps are different for each browser. For example, in Internet Explorer 5 you can access the Proxy Settings dialog box by clicking Tools > Internet Options > Connections > LAN Settings > Advanced.

You can also set the appliance as a streaming media forward proxy server. Users can then configure their streaming media players to point to the appliance IP address and port configured for the streaming media forward proxy service you have created.

Figure 4 illustrates how forward proxy services work.

Figure 4



- 1 The browser is configured to use the appliance as a proxy server.
- 2 All browser requests are sent to the appliance.
- 3 The appliance gets objects not in cache from the Web.
- 4 The appliance sends cached objects to the requesting browser.

For more information about forward proxy services, see [Chapter 8, “Accelerating Individual Browsers,”](#) on page 35, and [Chapter 13, “Accelerating Streaming Media to Individual Media Players,”](#) on page 81.

3

Network-Wide Browser Acceleration

Another method for accelerating content delivery to browsers is to configure a router or switch on the network to route all HTTP traffic to a transparent proxy service you have set up on the appliance.

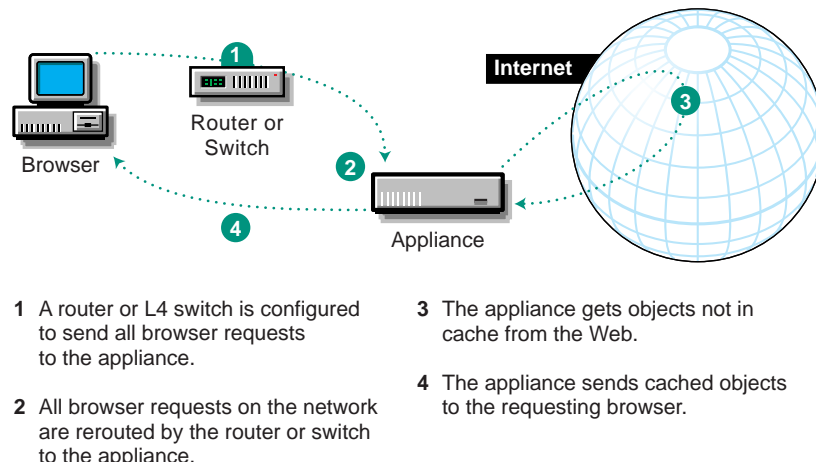
After this has been done, all network users automatically get accelerated content.

The specific configuration steps are different for each router or switch, but the basic concept is that the switch or router sends all browser requests that use the network's HTTP port number (port 80 in most cases) to the appliance.

You can also set the appliance as a streaming media transparent proxy server and have a router or switch send all streaming content requests to the appliance's streaming media transparent service.

Figure 5 on page 19 illustrates how transparent proxy services work.

Figure 5



For more information about transparent proxy services, see [Chapter 9, “Accelerating All the Browsers on a Network,”](#) on page 39, and [Chapter 14, “Accelerating Streaming Media to All Media Players on the Network,”](#) on page 87.

4

Web Server Acceleration


Your Excelerator appliance can dramatically improve the performance and response time of your Web site by offloading the burden of handling redundant content requests from the Web server.

After this has been done, the Web server can devote its bandwidth to handling requests for specific content and services and to supplying uncached and/or updated content to the appliance for subsequent caching.

You can also set the appliance to offload the request burden from a streaming media server.

Figure 6 on page 21 illustrates how Web server acceleration services work.

Figure 6



1 DNS is configured to resolve object requests to the appliance's IP address rather than to the Origin Web Server's address.


2 The appliance is configured as a Web server accelerator.

3 Browser requests are sent to the appliance rather than to the origin Web server.

4 The appliance gets objects not in cache from the origin Web server.

5 The appliance sends cached objects to the requesting browser.

For more about transparent proxy services, see [Chapter 10, "Accelerating Web Servers,"](#) on [page 51](#), and [Chapter 15, "Accelerating Streaming Media Servers,"](#) on [page 95](#).



Integrating Excelerator 2.3 into Your Network

You should have received a *Getting Started* guide with your Excelerator appliance. It is designed to help you quickly connect the appliance to your network and then test it to ensure it is configured correctly.

IMPORTANT: You should complete the initial setup before proceeding with the instructions in this guide.

The following table summarizes the tasks you can accomplish using the information in this section.

To	See
Learn about your options for managing the appliance	Chapter 5, “Managing the Appliance,” on page 25
Install the appliance on your network and prepare your network workstations to use appliance services	Chapter 6, “Preparing the Network,” on page 27
Troubleshoot any initial setup problems you encounter	Chapter 7, “Troubleshooting the Initial Appliance Setup,” on page 31

5

Managing the Appliance

The Excelerator appliance can be configured and managed in the following ways:

- ♦ Using the browser-based management tool from a workstation on the network.
- ♦ From the command line through a Telnet or null-modem connection. (You can also use an attached keyboard and monitor if your appliance has the required connections.)

The Browser-Based Management Tool

The browser-based management tool is unlike other management utilities because its interface appears in your browser as an HTML page originating from the appliance. The only programs associated with this tool that run on your workstation are a Java-compatible Web browser and the Java* components required by the HTML page.

If you experience problems with the interface, such as the page freezing, you can usually solve the problem by re-clicking any icon in the tool or refreshing the page.

For more information about using the browser-based management tool, see [Appendix 42, “Using the Browser-Based Management Tool,” on page 293](#).

The Command Line

Although it is possible to configure and monitor an appliance using only the command line interface, we strongly recommend that you use the browser-based tool for all administrative tasks whenever possible.

The browser-based tool includes extensive cross-checking, helpful messages, and other program features to ensure that Excelerator is configured correctly for optimal performance. The command line interface does not include these features. Even the most expert users can overlook critical steps in configuring Excelerator from the command line.

For more information about using the command line, see [Appendix A, “Command Line Reference,” on page 431](#).

6

Preparing the Network

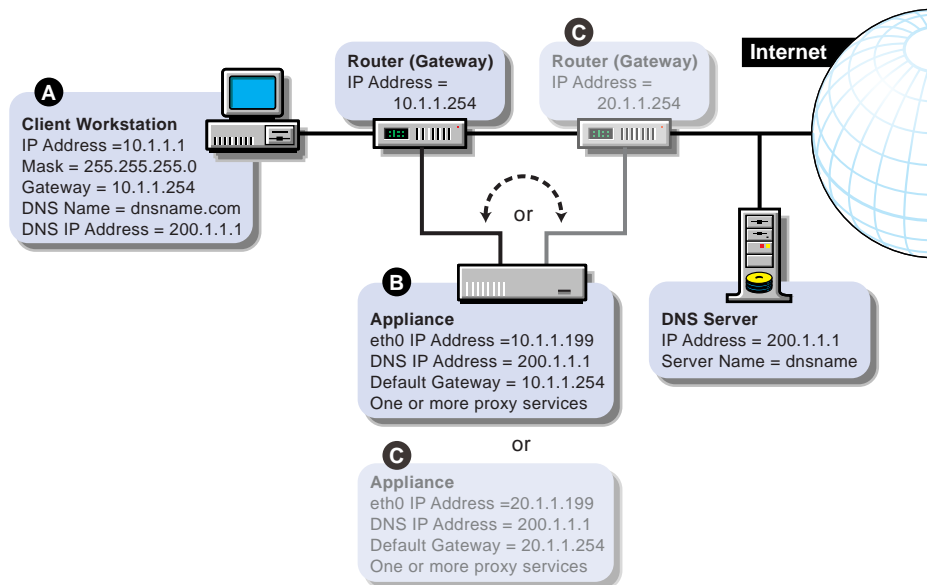
After you complete the initial appliance setup and test, review this chapter to ensure that all your network components are properly configured.

Basic Network Configuration Setup

Figure 7 on page 27 provides a visual map for the information in this chapter.

NOTE: The letters in Figure 7 on page 27 are referenced in the tables that follow. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 7



Configuring the Client Workstation

In most cases, client workstations on the network are already configured with IP address information to use the network. If that is the case with your client workstations, you can skip this section.

The workstation of each browser that will use appliance proxy services must be configured with the IP address information listed below. (List items marked with asterisks [*] must be on the same subnet.)

- ♦ A numeric IP address on the subnet *
- ♦ The subnet mask *

- ♦ The numeric IP address of the default gateway for the subnet *
- ♦ The numeric IP address of the DNS server the browser will use to resolve DNS names
- ♦ The domain name for the DNS server the client will use (optional)

Configuration procedures vary for each platform. Refer to the workstation documentation for specific instructions.

Configuration Requirements	Do This	Notes
A numeric IP address on the subnet	Refer to setup instructions for the system.	See A in Figure 7 on page 27 .
The subnet mask	The procedure is different for each platform. On a Windows* 95/98 or Windows NT* workstation, for example, right-click the Network Neighborhood icon on the desktop.	The IP address, subnet mask, and gateway address must all be on the same subnet.
The numeric IP address of the default gateway for the subnet		
The numeric IP address of the DNS server the browser will use to resolve DNS names		
The domain name for the DNS server the client will use (optional)		

Configuring the Excelerator Appliance

NOTE: If you used the Getting Started to set up your appliance, you have already completed most of the following steps.

IMPORTANT: When possible, connect the network cable to the network card on the appliance before assigning an IP address to the card. If this is not possible, you might need to restart the appliance after the cable is attached so the IP address assignment can take effect.

Configure the appliance following the steps below:

To Configure	Do This	Notes
IP addresses and subnet masks for the network connections (eth0, eth1, etc.) that will handle proxy services	<ol style="list-style-type: none"> 1. In the browser-based tool, click Network > IP Addresses > Add Address. 2. Enter the addresses in the appropriate fields > click the Assign to Adapter drop-down list > select the appropriate adapter. 3. Click Apply. 	<p>See <i>B</i> and <i>C</i> in Figure 7 on page 27.</p> <p>The appliance does not need to be on the same subnet as the browser. If the appliance is on a different subnet, its IP address will reflect a different subnet.</p> <p>Also, eth0, eth1, etc., can be on different subnets.</p> <p>IMPORTANT: If you plan to use an appliance cluster, do not assign the IP addresses for clustered services to appliance network adapters. IP addresses used for clustered services are assigned during the clustered service creation process. Addresses are dynamically bound to network adapters only while appliances are hosting the clustered services associated with them.</p>
At least one DNS server IP address	<ol style="list-style-type: none"> 1. In the browser-based tool, click Network > DNS. 2. Enter the addresses in the appropriate fields. 3. Click Apply. 	
The numeric IP address for a gateway (router) on the same subnet as the appliance	<ol style="list-style-type: none"> 1. In the browser-based tool, click Network > Gateway/Firewall. 2. Enter the address in the Default Gateway IP Address field. 3. Click Apply. 	<p>See <i>B</i> in Figure 7 on page 27.</p> <p>If the appliance is on the same subnet as the client workstation, the appliance and the workstation will have the same gateway address.</p> <p>If the appliance is on a different subnet than the browser, its gateway address will be the IP address of the router on the other subnet. See <i>C</i> in Figure 7 on page 27.</p>

To Configure	Do This	Notes
Passwords for Config and View users	<ol style="list-style-type: none"> 1. In the browser-based tool, click System > Actions > Password. 2. Click the User drop-down list > select Config. 3. Enter the password information > click Change. 4. Repeat Step 2 and Step 3 for the View user. 5. Click Apply. 	<p>For information about Config and View users, see “Change Password Dialog Box” on page 305.</p> <p>NOTE: Telnet is not secure unless a password is set.</p> <p>We strongly recommend you set system passwords as part of the initialization process. For more information, see “Change Password Dialog Box” on page 305.</p>
One or more proxy services	See “Accelerating HTML Content” on page 33 .	

IMPORTANT: If you are reinitializing the system, you should remove the CD, shut down Excelerator, turn the appliance off, and then restart it.

7

Troubleshooting the Initial Appliance Setup

This section covers troubleshooting the initial appliance setup.

Appliance Problems

My appliance isn't working

- ☐ Most problems are caused by invalid IP address configurations. Four things are critical:
 - ♦ A numeric IP address with a subnet mask
 - ♦ A valid gateway address on the same subnet as the IP address
 - ♦ A valid DNS server IP address
 - ♦ A valid DNS domain name

I can't ping the appliance from my client

- ☐ The IP address for the client must be 10.1.1.2 (or any other valid 10.1.1 subnet address other than 10.1.1.1) and the subnet mask must be 255.255.255.0. The gateway must be the address of the appliance; in the original configuration, that address is 10.1.1.1. DNS on the client must also be set to the IP address of the appliance.

All the numbers are correct and the appliance still won't ping

- ☐ Some Ethernet cards under Windows NT* or Windows* 95/98 do not allow a crossover cable. If you suspect this is the problem, try connecting the two machines with a standard Ethernet cable running through a hub.
- ☐ Windows 2000 requires modification of its registry to work with a cross-over cable.

To initialize an appliance from a Windows 2000 workstation, you must complete the instructions in “[How to Disable Media Sense for TCP/IP in Windows 2000](http://support.microsoft.com/support/kb/articles/Q239/9/24.ASP?LN=EN-US&SD=gn&FR=0)” (<http://support.microsoft.com/support/kb/articles/Q239/9/24.ASP?LN=EN-US&SD=gn&FR=0>) on the Web.

Browser Problems

My browser can't find the application

- ☐ The correct URL is <http://10.1.1.1:1959/appliance/config.html>.
- ☐ Make sure you specify <http://> in the URL window. Typing the address of the application without <http://> doesn't work.

Nothing ever comes up on my browser

- ☐ You must be using Netscape* Navigator* 4.07 (or higher), Netscape Communicator* 4.5 (or higher), or Internet Explorer 4.01 (or higher) with the appliance. Also, the Excelerator release notes might contain more information regarding browser compatibility.
- ☐ You must have a JVM* (Java* Virtual Machine) installed. It comes with Netscape but might not be installed.
- ☐ Check to see if a copy of Netscape is still in memory. (Open the task list by pressing Ctrl+Alt+Del.) Sometimes Netscape does not exit correctly from a previous session.
- ☐ Try exiting from your client OS and restarting.
- ☐ Try the SHUTDOWN command from a Telnet or command line session on the appliance. Then turn the appliance off and on again, and wait for it to come up.
- ☐ If you just started the appliance, you might be trying to start before the server is up. When the appliance starts, you will hear the startup beep pattern (two longs and four shorts) repeated four times. When the beeping stops, the appliance is ready.

None of the changes I made in the browser application are taking effect

- ☐ After making the changes, you must click Apply to make the changes effective.



Accelerating HTML Content

After you have connected the appliance to your network and completed the other basic setup instructions in [“Integrating Excelsior 2.3 into Your Network” on page 23](#), use the instructions in this section to learn more about the different acceleration services the Excelsior appliance offers and to set up the HTML acceleration services your Web content strategy requires.

Each chapter in this section provides:

- ♦ An overview of the proxy service
- ♦ Instructions for setting up the proxy service

The following table summarizes the tasks you can accomplish using the information in this section.

To	See
Accelerate Web content for individual browsers on your network	Chapter 8, “Accelerating Individual Browsers,” on page 35
Accelerate Web content for all the browsers on your network	Chapter 9, “Accelerating All the Browsers on a Network,” on page 39
Accelerate content delivery from your Web servers to the Web	Chapter 10, “Accelerating Web Servers,” on page 51

8

Accelerating Individual Browsers

This section contains information about accelerating individual browsers.

Overview of Forward Proxy

This section presents a conceptual overview of Excelerator forward proxy services.

Setup instructions are in [“Forward Proxy Setup” on page 36](#).


Key Functionality

You can configure browsers with the IP address of a forward proxy service or you can set up Web Proxy Auto-Discovery (WPAD) services on your network so that configured browsers can obtain proxy address information automatically.

After they are properly configured, browsers send requests directly to an appliance IP address configured for forward proxy services. The forward proxy service obtains the objects and forwards copies back to the browsers.

How Forward Proxy Works

Figure 8



- 1 A browser requests an origin Web server's Web page from its forward proxy server (the appliance).
- 2 The forward proxy service obtains the numeric IP from DNS.
- 3 The service obtains objects from the origin Web server.
- 4 The service forwards copies of the retrieved objects to the browser.
- 5 The forward proxy service handles subsequent requests for the same Web page objects without accessing DNS or the origin Web server.

Benefits of Forward Proxy

- ◆ Forward proxy doesn't require a special router configuration.
- ◆ Forward proxy provides an immediate improvement in browser performance.
- ◆ Forward proxy allows users to decide whether to use the proxy service.


For tips and guidelines on setting up forward proxy services, see [“Forward Proxy Setup” on page 36](#) and [“Setting Up Forward Proxy with WPAD” on page 287](#).

Forward Proxy Setup

Figure 9 provides a visual map for the information in this section.

NOTE: The letters in Figure 9 are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 9



Set up forward proxy services as follows:

To	Do This	Notes
Ensure your basic network configuration is complete	1. See “Basic Network Configuration Setup” on page 27 .	

To	Do This	Notes
Enable forward proxy services on the appliance	<ol style="list-style-type: none"> 1. In the browser-based tool, click Cache > Client Accelerator > Enable Client Acceleration (Forward Proxy). 2. In the Proxy IP Addresses list, check the IP addresses that forward proxy services will be available on. 3. Enter the port that Accelerator will receive and process forward proxy requests on. (The default is 8080.) 4. Click Apply. 	<p>See A in Figure 9 on page 36.</p> <p>For more information, see “Client Accelerator Tab” on page 333.</p>
Enable the client browsers to use proxy services	See the software vendor’s documentation for more information.	<p>See B in Figure 9 on page 36.</p> <p>Use one of the checked IP addresses as the address for the forward proxy server.</p> <p>Be sure to use the same port number as configured on the appliance.</p>

9

Accelerating All the Browsers on a Network

This sections contains information about accelerating all the browsers on a network.

Overview of Transparent Proxy

Transparent proxy services require browser requests to be routed to the appliance from a network router or switch. This chapter reviews four different router/switch configurations and contains setup instructions for each configuration type.

The four router/switch configurations are:

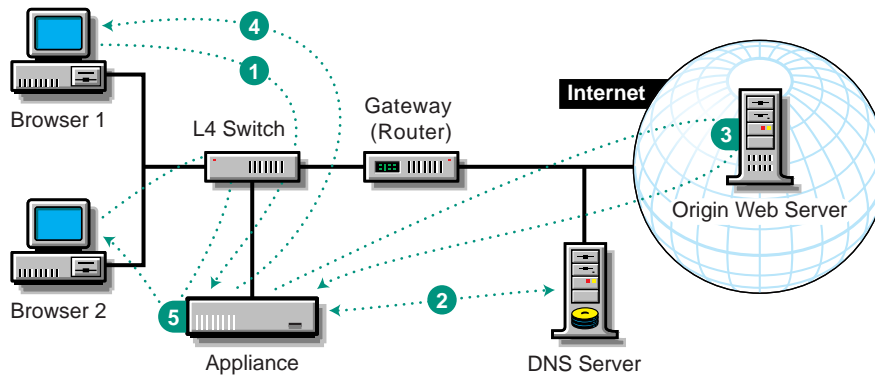
- ♦ An L4 switch
- ♦ A WCCP-capable network router
- ♦ The appliance's internal routing service, acting as the default gateway for the network subnet being accelerated
- ♦ The appliance's internal routing service, acting as an inline (main) network router

Transparent Proxy with an L4 Switch

An L4 switch on the same network as the client workstation intercepts browser requests from the client and sends them to the appliance. The transparent proxy service processes the request for the browser.

How Transparent Proxy Works with an L4 Switch

Figure 10



- 1 A browser requests a Web page from an origin Web server. The L4 switch detects that the request is on port 80, intercepts it, and sends it to the appliance's transparent proxy service.
- 2 The service obtains the numeric IP address from DNS.
- 3 The service gets the Web page objects from the origin Web server.
- 4 The service forwards copies of the retrieved objects to the browser.
- 5 The appliance's transparent proxy service handles subsequent requests for the same Web page objects without accessing DNS or the origin Web server.

Benefits of Transparent Proxy with an L4 Switch

Transparent proxy doesn't require browser configuration. After the switch and the appliance are configured, proxy services are transparent to the browser.


For tips and guidelines on setting up transparent proxy services using an L4 switch, see ["Transparent/L4 Proxy Setup" on page 43](#).

Transparent Proxy with a WCCP-Capable Router

A WCCP-capable router, which is configured as the default gateway for the client workstation, intercepts browser requests from the client and routes them to the appliance. The transparent proxy service processes the request for the browser.

How Transparent Proxy Works with a WCCP-Capable Router

Figure 11



- 1 A browser requests a Web page from an origin Web server. The WCCP-capable router detects that the request is on port 80 and routes it to the appliance's transparent proxy service
- 2 The appliance obtains the IP address of the origin Web server from DNS.
- 3 The proxy service obtains Web page objects from the origin Web server.
- 4 The service forwards copies of the retrieved objects to the browser.
- 5 The appliance's transparent proxy service handles subsequent requests for the same Web page objects without accessing DNS or the origin Web server.

Benefits of Transparent Proxy with a WCCP-Capable Router

Transparent proxy doesn't require browser configuration. After the router and the appliance are configured, proxy services are transparent to the browser.

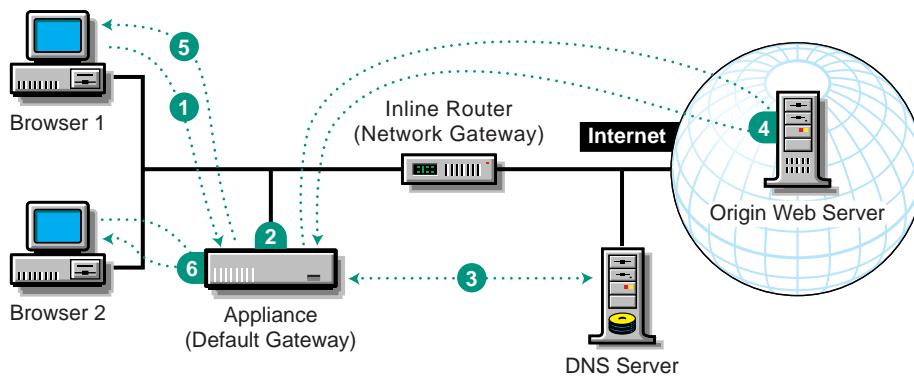
For tips and guidelines on setting up transparent proxy services using WCCP-capable routers, see [“Transparent/WCCP Proxy Setup” on page 44](#).

Transparent Proxy as a Default Gateway (Router)

An appliance's IP address is specified as the default gateway for the client workstation that the browser resides on. The appliance provides both routing and transparent proxy services.

How Transparent Proxy Works as a Default Gateway

Figure 12



- 1 A browser requests a Web page from an origin Web server through the client workstation's default gateway (the appliance).
- 2 The appliance detects that the request is to port 80 and routes it to its transparent proxy service.
- 3 The service obtains the numeric IP from DNS.
- 4 The service obtains Web page objects from the origin Web server.
- 5 The service forwards copies of the retrieved objects to the browser.
- 6 The transparent proxy service handles subsequent requests for the same Web page objects without accessing DNS or the origin Web server.

Benefits of Transparent Proxy as a Default Gateway

- ◆ Transparent proxy doesn't require browser configuration.
- ◆ This configuration might be useful for small business networks where the transparent proxy server is on the same subnet as the browsers.
- ◆ Using the appliance as a low-cost router might be sufficient for some small businesses.


For tips and guidelines on setting up the appliance as a default gateway and a transparent proxy server, see [“Transparent/Default Gateway Setup” on page 47](#).

Transparent Proxy as an Inline Router (Network Gateway)

An appliance's eth0 IP address is specified as the default gateway for the client workstations that the browsers reside on. The eth1 port connects to the Internet. The appliance provides routing and caching services and serves as the gateway to the Internet.

How Transparent Proxy Works as an Inline Router

Figure 13



- 1 A browser requests a Web page from an origin Web server through the client workstation's default gateway (the appliance).
- 2 The appliance detects that the request is on port 80 and routes it to its transparent proxy service.
- 3 The service obtains the numeric IP address from DNS.
- 4 The service obtains Web page objects from the origin Web server.
- 5 The service forwards copies of the retrieved objects to the browser.
- 6 The transparent proxy service handles subsequent requests for the same Web page objects without accessing DNS or the origin Web server.

Benefits of Transparent Proxy as an Inline Router

- ♦ Transparent proxy doesn't require browser configuration.
- ♦ This configuration might be useful for small business networks where all the client workstations and the transparent proxy server are on the same network.
- ♦ Using the appliance as a low-cost router might be sufficient for some small businesses.

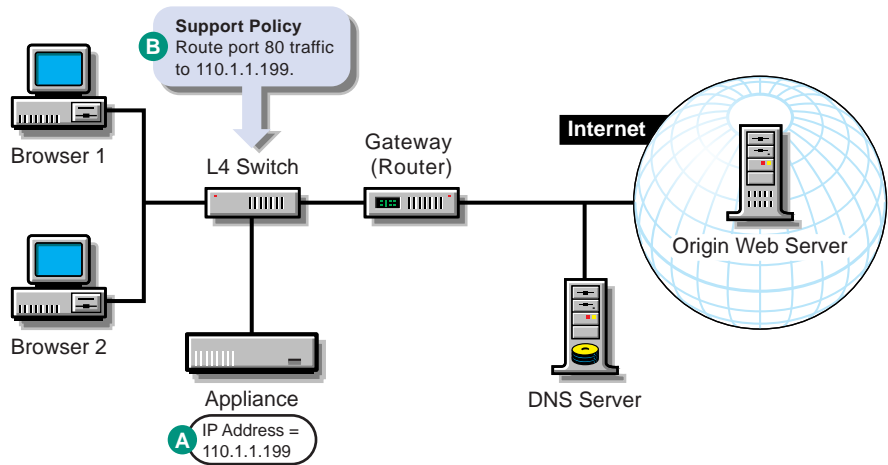
For tips and guidelines on setting up transparent proxy services with the appliance as an inline router, see [“Transparent/Inline Router Setup” on page 49](#).

Transparent/L4 Proxy Setup

Figure 14 provides a visual map for the information in this section.

NOTE: The letters in **Figure 14** are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 14



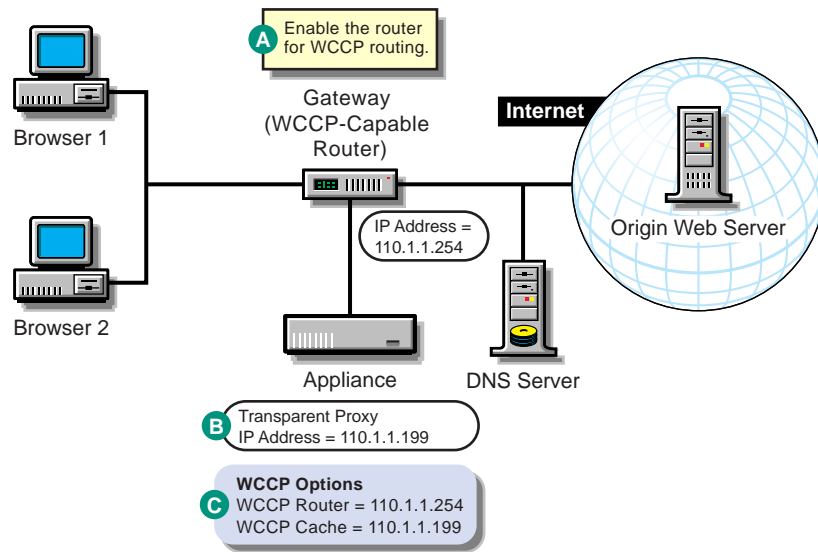
To	Do This	Notes
Ensure your basic network configuration is complete	1. See “Basic Network Configuration Setup” on page 27.	
Set up transparent proxy services on the appliance	<ol style="list-style-type: none">1. In the browser-based tool, click Cache > Transparent Handling > Enable Transparent Client Acceleration (Transparent Proxy L4 Switch Support).2. Insert one or more ports. (The default is 80.)3. Check the IP addresses you want to service transparent proxy requests. Only check one address for any given network card.4. Click Apply.	<p>See A in Figure 14 on page 44.</p> <p>When transparent proxy is enabled, it is active for all IP addresses on the appliance, except those configured for origin Web server acceleration services on the same port.</p> <p>No additional configuration is necessary for Excelerator to work with an L4 switch.</p> <p>For more information, see “Transparent Handling Tab” on page 342.</p>
Set up your L4 switch to route browser requests (port 80 traffic) to the appliance	1. Configure a support policy to redirect traffic to a transparent proxy address on the appliance. Refer to the documentation for your switch.	See B in Figure 14 on page 44 .

Transparent/WCCP Proxy Setup

[Figure 15](#) provides a visual map for the information in this section.

NOTE: The letters in [Figure 15](#) are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 15



To	Do This	Notes
Ensure your basic network configuration is complete	1. See “Basic Network Configuration Setup” on page 27.	
Enable WCCP routing on the router	1. Enable the router for WCCP routing. Follow the router manufacturer’s directions.	See A in Figure 15 . A WCCP-capable router can service more than one transparent proxy server.
Set up transparent proxy services on the appliance	<ol style="list-style-type: none"> 1. In the browser-based tool, click Cache > Transparent Handling > Enable Transparent Client Acceleration (Transparent Proxy L4 Switch Support). 2. Insert one or more ports. (The default is 80.) 3. Check the IP addresses you want to service transparent proxy requests. Only check one address for any given network card. 4. Click Apply. 	See B in Figure 15 on page 45. When transparent proxy is enabled, it is active for all IP addresses on the appliance, except those addresses configured for origin Web server acceleration services on the same port. For more information, see “Transparent Handling Tab” on page 342.

To	Do This	Notes
Register the appliance with a WCCP version 1 router if the routers on your network use WCCP version 1	<ol style="list-style-type: none"> 1. After you have enabled transparent proxy services, check Enable WCCP. 2. Click WCCP V1 Options. 3. Enter a proxy name and a farm name in their respective fields. (See Notes column for more information.) 4. In the WCCP Router field, enter the IP address of the WCCP router. 5. In the WCCP Cache field, enter an appliance IP address configured for transparent proxy service. 6. Click Apply. 	<p>See C in Figure 15 on page 45.</p> <p>The Proxy Name and Farm Name fields on the WCCP options form are text strings for your reference. They have no other function.</p> <p>The appliance needs the WCCP Router IP address in order to register with the router.</p> <p>The router needs the WCCP Cache IP address to know where to send browser requests.</p> <p>An appliance can register with only one WCCP-capable router.</p> <p>If WCCP routing isn't working, try entering the <code>get stats wccp</code> command on the command line. Check the configuration for problems.</p> <p>For more information, see "WCCP Version 1.0 Options Dialog Box" on page 345.</p>

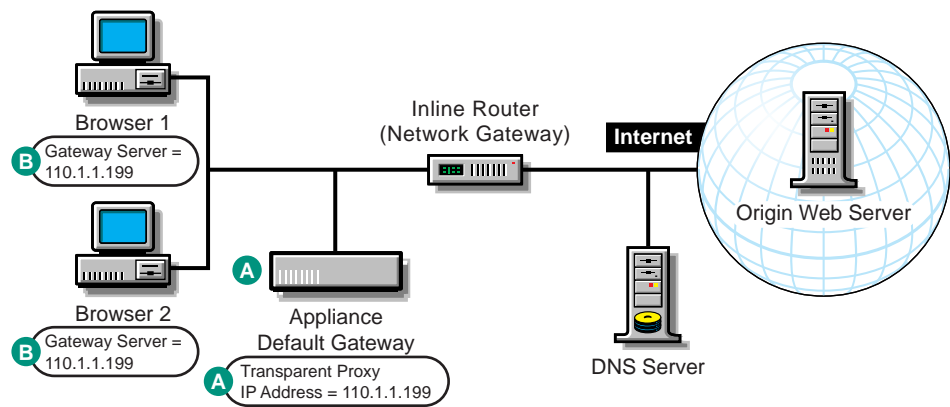
To	Do This	Notes
Register the appliance with the WCCP version 2 routers if the routers on your network use WCCP version 2	<ol style="list-style-type: none"> 1. After you have enabled transparent proxy services, check Enable WCCP. 2. Click WCCP V2 Options. 3. Click the Cache IP Address drop-down list > select an appliance IP address configured for transparent proxy service. This IP address should also be the main appliance connection to the Internet or network containing the origin Web servers. 4. Check either Use Unicast or Use Multicast. 5. If you checked Use Multicast, enter the multicast IP address > click OK > click Apply. Excelsior validates the address. 6. If you checked Use Unicast, insert one or more WCCP version 2 router IP addresses. 7. If you want the appliance to use signed packets for WCCP version 2 communications, check Enable Secure WCCP Communications > type a password > click OK. 8. Click Apply. 	<p>See C in Figure 15 on page 45.</p> <p>The router needs the WCCP Cache IP address to know where to send browser requests.</p> <p>Excelsior needs one or more WCCP router IP addresses in order to register with the routers.</p> <p>An appliance can register with multiple WCCP version 2 routers.</p> <p>If WCCP routing isn't working, try entering the <code>get stats wccp</code> command on the command line. Check the configuration for problems.</p> <p>For more information, see "WCCP Version 2.0 Options Dialog Box" on page 345.</p>

Transparent/Default Gateway Setup

[Figure 16](#) provides a visual map for the information in this section.

NOTE: The letters in [Figure 16](#) are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 16



To	Do This	Notes
Ensure your basic network configuration is complete	<ol style="list-style-type: none">1. See “Basic Network Configuration Setup” on page 27.	
Set up transparent proxy services on the appliance	<ol style="list-style-type: none">1. In the browser-based tool, click Cache > Transparent Handling > Enable Transparent Client Acceleration (Transparent Proxy L4 Switch Support).2. Insert one or more ports. (The default is 80.)3. Check the IP addresses you want to service transparent proxy requests. Only check one address for any given network card.4. Click Apply.	<p>See A in Figure 16.</p> <p>For more information, see “Transparent Handling Tab” on page 342.</p>
Set up the appliance as a router	<ol style="list-style-type: none">1. After you have enabled Exceleator for transparent proxy services, check Router Options.2. Click Router Options.3. Set up appliance routing.4. Click Apply.	<p>See A in Figure 16 on page 48.</p> <p>For more information, see “Router Capabilities” on page 277 and “Router Options Dialog Box” on page 347.</p>

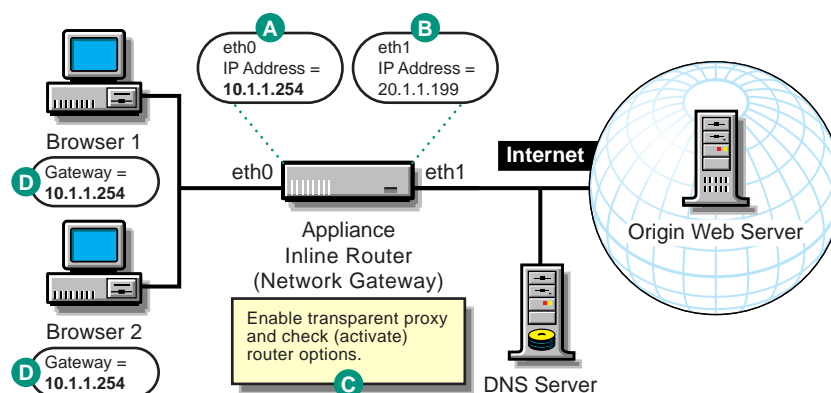
To	Do This	Notes
Configure the client workstations to use the appliance as their gateway	Refer to setup instructions for the system. The procedure is different for each platform. On a Windows 95/98/NT workstation, for example, right-click the Network Neighborhood icon on the desktop.	See <i>B</i> in Figure 16 on page 48 . Use the transparent proxy IP address as the gateway IP address.

Transparent/Inline Router Setup

[Figure 17](#) provides a visual map for the information in this section.

NOTE: The letters in [Figure 17](#) are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 17



To	Do This	Notes
Ensure your basic network configuration is complete	1. See “Basic Network Configuration Setup” on page 27 .	
Configure the network adapters for inline routing	<ol style="list-style-type: none"> 1. In the browser-based tool, click Network > IP Addresses. 2. Configure eth0 with an IP address on the same subnet as the client workstations. 3. Configure eth1 with an IP address on the network. 4. Click Apply. 	<p>See <i>A</i> in Figure 17 on page 49.</p> <p>The eth0 IP address handles transparent proxy services and doubles as the gateway address for all client workstations on the subnet.</p> <p>See <i>B</i> in Figure 17 on page 49.</p> <p>The eth1 IP address provides access to the network.</p>

To	Do This	Notes
Set up transparent proxy services on the appliance	<ol style="list-style-type: none"> 1. In the browser-based tool, click Cache > Transparent Handling > Enable Transparent Client Acceleration (Transparent Proxy L4 Switch Support). 2. Insert one or more ports. (The default is 80.) 3. Check the IP addresses you want to service transparent proxy requests. Only check one address for any given network card. 4. Click Apply. 	<p>See C in Figure 17 on page 49.</p> <p>For more information, see “Transparent Handling Tab” on page 342.</p>
Set up the appliance as a router	<ol style="list-style-type: none"> 1. After you have enabled Excelerator for transparent proxy services, check Router Options. 2. Click Router Options. 3. Set up appliance routing. 4. Click Apply. 	<p>See C in Figure 17 on page 49.</p> <p>For more information, see “Router Capabilities” on page 277 and “Router Options Dialog Box” on page 347.</p>
Configure the client workstations to use the appliance as their gateway	<p>Refer to setup instructions for the system.</p> <p>The procedure is different for each platform. On a Windows 95/98/NT workstation, for example, right-click the Network Neighborhood icon on the desktop.</p>	<p>See D in Figure 17 on page 49.</p> <p>Use the transparent proxy IP address as the gateway IP address.</p>

10 Accelerating Web Servers

This section contains information about accelerating Web servers.

Overview of Web Server Acceleration


The appliance's origin Web server accelerator relies on DNS to cause the accelerator to receive requests originally targeted at the origin Web server. The Web server accelerator handles the requests, accessing the origin Web server only when needed objects are not cached.

How Origin Web Server Acceleration Works

The mechanism for routing browser requests to the Web server accelerator instead of the Web server can be summarized as follows:

- ♦ Without acceleration, DNS resolves the origin Web server's DNS name to the origin server's IP address.
- ♦ With acceleration, DNS resolves the server's name to the IP address of an appliance Web server accelerator (reverse proxy) service.

Figure 18



- 1 A browser on the Web requests an origin Web server Web page. This generates a request to DNS for the numeric IP address of the Web server.
- 2 Instead of returning the origin Web server's numeric IP address, DNS returns the numeric IP address of the accelerator service on the appliance.
- 3 The browser requests the Web page using the numeric IP address of the accelerator service.
- 4 The accelerator service obtains the Web page objects from the origin Web server.
- 5 The accelerator returns copies of the objects to the browser.

Benefits of Origin Web Server Acceleration

- ◆ A Web server accelerator reduces response time to browser requests and frees up origin Web server bandwidth, allowing it to handle requests for less frequently requested, uncached data much more quickly.
- ◆ The appliance can accelerate origin Web servers at remote locations that don't offer broadband connections. The Web server accelerator can be located close to the Internet backbone, delivering high-speed access to browsers for all cached objects. The connection to the origin Web server is then used for transporting only those objects not already in cache.

For tips and guidelines on setting up origin Web server accelerators, see “Web Server Accelerator Setup” on page 52.

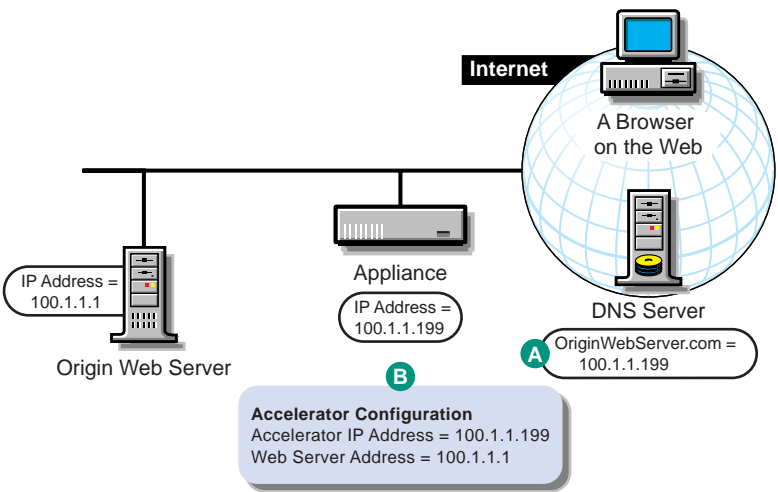
The procedure for configuring DNS to work with Web server accelerators is explained in “Working with DNS” on page 53.

Web Server Accelerator Setup

Figure 19 provides a visual map for the information in this section.

NOTE: The letters in Figure 19 are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 19



To	Do This	Notes
Ensure your basic network configuration is complete for each appliance	1. See “Configuring the Excelerator Appliance” on page 28.	
Ensure that DNS resolves browser requests to the appliance IP addresses configured for the Web server accelerator services	1. See “Working with DNS” on page 53.	See A in Figure 19.

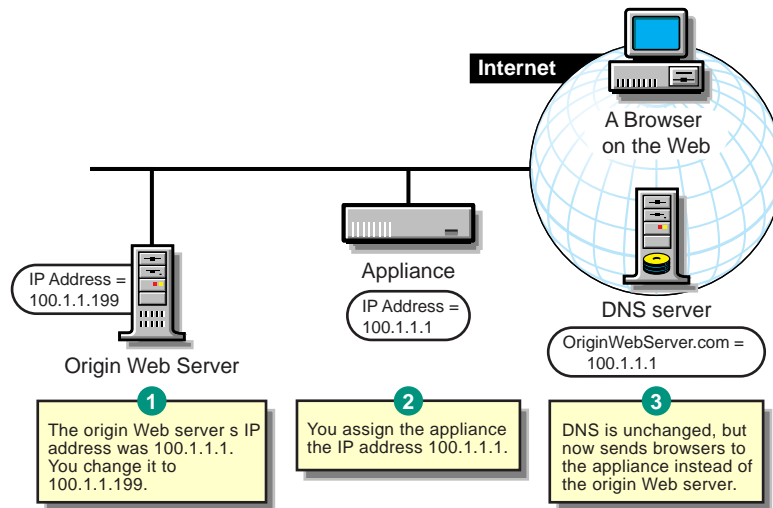
To	Do This	Notes
Set up one or more Web server accelerator services	<ol style="list-style-type: none"> 1. In the browser-based tool, click Cache > Web Server Accelerator > Insert. 2. For your tracking purposes, enter a name for the Web server accelerator. 3. Enter a DNS name. 4. In the Accelerator Proxy Port field, enter the port on which the Web server accelerator will receive requests and vend data. 5. In the Accelerator IP Addresses list, check one or more addresses on which the Web server accelerator will receive requests and vend data. (DNS resolves requests to these addresses.) 6. In the Web Server Port field, enter the port on which the appliance and origin Web server will communicate. 7. In the Web Server Addresses list, insert one or more IP addresses (or DNS names). The Web server accelerator will fill its cache from these addresses or DNS names. (Excelerator must be able to fill all requests through any of these addresses or names.) 8. To activate the Web server accelerator, check Enable This Accelerator. 9. Click OK > Apply. 	<p>See <i>B</i> in Figure 19 on page 52.</p> <p>If server persistence is enabled on the Web Server Accelerator tab, Excelerator will use the same Web server to fill browser requests throughout a session. This setting affects all accelerators on the appliance and saves e-business users from having to log in multiple times. See “Web Server Accelerator Tab” on page 348.</p> <p>If logging is enabled, accelerator log files for the Web server accelerator will have the same name as the Web server accelerator.</p> <p>The DNS name is required when:</p> <ul style="list-style-type: none"> ♦ You are accelerating multiple Web servers on the same IP address. (Multiple accelerator services use the same IP address.) See “Standard Multihoming for Multiple Web Sites” on page 122. ♦ You are accelerating a single Web site using path-based multihoming. See “Multihoming and Path-Based Support” on page 123. ♦ The appliance is part of an ICP hierarchy that needs to resolve relative URLs. <p>If you enter DNS names in the Web Server Addresses list, make sure they are not the names that now resolve to appliance numeric IP addresses. That would create an endless loop.</p>

Working with DNS

The steps you take for having DNS resolve requests to the appliance rather than to the origin server depend on whether the appliance and the origin Web server are on the same subnet.

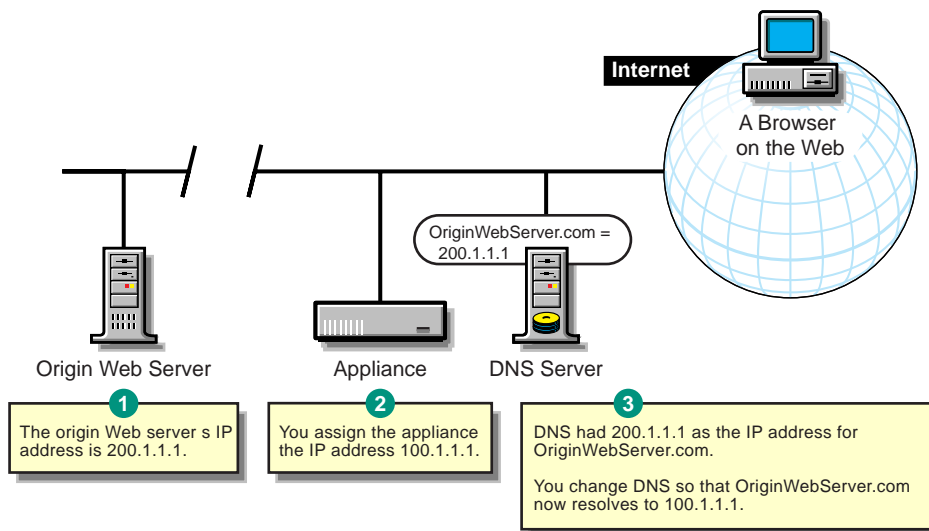
If the appliance and the origin Web server are on the same subnet, you can swap IP addresses as shown in [Figure 20](#).

Figure 20



If the origin Web server is on a remote network, you need to alter DNS as shown in [Figure 21 on page 54](#).

Figure 21



Configuration Considerations When Using Appliance Multihoming Features

Appliance multihome capabilities are explained in [“Standard Multihoming for Multiple Web Sites” on page 122](#) and [“Multihoming and Path-Based Support” on page 123](#). Keep the following points in mind when configuring multihomed support on the appliance:

- ◆ **Support for SSL is restricted.** Although Excelerator allows multiple Web server accelerators to use the same IP address and port combination, this is not supported for Web servers using SSL.

You must ensure that Web servers using SSL are each accelerated using a unique IP address and port combination. Attempts to do otherwise will cause Excelerator to report a TCP bind error.

- ♦ **DNS names must be unique when accelerating multiple sites.** If you are accelerating multiple Web sites on the same IP address, the DNS names in the Web server accelerator definitions must exactly match the DNS names that are used in browser requests. Each accelerator definition must use a unique combination of DNS name, IP address, and port number for Excelerator to properly route browser requests.
- ♦ **DNS names must be the same when accelerating a single Web site.** If you are accelerating multiple Web servers with different IP addresses as a single Web site, the DNS names, accelerator IP addresses, and the accelerator proxy port in the multihoming master and child accelerator definitions must exactly match each other.

The first accelerator must be multihoming master and it must fill from the Web server you want Excelerator to contact for all non-specific data requests. Subsequent child accelerators can then be defined for each server you want Excelerator to contact for specific data requests.

With graphics requests, for example, the multihoming master uses the path-based rules you define to determine which child accelerator to route requests to and, therefore, which Web server to fill requests from.

IV

Accelerating Streaming Media

After you have connected the appliance to your network and completed the other basic setup instructions in [“Integrating Excelerator 2.3 into Your Network” on page 23](#), use the information in this section to learn more about Excelerator streaming media acceleration services and to set up these services on your network.

Excelerator has always supported HTTP streaming wherein a streaming file is delivered through HTTP and played in a browser plug-in.

Starting with the 2.0 release, the appliance includes support for non-proprietary RTSP/RTP streaming media content and tunneling of an RTSP/RTP session inside an HTTP connection as implemented in the QuickTime* and Darwin* products from Apple Computer, Inc.

The following table summarizes the tasks you can accomplish using the information in this section.

To	See
Learn about Excelerator streaming media support	Chapter 11, “Streaming Media Overview,” on page 59
Prepare your network and Excelerator appliance for streaming media acceleration	Chapter 12, “Preparing to Cache Streaming Media Data,” on page 67
Configure your appliance to provide caching services to individual browsers	Chapter 13, “Accelerating Streaming Media to Individual Media Players,” on page 81
Configure your network and appliance to provide streaming media content to all the browsers on a network	Chapter 14, “Accelerating Streaming Media to All Media Players on the Network,” on page 87
Configure your appliance to accelerate streaming media content delivery from your streaming media servers to the Web	Chapter 15, “Accelerating Streaming Media Servers,” on page 95
Configure your appliance as an upstream proxy	Chapter 16, “Configuring an Upstream Proxy for the Appliance,” on page 101
Configure the QuickTime players on your network to use the streaming media caching services you have created	Chapter 17, “Configuring QuickTime Media Players to Use Proxy Services,” on page 103

11

Streaming Media Overview

This chapter is informational. If you are already familiar with streaming media requirements, players, and protocols, you might want to skip to [Chapter 12, “Preparing to Cache Streaming Media Data,”](#) on page 67.

To understand the requirements for delivering streaming content on your network, it is important to consider key differences and similarities between the HTML-based Web content that is displayed in Web browsers and streaming media Web content that is viewed using streaming media players or browser plug-ins.

The following table contrasts HTML-based content with streaming content.

Table 10 **HTML-Based Versus Streaming Content**

Issue	HTML-Based Content	Streaming Content
User Access	Most Web browsers can display most HTML-based content.	A media-specific player is required. For information on configuring the QuickTime player, see Chapter 17, “Configuring QuickTime Media Players to Use Proxy Services,” on page 103.
File Size	Most HTML content files are relatively small.	Most media files are relatively large. This has implications for ensuring you have adequate caching disk space on your appliance and that you have configured the appliance for streaming media objects. For more information, see “Managing Disk Space and Streaming Objects” on page 74.

Issue	HTML-Based Content	Streaming Content
Transmission Rate and Time Constraints	<p>Relatively unimportant</p> <p>If users are patient, they can eventually view the objects.</p>	<p>The transmission rate is critically important.</p> <p>Media data is time-oriented. Each file has a natural duration; each byte in the file has a predefined time at which it will be played relative to the other bytes in the file.</p> <p>Movie frames and the accompanying soundtrack must be delivered smoothly, at the correct time, and in sync with each other. For help determining whether your network can accommodate streaming media requirements, see “Assessing Your Network Bandwidth Capacity” on page 67, “Managing Streaming Bandwidth (Admission Control)” on page 72, and “Managing the Total Sessions Allowed” on page 77.</p>
Downloading	<p>Each object is downloaded as a whole and viewed only after the download is completed.</p>	<p>Objects are viewed as they download.</p> <p>Complete downloads are not required. Users might be interested in only the middle 10 minutes of the file, resulting in only part of the stream being cached. For information on how streaming objects are cached, see “Managing the Caching of Streaming Objects” on page 75.</p>
Time of Creation	<p>HTTP content might have been created in the past, or it might be generated based on the browser request. For example, stock quotes are never stored as objects on an origin Web server but are generated on the fly based on specific browser request parameters.</p>	<p>Streaming objects might be files created in the past, or they might be live transmissions.</p> <p>If you are concerned about freshness of streams in cache, see “Managing Streaming Object Cache Freshness” on page 74.</p>
Coding	<p>There is one version of each object.</p>	<p>There are usually multiple versions of each object, one for each playback bandwidth.</p>

Issue	HTML-Based Content	Streaming Content
Media and User Impacts	After the objects are displayed, the browser and origin Web server do not interact until the user requests new objects.	<p>Displaying objects requires constant interaction between the media player and the streaming media server, including handling user requests for fast forwarding, pausing, etc.</p> <p>For information on how the appliance handles interruptions and other user requests, see “Managing the Caching of Streaming Objects” on page 75.</p>

Streaming Media Protocols

RTSP and RTP are the TCP/IP protocols for streaming media content. RTSP handles all aspects of transport control for the stream and RTP carries the media stream packets. Further information on the RTSP/RTP protocol is available on the [Web \(http://www.cs.columbia.edu/~hgs/rtsp/\)](http://www.cs.columbia.edu/~hgs/rtsp/).

The QuickTime* streaming media product from Apple Computer, Inc. conforms exactly with the RTSP/RTP standard. The other major vendors of products that create streaming media generally either include proprietary extensions in their RTSP/RTP packets, or they use their own proprietary protocols for streaming media content delivery.

The base Excelerator appliance supports caching of QuickTime streaming objects and the session management functionality required by RTSP/RTP-compatible media players.

The following sections contain tables that list the streaming media protocols associated with streaming products from the major vendors and summarize the advantages and disadvantages associated with each protocol.

QuickTime

The following table summarizes the protocols supported by the QuickTime Player* from Apple Computer, Inc., the advantages and disadvantages associated with each protocol, and Excelerator support of the protocol.

Table 11 QuickTime Protocol Support

Name	Protocol Details	Advantages	Disadvantages	Base Appliance Streaming Support
HTTP Streaming	<p>HTTP in TCP: Media data only (no control data).</p> <p>Transparent and reverse services use 80 as their standard port.</p> <p>Forward services most commonly use port 8080.</p> <p>TCP delivers the data as quickly as possible. If the buffer slider stays ahead of the player slider, the stream plays normally. If the player reaches the buffer, the stream pauses until adequate content buffering is restored.</p>	<ul style="list-style-type: none"> ♦ Uses the Web's most standard protocol. ♦ Setup is easy. ♦ Passes through firewalls with other Web data. 	<ul style="list-style-type: none"> ♦ No advanced functionality. Players cannot fast forward, pause, etc. ♦ Workstation resources must accommodate buffering of incoming content. 	Version 1.x and later
RTSP/RTP	<p>RTSP in TCP: Control data.</p> <p>RTP in UDP: Media data.</p> <p>Transparent and reverse services use 554 as their standard port.</p> <p>Forward services most commonly use port 9090.</p> <p>Data is delivered in real time exactly when needed</p>	<ul style="list-style-type: none"> ♦ Supports real-time delivery of media data. ♦ Lets players fast forward, pause, etc. 	<ul style="list-style-type: none"> ♦ Hard to get through firewalls. ♦ Unreliable delivery of media data. 	Version 2.0

Name	Protocol Details	Advantages	Disadvantages	Base Appliance Streaming Support
HTTP Tunneling	<p>RTP or RTSP in HTTP in TCP: Both control and media data.</p> <p>Transparent and reverse services use 80 as their standard port.</p> <p>Forward services most commonly use port 8080.</p> <p>Data is delivered in real time exactly when needed.</p>	<ul style="list-style-type: none"> Supports real-time delivery of media data. Lets players fast forward, pause, etc. Provides reliable delivery of media data. Passes through most firewalls. 	Subject to TCP congestion control, which might not provide adequate bandwidth for the stream.	Version 2.0

Windows* Media Player*

The following table summarizes the protocols supported by the Windows Media Player and the advantages and disadvantages associated with each protocol.

Table 12 Windows Media Player Protocol Support

Name	Protocol Details	Advantages	Disadvantages	Base Appliance Streaming Support
HTTP Streaming	<p>HTTP in TCP: Media data only (no control data).</p> <p>Transparent and reverse services use 80 as their standard port.</p> <p>Forward services most commonly use port 8080.</p> <p>TCP delivers the data as quickly as possible. If the buffer slider stays ahead of the player slider, the stream plays normally. If the player reaches the buffer, the stream pauses until adequate content buffering is restored.</p>	<ul style="list-style-type: none"> Uses the Web's most standard protocol. Setup is easy. Passes through firewalls with other Web data. 	<ul style="list-style-type: none"> No advanced functionality. Players cannot fast forward, pause, etc. Workstation resources must accommodate buffering of incoming content. 	Version 1.x and later

Name	Protocol Details	Advantages	Disadvantages	Base Appliance Streaming Support
MMS in UDP	<p>MMS in TCP: Control data.</p> <p>MMS in UDP: Media data.</p> <p>Transparent, reverse, and forward services use 1755 as their standard port.</p>	<ul style="list-style-type: none"> ♦ UDP provides the most efficient network throughput. ♦ Supports real-time delivery of media data. ♦ Lets players fast-forward, pause, etc. 	<ul style="list-style-type: none"> ♦ Many network administrators close their firewalls to UDP traffic, limiting the potential audience of UDP-based streams. ♦ UDP packet delivery is unreliable. 	
MMS in TCP	<p>MMS in TCP: Both control and media data.</p> <p>Transparent, reverse, and forward services use 1755 as their standard port.</p>	<ul style="list-style-type: none"> ♦ Guaranteed delivery of media delivery. ♦ Supports real-time delivery of media data. ♦ Lets players fast forward, pause, etc. 	<ul style="list-style-type: none"> ♦ Might be blocked by firewall. ♦ Slight protocol overhead for guaranteed delivery. 	
MMS in HTTP	<p>MMS in HTTP in TCP: Both control and media data.</p> <p>Transparent and reverse services use 80 as their standard port.</p> <p>Forward services most commonly use port 8080.</p>	<ul style="list-style-type: none"> ♦ Uses the Web's most standard protocol. ♦ Setup is easy. ♦ Guaranteed delivery of media data. ♦ Supports real-time delivery of media data. ♦ Lets players fast-forward, pause, etc. ♦ Passes through firewalls with other Web data. 	<ul style="list-style-type: none"> ♦ Slight protocol overhead for guaranteed delivery. 	Through an add-on

Real

The following table summarizes the protocols supported by Real* and the advantages and disadvantages associated with each protocol.

Table 13 Real Protocol Support

Name	Protocol Details	Advantages	Disadvantages	Base Appliance Streaming Support
HTTP Streaming	<p>HTTP in TCP: Media data only (no control data).</p> <p>Transparent and reverse services use 80 as their standard port.</p> <p>Forward services most commonly use port 8080.</p> <p>TCP delivers the data as quickly as possible. If the buffer slider stays ahead of the player slider, the stream plays normally. If the player reaches the buffer, the stream pauses until adequate content buffering is restored.</p>	<ul style="list-style-type: none"> ♦ Uses the Web's most standard protocol. ♦ Setup is easy. ♦ Passes through firewalls with other Web data. 	<ul style="list-style-type: none"> ♦ No advanced functionality. Players cannot fast forward, pause, etc. ♦ Workstation resources must accommodate buffering of incoming content. 	Version 1.x and later
HTTP Cloaking	<p>RDT in HTTP in TCP: Media and control data.</p> <p>Transparent and reverse services use 80 as their standard port.</p> <p>Forward services most commonly use port 8080.</p>	<ul style="list-style-type: none"> ♦ Supports real-time delivery of media data. ♦ Lets players fast forward, pause, etc. ♦ Provides reliable delivery of media data. ♦ Passes through most firewalls. 	Subject to TCP congestion control, which might not provide adequate bandwidth for the stream.	

Name	Protocol Details	Advantages	Disadvantages	Base Appliance Streaming Support
RTSP/RDT	<p>RDT in UDP: Media data.</p> <p>RTSP in TCP: Control data.</p> <p>Transparent and reverse services use 554 as their standard port.</p> <p>Forward services most commonly use port 9090.</p> <p>Data is delivered in real time exactly when needed.</p>	<ul style="list-style-type: none"> Supports real-time delivery of media data. Lets players fast forward, pause, etc. 	<ul style="list-style-type: none"> Hard to get through firewalls. Unreliable delivery of media data. 	
RTSP/RTP	<p>RTP in UDP: Media data.</p> <p>RTSP in TCP: Control data.</p> <p>Transparent and reverse services use 554 as their standard port.</p> <p>Forward services most commonly use port 9090.</p> <p>Data is delivered in real time exactly when needed.</p>	<ul style="list-style-type: none"> Supports real-time delivery of media data. Lets players fast forward, pause, etc. 	<ul style="list-style-type: none"> Hard to get through firewalls. Unreliable delivery of media data. 	

12

Preparing to Cache Streaming Media Data

This chapter contains information on assessing and preparing your network and appliance to handle streaming media.

Excelerator 2.3 supports caching of QuickTime streaming content by default.

Windows Media and RealMedia require that you purchase add-on products.

Instructions for creating caching services for QuickTime streaming content begin with [Chapter 13, “Accelerating Streaming Media to Individual Media Players,”](#) on page 81.

Instructions for creating Windows Media and RealMedia caching services are contained in *Volera Media Excelerator 1.2 for Windows Media Administration Guide* and *Volera Media Excelerator 1.2 for RealSystem Proxy 8 Startup Guide*, respectively.

Identifying Your Streaming Media Types

To plan and create streaming media caching services on your network, you must know the types of streaming media that network users are accessing and understand which of these types the base appliance can cache.

An overview of the standard media players, the protocols they handle, and Excelerator support for these protocols is found in [“Streaming Media Protocols”](#) on page 61.

Assessing Your Network Bandwidth Capacity

As noted in [Table 10, “HTML-Based Versus Streaming Content,”](#) on page 59, streaming media requires considerably more network bandwidth (both capacity and speed) than HTML-based traffic.

As you assess your network bandwidth capacity, the decisions you make regarding upgrading or modifying your network will depend largely on which of the following describes your main focus:

- ♦ If you are accelerating content to streaming media players in an enterprise environment, see [“Assessing Bandwidth for Forward and/or Transparent Streaming Proxy Services”](#) on page 68.
- ♦ If you are accelerating content from streaming servers to players on the Web, see [“Assessing Bandwidth for a Streaming Server Accelerator”](#) on page 70.

Upstream Versus Downstream Bandwidth

As you read the following explanations and configure your appliance, you will want to clearly understand the distinction between upstream bandwidth and downstream bandwidth.

One simple way to remember the difference between upstream and downstream is to think of the media stream as being similar to a stream of water. The media comes toward the appliance from upstream (the origin streaming server), passes through the appliance, and flows downstream to the players that have requested it.

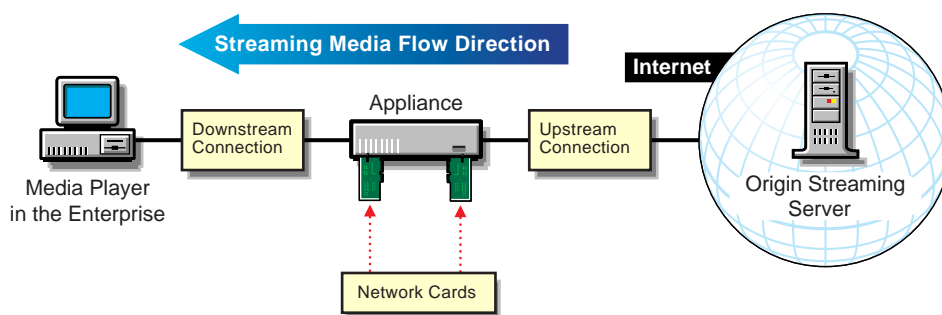
The players and the origin streaming server might be on your local network, or they might be on the Internet. Where they are doesn't matter. Upstream and downstream refer to the direction of the stream, which is always from the server to the player. The location of these components on the network is irrelevant.

Assessing Bandwidth for Forward and/or Transparent Streaming Proxy Services

Enterprises often cache streaming content from the Web and deliver it to players within the enterprise, thus reducing their Internet connection overhead.

Figure 22 on page 68 shows a simplified view of the components that affect bandwidth in enterprise deployments of streaming media caching.

Figure 22



NOTE: The placement of network components in the graphic is not literal. For example, it does not imply that the appliance is acting as an inline router on the enterprise backbone.

Assessing the Downstream Connection (Enterprise Network Backbone)

To assess the capacity of your network's downstream connection to deliver streaming content to players, complete the following steps:

- 1 Determine the bandwidth capacity of your network backbone.

To ensure that the players on your network can effectively view cached streams, the enterprise's network backbone must be capable of handling multiple streaming connections.

The recommended minimum capacity is 100 Mbps.

- 2 Determine the bandwidth of connections to players on the network.

Connections to players must also be capable of delivering a cached stream in addition to other standard network traffic.

The minimum is 10 Mbps; 100 Mbps is recommended.

- 3 Determine how much bandwidth you want consumed by streaming content.

It is wise to limit the amount of network bandwidth that can be consumed by cached streaming content. A good starting point is 50 percent of total available bandwidth.

For example, if you have a 100 Mb backbone (100,000 Kbps), you might want to start by setting the Max Downstream Bandwidth parameter in the **Streaming Management Configuration Dialog Box** to 50,000 Kbps (50 Megabits per second). You can then adjust this parameter up or down depending on your network requirements.

- 4** (Optional) Calculate how many streams the appliance is allowed to send simultaneously on your network.

For example, if you have limited the downstream bandwidth to 50,000 Kbps and each stream consumes 128 Kbps (which is standard for most streaming content on the Web today), the appliance will send up to 390 simultaneous streams on your network ($50,000 / 128 = 390.625$).

For more information on setting parameters after your bandwidth evaluation, see **“Configuring Exceleator for Streaming Bandwidth Management” on page 73**.

Assessing the Upstream Connection (Internet Connection)

To assess the capacity of your network’s Internet connection to retrieve streaming content for caching, complete the following steps:

- 1** Determine the bandwidth capacity of your Internet connection.

To process initial player requests effectively, the connection must be able to support real time interaction between players and origin streaming servers on the Web.

- 2** Determine how much Internet bandwidth you want consumed by initial requests for streaming content. A good starting point might be 25 percent of total available bandwidth.

For example, if you have a 10 Mb (10,000 Kbps) Internet connection and you are allowing up to 25 percent for retrieving streaming objects, you would set the Max Upstream Bandwidth parameter in the **Streaming Management Configuration Dialog Box** to 2500 Kbps (2.5 Mb per second).

- 3** Weigh your available Internet streaming bandwidth (determined in **Step 2**) against the initial stream caching requirements for your network and plan to preload streams into cache if required.

Just as business connections to the Internet vary widely from T-1 connections or less to multiple T-3 connections and more, the requirements for initial stream caching also vary according to business requirements.

For example, you might need to make specific streams available to multiple network users but have limited upstream bandwidth for initial stream requests. If all your network users arrive at work at 8 a.m. and begin requesting the same streams from the Internet at about the same time, your Internet connection won’t be able to handle the load, users will become extremely frustrated, and productivity will decrease.

To prevent this from occurring, you might consider downloading streams to your appliance during off hours while your Internet connection is uncongested.

NOTE: The Download tab does not currently support stream downloads. You must therefore use player requests to preload the required streams.

On the other hand, if you have ample upstream bandwidth, you can probably afford to let the players on your network request streams as they are needed.

For information on setting parameters after your bandwidth evaluation, see **“Configuring Exceleator for Streaming Bandwidth Management” on page 73**.

Assessing the Capacity of the Appliance's Network Cards

The appliance's network cards are the third critical component to consider in your network bandwidth assessment.

Setting upstream and downstream bandwidth limits will help avoid overloading the appliance's network cards.

If your appliance has multiple network cards, you can tune the flow of streaming traffic by devoting more network cards to delivering cached content to players than to retrieving content from the Internet. You can do this by assigning the cards IP addresses on specific subnets on your network.

For example, if your appliance has three network cards, you could assign two cards IP addresses on subnets on the downstream portion of your network backbone and one card an IP address on the upstream connection to the Internet.

IMPORTANT: The appliance doesn't provide load balancing across network cards. The load on each card is determined by which services use the IP addresses bound to the card.

If your appliance has multiple network cards and you have estimated the total downstream or upstream capacity based on the combined capacities of multiple cards, you should track which IP addresses are assigned to which services and estimate the load per IP address.

Next you should add together the loads from all IP addresses on each network card. This will tell you the estimated load for the card and let you judge whether the card can handle the load assigned to it.

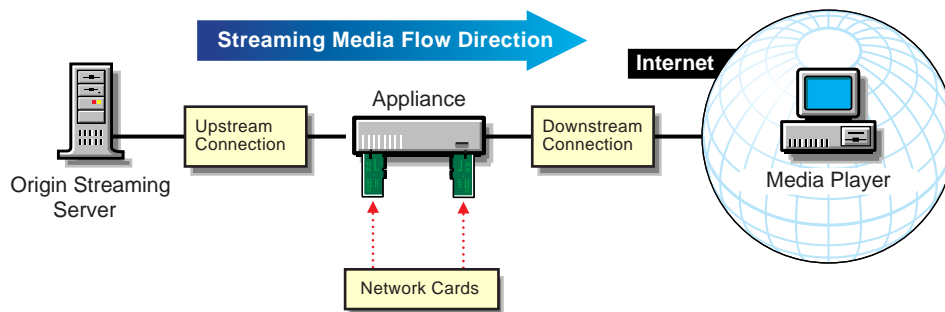
If you don't perform this check, your configuration might result in one of the network cards handling a disproportionate amount of traffic. This could cause the card to fail and the system to crash.

Assessing Bandwidth for a Streaming Server Accelerator

Content providers often cache streaming content prior to delivery to players on the Web, thus off loading request processing from their streaming servers and dramatically increasing their content delivery capacity.

Figure 23 shows a simplified view of the three main components that affect bandwidth in content provider deployments of streaming media caching.

Figure 23



NOTE: The placement of network components in the graphic is not literal. For example, it does not imply that the appliance is acting as an inline router.

Assessing the Upstream Connection (Your Network's Backbone)

To assess the capacity of your network's connection to streaming servers, complete the following steps:

1 Determine the bandwidth capacity of your network backbone.

For initial player requests to be cached properly, the backbone must support real-time interaction with players through the caching appliance as well as through caching of the stream on the appliance.

The recommended minimum capacity is 100 Mbps.

2 Determine how much network bandwidth you want consumed by stream caching.

As you decide on a percentage, consider what other network activities consume backbone bandwidth.

It might also be useful to consider how many streams you can simultaneously cache based on given percentages.

For example, if you have a 100 Mb backbone and delivering streaming content is a main focus of your business, you might want to start by devoting 50 percent of network bandwidth to handling streaming requests. To do this, you would set the Max Upstream Bandwidth parameter in the **Streaming Management Configuration Dialog Box** to 50,000 Kbps (50 Mb per second), meaning that caching streams on your appliance can consume up to 50 percent of your total network's bandwidth.

In this example, if your streams each consume 128 Kbps, your appliance will accept up to 390 different streams for caching at the same time ($50,000 / 128 = 390.625$). This is probably more streams than most Web sites would need to cache simultaneously. Keep in mind, however, that setting the appliance's bandwidth parameters only affects the bandwidth available to other network traffic when streams are actually being cached.

For information on setting parameters after your bandwidth evaluation, see **“Configuring Exceleator for Streaming Bandwidth Management” on page 73**.

Assessing the Downstream Connection (Players on the Internet)

To assess the capacity of your Internet connection to deliver cached streaming content to requesting players, complete the following steps:

1 Determine the bandwidth capacity of your Internet connection.

In order for players on the Web to effectively view cached streams, the Internet connection must be capable of handling multiple streaming connections.

The minimum recommended bandwidth is 10 Mbps. Large content providers will require many times this amount.

2 Determine how much of your available Internet bandwidth you want consumed by streaming content.

Although you will want to limit the amount of bandwidth that can be consumed by cached streaming content, you will want to provide enough bandwidth to avoid network congestion and ensure the highest possible transmission quality.

A good starting point might be 80 percent of total available Internet bandwidth. This would mean that, for a 10 Mbps connection, you would set the Max Downstream Bandwidth parameter in the **Streaming Management Configuration Dialog Box** to 8000 Kbps (8 Mbps). If your streams each consume 128 Kbps, you could service up to 62 stream requests simultaneously ($8000 / 128 = 62.5$).

By the same token, if you have T-3 (45 Mbps) connections to the Internet and you want 80 percent of total bandwidth available for streaming, you would set the Max Downstream Bandwidth to 36,000 Kbps for each T-3 connection. ($45,000 \times 8 = 36,000$) Each T-3

connection could then service up to 281 requests for 128 Kbps streams. (36,000 / 128 = 281.25)

For information on setting parameters after your bandwidth evaluation, see [“Configuring Excelsator for Streaming Bandwidth Management” on page 73](#).

Assessing the Capacity of the Appliance’s Network Cards

The appliance’s network cards are the third critical component to consider in your network bandwidth assessment.

Setting upstream and downstream bandwidth limits will help avoid overloading the appliance’s network cards.

If your appliance has multiple network cards, you can tune the flow of streaming traffic by devoting more network cards to delivering cached content to players on the Internet than to retrieving content from your streaming servers. You can do this by assigning the cards IP addresses on specific subnets on your network.

For example, if your appliance has three network cards, you could assign one card an IP address on a subnet of your upstream network backbone, and two cards IP addresses on the downstream connection to the Internet.

IMPORTANT: The appliance doesn’t provide load balancing across network cards. The load on each card is determined by which services use the IP addresses bound to the card.

If your appliance has multiple network cards and you have estimated the total downstream or upstream capacity based on the combined capacities of multiple cards, you should track which IP addresses are assigned to which services and estimate the load per IP address.

Next, you should add together the loads from all IP addresses on each network card. This will tell you the estimated load for the card and let you judge whether the card can handle the load assigned to it.

If you don’t perform this check, your configuration might result in one of the network cards handling a disproportionate amount of traffic. This could cause the card to fail and the system to crash.

Configuring the Appliance to Match Your Requirements

The [Streaming Management Configuration Dialog Box](#) contains all the parameters you can use to configure the appliance to match your network capacity and streaming content delivery requirements. The following sections explain various aspects of appliance management and tuning. For reference information on the fields in the Policy Management dialog box, see [“Streaming Management Configuration Dialog Box” on page 358](#).

Managing Streaming Bandwidth (Admission Control)

To provide streaming media support, you must ensure that your network bandwidth can handle the increased load that streaming media requires.

Because the Excelsator appliance can deliver extremely high volumes of data to the network, its delivery of cached streaming content on your network can easily exceed the capacity of your network hardware, including the appliance’s network cards.

When this happens, the resulting backlog can result in lost packets and other more severe problems, including the Excelsator system crashing.

This section explains the controls built into the appliance for managing the streaming load on your network bandwidth, and it provides basic suggestions for ensuring optimal tuning of your appliance's ability to provide streaming content to network users.

How the Appliance Manages Bandwidth

Each time a media player sets up an RTSP/RTP streaming connection (either native RTSP/RTP or tunneled RTSP/RTP), the player must negotiate with the origin streaming server or Excelsior appliance providing the stream. One item of negotiation is reserving the average bandwidth expected to be required during the time the stream is playing.

When the appliance gets a request for an RTSP/RTP or tunneled streaming media connection, it must decide whether to accept (admit) the connection or reject it. The key to this decision is determining whether there is enough network bandwidth available to accommodate the request.

Excelsior uses the values specified for the following parameters to make bandwidth-based admission control decisions:

- ♦ **Max Bandwidth Per Stream:** This parameter sets a limit on the network bandwidth an individual stream can consume.
- ♦ **Max Upstream Bandwidth:** This parameter limits the network bandwidth that connections with origin streaming servers can consume.
- ♦ **Max Downstream Bandwidth:** This parameter limits the network bandwidth that connections with media players can consume.

For more information on these parameters, see [“Streaming Management Configuration Dialog Box” on page 358](#).

As new connections are admitted, the bandwidth they are estimated to consume is deducted from the applicable bandwidth limits configured for the system.

If a requested stream would cause one of the configured bandwidth limits to be exceeded, the connection is denied and the requesting player receives a 453 Not Enough Bandwidth error.

Configuring Excelsior for Streaming Bandwidth Management

By default, the maximum bandwidth parameters are set to take advantage of unlimited network bandwidth. However, every network has upper bandwidth limitations. You should, therefore, generally set all three parameters to ensure optimal delivery of cached streaming content.

IMPORTANT: As explained in [“Managing Streaming Bandwidth \(Admission Control\)” on page 72](#), you must, at the very least, set the Max Downstream Bandwidth parameter to match the capacity of your network hardware, including the appliance's network adapters.

Based on the decisions you made in [“Assessing Your Network Bandwidth Capacity” on page 67](#), set your appliance's streaming bandwidth parameters by completing the following steps.

- 1** In the browser-based management tool, click Cache > Media Cache > Policy Management Options.
- 2** In the Max Bandwidth Per Stream field, type the bandwidth of the highest-coded bandwidth you want players to be able to access.

The most commonly used bandwidth on the Internet is 128 Kbps.

Other common values are 256 Kbps, 1024 Kbps (1 Mb), and 1536 Kbps (1.5 Mb).

- 3** In the Max Upstream Bandwidth field, type the upstream bandwidth limit that you determined in [“Assessing Your Network Bandwidth Capacity” on page 67](#).

If you are a content provider, you might be reluctant to change the Maximum Upstream Bandwidth parameter from its default (unlimited) setting because you want to cache all the content you have created as quickly as possible. However, in most situations you will get optimal results by limiting the upstream bandwidth to 80% of your network's upstream capacity. This will ensure that the network remains in an uncongested state and maintains the highest transmission quality.

4 In the Max Downstream Bandwidth field, type the downstream bandwidth limit you determined in [“Assessing Your Network Bandwidth Capacity” on page 67](#).

5 Click OK > Apply.

NOTE: If you are closing the browser-based tool for some reason, make sure you click Apply before exiting. Otherwise, your bandwidth parameter settings will be lost.

Managing Disk Space and Streaming Objects

Excelerator uses the following parameters in the [Streaming Management Configuration Dialog Box](#) to limit streaming object size and the amount of disk space consumed by streaming objects:

- ♦ Max Object Size
- ♦ Max Object Duration
- ♦ Max Disk Usage

Normally it is not necessary to adjust the first two object-related parameters because the system calculates and enforces a maximum object size that is approximately one-fourth of the appliance's smallest hard disk size.

You cannot set an object size or duration value greater than the system-calculated size. And unless you want to arbitrarily limit the size of streaming objects that can be cached, there is no reason to set a lesser value than calculated by the system.

The Max Disk Usage parameter should usually be set at about 50 percent of caching disk space. To do this, complete the following steps:

1 In the browser-based management tool, click Monitoring > Summary.

2 Write down the Cache Disk Space field value in megabytes.

3 Calculate 50 percent of the cache disk space in megabytes and convert this to gigabytes by dividing by 1,000.

4 Click Cache > Media Cache.

5 In the Max Disk Usage field, enter the result obtained in [Step 3](#).

IMPORTANT: Remember to convert the Summary tab value from megabytes to gigabytes.

6 Click OK > Apply.

Managing Streaming Object Cache Freshness

The following fields in the [Streaming Management Configuration Dialog Box](#) control the freshness of streaming objects in cache:

- ♦ Max TTL
- ♦ Min TTL
- ♦ Default TTL

The discussion on cache freshness found in [Chapter 26, “Cache Freshness,” on page 185](#) generally applies to streaming media cache freshness.

As you read that discussion, keep the following points in mind:

- ♦ Some streaming servers set the object expiration time to 0 by default.

Because Excelerator honors all object header values, streams with an expiration time of 0 are not cached unless the Min TTL value is set to a larger value to force caching. By default, the Min TTL value is set to 3600 seconds (one hour).

The minimum value you set should be no less than the playing time of the longest streaming object you plan to cache. Otherwise, the object will be removed from cache before it is finished playing.

In most situations you will want to set the Min TTL value even higher, as explained in the next point.

- ♦ Streaming objects tend not to be particularly dynamic.

For example, movie trailers don’t change very often once they are posted.

Because of this, you can probably set the Min and Max TTL values higher than you set the corresponding values for HTTP objects.


Managing the Caching of Streaming Objects

When a player requests a streaming media object, the appliance caches the media stream exactly as it comes to the player. Because of fast forward requests and other user actions, only part of the stream may be cached.

How Streaming Objects Are Cached

The following graphic summarizes how Excelerator caches streaming content when interruptions occur and fills subsequent requests for the same object using the cached portions of the object.

Figure 24



- 1 A player requests a stream from the appliance.
- 2 The service begins to fill the stream from the origin streaming server.
- 3 The service transmits the stream in real time to the player.
- 4 The player executes various fast forward/rewind actions.
- 5 The appliance handles these requests either from cache or by initiating new fill requests as necessary.
- 6 The player closes the connection prior to the appliance fully caching the stream.
- 7 The appliance continues to fill the stream based on the Continue Fill Time parameter. When the time expires, if no other players are requesting the same stream, the appliance stops the fill process.
- 8 The appliance handles subsequent requests for the same stream using the portions that have been cached and initiating new fill requests to the origin streaming server as required.

Ensuring that Stream Filling Continues

You cannot control whether a user fast forwards through a streaming object while the object is being cached. Such control is not necessary because Excelerator automatically fills gaps in the streaming object during subsequent requests.

However, you can control what happens when a user terminates a streaming request before the stream has been fully cached.

Continued filling of streaming objects is controlled by the Continue Streaming Fill Time option in the **Streaming Management Configuration Dialog Box**. By default this parameter is set for 0 minutes, which means that caching of the stream ceases immediately.

You can have Excelerator continue to fill streaming objects for up to 10 hours (600 minutes) after requests are terminated. If, during that time another user requests the same stream, the connection with the origin streaming server will already be established and caching will continue uninterrupted. This will provide the highest possible quality for the playback of the cache stream.

As you set the Continue Streaming Fill Time option, consider the following general guidelines:

- ♦ If you are accelerating a streaming server as a content provider, you should use a relatively high setting.

The field value should be no less than the playing time of the longest object on the streaming server.

- ♦ If you are setting up forward and/or transparent streaming proxy services in an enterprise, you will probably want to set a relatively low fill time to avoid wasting bandwidth on streaming objects that might not be requested again.

Managing the Total Sessions Allowed


The Max Sessions parameter in the **Streaming Management Configuration Dialog Box** provides an alternate way of managing bandwidth.

For most streaming cache scenarios, the Maximum Bandwidth parameters provide acceptable bandwidth management control. However, ISPs who host dial-up Web connections might find the Max Sessions parameter more useful in managing network resources.

Getting QuickTime Streaming Content Through Firewalls

Most firewalls block UDP/RTSP traffic, as shown in **Figure 25**.

Figure 25 UDP/RTSP Traffic Is Blocked by Most Firewalls



HTTP-tunneled requests, on the other hand, pass through most firewalls as shown in **Figure 26**.

Figure 26 HTTP-Tunneled Traffic Passes Through Most Firewalls




Figure 26 illustrates why HTTP tunneling is widely used for QuickTime streaming on the Web. The HTTP packets pass through most firewalls, and the RTSP data embedded in the packets lets the QuickTime players communicate with the streaming servers.

The Appliance and HTTP Tunneling

Exceleator can be configured to either process and actively fill HTTP tunneled requests or to simply pass them through to the origin streaming servers. **Table 14** summarizes the differences and trade-offs of having HTTP tunnel support enabled and disabled.

Table 14 QuickTime HTTP Tunnel Options

	QT HTTP Tunnel Support Enabled	QT HTTP Tunnel Support Not Enabled (HTTP pass through)
How player requests are handled	The appliance acts as a proxy for the player.	The appliance simply forwards player requests.
Which protocols are used	UDP/RTSP	HTTP
What a firewall does	The firewall blocks communication with the server.	Most firewalls let packets through
Is the stream cached?	Yes	No

Passing Streaming Content Through the Appliance Without Caching It

If you don't want to cache QuickTime HTTP tunneled streams, appliance setup is very simple. You must only ensure that the QuickTime HTTP Tunnel Enable option is not checked in the **Streaming Tab**.

After that, any HTTP tunneled requests that come through a forward and/or transparent streaming proxy service defined on the appliance will be passed through to the origin streaming server and the returned streaming data will not be cached.

Because HTTP-tunneled traffic generally passes through firewalls, the location of the appliance in relation to the players is not a critical issue.

On the other hand, if you want to cache QuickTime HTTP tunneled streams, you have several options for overcoming firewall limitations.


Setting Up Your Appliance to Work with Firewalls

This section outlines three simplified configuration scenarios for getting QuickTime streaming content through firewalls. The first scenario works only for HTTP tunneled player requests. The other two scenarios work for both HTTP tunneled requests and RTSP requests.

Placing the Appliance Outside the Firewall

If you place the appliance outside the firewall, the QuickTime players can use HTTP tunneled requests to go through the firewall to the appliance. The appliance can then use UDP/RTSP to communicate with the origin streaming server, as shown in **Figure 27**.

Figure 27



In this scenario, all appliance IP addresses, the default gateway, and the DNS server would be on subnets outside the firewall.

NOTE: This approach does not apply to players configured to use the UDP/RTSP transport option.

Going Through an RTSP Proxy on the Firewall Network Address Translator (NAT)


If you place the appliance inside a firewall or behind a network address translator, you will need a way to get the UDP/RTSP packets through the firewall or translator to the origin streaming server.

Some firewalls and translators have an RTSP proxy feature that is designed for this purpose.

If your firewall or translator has an RTSP proxy feature, you can configure the IP address and port number of the RTSP proxy service on the firewall or translator as an upstream proxy to the appliance. (For more information, see [“Configuring an Upstream Proxy for the Appliance” on page 101.](#))

The appliance can then use the UDP and RTSP protocols to communicate with the origin streaming server through the firewall or translator, as shown in [Figure 28 on page 79.](#)

Figure 28



In this scenario, the appliance IP addresses, the default gateway, the DNS server, and the RTSP proxy service on the firewall or network address translator would all be on subnets inside the firewall.


Using the Appliance as a Component in the Firewall

If neither of the previous two options is possible for your network, you can configure the appliance as a firewall component by doing the following:

- ♦ Using IP addresses on network cards that are on subnets inside the firewall to receive forward and/or transparent streaming requests
- ♦ Using IP addresses on other network cards that are outside the firewall to communicate with origin streaming servers on the Web

[Figure 29](#) illustrates this scenario.

Figure 29



In this scenario, the appliance IP addresses would be both inside and outside the firewall, depending on the network cards to which they were assigned. The default gateway and the DNS server would be on subnets outside the firewall.

Logging Streaming Media Transactions

Logging of streaming caching activity can be useful for a number of reasons, such as billing for services rendered. Excelsior lets you specify how often a new log file will be started (rolled over), how long old log files will be retained, and how the log files will be formatted.

Your appliance offers the following streaming logging services:

- ◆ You can turn on logging for each forward, transparent, and reverse streaming proxy service you create.
- ◆ You can have the appliance automatically download streaming log files to an FTP server and automatically delete downloaded files.
- ◆ You can control the deleting of old log files based on an older-than-*x* time period or the number of log files in the system.
- ◆ You can customize what transaction information is logged for each streaming service.

For a general discussion of logging issues and specific suggestions on developing a logging strategy, see [Chapter 34, “Logging,” on page 237](#).

For information on streaming-specific logging options, see [“Streaming Media Log Options Dialog Box \(MMS\)” on page 369](#).

13

Accelerating Streaming Media to Individual Media Players

This chapter contains instructions for creating QuickTime caching services.

Instructions for creating Windows Media and RealMedia caching services are contained in *Volera Media Excelerator 1.2 for Windows Media Administration Guide* and *Volera Media Excelerator 1.2 for RealSystem Proxy 8 Startup Guide*, respectively.

Overview of Forward Streaming Proxy

This section presents a conceptual overview of Excelerator streaming forward proxy services.

Setup instructions are in “[Forward Streaming Proxy Setup](#)” on page 82.


Key Functionality

You can configure RTSP/RTP-compliant media players with the IP address of a streaming forward proxy service you have created on the appliance.

After the players are properly configured, they send streaming requests directly to an appliance IP address configured for streaming forward proxy services. The streaming forward proxy service obtains the media streams, caches the streams, and sends copies of the streams back to the players in the same way the players would receive the streams from origin streaming servers.

How Forward Streaming Proxy Works

Figure 30



- 1 A player requests a stream from its forward proxy server (the appliance).
- 2 The forward proxy service obtains the numeric IP from DNS.
- 3 The service begins to fill the stream from the origin streaming server.
- 4 The service transmits the stream in real time to the player.
- 5 The streaming forward proxy service handles subsequent requests for the same stream without accessing DNS or the origin streaming server.

Benefits of Forward Streaming Proxy

- ◆ Doesn't require a special router configuration
- ◆ Provides an immediate improvement in player performance
- ◆ Lets users decide whether to use the proxy service

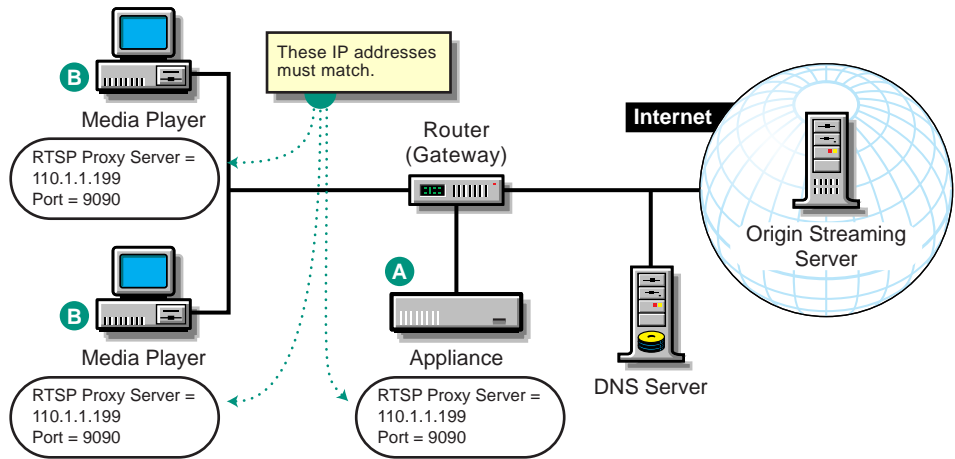
For tips and guidelines on setting up forward proxy services, see [“Forward Streaming Proxy Setup” on page 82](#).

Forward Streaming Proxy Setup

[Figure 31](#) provides a visual map for the information in this section.

NOTE: The letters in [Figure 31](#) are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 31



Set up forward proxy services as follows:

To	Do This	Notes
Ensure your basic network configuration is complete	1. See “Basic Network Configuration Setup” on page 27.	
Set up the appliance to work around firewall limitations	For guidelines and more information, see “Getting QuickTime Streaming Content Through Firewalls” on page 77.	The choices you have depend on whether the players on your network will use HTTP tunneling or UDP, RTSP.
Enable streaming forward proxy services on the appliance	<div>1. In the browser-based tool, click Cache > Media Cache > Forward > Enable.</div> <div>2. In the Name field, type a name for the forward service.</div> <div>3. Ensure the port is set to 9090.</div> <div>4. Check the appliance IP address for the service.</div> <div>5. Click Apply.</div>	<div>See A in Figure 31 on page 83.</div> <div>IMPORTANT: Enabling streaming forward proxy services is required even if HTTP tunneling is employed.</div> <div>For more information, see “Forward Streaming Services (RTSP)” on page 361.</div>

To	Do This	Notes
(Optional) Enable HTTP tunneling support	<ol style="list-style-type: none"> 1. In the browser-based tool, click Cache > click Media Cache > check QuickTime HTTP Tunnel Enable. 2. Click Cache > click Client Accelerator > check Enable Client Acceleration > check the IP address that you checked for the streaming forward proxy service you created. 3. Click Apply. 	When using HTTP tunneling, the appliance uses RTSP to communicate with the origin streaming server. Because RTSP traffic can't pass through most firewalls, you must configure the appliance to work around this issue. For more information, see "Getting QuickTime Streaming Content Through Firewalls" on page 77 .
<p>Enable the client players to use RTSP unless a firewall separates the player from the appliance</p> <p>(See the Important in the Notes column.)</p>	<p>Player setup procedures vary slightly, depending on the version of QuickTime you are using. For example, in QuickTime 5, you do the following:</p> <ol style="list-style-type: none"> 1. In the player, click Edit > click Preferences > click Streaming Transport > select Use UDP, RTSP. 2. Select the optional port field > type 9090 in the field. 3. In the drop-down list, select Streaming Proxy > check RTSP Proxy Server. 4. In the field below RTSP Proxy Server, type the IP address of the appliance's forward streaming proxy service. 5. In the Port ID field type 9090. 6. Close the QuickTime Settings dialog box. 	<p>See <i>B</i> in Figure 31 on page 83.</p> <p>Use the appliance's checked IP address as the address for the forward proxy server.</p> <p>Be sure to specify port 9090, as configured on the appliance.</p> <p>IMPORTANT: If the appliance is on the other side of a firewall, you must configure the player to use HTTP.</p>

To	Do This	Notes
Enable the client players to use HTTP	<p>Player setup procedures vary slightly, depending on the version of QuickTime you are using. For example, in QuickTime 5, you do the following:</p> <ol style="list-style-type: none"> 1. In the player, click Edit > click Preferences > click Streaming Transport > select Use HTTP. 2. Select the optional port field > type 8080 in the field. 3. In the drop-down list, select Streaming Proxy > check HTTP Proxy Server. 4. In the field below HTTP Proxy Server, type the IP address of the appliance's forward streaming proxy service. 5. In the Port ID field type 8080. 6. Close the QuickTime Settings dialog box. 	<p>See <i>B</i> in Figure 31 on page 83.</p> <p>Use the appliance's checked IP address as the address for the forward proxy server.</p> <p>Be sure to specify port 8080, as configured on the appliance.</p> <p>IMPORTANT: This option is required when the player and appliance are separated by a firewall.</p>

14

Accelerating Streaming Media to All Media Players on the Network

This chapter contains instructions for creating QuickTime caching services.

Instructions for creating Windows Media and RealMedia caching services are contained in *Volera Media Excelsior 1.2 for Windows Media Administration Guide* and *Volera Media Excelsior 1.2 for RealSystem Proxy 8 Startup Guide*, respectively.

Overview of Transparent Streaming Proxy

Transparent streaming proxy services require that media player requests be routed to the appliance from a network router or switch. This chapter reviews two different router/switch configurations and contains setup instructions for each configuration type.

The following router/switch configurations are currently supported:

- ♦ An L7 switch routing streaming requests to the transparent streaming service
IMPORTANT: This is the only viable option if the network has both QuickTime and Real RTSP traffic.
- ♦ An L4 switch routing streaming requests to the transparent streaming service

The following options, which are supported for transparent HTTP proxy, are not currently supported for transparent streaming proxy:

- ♦ WCCP
- ♦ Appliance internal routing services acting as the default gateway
- ♦ Appliance internal routing services acting as an inline (main) network router

Transparent/L7 Streaming Proxy

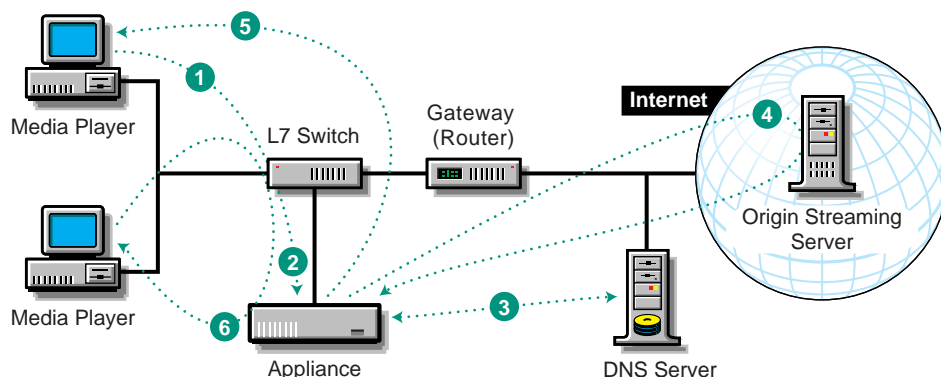
IMPORTANT: If your network has both QuickTime and Real RTSP traffic, you must use this option for any transparent streaming services you create.

Key Functionality

An L7 switch on the same network as the client workstation intercepts media player requests from the client, identifies which requests are for QuickTime content, and sends only those requests to the appliance. (Other streaming media requests are sent to the target origin streaming server.) The transparent streaming proxy service processes the QuickTime requests for the players.

How Transparent Streaming Proxy Works with an L7 Switch

Figure 32



- 1 A player requests a stream from an origin streaming server.
- 2 The L7 switch detects that the request is on port 554, intercepts it, checks to verify it is for QuickTime content, and sends it to the appliance's transparent streaming proxy service.
- 3 The service obtains the numeric IP address of the origin streaming server from DNS.
- 4 The service begins to fill the stream from the origin streaming server.
- 5 The service transmits the stream in real time to the player.
- 6 The transparent streaming forward proxy service handles subsequent requests for the same stream without accessing DNS or the origin streaming server.

Benefits of Transparent Streaming Proxy with an L7 Switch

Assuming the default streaming transport settings (UDP, RTSP) have not been changed, transparent streaming proxy doesn't require player configuration.


After the switch and the appliance have been configured, proxy services are transparent to the player.

Setting Up Transparent Streaming Proxy with an L7 Switch

Figure 33 provides a visual map for the information in this section.

NOTE: The letters in Figure 33 are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 33



To	Do This	Notes
Ensure your basic network configuration is complete	1. See “Basic Network Configuration Setup” on page 27.	
Set up the appliance to work around firewall limitations	For guidelines and more information, see “Getting QuickTime Streaming Content Through Firewalls” on page 77.	
Set up transparent streaming proxy services on the appliance	<ol style="list-style-type: none"> 1. In the browser-based tool, click Cache > Media Cache > Forward > Enable. 2. Type a name for the transparent streaming proxy service. 3. Ensure the Port field has 554 as its value. 4. Click OK. 5. Click Apply. 	<p>See A in Figure 33 on page 89.</p> <p>If you enable logging, log files for the transparent streaming service will have the same name as the streaming service.</p> <p>For more information, see “Client Accelerator Tab” on page 333 and “Transparent Streaming Service Dialog Box (MMS)” on page 367.</p>
Set up your L7 switch to route QuickTime player requests (port 554 traffic) to the appliance	1. Referring to the documentation for your switch, configure a support policy to redirect all port 554 traffic for URLs with .MOV extensions to a transparent proxy address on the appliance.	See B in Figure 33 on page 89.

To	Do This	Notes
Set up players	Assuming the default player configuration settings have not been changed, transparent L7 streaming proxy doesn't require player configuration.	Streaming transport settings should be UDP, RTSP. Streaming proxy settings should not be enabled for an RTSP proxy server.

Transparent/L4 RTSP Streaming Proxy

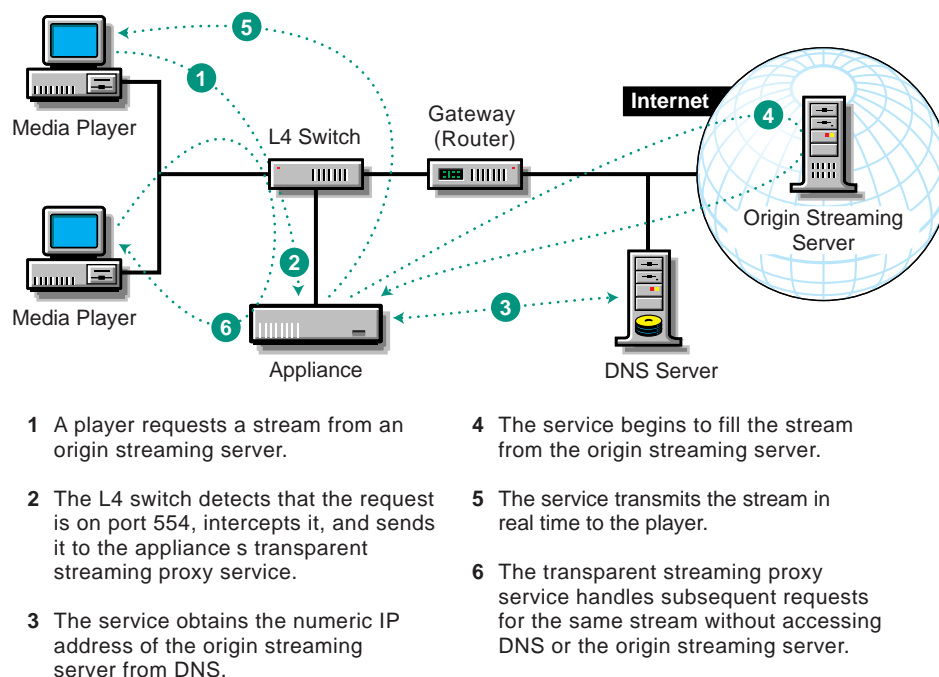
IMPORTANT: If your network has both QuickTime and Real RTSP traffic, you must use the L7 switch option. For more information, see [“Transparent/L7 Streaming Proxy” on page 87](#).

Key Functionality

An L4 switch on the same network as the client workstation intercepts media player requests from the client and sends them to the appliance. The transparent streaming proxy service processes the QuickTime requests for the players.

How Transparent RTSP Streaming Proxy Works with an L4 Switch

Figure 34



Benefits of Transparent RTSP Streaming Proxy with an L4 Switch

Assuming the default streaming transport settings have not been changed, transparent streaming proxy doesn't require player configuration.

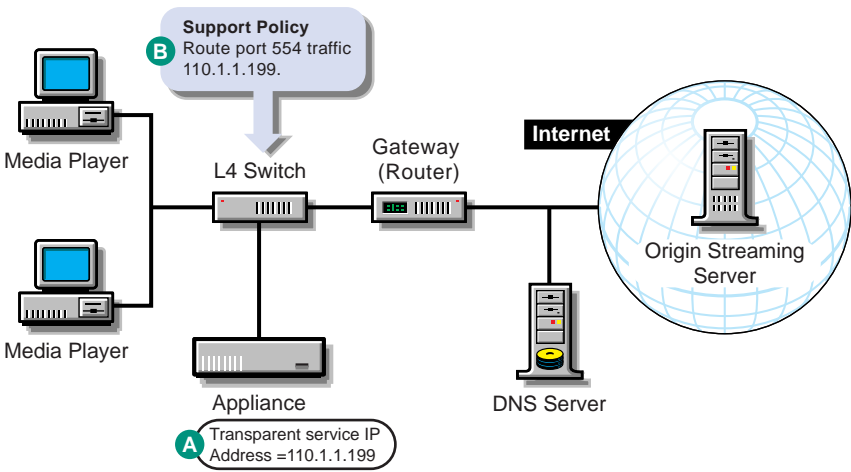
After the switch and the appliance have been configured, proxy services are transparent to the player.

Setting Up Transparent RTSP Streaming Proxy with an L4 Switch

Figure 35 provides a visual map for the information in this section.

NOTE: The letters in Figure 35 are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 35



To	Do This	Notes
Ensure your basic network configuration is complete	1. See “ Basic Network Configuration Setup ” on page 27.	
Set up the appliance to work around firewall limitations	For guidelines and more information, see “ Getting QuickTime Streaming Content Through Firewalls ” on page 77.	The choices you have depend on whether the players on your network will use HTTP tunneling or UDP, RTSP.
Set up transparent streaming proxy services on the appliance	<ol style="list-style-type: none">1. In the browser-based tool, click Cache > click Media Cache > Transparent > Enable.2. Type a name for the transparent streaming proxy service.3. Ensure the Port field has 554 as its value.4. Click OK.5. Click Apply.	<p>See A in Figure 35 on page 91.</p> <p>If you enable logging, log files for the transparent streaming service will have the same name as the streaming service.</p> <p>For more information, see “Client Accelerator Tab” on page 333 and “Transparent Streaming Service Dialog Box (MMS)” on page 367.</p>

To	Do This	Notes
Set up your L4 switch to route QuickTime player requests (port 554 traffic) to the appliance	1. Referring to the documentation for your switch, configure a support policy to redirect port 554 traffic to a transparent proxy address on the appliance.	See <i>B</i> in Figure 35 on page 91 .
Set up players	Assuming the default player configuration settings have not been changed, transparent L4 RTSP streaming proxy doesn't require player configuration.	Streaming transport settings should be UDP, RTSP. Streaming proxy settings should not be enabled for an RTSP proxy server.

Transparent/L4 HTTP Tunneled Streaming Proxy

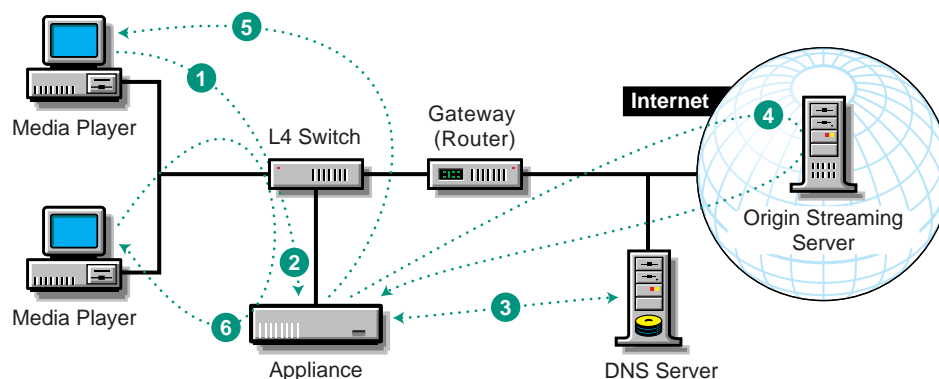
IMPORTANT: If your network has both QuickTime and Real RTSP traffic, you must use the L7 switch option. For more information, see [“Transparent/L7 Streaming Proxy” on page 87](#).

Key Functionality

An L4 switch on the same network as the client workstation intercepts media player requests from the client and sends them to the appliance. The transparent streaming proxy service processes the QuickTime requests for the players.

How Transparent HTTP Tunneled Streaming Proxy Works with an L4 Switch

Figure 36



- 1 A player requests a stream from an origin streaming server.
- 2 The L4 switch detects that the request is on port 80, intercepts it, and sends it to the appliance's transparent streaming proxy service.
- 3 The service obtains the numeric IP address of the origin streaming server from DNS.
- 4 The service begins to fill the stream from the origin streaming server.
- 5 The service transmits the stream in real time to the player.
- 6 The transparent streaming proxy service handles subsequent requests for the same stream without accessing DNS or the origin streaming server.

Benefits of Transparent HTTP Tunneled Streaming Proxy with an L4 Switch

Assuming the default streaming transport settings have not been changed, transparent streaming proxy doesn't require player configuration.

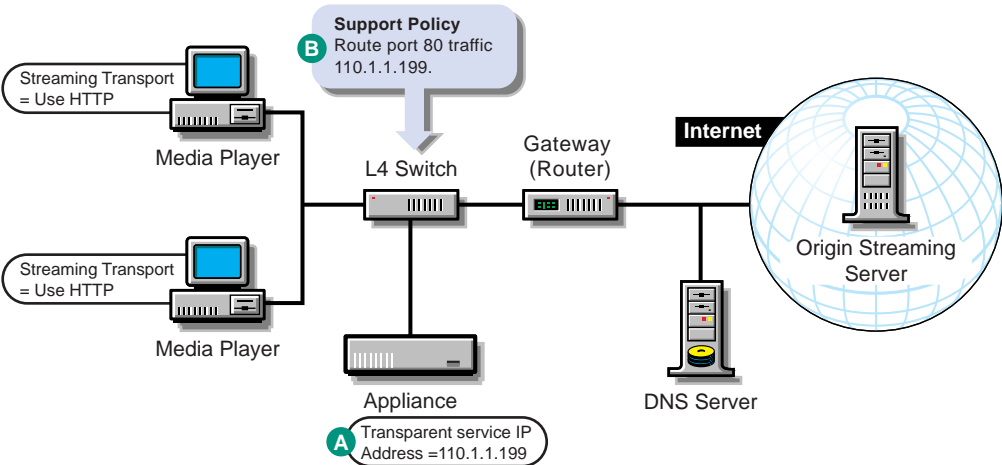
After the switch and the appliance have been configured, proxy services are transparent to the player.

Setting Up Transparent HTTP Tunneled Streaming Proxy with an L4 Switch

Figure 37 provides a visual map for the information in this section.

NOTE: The letters in Figure 37 are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 37



To	Do This	Notes
Ensure your basic network configuration is complete	<div>1. See “Basic Network Configuration Setup” on page 27.</div> <div>2. Set up HTTP transparent services as described in “Transparent/L4 Proxy Setup” on page 43.</div>	

To	Do This	Notes
Set up transparent streaming proxy services on the appliance	<ol style="list-style-type: none"> 1. In the browser-based tool, click Cache > click Media Cache > Transparent > Enable. 2. Type a name for the transparent streaming proxy service. 3. Ensure the Port field has 80 as its value. 4. Click OK. 5. Ensure the QuickTime HTTP Tunnel Enable option is checked. 6. Click Apply. 	<p>See A in Figure 37 on page 93.</p> <p>If you enable logging, log files for the transparent streaming service will have the same name as the streaming service.</p> <p>For more information, see “Client Accelerator Tab” on page 333 and “Transparent Streaming Service Dialog Box (MMS)” on page 367.</p>
Set up your L4 switch to route QuickTime player requests (port 80 traffic) to the appliance	<ol style="list-style-type: none"> 1. Referring to the documentation for your switch, configure a support policy to redirect port 80 traffic to a transparent proxy address on the appliance. 	<p>See B in Figure 37 on page 93.</p> <p>This will route all port 80 traffic to the appliance.</p>
Enable the client players to use HTTP	<p>Player setup procedures vary slightly, depending on the version of QuickTime you are using. For example, in QuickTime 5, you do the following:</p> <ol style="list-style-type: none"> 1. In the player, click Edit > click Preferences > click Streaming Transport > select Use HTTP. 2. Select the optional port field > type 80 in the field. 3. Close the QuickTime Settings dialog box. 	<p>Be sure to specify port 80 as configured on the appliance.</p>

15 Accelerating Streaming Media Servers

This chapter contains instructions for creating QuickTime caching services.

Instructions for creating Windows Media and RealMedia caching services are contained in *Volera Media Excelsator 1.2 for Windows Media Administration Guide* and *Volera Media Excelsator 1.2 for RealSystem Proxy 8 Startup Guide*, respectively.

Overview of Streaming Media Server Acceleration

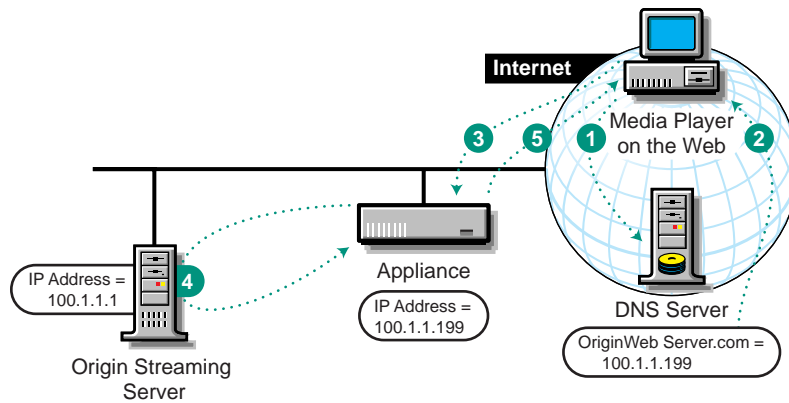
The appliance's streaming media accelerator relies on DNS to cause the accelerator to receive requests originally targeted at the origin streaming server. The appliance's streaming media accelerator handles the requests, accessing the origin streaming server only when needed objects are not cached.

How Origin Streaming Server Acceleration Works

The mechanism for routing player requests meant for origin streaming servers to the streaming media accelerator instead can be summarized as follows:

- ♦ Without acceleration, DNS resolves the origin streaming server's DNS name to the origin streaming server's IP address.
- ♦ With acceleration, DNS resolves the server's name to the IP address of an appliance streaming media accelerator (reverse proxy) service.

Figure 38



- 1 A player requests a stream from an origin streaming server. This generates a request to DNS for the numeric IP address of the origin streaming server.
- 2 Instead of returning the origin streaming server's numeric IP address, DNS returns the numeric IP address of the accelerator service on the appliance.
- 3 The player requests the stream using the numeric IP address of the accelerator service.
- 4 The accelerator service begins filling the stream from the origin streaming server.
- 5 The accelerator service transmits the stream in real time to the player.

Streaming Accelerator Services Are Created in Matching Pairs

As shown in [Figure 38](#), the streaming accelerator service might receive both RTSP and HTTP tunneled requests. For this reason you need to create both an RTSP accelerator and a matching HTTP accelerator for each streaming accelerator service. These are referred to in this document as the streaming accelerator pair.

Guidelines and steps for creating the streaming accelerator pair are described in [“Streaming Server Accelerator Setup” on page 97](#) and [“Deploying Multiple Streaming Accelerators on the Appliance” on page 98](#).

Benefits of Origin Streaming Server Acceleration

- ♦ A streaming server accelerator reduces response time to browser requests and frees up origin streaming server bandwidth, allowing it to handle requests for less frequently requested, uncached streams much more quickly.
- ♦ The appliance can accelerate origin streaming servers at remote locations that don't offer high-bandwidth connections. The streaming server accelerator can be located close to the Internet backbone, delivering high-speed access to browsers for all cached objects. The connection to the origin streaming server is then used for transporting only those portions of streams not already in cache.

For tips and guidelines on setting up origin streaming server accelerators, see [“Streaming Server Accelerator Setup” on page 97](#).

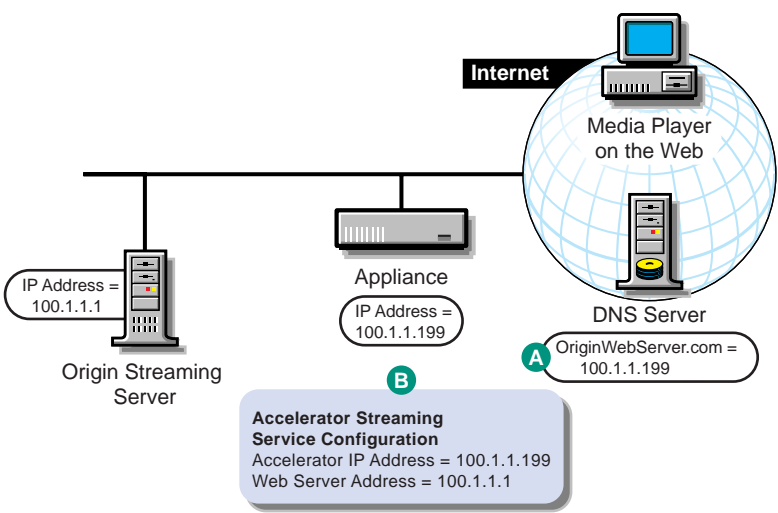
The procedure for configuring DNS to work with Web server accelerators is explained in [“Working with DNS” on page 53](#).

Streaming Server Accelerator Setup

Figure 39 provides a visual map for the information in this section.

NOTE: The letters in Figure 39 are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 39



To	Do This	Notes
Ensure your basic network configuration is complete for each appliance	1. See “Configuring the Exceleator Appliance” on page 28.	
Ensure that DNS resolves browser requests to the appliance IP addresses configured for the streaming server accelerator services	1. See “Working with DNS” on page 53.	See A in Figure 39.

To	Do This	Notes
Set up the streaming (RTSP) server accelerator service member of the streaming accelerator pair	<ol style="list-style-type: none"> 1. In the browser-based tool, click Cache > Media Cache > Accelerator. 2. Check Enable. 3. For your tracking purposes, enter a name for the streaming server accelerator. 4. Enter a DNS name, if desired. 5. Enable logging and set up logging options, if desired. 6. Ensure the Web Server and Accelerator Port field values are both 554. 7. In the Accelerator Address list, check one or more addresses from which the streaming server accelerator will receive requests and vend data to. (DNS resolves requests to these addresses.) 8. In the Web Server Address list, insert one IP address (or DNS name) from which the streaming server accelerator will fill its cache. 9. Click OK. 10. Check QuickTime HTTP Tunnel Enable. 11. Click Apply. 	<p>See B in Figure 39 on page 97.</p> <p>Streaming accelerators are created in pairs that consist of an RTSP accelerator (created in this step) and an HTTP accelerator (created in the following step).</p> <p>If server persistence is enabled in the Web Server Accelerator tab, Excelerator will use the same Web server to fill browser requests during a session. This setting affects all accelerators on the appliance and saves e-business users from having to log in multiple times. See “Web Server Accelerator Tab” on page 348.</p> <p>If logging is enabled, accelerator log files for the streaming server accelerator will have the same name as the streaming server accelerator.</p> <p>IMPORTANT: The accelerator IP addresses must be used only by this streaming accelerator service and the matching HTTP accelerator service that you create in the next step.</p> <p>If you enter DNS names in the Web Server Addresses list, make sure they are not the names that now resolve to appliance numeric IP addresses. That would create an endless loop.</p>
Set up an HTTP server accelerator to support HTTP tunneled traffic coming to the accelerator pair	<p>Create a corresponding Web server accelerator by completing the instructions found in “Web Server Accelerator Setup” on page 52.</p> <p>Use the same IP addresses, Web server address, and DNS name as entered for the streaming server accelerator you just created.</p>	<p>IMPORTANT: Do not check the Act as a Tunnel option when creating the Web server accelerator service.</p> <p>IMPORTANT: The accelerator IP addresses must be used only by this HTTP server accelerator service and its matching streaming accelerator service created in the previous step.</p>

Deploying Multiple Streaming Accelerators on the Appliance

If multiple accelerator services use the same IP addresses and port numbers, the appliance cannot map streaming requests to the appropriate accelerators and therefore cannot process the requests.

You must therefore ensure that each streaming accelerator service pair (RTSP and HTTP) complies with the following requirements as specified in “**Streaming Server Accelerator Setup**” on page 97:

- ◆ Specify only one origin streaming server per accelerator service pair.

Each accelerator in the pair (RTSP and HTTP) must have the same origin streaming server in its Web Server Address list.

- ◆ Ensure that the IP addresses in the Accelerator Address list of each accelerator pair are used only by the members of the pair and not by any other accelerator service.

If your appliance has only one network card, you can add unique addresses for each accelerator service to the network card.

NOTE: QuickTime players do not provide host header information.

16

Configuring an Upstream Proxy for the Appliance

The appliance can be configured to use another proxy server for filling RTSP requests. The relationship created between the appliance and the other server is similar to the CERN hierarchical relationship for HTTP requests. However, only one parent is supported.

This feature is useful for getting streaming requests through some firewalls as explained in [“Going Through an RTSP Proxy on the Firewall Network Address Translator \(NAT\)” on page 79](#).

It can also be used to create simple hierarchies of appliances or to enable the appliance to leverage streaming proxies on the Web.

To create an upstream proxy for the appliance, see Enable Upstream Proxy under [“Streaming Tab” on page 357](#).

17

Configuring QuickTime Media Players to Use Proxy Services

Two media player settings must match your network setup:

- ♦ **Streaming Proxy Options:** Lets you designate a forward proxy server for RTSP- and/or HTTP-based streaming media requests.
- ♦ **Streaming Transport Options:** Lets you specify whether streaming media packets are transported in the UDP and RTSP protocols or wrapped in HTTP packets.

Setting Streaming Proxy Options

For media players to use forward streaming proxy services, they must be configured with the service's IP address and port number.

You can configure players to send both HTTP (port 8080) and/or RTSP (port 9090) requests to an appliance forward streaming proxy service IP address. The options you enable depend on the type of requests you want handled by the forward service.

The location of the Streaming Proxy options varies slightly, depending on the version of the player being configured. For example, in the QuickTime 5 player, the path is Edit > Preferences > Streaming Proxy.

Setting Streaming Transport Options

QuickTime players offer two transport options:

- ♦ UDP, RTSP
- ♦ HTTP

The main consideration when selecting a transport option for proxy requests is whether the requests must pass through a firewall to reach the appliance:

- ♦ If the appliance and the player are not separated by a firewall, the UDP, RTSP option is best.
- ♦ If the appliance and the player are separated by a firewall, HTTP is probably required.

NOTE: Issues dealing with the position of the appliance relative to a firewall are discussed in [“Setting Up Your Appliance to Work with Firewalls” on page 78](#).

Specific player configuration steps are included in each streaming proxy configuration chapter. For more information, see [Chapter 13, “Accelerating Streaming Media to Individual Media Players,” on page 81](#); [Chapter 14, “Accelerating Streaming Media to All Media Players on the Network,” on page 87](#); and [Chapter 15, “Accelerating Streaming Media Servers,” on page 95](#).

V

Hierarchies, Clusters, and Multihoming

Large Web sites require multiple Web servers in rather complex configurations. The Excelerator appliance is designed to accelerate most complex configurations, including Web server farms and various multihoming configurations.

Appliances can be configured in hierarchical relationships for efficiency in obtaining and caching Web content. Appliances can also be clustered together to provide multiple proxy services and failover protection of the services.

The chapters in this section provide overviews of the advanced configurations your Web content strategy might need, as well as setup and troubleshooting help for each configuration.

To	See
Set up appliance hierarchies and/or configure the appliance to get and share content in hierarchical relationships with other proxy servers	Chapter 18, “Hierarchical Caching,” on page 107
Set up a cluster of appliances to provide multiple proxy services and ensure failover protection of the services	Chapter 19, “Clustering,” on page 113
Set up multihoming configurations	Chapter 20, “Appliance Groups and Multihomed Configurations,” on page 121

18 Hierarchical Caching

This section contains information about hierarchical caching.

Overview

Caching hierarchies extend an appliance's cache performance by enabling it to get uncached objects from other appliances on the network rather than from the origin Web servers on the Internet. The mechanism for achieving communication between proxy servers is the caching hierarchy.

Caching hierarchies make it possible for a high percentage of browser requests to be filled from within the network.

The Excelerator appliance supports two types of caching hierarchies: CERN and ICP.

CERN Hierarchies

Key Functionality

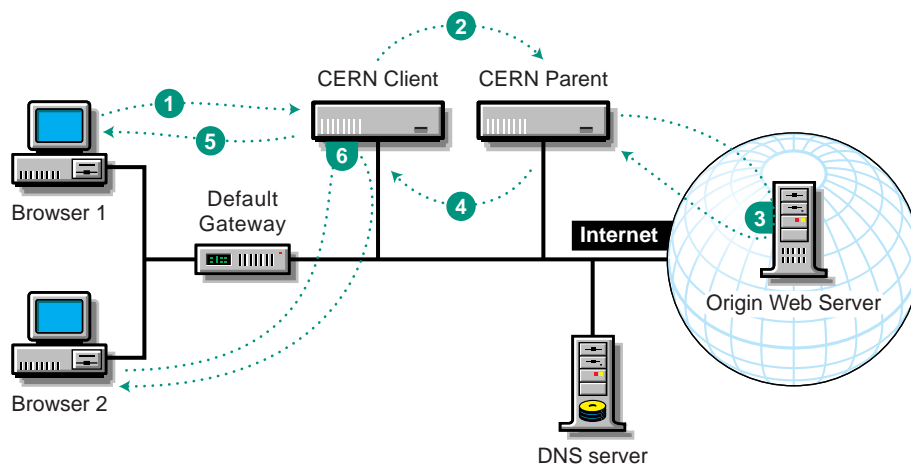
CERN hierarchies consist of CERN clients and CERN parents. The client is basically an appliance that is configured to use one or more proxy servers as forward proxy servers (CERN parents).

CERN hierarchies are very similar to the relationship between a client workstation browser and a forward proxy server. The difference is that the CERN client is an intermediary and the browser is a user agent.

CERN parents can also designate other proxy servers as their CERN parents.

How CERN Hierarchies Work

Figure 40



- 1 The CERN client receives a browser request which it can't fill from cache.
- 2 The client sends a forward proxy request to its CERN parent.
- 3 The parent obtains the objects (either through its own hierarchy as a client or directly from the origin Web server).
- 4 The CERN parent sends copies of the objects to the client.
- 5 The client caches the objects and sends copies to the requesting browser.
- 6 The client fills subsequent requests for the same objects from cache.

Benefits of CERN Hierarchies

- ♦ A higher percentage of browser requests can be filled without accessing origin Web servers.
- ♦ Because CERN hierarchies use static routing without queries, they generate less packet traffic on the network than ICP hierarchies do.

For tips and guidelines on setting up CERN hierarchies, see [“CERN Hierarchy Setup” on page 109](#).

ICP Hierarchies

Key Functionality

ICP hierarchies consist of ICP clients, parents, and peers. The ICP client is an appliance that is configured to communicate with one or more proxy servers that have been configured as ICP servers. These ICP servers are designated as either peers or parents of the ICP client.

Servers in an ICP hierarchy use:

- ♦ The ICP protocol to locate URLs
- ♦ The forward proxy service to send data

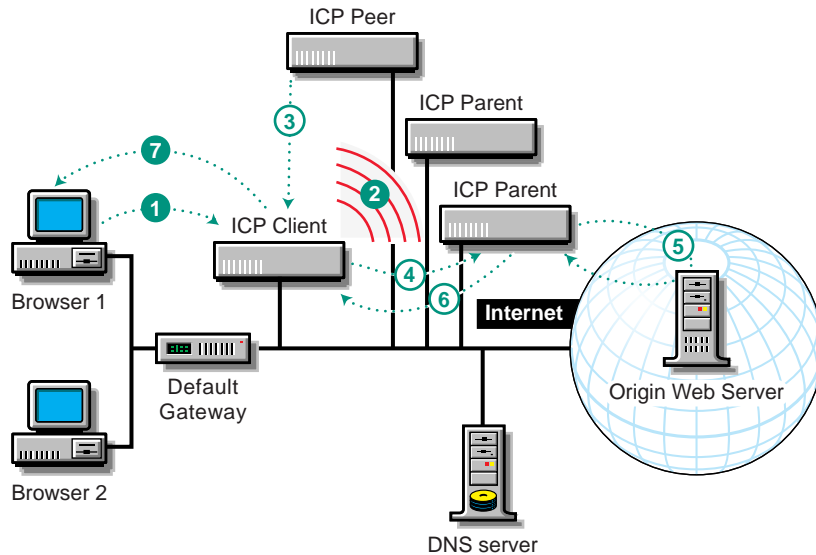
ICP hierarchies are more complex than CERN hierarchies. The ICP protocol offers intelligence for deciding which routes to pursue first when seeking data.

ICP parents can also specify other proxy servers as their ICP parents or peers.

How ICP Hierarchies Work

NOTE: Steps shown in reverse color are conditional.

Figure 41



- 1 The ICP client receives a browser request which it can't fill from cache.
- 2 The client sends a query request to its peers and parents.
- 3 If a member of the hierarchy has the objects, it sends copies to the appliance.
- 4 If all peers and parents respond miss, or the request times out, the appliance selects a parent and issues a direct request for the objects.
- 5 The parent gets the objects, either through its hierarchy or from the origin Web server.
- 6 The parent returns the objects to the ICP client.
- 7 The client caches the objects and returns them to the browser.

Benefits of ICP Hierarchies

- ♦ A higher percentage of browser requests can be filled without accessing origin Web servers.
- ♦ ICP hierarchies are more powerful than CERN hierarchies because they include peer relationships and allow hierarchy members to communicate regarding the best routes for obtaining objects.

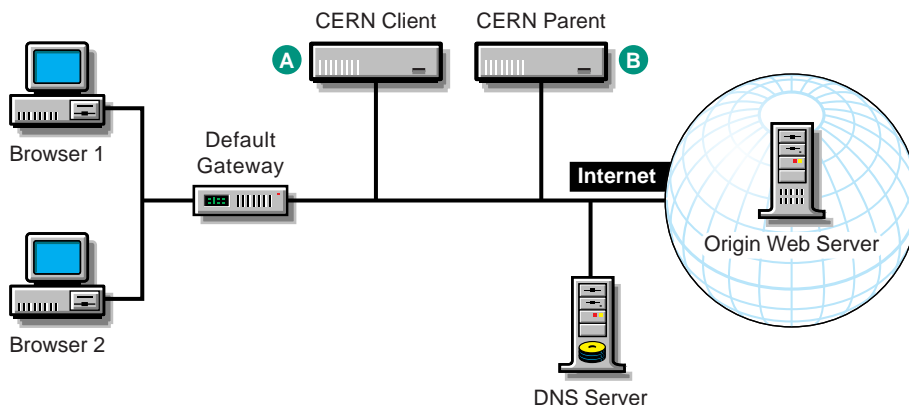
For tips and guidelines on setting up ICP hierarchies, see [“ICP Hierarchy Setup” on page 111](#).

CERN Hierarchy Setup

Figure 42 provides a visual map for the information in this section.

NOTE: The letters in Figure 42 are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 42



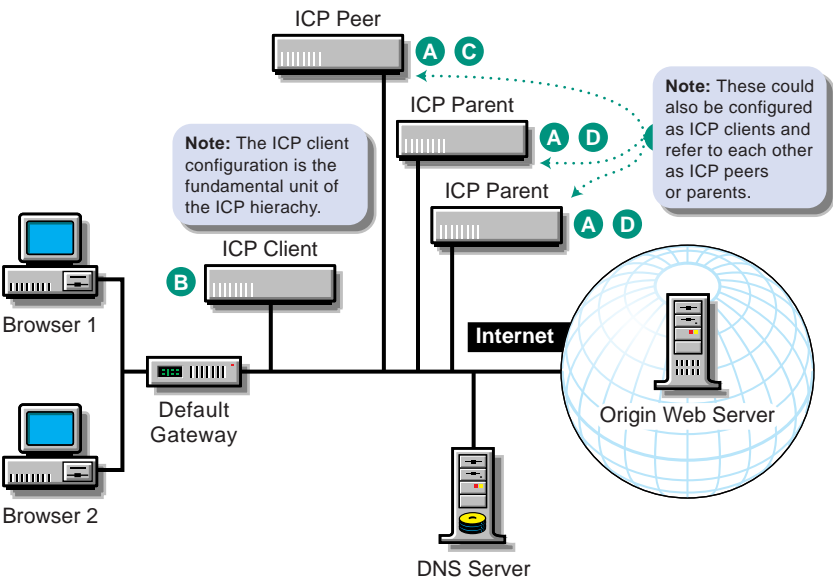
To	Do This	Notes
Ensure your basic network configuration is complete for each appliance in the CERN hierarchy	1. See “Configuring the Excelerator Appliance” on page 28.	
Enable the appliance to request and receive data through the hierarchy	1. In the browser-based tool, click Hierarchy > ICP/ CERN Configuration. 2. Check Enable ICP/CERN Client. 3. Click Apply.	See A in Figure 42 . For more information, see “ICP/ CERN Configuration Tab” on page 405.
Designate one or more CERN parents for the appliance	1. In the ICP/CERN Configuration tab, click CERN Parent. 2. In the Hostname field, enter the numeric IP address or DNS name of the proxy server that you want to serve as a parent for the appliance. 3. In the HTTP Proxy Port field, enter the port number that the parent will receive and transmit data on. 4. In the Priority field, enter a priority for the parent. 5. Click Apply.	See A in Figure 42 on page 110 . If the appliance has multiple CERN parents, one parent must have a priority number lower than the others. (This parent will be accessed first.) For more information, see “CERN Parent Dialog Box” on page 409.
Ensure each CERN parent is configured as a forward proxy server	1. See “Forward Proxy Setup” on page 36.	See B in Figure 42 on page 110 . A CERN parent is a proxy server that is enabled for forward proxy services. See “Client Accelerator Tab” on page 333.

ICP Hierarchy Setup

Figure 43 provides a visual map for the information in this section.

NOTE: The letters in Figure 43 are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 43



To	Do This	Notes
Ensure your basic network configuration is complete for each appliance in the ICP hierarchy	1. See “Configuring the Excelerator Appliance” on page 28.	
Configure all appliances that will receive and process requests from other members of the hierarchy as ICP servers	1. In the browser-based tool, click Hierarchy > ICP/ CERN Configuration. 2. Check Enable ICP Server. 3. Make sure each appliance is also configured as a forward proxy server. See “Forward Proxy Setup” on page 36. 4. Click Apply.	See A in Figure 43 on page 111. For more information on ICP server options, see “ICP/ CERN Configuration Tab” on page 405.
Configure all appliances that will request and receive data through the hierarchy as ICP/ CERN clients	1. In the browser-based tool, click Hierarchy > ICP/ CERN Configuration. 2. Check Enable ICP/ CERN Client. 3. Click Apply.	See B in Figure 43 on page 111. For more information on ICP/ CERN client options, see “ICP/ CERN Configuration Tab” on page 405.

To	Do This	Notes
Designate one or more ICP peers for the appliance	<ol style="list-style-type: none"> 1. In the ICP/CERN Configuration tab, click ICP Peer. 2. In the Hostname field, enter the numeric IP address or DNS name of a proxy server that you want to serve as an ICP peer for the appliance. 3. In the HTTP Proxy Port field, enter the port number that the peer will receive and transmit data on. 4. In the ICP Port field, enter the port number that members of the hierarchy will communicate on. (The default is 3130.) 5. Click Apply. 	<p>See <i>C</i> in Figure 43 on page 111.</p> <p>As you define peers and parents for the appliance, you are building the ICP hierarchy.</p> <p>Having ICP peers is optional. An appliance could have only ICP parents.</p> <p>For more information, see “ICP Peer Dialog Box” on page 408.</p>
Designate one or more ICP parents for the appliance	<ol style="list-style-type: none"> 1. In the ICP/CERN Configuration tab, click ICP Parent. 2. In the Hostname field, enter the numeric IP address or DNS name of a proxy server that you want to serve as an ICP parent for the appliance. 3. In the HTTP Proxy Port field, enter the port number that the parent will receive and transmit data on. 4. In the ICP Port field, enter the port number that members of the hierarchy will communicate on. (The default is 3130.) 5. In the Priority field, enter a priority for the parent. 6. Click Apply. 	<p>See <i>D</i> in Figure 43 on page 111.</p> <p>Having ICP parents is optional. An appliance could have only ICP peers.</p> <p>If an appliance has only ICP peers and a request can't be filled by its peers, the appliance will go directly to the origin Web server to get the objects (which it would then share with the hierarchy on future requests).</p> <p>For more information, see “ICP Parent Dialog Box” on page 408.</p>

19 Clustering

This section contains information about clustering.

Overview

You can provide failover protection of caching services on your network by using the following:

- ♦ Web server accelerator groups
- ♦ Appliance clusters

Web Server Accelerator Groups

Key Functionality

DNS Round-Robin dispenses the multiple IP addresses assigned to a single origin Web server DNS name in a rotating sequence, achieving load balancing and automatic failover protection for the acceleration service.

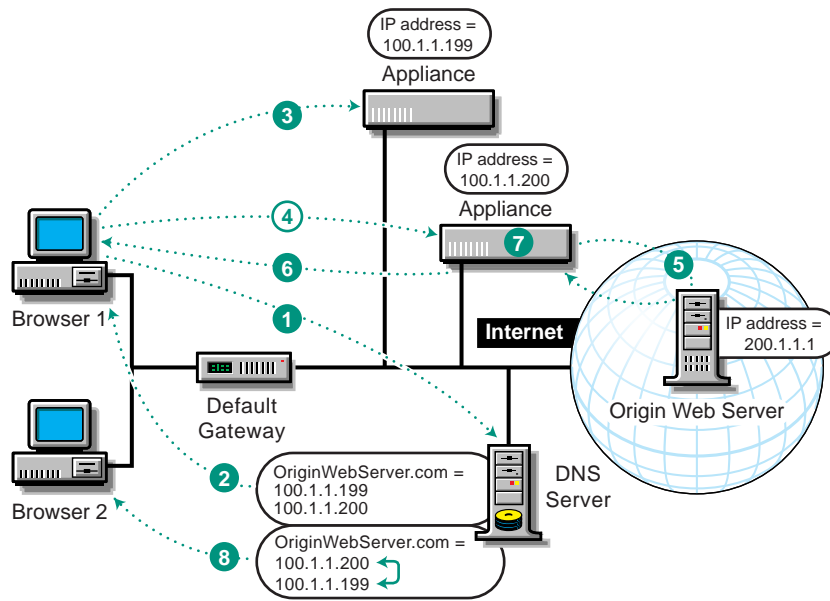
If the appliances are configured as ICP peers of each other, they automatically share information within the group. See [“How ICP Hierarchies Work” on page 109](#) for more information.

Although the following illustration shows the origin Web server on the Internet, it could be on the same network as the DNS server and the Web server accelerator, and the browsers could be on the Web, as shown in [“How Origin Web Server Acceleration Works” on page 51](#).

How a Web Server Accelerator Group Works

NOTE: Steps shown in reverse color are conditional.

Figure 44



- 1 A browser requests an origin Web server's Web page, resulting (as always) in a DNS request.
- 2 DNS dispenses the IP addresses of the accelerator services on both appliances.
- 3 The browser requests the Web page using the first IP address returned by DNS.
- 4 If the request times out, the browser reissues the request using the next IP address. This provides fault tolerance.
- 5 The accelerator service on the appliance obtains any uncached objects from the origin Web server.
- 6 The accelerator service returns copies of the objects to the browser.
- 7 The accelerator service handles subsequent requests for the same objects without accessing the origin Web server.
- 8 With each request for the same domain name, DNS rotates the order of the IP addresses. This provides load balancing within the appliance group.

Benefits of Web Server Accelerator Groups

- ◆ Deploying multiple Web server accelerators, as described in [Figure 44](#), provides load balancing, thus increasing throughput and performance.
- ◆ Having multiple Web server accelerators for a single origin Web server provides failover protection and increases reliability.

For tips and guidelines on setting up Web server accelerator groups, see [“Appliance Web Server Accelerator Group Setup” on page 121](#).

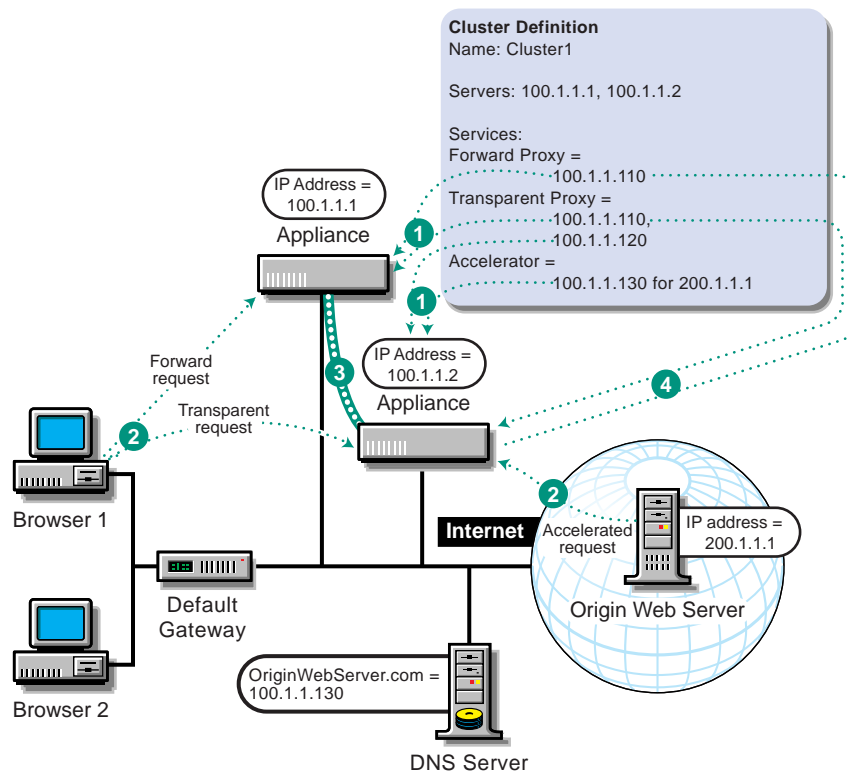
Appliance Clusters

Key Functionality

An appliance cluster is a group of appliances that have been configured with the same cluster definition. The cluster dynamically handles forward, transparent, and reverse proxy services that are defined for the cluster.

How Clusters Work

Figure 45



- 1 When the cluster is enabled in the Cluster tab, caching service IP addresses are allocated to members of the cluster.
- 2 Forward, transparent, and accelerator services are handled exactly as if configured on individual appliances.
- 3 Cluster members communicate their status with each other through heart-beat packets.
- 4 If a member of the cluster ceases to function, the services it was handling are reassigned to other members of the cluster (in this example, the appliance with IP address 100.1.1.1 fails).

Key Benefits of Clusters

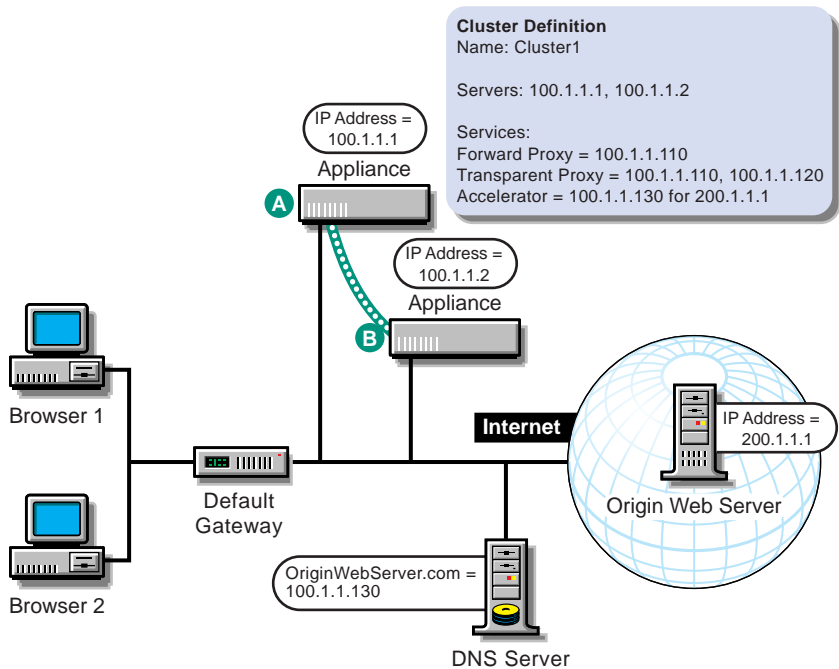
- ◆ Assigning services to a cluster provides failover support for all defined services.
- ◆ Distribution of service IP addresses among the appliances is automatically negotiated within the cluster.
- ◆ If one of the appliances shuts down for any reason, the service IP addresses that were previously assigned to it are redistributed among the remaining appliances.

Cluster Setup

Figure 46 provides a visual map for the information in this section.

NOTE: The letters in Figure 46 are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 46



Be sure to observe the following configuration rules:

- ◆ Configure at least one service on each appliance.
- ◆ Include all cluster members in the Servers list on each appliance.
- ◆ Ensure that the last appliance configured has the lowest IP address.
- ◆ Include all services on the last appliance configured.

To	Do This	Notes
Ensure your basic network configuration is complete for each appliance	1. See “Configuring the Excelerator Appliance” on page 28.	<p>Each appliance in the cluster must have an IP address on the same subnet as the cluster.</p> <p>If you are defining a transparent service in the cluster and you plan to use the clustered appliances in an inline routing configuration, you must configure each client to use the transparent service IP address as its default gateway.</p> <p>Also, each client’s DNS server definition must include an IP address that is assigned to a network card in each appliance. The IP addresses you list must be on the same subnet as the client.</p>

To	Do This	Notes
Create a cluster definition on each appliance	<ol style="list-style-type: none"> 1. In the browser-based tool, click Cache > Cluster. 2. Check Enable Cluster. 3. Enter a cluster name. 4. Click the Subnet drop-down list > select the cluster subnet address. 5. In the Servers list, insert the IP address, name, role, and capacity of each appliance that will participate in the cluster. 6. When configuring the last appliance, skip to the next entry in this table, "To include the complete list of services in the cluster definition on the last appliance configured." 7. Insert at least one service into the Services list. 8. Click Apply. 9. Go to Step 1 and configure the next appliance in the cluster. 	<p>See A in Figure 46 on page 116.</p> <p>IMPORTANT: Configure the appliance with the lowest IP address last.</p> <p>An appliance can belong to only one cluster. You assign each cluster a unique name.</p> <p>The cluster name and subnet address must be the same for each appliance in the cluster.</p> <p>The IP addresses you enter must all fit within the cluster's subnet address range.</p> <p>IMPORTANT: IP addresses assigned to clustered services must not be assigned to network adapters in Network > IP Addresses.</p> <p>The cluster definition for each appliance must include all the appliances in the cluster's Servers list.</p> <p>The name, role, and capacity of each appliance must be entered exactly the same in each appliance's cluster definition. To learn more about the role of the capacity field, see "About Capacity and Weight" on page 118.</p> <p>For more information on appliance clusters, see "Cluster Tab" on page 371.</p>
Include the complete list of services in the cluster definition on the last appliance configured	<ol style="list-style-type: none"> 1. On the last appliance only, insert all the services the cluster will handle in the Services list. 2. Click Apply. 	<p>See A in Figure 46 on page 116.</p> <p>IMPORTANT: The last appliance configured must be the one with the lowest IP address.</p> <p>The cluster definition for the last appliance added to the cluster must also include all the services the cluster will handle.</p> <p>The list of services is automatically propagated to other appliances in the cluster.</p>

About Capacity and Weight

The appliance employs complex formulas to ensure that clustered service loads are distributed among cluster members in such a way that cluster resources are fully leveraged.

It is important that the capacity and weight numbers are correctly set for the service distribution formulas to work properly.

Server Capacity

The Servers lists you create on cluster members must each contain all cluster members. The contents of each Servers list should be identical on each cluster member. It is especially important that Capacity field values accurately reflect the relative capacity of each appliance in the cluster.

Follow these rules when assigning capacity numbers:

- ◆ Ensure that numbers accurately reflect relative appliance capacity.

If appliances are identical, assign the same capacity number to each appliance. If appliances are not identical, you can check with your appliance vendor for recommended capacity ratios. Ultimately, however, deciding relative capacities is a subjective process.

- ◆ Keep numbers as small as possible.

Assign identical appliances a capacity number of 1. If you determine that one appliance is fifty percent more powerful than another, assign the first a capacity number of 3 and the second a capacity number of 2.

Service Weight

Each appliance's cluster definition must list at least one clustered service in its Services list. The definition on the last appliance configured must contain all clustered services.

The best metric to use for weight assignment is service traffic or the number of requests handled by each service during a specific period of time.

For example, a cluster might contain two services: a Web accelerator that handles 100,000 hits per day and a forward service that handles 20,000 requests per day.

To have the weight of each service reflect its number of requests, you could assign the accelerator service a weight of 10 and the forward service a weight of 2.

Determining an appropriate weight number requires some investigation on your part. Initially you might need to estimate the number of requests for each service. After the cluster has been running for a while, you can use statistics from the Monitoring panel tabs and from log files to refine your weight numbers.

Allocation of Clustered Service IP Addresses

As previously mentioned, the formulas Excelsator uses to ensure load balancing are complex. The results, on the other hand, are fairly straightforward.

The following discussion is provided for those who want to understand how the caching system allocates clustered service IP addresses among cluster members. It also explains why service log files are often distributed among cluster members.

Excelsator distributes services among cluster members by clustered service IP address as follows:

1. The system divides the weight you've assigned to each service by the number of IP addresses assigned to the service.

For example, the Web server accelerator service mentioned above has a weight of 10. If you've assigned two IP addresses to the service, Excelerator assigns each IP address a weight of 10 divided by 2, or 5.

By the same token, if you've assigned one IP address to the forward service, Excelerator assigns this IP address a weight of 2.

2. The system assigns the IP addresses with the highest weights first. If an IP address is handling two or more services, Excelerator uses the total of the combined weights for all services for that IP address's weight.

Let's consider two possible scenarios:

Scenario 1: You've assigned two IP addresses to the Web server accelerator service and a third IP address to the forward service.

Scenario 2: You've assigned one IP address to both the Web server accelerator and the forward service. You've assigned a second IP address to the Web server accelerator. The first address has a combined total weight of 5 plus 2, or 7. The second address has a weight of 5.

3. The system assigns each IP address to only one cluster member. If the IP address is handling multiple services, the services are all assigned with the IP address to a single appliance.

Scenario 1: The system has three IP addresses to assign, with service weights of 5, 5, and 2.

Scenario 2: The system has two IP addresses to assign, with service weights of 7 and 2.

4. The system subtracts an amount of capacity from each appliance as services are allocated to it. The amount of capacity subtracted reflects the relative combined service weight of the IP address.

Scenario 1: You have two appliances of equal capacity in your cluster. The system assigns the first 5-weight IP address to the first appliance and subtracts a calculated amount of capacity from the first appliance. Since the second appliance now has the most current capacity, Excelerator assigns it the second 5-weight IP address, making the capacities equal once more. The system then assigns the 2-weight IP address to the first appliance.

All services have been allocated. The first appliance has two IP addresses, one handling a Web server accelerator and the other handling a forward proxy service. The second appliance has one IP address, also handling the Web server accelerator service.

Scenario 2: You have two appliances of equal capacity in your cluster. The system assigns the first 7-weight IP address to the first appliance and subtracts a calculated amount of capacity. Since the second appliance now has the most current capacity, Excelerator assigns it the 5-weight IP address.

All services have been allocated. The first appliance has one IP address handling a Web server accelerator and a forward proxy service. The second appliance has one IP address handling the Web server accelerator service.

20

Appliance Groups and Multihomed Configurations

This section contains information about appliance groups and multihomed configurations.

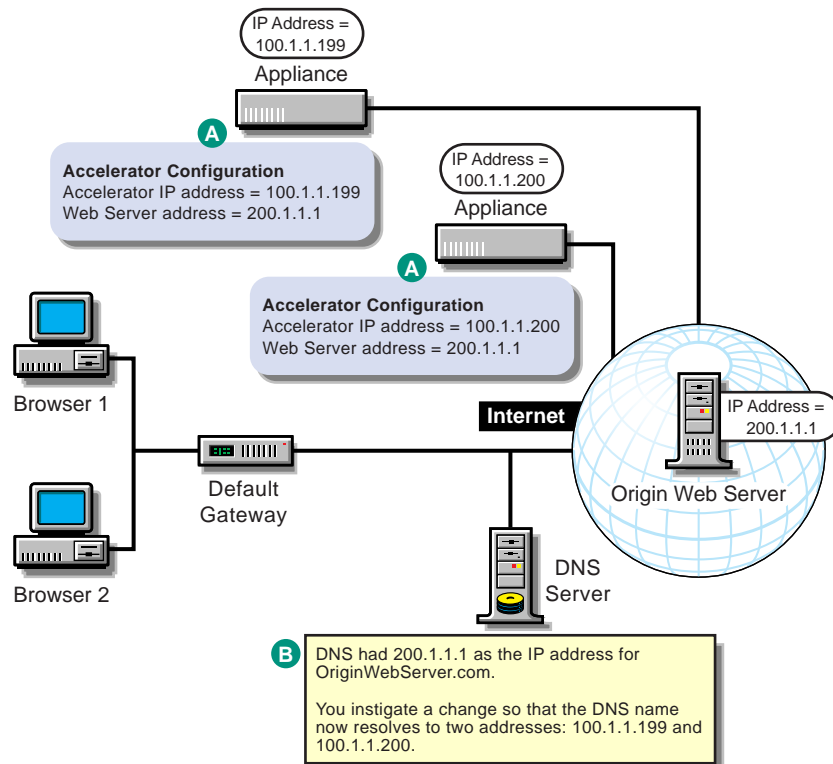
Appliance Web Server Accelerator Group Setup

Figure 47 on page 121 provides a visual map for the information in this section.

NOTE: The letters in Figure 47 are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Although Figure 47 shows the origin Web server on the Internet, it could be on the same network as the DNS server and the Web server accelerators, and the browsers could be on the Internet, as shown in “Web Server Accelerator Setup” on page 52.

Figure 47



To	Do This	Notes
Ensure your basic network configuration is complete	1. See “Basic Network Configuration Setup” on page 27.	
Set up two or more appliances as Web server accelerators to the same origin Web server	1. See “Web Server Accelerator Setup” on page 52.	See A in Figure 47 on page 121. If you want the appliances to share cached objects, configure them as ICP peers. See “ICP Hierarchy Setup” on page 111.
Modify DNS	1. Take the necessary steps to have DNS resolve requests for the origin Web server to the numeric IP addresses of appliances configured as accelerators.	See B in Figure 47 on page 121. Round-Robin DNS is a standard DNS function.

Standard Multihoming for Multiple Web Sites

Multiple Web sites can be hosted through a single IP address and port combination in two ways:

- ♦ **Multiple Web Servers Through a Single IP Address and Port:** This requires that each DNS name in the Web server accelerator definition is unique and that none of the Web servers uses SSL. (See [“Configuration Considerations When Using Appliance Multihoming Features” on page 54.](#)) To do this, you simply configure Excelsior with a Web accelerator for each DNS name.
- ♦ **Multihomed Web Servers:** You can also accelerate multihomed Web servers, which are Web servers that host multiple virtual Web servers with different DNS identities.

The following sections contain simplified examples for each scenario.

Example: Accelerating Multiple Web Servers on a Single IP Address

A company named Server Consolidation, Inc., offers acceleration services to several small companies, each of which has its own Web server. Server Consolidation knows that it can provide enough bandwidth to handle the combined Web traffic on one IP address, 30.1.1.1.

Server Consolidation installs an appliance and configures it with multiple Web server accelerators that each use the same IP address, 30.1.1.1, and fill from different Web servers.

Each company then arranges to have its DNS name (or names) resolve to 30.1.1.1.

After the DNS changes are complete, Server Consolidation’s appliance is accelerating multiple Web servers through IP address 30.1.1.1.

Because Excelsior uses the DNS name to determine which Web server to fill a request from, one accelerator definition is required for each DNS name accelerated by the appliance.

NOTE: In this example, none of the Web servers could use SSL. See [“Configuration Considerations When Using Appliance Multihoming Features” on page 54](#) for more information.

Example: Accelerating a Multihomed Web Server

A company named Web Host, Inc., provides hosting services for 50 different companies on a single Web server with an address of 10.1.1.1.

DNS resolves browser requests to each of the 50 DNS names to the same IP address, 10.1.1.1. The Web server routes each request to the appropriate area on its hard disks, based on the DNS name in the request.

Web Host decides to accelerate its Web server and installs an appliance. The company assigns 10.1.1.1 to the appliance and 10.1.1.2 to the origin Web server. It then defines an accelerator service on 10.1.1.1 that fills from 10.1.1.2.

DNS resolves requests to each of the 50 Web sites to 10.1.1.1 just as before, but the appliance now receives and services the requests, obtaining uncached objects from the origin Web server when necessary.

Because the appliance has only one accelerator defined on 10.1.1.1 and it is accelerating a single Web server on that address, it does not need the DNS name to resolve requests.

IMPORTANT: When an IP address is used for accelerating a *single Web server*, you need only one Web server acceleration service, even if the server is a multihomed server servicing browser requests to multiple DNS names.

However, if you configure Excelerator to accelerate more than one Web server on any given IP address, you must create a separate accelerator service for each DNS name that might appear in browser requests, including all the DNS names used on any multihomed servers.

We recommend you avoid the overhead of creating multiple accelerator services for a multihomed server by ensuring that its accelerator service has a dedicated IP address and port combination—one that is not used by another accelerator service.

Multihoming and Path-Based Support

The Excelerator appliance can be configured as a Web server accelerator:

- ♦ The appliance can support multiple Web server accelerators, each of which has a unique IP address and/or port number.
- ♦ The appliance can support various configurations on each of these accelerators, as summarized below.

NOTE: In the following descriptions, the term *Web server farm* refers to either a single Web server or a bank of mirrored Web servers.

No Multihoming: The Excelerator appliance accelerates a single, non-multihoming Web server farm that has one DNS name.

Host-Based Multihoming: The Excelerator appliance accelerates a single, multihoming Web server farm.

The Web server farm is configured to support multihoming (multiple DNS names). Both the appliance and the Web server farm use browser host headers to select the correct content. The appliance routes all fill traffic to this single Web server farm.

Appliance-Based Multihoming: The Excelerator appliance accelerates multiple, non-multihoming Web server farms, each of which has one unique DNS name.

The appliance uses browser host headers to select the correct cached content and to route fill requests to the correct Web server farm.

Path-Based Multihoming: The Excelerator appliance accelerates multiple, non-multihoming Web server farms from a single host name.

The appliance uses the path portion of the URL to select the correct cached content and to determine which Web server farm to fill content from.

Appliance-Based Path-Based Multihoming: The Excelerator appliance accelerates multiple, non-multihoming Web server farms.

The appliance uses both the hostname sent from the browser and the URL path to select the correct cached content and to determine which Web server farm to fill content from.

Path-based multihoming can be configured to use either the beginning or the ending portion of the path to select cached content and determine which Web server farm to fill content from.

Path-Based Multihoming Examples

Example One

The ZXY Company wants to accelerate its support and sales Web sites as a single external Web site.

The company set up two accelerators for www.zxy.org on the same IP address and port number, and configured them with path-based multihoming rules.

One accelerator has a rule for paths that start with `/sales`; the other has a rule for paths that start with `/support`.

Customers can now access the single www.zxy.org Web site. All requests starting with www.zxy.org/sales are directed to the sales Web server farm, and all requests starting with xxx.zxy.org/support are directed to the support Web server farm.

The ZXY Company can decide whether a URL such as www.zxy.org/sales/newproducts.html gets sent to the Web server with sales included in the path (www.zxy.org/sales/newproducts.html) or without sales included in the path (www.zxy.org/newproducts.html) by selecting whether the matching starting substring gets removed from the path.

Example Two

CAB Unlimited has a Web site consisting of the following components:

- ◆ A Web server farm that processes URLs ending in ASP
- ◆ A Web server farm that processes URLs ending in JPEG and GIF
- ◆ A Web server farm that processes all other URLs

CAB Unlimited sets up four accelerator services for www.cab.org on the same IP address and port, and configures them with path-based multihoming rules.

One accelerator has a rule for paths that end with ASP. The second has a rule for paths that end with JPEG. The third accelerator has a rule for paths that end with GIF. And the fourth accelerator is configured as the default for all other paths.

Browsers can now access the single www.cab.org Web site. Requests for www.cab.org/main.asp are directed to the ASP Web server farm, requests for www.cab.org/logo.gif and www.cab.com/

photo.jpg are directed to the graphics Web server farm, and requests for www.cab.com/directory.html are directed to the third Web server farm.

Domain-Based Acceleration

In Excelerator 2.3, appliance-based multihoming functionality has been extended to support Domain acceleration.

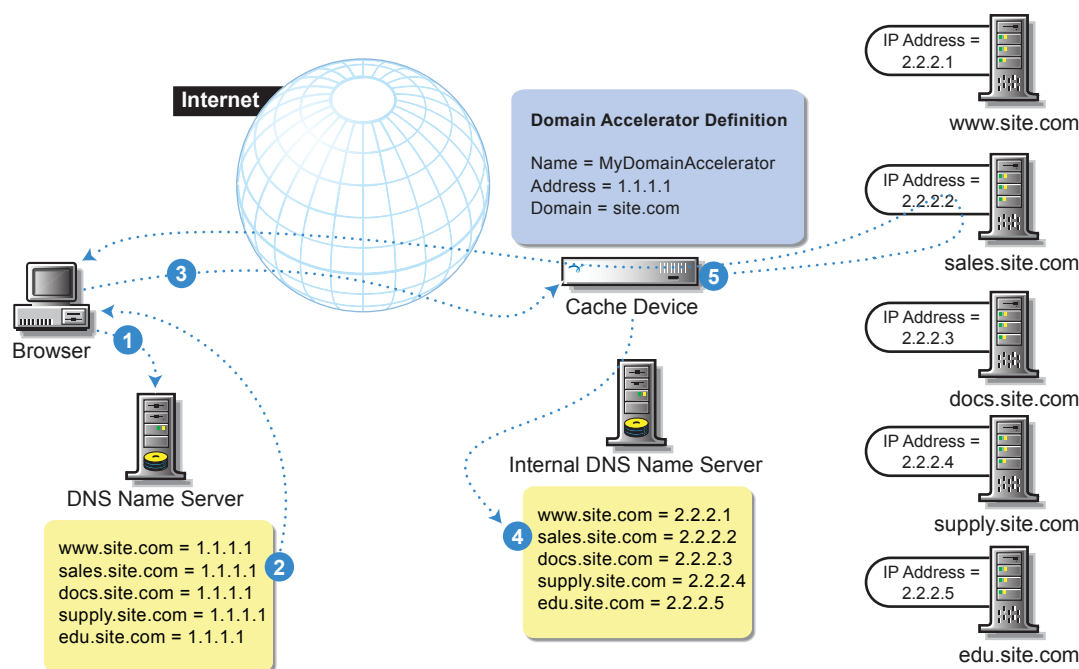
Understanding Domain-Based Acceleration

In the past, a separate accelerator service had to be created for each Web server being accelerated. Now, domain-based multihoming lets you create a single domain-based accelerator to accelerate all the Web servers in a single domain.

Graphical Overview

Figure 48 presents a basic overview of a single Excelerator cache device accelerating five Web servers in a single domain through a single domain-based acceleration service.

Figure 48



- 1 A browser requests an IP address for sales.site.com.
- 2 The DNS Name Server resolves sales.site.com to IP address 1.1.1.1 and returns the address to the browser.
- 3 The browser requests the objects from sales.site.com at IP address 1.1.1.1.
- 4 The cache device receives the request, and using its internal DNS server obtains the IP address for the sales.site.com Web server.
- 5 The cache device obtains the objects from the Web server (if they aren't already in its cache) and returns them to the browser.

Authentication Profiles and Secure Excelerator Are Supported

Domain-based multihoming also works with authentication profiles and Secure Excelerator implementations.

URL Overrides Are Not Required

Domain-based accelerators automatically change cross-referenced URLs between Web servers within a domain. This eliminates the need to create URL overrides.

Using Authentication Profiles

We recommend you create authentication profiles using the browser-based tool. Once the profiles have been created, you can assign them to domain-based accelerators using the CLI. See [Table 2 on page 127](#) for a list of CLI commands.

Once assigned, the profiles work as described for Web acceleration services.

Using Secure Excelerator

Secure Excelerator interacts with domain-based acceleration services in basically the same way as with Web server-based acceleration services. For an overview of Web server acceleration functionality, see [Chapter 10, “Accelerating Web Servers,” on page 51](#).

In addition, there are two important differences:

- ◆ Domain-based acceleration services do not require the creation of URL overrides for links between servers within the domain being accelerated
- ◆ Domain-based acceleration services can only use one trusted root for encryption. Therefore, to use the secure fill feature with a domain acceleration service, you must ensure that every Web server in the domain is using the same trusted root as the acceleration service.

For more information on certificate maintenance, see [Chapter 24, “Managing Appliance Certificates,” on page 165](#).

The Domain Name and Nodes

Domain-based accelerator services can handle up to five nodes before the domain name.

For example, an internal Web server with the domain name *svc.class1.edu.utah.provo.volera.com* has five nodes before the domain name *volera.com*.

Logging

You must use the CLI to configure logging for domain-based acceleration services. However, once the log files are being created, you can use the browser-based tool to manage the downloading and pushing of log files. See [Table 2 on page 127](#) for a list of CLI commands.

NOTE: Log files for domain-based acceleration services are listed in the REVERSE log grouping in the browser-based tool along with other accelerator service logs.

Using the CLI

[Table 2 on page 127](#) summarizes all the commands associated with domain-based acceleration.

The following sections outline some general tips for working with this feature.

Creating a Domain-Based Accelerator

Complete the following steps:

- 1

At the system prompt, create a domain accelerator using the following command:

`add domainaccelerator = name`

where *name* is the name you will use for managing the accelerator service.
- 2

Designate an IP address for handling incoming browser requests using the following command:

`set domainaccelerator name address = iii.iii.iii.iii`

where *name* is the accelerator name you specified in **Step 1** and the *i*'s are a dot-delimited IP address.
- 3

Set one or more port numbers for the service using the following command:

`set domainaccelerator name port = port`

where *name* is the accelerator name you specified in **Step 1** and *port* is the port number the service will listen on for incoming browser requests. If you are using Secure Excelerator, you will need to specify both ports for both HTTP and HTTPS traffic (for example, port 80 and port 443).

IMPORTANT: You must ensure that no other services on the cache device are using the same IP address and port combinations.
- 4

Specify the internal domain name from which the accelerator service will fill incoming browser requests, by entering the following command:

`set domainaccelerator name domainname = domain`

where *name* is the accelerator name you specified in **Step 1** and *domain* is the name of the domain from which the service will fill browser requests.
- 5

Save the settings by entering

`apply`

CLI Summary

The following table summarizes the commands used to set up and configure domain-based authentication. For detail on individual commands, refer to the CLI help.

Table 2

Feature	CLI Command
Basic Service	<code>add domainaccelerator = <i>name</i></code>
	<code>set domainaccelerator <i>name</i> enable = yes no</code>
	<code>set domainaccelerator <i>name</i> domainname = <i>domain</i></code>
	<code>set domainaccelerator <i>name</i> port = <i>port</i></code>
	<code>set domainaccelerator <i>name</i> address = <i>address</i></code>
	<code>set domainaccelerator <i>name</i> sendxforwardedforheader = yes no</code>

Feature	CLI Command
	set domainaccelerator <i>name</i> initialtcpreceivewindowsize = <i>size</i> set domainaccelerator <i>name</i> limitfillbandwidth = yes no set domainaccelerator <i>name</i> cacheobjectswithnovalidatororexpirationdate = yes no set domainaccelerator <i>name</i> browsernocacherequests = revalidate ignore refill set domainaccelerator <i>name</i> sslport = <i>port</i> set domainaccelerator <i>name</i> sslkeyid = <i>cert_name</i>
Authentication	set domainaccelerator <i>name</i> authentication enable = yes no set domainaccelerator <i>name</i> authentication profile = <i>profile</i> set domainaccelerator <i>name</i> authentication sslmaxidletime = <i>time</i> set domainaccelerator <i>name</i> authentication cookiedomain = <i>domain</i> set domainaccelerator <i>name</i> authentication profilerule = and or set domainaccelerator <i>name</i> authentication whenaclfails = no yes set domainaccelerator <i>name</i> authentication requiredonpost = no yes set domainaccelerator <i>name</i> authentication requiredonput = no yes set domainaccelerator <i>name</i> authentication requiredondelete = no yes set domainaccelerator <i>name</i> authentication requiredontrace = no yes set domainaccelerator <i>name</i> authentication requiredonoptions = no yes set domainaccelerator <i>name</i> authentication requiredonconnect = no yes set domainaccelerator <i>name</i> authentication requiredonhttp = no yes set domainaccelerator <i>name</i> authentication requiredonother = no yes
Referer Header Control	set domainaccelerator <i>name</i> referercontrol enable = yes no set domainaccelerator <i>name</i> referercontrol url = <i>url</i>
Access Control	set domainaccelerator <i>name</i> accesscontrol enable = yes no set domainaccelerator <i>name</i> accesscontrol policy = <i>policy</i>
Secure Exclerator	set domainaccelerator <i>name</i> securex enable = yes no set domainaccelerator <i>name</i> securex secureaccesstowebserver = yes no set domainaccelerator <i>name</i> securex trustedrootsfilelist = <i>list</i> set domainaccelerator <i>name</i> securex interactivehandshake = yes no set domainaccelerator <i>name</i> securex acceptanyroot = yes no

Feature	CLI Command
	set domainaccelerator <i>name</i> securex allowcacheatbrowser = yes no
Custom Cache Header	set domainaccelerator <i>name</i> customcacheheaderenable = yes no set domainaccelerator <i>name</i> customcacheheader = <i>header</i> set domainaccelerator <i>name</i> customcacheheaderonobjexpiry = checkmodified getnew
Logging - Common Format	set domainaccelerator <i>name</i> comlog = yes no set domainaccelerator <i>name</i> comlogrolloption = bysize bytime set domainaccelerator <i>name</i> comlogrollsize = <i>size</i> set domainaccelerator <i>name</i> comlogrollperiod = <i>minutes</i> set domainaccelerator <i>name</i> comlogrollday = 1stofmonth monday tuesday . . . set domainaccelerator <i>name</i> comlogrolltime = 12am 1 am 2 am . . . set domainaccelerator <i>name</i> comlogrolltimezone = local gmt set domainaccelerator <i>name</i> comlogdeleteoption = nodelete bytime byfiles set domainaccelerator <i>name</i> comlogdeletemaxtime = <i>hours</i> set domainaccelerator <i>name</i> comlogdeletemaxfile = <i>files</i>
Logging - Extended Format	set domainaccelerator <i>name</i> extlog = yes no set domainaccelerator <i>name</i> extlogrolloption = bysize bytime set domainaccelerator <i>name</i> extlogrollperiod = <i>minutes</i> set domainaccelerator <i>name</i> extlogrollday = 1stofmonth monday tuesday . . . set domainaccelerator <i>name</i> extlogrolltime = 12am 1 am 2 am . . . set domainaccelerator <i>name</i> extlogrolltimezone = local gmt set domainaccelerator <i>name</i> extlogrollsize = <i>size</i> set domainaccelerator <i>name</i> extlogdeleteoption = nodelete bytime byfiles set domainaccelerator <i>name</i> extlogdeletemaxtime = <i>hours</i> set domainaccelerator <i>name</i> extlogdeletemaxfile = <i>files</i> set domainaccelerator <i>name</i> extlogusername = yes no set domainaccelerator <i>name</i> extlogservername = yes no set domainaccelerator <i>name</i> extlogserverip = yes no

Feature	CLI Command
	set domainaccelerator <i>name</i> extlogsitename = yes no
	set domainaccelerator <i>name</i> extlogmethod = yes no
	set domainaccelerator <i>name</i> extloguri = yes no
	set domainaccelerator <i>name</i> extloguristem = yes no
	set domainaccelerator <i>name</i> extloguriquery = yes no
	set domainaccelerator <i>name</i> extloghttpversion = yes no
	set domainaccelerator <i>name</i> extloghttpstatus = yes no
	set domainaccelerator <i>name</i> extlogbytessent = yes no
	set domainaccelerator <i>name</i> extlogbytesreceived = yes no
	set domainaccelerator <i>name</i> extlogtimetaken = yes no
	set domainaccelerator <i>name</i> extloguseragent = yes no
	set domainaccelerator <i>name</i> extlogcookie = yes no
	set domainaccelerator <i>name</i> extlogreferer = yes no
	set domainaccelerator <i>name</i> extlogcachedstatus = yes no
	set domainaccelerator <i>name</i> extlogfillproxyip = yes no
	set domainaccelerator <i>name</i> extlogoriginip = yes no
	set domainaccelerator <i>name</i> extlogxwdfor = yes no
	set domainaccelerator <i>name</i> extlogcontentrange = yes no
	set domainaccelerator <i>name</i> extlogrange = yes no
	set domainaccelerator <i>name</i> extlogifrange = yes no
	set domainaccelerator <i>name</i> extlogcontentlength = yes no
	set domainaccelerator <i>name</i> extlogetag = yes no
	set domainaccelerator <i>name</i> extlogrequestpragma = yes no
	set domainaccelerator <i>name</i> extlogreplypragma = yes no
	set domainaccelerator <i>name</i> extlogcompleted = yes no
	set domainaccelerator <i>name</i> extlogheadersize = yes no
	set domainaccelerator <i>name</i> extlogcacheinfo = yes no

VI

Managing and Leveraging Excelerator's Advanced Features

As you set up your appliance and fine-tune its installation, you should be aware of the many supporting functions the appliance offers.

We recommend you review the chapters in this section and use the information in them to ensure your appliance is providing exactly the services your Web content delivery strategy requires.

The following table summarizes the tasks you can accomplish using the information in this section.

To	See
Install and manage product licenses	Chapter 21, "Installing and Upgrading Licenses," on page 133
Use appliance authentication services	Chapter 22, "Authentication Services," on page 135
Set up and implement access control policies	Chapter 23, "Access Control," on page 161
Manage appliance certificates	Chapter 24, "Managing Appliance Certificates," on page 165
Prepare internal content for Web access	Chapter 25, "Transforming Content for Internet Delivery," on page 173
Tune the appliance to meet your cache freshness requirements	Chapter 26, "Cache Freshness," on page 185
Learn about managing appliance security	Chapter 27, "Managing Appliance Security Features," on page 191.
Learn how the appliance stores various configuration settings, including those you make	Chapter 28, "Automatic Configuration Mechanisms," on page 195
Learn important information about re-imaging and restoring appliance configurations	
Learn about appliance content filtering capabilities, how they function, and how to use them	Chapter 29, "Content Filtering," on page 203
Configure the appliance so that DNS names in browser requests resolve as expected	Chapter 30, "DNS Name Resolution," on page 215
Control referred access to accelerated content by other Web sites	Chapter 31, "Controlling Referred Access to Content," on page 219

To	See
Prevent caching of errors from origin Web servers	Chapter 32, “Dynamic Bypass,” on page 231
Customize appliance error messages and add error messages for additional languages	Chapter 33, “Appliance Error Messages,” on page 233
Plan and implement a logging strategy so log files are downloaded before appliance logging space fills up	Chapter 34, “Logging,” on page 237
Use the appliance's FTP services	Chapter 35, “FTP Services,” on page 261
Keep specific objects in cache and prevent their being replaced in cache by more recently requested objects	Chapter 36, “Object Pinning,” on page 269
Set up and use appliance routing in small network installations	Chapter 37, “Router Capabilities,” on page 277
Shut down and restart the appliance	Chapter 38, “Shutting Down and Restarting,” on page 279
Use appliance SOCKS client services	Chapter 39, “SOCKS Client Services,” on page 281
Ensure appliance time is synchronized with the network	Chapter 40, “Time Synchronization,” on page 283
Use appliance WPAD capabilities to automatically configure network workstations to use appliance forward proxy services	Chapter 41, “Web Proxy Auto-Discovery (WPAD),” on page 285

21

Installing and Upgrading Licenses

Each cache device service requires that a corresponding license be installed.

If you are upgrading to from a previous version of Excelerator, you must remove all previously installed product licenses and install new licenses before you can use the product.

Also, if you've added a license to allow for a higher MB capacity on your cache device, you must remove all the lower MB licenses from the appliance before you install the new Excelerator license. For example, if you have two 256 MB licenses on your cache device, you must remove both 256 MB licenses and install a new 512 MB license to use the product.

Obtaining Product Licenses

For information on obtaining Excelerator product licenses, see “[Understanding Product Activation Licenses](#)” in the *Volera Excelerator 2.3 Getting Started Guide*.

Listing Currently Installed Licenses

To display a list of licenses currently installed on a cache device, complete the following steps:

- 1 At the System Console prompt, enter the following command:

```
get license
```
- 2 Note the name of the licenses you want to get information for or remove from the cache device.

Viewing License Information

To view the information for an installed license, complete the following steps:

- 1 At the System Console prompt, enter the following command:

```
get license license_name
```


where *license_name* is the name of a currently installed license.

Removing Licenses

To remove or uninstall a license, complete the following steps:

- 1 Enter the following command for each license file you want to remove:

```
remove license = licensename
```


where *licensename* is the name of the license file you want to remove.

- 2** If you are removing a license file as part of an upgrade, *do not restart the device* even though the system prompts you to. Restarting will cause all services that no longer have valid licenses installed to be automatically removed from the device.
- 3** If applicable, install new licenses to replace those you have removed by continuing with the next sections.

If you are not replacing the licenses you have removed, you must restart the device to complete the removal process. Any services that are dependent on the licenses you have removed are also removed from the device.

Installing Licenses Using a Floppy Disk

To install a license using a floppy disk, complete the following steps:

- 1** Copy the license files (.LIC files) to a floppy disk.
- 2** Insert the floppy disk in the cache device's disk drive.
- 3** At the System Console prompt, enter the following command:

```
importlicense floppy
```

- 4** Restart the cache device by entering the following command:

```
restart
```

Installing Licenses Using FTP

To install a license using FTP, complete the following steps:

- 1** Start an FTP client and point it to an IP address on the cache device.
For more information, see [“Starting an FTP Session with the Appliance” on page 263](#).

- 2** Copy (put) the license files (.LIC files) to the default FTP directory on the appliance.

- 3** At the System Console prompt, enter the following command:

```
importlicense
```

- 4** Restart the cache device by entering the following command:

```
restart
```

22 Authentication Services

You can control access to Excelerator proxy services by creating authentication profiles and assigning one or two of them to each service.

The following sections will help you to match Excelerator’s authentication features to your security infrastructure and requirements, and to create authentication profiles for your proxy services.

Matching Authentication Profiles to Your Requirements

Chances are good that your network already requires authentication of those seeking access to network services through a database of some kind (LDAP-compliant, RADIUS, Novell Directory Services, or NTLM). Excelerator 2.x lets you extend your authentication infrastructure to include access to proxy services.

The following sections provide information to help you match Excelerator’s authentication profile types with your network requirements.

Understanding How Profiles Work

Each profile type has different authentication mechanisms as indicated in [Table 3](#):

Table 3

Profile Type	Summary of Authentication Mechanisms
Mutual (certificate-based)	<p>To gain access to a proxy service, browsers must present information to the cache device for a certificate that has been signed by the Certificate Authority (CA) assigned to the profile.</p> <p>This method is much less secure if the profile is used alone and uses a well known trusted root.</p>

Profile Type	Summary of Authentication Mechanisms
LDAP	<p>To gain access to a proxy service, users must enter the information required by the profile, normally a valid username and password.</p> <p>The exchange of the username and password between the browser and the cache device is encrypted.</p> <p>The exchange of information between the cache device and the LDAP server can be encrypted or non-encrypted depending on whether the LDAP server's Trusted Root has been imported to the cache device.</p> <p>IMPORTANT: NDS Single-sign-on functionality is available if the LDAP-compatible database is NDS. The service is configurable only from the command line.</p>
RADIUS	<p>To gain access to a proxy service, users must enter their RADIUS username and password.</p> <p>The exchanges of username and password information between the browser, cache device, and RADIUS server are all encrypted.</p> <p>Single-sign-on functionality is configurable only from the command line.</p>
NDS	<p>To gain access to a proxy service, users must enter their Novell Directory Services username and password.</p> <p>All exchanges of username and password information between the browser, cache device, and NDS server are encrypted.</p> <p>IMPORTANT: NDS Single-sign-on functionality is configurable only from the command line.</p>
Basic	<p>To gain access to a proxy service, users must enter the information required by the profile on which the Basic profile is based (LDAP, RADIUS, or NDS).</p> <p>The exchange of username and password information between the browser and the cache device is lightly encrypted.</p> <p>The exchange between the cache device and the directory or database is the same as the profile on which the basic profile is based.</p>
NTLM	<p>To gain access to a proxy service, user must enter their NTLM username and password.</p> <p>All exchanges of user credentials are encrypted.</p> <p>Only <i>Global</i> NTLM users can authenticate using this service.</p>

A Summary of Authentication Method Pros and Cons

Excelerator 2.3 provides support for various authentication sources as summarized in [Table 4](#):

Table 4

Profile Type	Pros	Cons
Mutual (certificate-based)	<ul style="list-style-type: none"> ♦ The browser must have a certificate signed by the Certificate Authority of the profile's Trusted Root. This protects against spoofing by unauthorized persons. ♦ Authentication to the service lasts for the entire session (there is no timeout). ♦ A username and password are not required. ♦ Combining a Mutual profile with an LDAP, RADIUS, or NDS profile creates the most secure authentication method. 	<ul style="list-style-type: none"> ♦ Management of certificates on browsers can be support-intensive ♦ If the workstation is not secure or is stolen, unauthorized persons can access the service. ♦ This method is much less secure if the profile is used alone and is using a well known trusted root.
LDAP	<ul style="list-style-type: none"> ♦ LDAP is a widely used directory service protocol. ♦ This method works with Secure LDAP servers. ♦ This implementation supports context-less login. ♦ This implementation supports group access. 	LDAP trees might be viewable by unauthorized persons.
RADIUS	<ul style="list-style-type: none"> ♦ RADIUS dial-in servers are widely used. ♦ The username and password are automatically encrypted. 	RADIUS requires maintenance of various access control lists.
NDS	<ul style="list-style-type: none"> ♦ If users are currently managed using NDS, this is easy to set up. ♦ This method supports context-less login. ♦ NDS groups are supported on NetWare 5 and NetWare 6 servers through an LDAP group implementation. 	<ul style="list-style-type: none"> ♦ eDirectory (NDS) is required. ♦ eDirectory (NDS) groups are not natively supported.
Basic	<ul style="list-style-type: none"> ♦ This works with streaming media services and other applications that don't support SSL. ♦ This is a good solution for those who don't need username and password encryption. ♦ This method doesn't involve cookies. 	The username and password are more easily decrypted than with other methods.

Profile Type	Pros	Cons
NTLM	<ul style="list-style-type: none"> ♦ This method takes advantage of Microsoft NT's security scheme. ♦ If users are currently managed using NT Domains, this is easy to set up. ♦ If the user has previously logged into the Domain, authentication is transparent. ♦ This method doesn't involve cookies. 	The Domain Controller with the username and password must be on the same IP subnet as the cache device.

Combining Authentication Profiles

Depending on the profile types you are using, you might be able to assign two authentication profiles to a service for added security and/or flexibility. Only the specific profile combinations indicated in [Table 5](#) are allowed.

Table 5

Profile Type	Can Combine with
Mutual (certificate-based)	One of the following: <ul style="list-style-type: none"> ♦ LDAP ♦ RADIUS ♦ NDS
LDAP	♦ Mutual
RADIUS	♦ Mutual
NDS	♦ Mutual
Basic	Cannot be paired with other authentication profiles.
NTLM	Cannot be paired with other authentication profiles.

AND and OR Relationships Between Profiles

When you configure a service to use two authentication profiles, you must specify whether the profiles have an *AND* relationship (the user must pass both profiles' criteria to access the service) or an *OR* relationship (the user must pass only one profile's criteria to access the service).

Combining Mutual (Certificate-Based) Profiles with Other Profiles

As indicated in [Table 5](#), you can combine Mutual profile types with LDAP, RADIUS, or NDS profiles.

If the service is configured so that the profiles have an *AND* relationship, the information in the user certificate used for the Mutual profile must exactly match the information in the directory or database for the associated profile. [Table 6](#) summarizes this requirement:

Table 6

Certificate Information	Must Match LDAP Context	Must Match RADIUS Information	Must Match NDS Context
Country: USA	c=USA	ignored	c=USA
Organization: Company Name, Inc.	o=Company Name, Inc.	ignored	o=Company Name, Inc.
Department: Sales	ou=Sales	ignored	ou=Sales
Name: John Doe	cn=John Doe (Netscape iPlanet uses uid=John Doe)	A valid username in the user list	cn=John Doe

If the fields shown in [Table 6](#) don't match exactly, the user receives a 403 error indicating that the contents of the certificate doesn't match the username and password. For example, if the Country information in [Table 6](#)'s certificate were United States of America (or even U.S.A.) the user would receive the 403 error.

Combining Authentication Profiles for Same-Domain Accelerators

If you combine authentication profiles for Web server accelerators that service origin Web servers in the same DNS domain, users of the accelerator services could receive username authentication mismatch errors.

To avoid this, check your configuration against the explanation in [Table 7](#). If your configuration matches all the points in the table, follow the suggestions provided.

Table 7

If	You Must
<ul style="list-style-type: none"> You have created multiple Web Server Accelerators for Web servers that are on the same DNS domain. For example: server1.foo.com, server2.foo.com, etc. And you have created multiple username/password profiles to assign to the accelerators And two or more of the username/password profiles are of the same type (LDAP, NDS, or RADIUS) For example, you might have profiles named LDAP1, LDAP2, etc. And you want to combine the username/password profiles with mutual (certificate-based) profiles 	<ol style="list-style-type: none"> Create a separate mutual authentication profile to pair with each of the username/password profiles that are of the same type (LDAP, NDS, or RADIUS). For example, you might name the mutual profiles CERT1, CERT2, etc. Assign the username/password profiles and the mutual profiles to the accelerator services in matching pairs. For example, you would assign CERT1 with LDAP1, CERT2 with LDAP2, etc.

NOTE: The mutual authentication profiles can each use the same trusted root if desired.

Understanding How Authentication Cookies Are Used

If you are concerned that some authentication methods use cookies, you should understand the following two points about Excelerator cookie functionality.

- ♦ Authentication cookies are session-based, meaning that they expire either when the session ends or when an inactivity timeout occurs.
- ♦ Because unique authentication cookies are sent to each browser, there is no risk of global service access in NAT IP installations. In other words, each NAT client must authenticate to use the service.

The usage of cookies by authentication profiles is summarized in [Table 3](#):

Table 3

Profile Type	Cookies Used	Cookie Effective Until
Mutual (certificate-based)	Yes	Current session ends
LDAP	Yes	Specified inactivity timeout occurs
RADIUS	Yes	Specified inactivity timeout occurs
NDS	Yes	Specified inactivity timeout occurs
Basic	Yes, if assigned to a transparent service No, if assigned to a forward or Web acceleration service	Specified inactivity timeout occurs (transparent only)
NTLM	No	

Setting Up Authentication Services (Overview)

To enable authentication services, you must complete the following steps:

- 1 Define one or more authentication profiles on the appliance using the Cache > Authentication tab.

See [“Authentication Tab” on page 377](#) and the Setting Up section below for your profile type.

- 2 Enable authentication for the proxy service (Client Accelerator, Transparent Handling, Web Server Accelerator).

See [“Client Accelerator Tab” on page 333](#), [“Transparent Handling Tab” on page 342](#), and [“Web Server Accelerator Dialog Box” on page 349](#).

IMPORTANT: Excelerator requires each SSL service (including authentication) to use a unique IP address and port combination. The default authentication port is 443. Attempts to enable authentication for more than one proxy service on the same IP address and port will result in a TCP bind error.

- 3 Select one or more profile(s) to use for each service.

See [“Add Authentication Profiles Dialog Box” on page 338](#).

- 4 Inform users regarding the following:

- ♦ Any workstation and/or browser preparation they must make, such as installing certificates
- ♦ The steps they will need to complete to log in and use the service
- ♦ What they should expect regarding inactivity timeouts, etc. while using the service

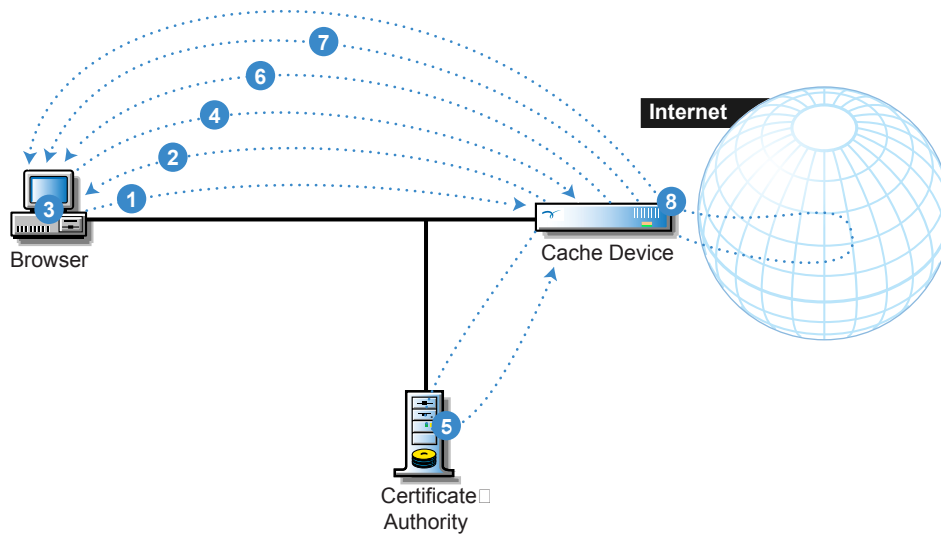
Using Mutual (Certificate-Based) Authentication

Use the information in this section to understand, create, and use mutual authentication profiles.

How Mutual Authentication Works

Figure 49 illustrates how mutual authentication can be used to control access to proxy services.

Figure 49



1. A browser requests access to a Web page through service on a cache device.
2. The cache device requests information for a Client certificate signed by the Certificate Authority (CA) listed in the profile assigned to the service.
3. The browser displays certificates that match the CA to the user.
4. The user chooses a certificate and the browser sends the certificate's public key to the cache device.
5. The cache device verifies the certificate with the CA.
6. If the certificate is verified, the process continues with Step 6. If the certificate is not verified, the cache device sends a 403 error to the browser indicating that the certificate cannot be verified.
7. The cache device sends a session-based cookie to the browser.
8. The cache device returns the requested Web page to the browser.
9. Because subsequent browser requests contain the session cookie, reauthentication is not required again during the session.

Platforms Requirements

The following table summarizes the platform requirements for mutual authentication:

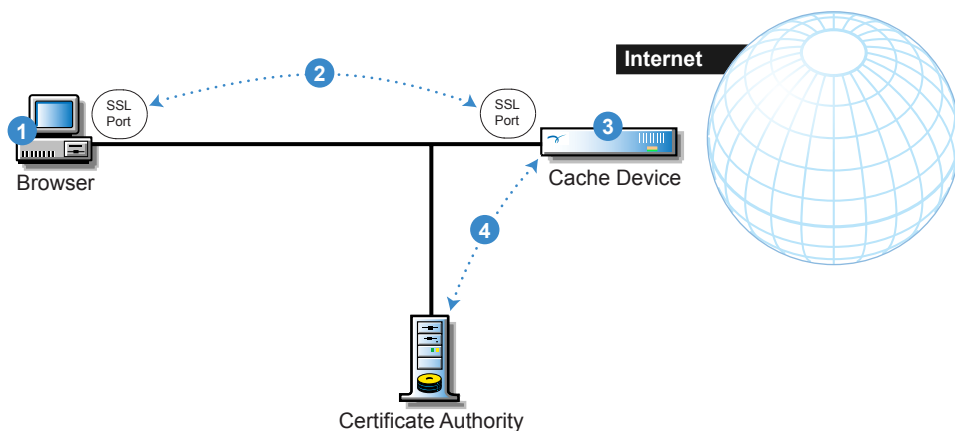
Table 4

Network Component	Software Requirements
Workstation	Any SSL-capable Internet browser
Cache Device	Excelerator 2.1 or later
Certificate Authority Server	The Certificate Authority that issued the Client certificate to the browser

Preparing Your Network for Mutual Authentication

Figure 50 summarizes the configuration requirements for mutual authentication:

Figure 50



1 ☐ The browser must have both the trusted root of the
☐ Certificate Authority (CA) and a user certificate
☐ issued by the CA.

2 ☐ The browser and the cache device must be able
☐ to communicate using an SSL connection.

3 ☐ The cache device must have the trusted root of
☐ the CA for the profile being used by the service.

4 ☐ The cache device must have network access to
☐ the CA.

NOTE: If the profile will be used in combination with another profile, make sure the information in each client certificate meets the requirements outlined in [“Combining Mutual \(Certificate-Based\) Profiles with Other Profiles”](#) on page 138.

Setting Up Mutual Authentication

After you have completed the steps in [“Preparing Your Network for Mutual Authentication”](#) on page 142, you can set up mutual authentication by completing the instructions in the following sections.

Creating a Mutual Authentication Profile

1 In the browser-based management tool, click Cache > Authentication > Insert.

2 Type a name for the profile in the Authentication Profile Name field.

IMPORTANT: Each profile name created on a cache device must be unique. Excelerator doesn't recognize case differences (MyProfile and myprofile are the same name to Excelerator) and it will

overwrite and concatenate previously created profiles without warning if a duplicate name is used. For more information, see [“Authentication Dialog Box” on page 378](#).

- 3** Check Mutual Authentication > click Options.
- 4** If the List of Trusted Roots contains the trusted root for the Client certificate the workstations will use, select the trusted root > click Insert.
- 5** If the List of Trusted Roots doesn’t contain the trusted root for the Client certificate the workstations will use, click Import Trusted Root and import the appropriate trusted root.
For information on importing trusted root files, see [Chapter 24, “Managing Appliance Certificates,” on page 165](#), specifically [“Importing a Trusted Root to a Cache Device” on page 169](#).
- 6** Click OK > OK.
- 7** Assign the profile to one or more proxy services as described in each service tab section in [“Using the Cache Panel” on page 333](#).

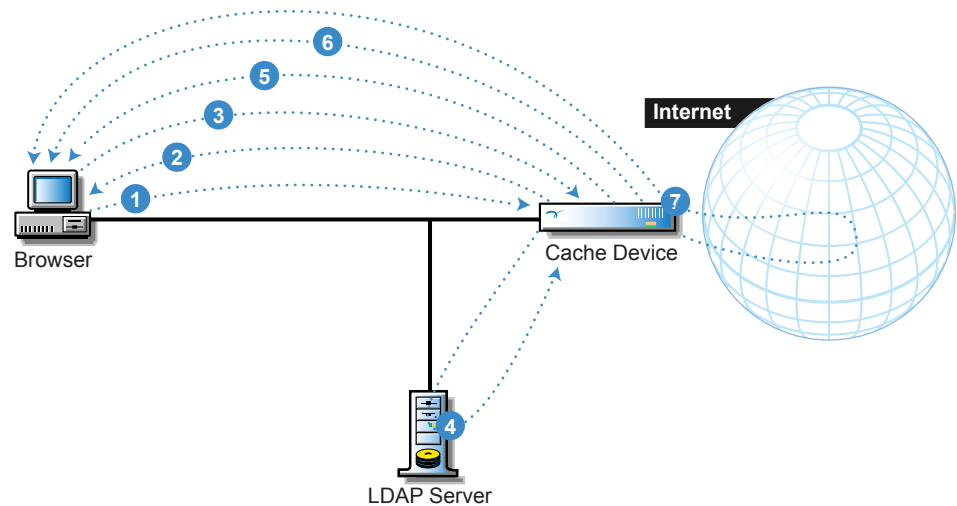
Using LDAP Authentication

Use the information in this section to understand, create, and use LDAP authentication profiles.

How LDAP Authentication Works

[Figure 51](#) illustrates how LDAP authentication can be used to control access to proxy services

Figure 51



- 1. A browser requests access to a Web page through service on a cache device.
- 2. The cache device sends the LDAP authentication form to the browser over a secure channel.
- 3. The user provides the username and password information.
- 4. The cache device verifies the username and password with the LDAP server.
 - If the information is verified, the process continues with Step 5.
 - If the information is not verified, the cache device resends the authentication form with a message that the login attempt failed.
- 5. The cache device sends a session cookie to the browser with an inactivity timeout limitation.
- 6. The cache device returns the requested Web page to the browser.
- 8. Subsequent browser requests contain the session cookie, and reauthentication is not required again unless an inactivity timeout occurs.

Platforms Supported

The following table summarizes the platforms supported for LDAP authentication:

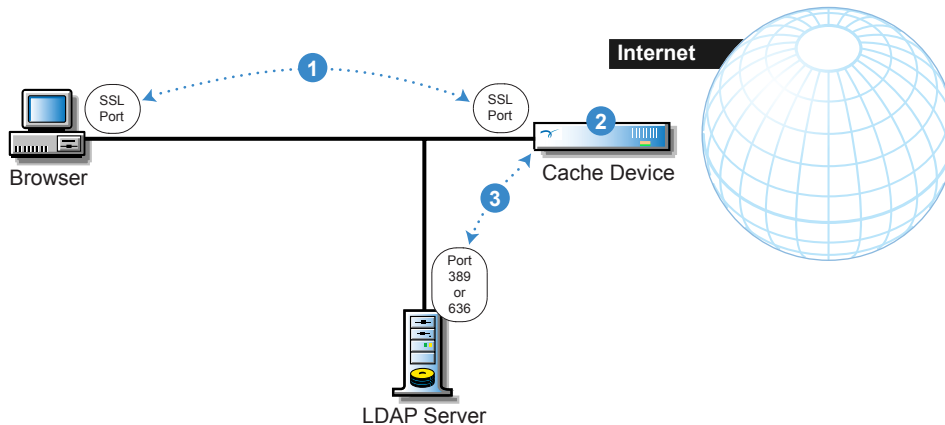
Table 5

Network Component	Software Requirements
Workstation	Any SSL-capable Internet browser (For NDS single sign-on, Windows 95, 98, NT, or 2000)
Cache Device	Excelerator 2.1 and later
LDAP server	LDAP-compliant database as specified in the profile

Preparing Your Network for LDAP Authentication

Figure 52 summarizes the configuration requirements for LDAP authentication:

Figure 52



- 1 ☐ The browser and the cache device must be able to communicate using an SSL connection.
- 2 ☐ If Secure LDAP is enabled, the cache device must have the trusted root of the LDAP server used by the service.
- 3 ☐ The cache device must be able to communicate with the LDAP server using port 389 (or port 636 if using Secure LDAP).

Setting Up LDAP Authentication

After you have completed the steps in [“Preparing Your Network for LDAP Authentication” on page 144](#), you can set up an LDAP authentication profile by completing the steps in the following sections.

Creating an LDAP Profile

Complete the following steps:

- 1 Define an authentication profile by clicking Cache > Authentication > Insert.
- 2 Type a name for the profile in the Authentication Profile Name field.
IMPORTANT: Each profile name created on a cache device must be unique. Excelerator doesn't recognize case differences (MyProfile and myprofile are the same name to Excelerator) and it will overwrite and concatenate previously created profiles without warning if a duplicate name is used. For more information, see [“Authentication Dialog Box” on page 378](#).
- 3 Select LDAP Authentication > click LDAP Options.
- 4 Specify the IP address of the server containing the LDAP-compliant directory in the LDAP Server Address field.
- 5 Type the port number on which the LDAP server will listen for requests from the cache device. The default ports are: 389 for non-secure access and 636 for secure (SSL) access.
- 6 If the cache device and the server will communicate using SSL, check Enable Secure Access to LDAP Server > click Import Trusted Root > complete the instructions in [“Importing a Trusted Root to a Cache Device” on page 169](#).

NOTE: Once you have imported a Trusted Root for one LDAP profile, you can use the same file for multiple LDAP profiles by typing the filename in the LDAP Server Trusted Root File field.

- 7 Select an LDAP Login Name Format and complete the instructions in the applicable following section: “Use User’s E-Mail” on page 146, “Use Distinguished Name” on page 146, or “Use Field Name” on page 146.

Use User’s E-Mail

This options lets users authenticate using their e-mail name stored in the LDAP database. You must specify the LDAP containers from which the e-mail name search should begin and the method the cache device should use to communicate with the LDAP server.

- 1 In the Use User’s E-mail dialog, create an LDAP search base by clicking Insert and typing an LDAP container from which the e-mail name search should begin.
- 2 Insert additional LDAP containers in the search base as required.
- 3 If the cache device can authenticate to the LDAP server using anonymous bind, click Use Anonymous Bind for LDAP Search.
- 4 If anonymous bind is not enabled on the LDAP server, click Use User Name/Password Bind for LDAP Search > enter the username and password pair through which the appliance will authenticate to the LDAP server before requesting the search.
- 5 If you plan to use LDAP groups, complete the instructions in “Enabling and Using LDAP Groups” on page 147. Otherwise, click OK to create the LDAP authentication profile.

Use Distinguished Name

This option lets users authenticate using their LDAP usernames. You can have users enter their fully distinguished (full LDAP context) usernames, or you can provide a list of LDAP contexts so they need only type their usernames.

- 1 In the Use Distinguished Name dialog, specify the field name your LDAP directory uses for username information.
NOTE: Netscape’s iPlanet directory stores usernames in the UID field. Most other LDAP-compliant directories use the CN field. If no value is entered, CN is used by default.
- 2 If you want users to be able to log in using only their usernames, insert each of the LDAP contexts of the users who will be authenticating through the authentication profile.
- 3 If you plan to use LDAP groups, complete the instructions in “Enabling and Using LDAP Groups” on page 147. Otherwise, click OK to create the LDAP authentication profile.
- 4 Click OK.

Use Field Name

This options lets users authenticate using a designated field name stored in the LDAP database. You must specify the field name to be used, the LDAP containers from which the field name search should begin, and the method the cache device should use to communicate with the LDAP server.

- 1 In the Use Field Name dialog, type the LDAP field name which users will use to authenticate
- 2 Create an LDAP search base by clicking Insert and typing an LDAP container from which the field name search should begin.
- 3 Insert additional LDAP containers in the search base as required.
- 4 If the cache device can authenticate to the LDAP server using anonymous bind, click Use Anonymous Bind for LDAP Search.

- 5 If anonymous bind is not enabled on the LDAP server, click Use User Name/Password Bind for LDAP Search > enter the username and password pair through which the appliance will authenticate to the LDAP server before requesting the search.
- 6 If you plan to use LDAP groups, continue with [Enabling and Using LDAP Groups](#). Otherwise, click OK to create the LDAP authentication profile.

Enabling and Using LDAP Groups

You can designate LDAP groups for authentication to Excelerator proxy services by including the LDAP context (parent container) for target groups. Users who are members of the groups will be able to authenticate using only their username.

Designating the Group Class and/or Attribute Name

Each LDAP-compliant directory uses a different mechanism for implementing group support. If you plan to set access control based on LDAP groups, you must also specify how the target directory's schema defines groups.

- 1 In the browser-based management tool, click Cache > Authentication > Insert > LDAP Authentication > Options.

The two fields, *LDAP Object Class Group Name* and *LDAP User Attribute Member*, tell Excelerator the mechanism the target directory's schema uses to designate an LDAP group.

For example, Active Directory uses the LDAP Object Class Group Name *group* and Novell Directory Services uses the name *groupofnames*.

- 2 If the LDAP group object class name is something other than *groupofnames* (the name used by Novell Directory Services), enter the object class name in the LDAP Object Class Group Name field.

For example, for Active Directory you must enter the name *group*.

- 3 Enter the user object attribute name designating group membership in this field.

For example, Active Directory uses *memberof* and Novell Directory Services uses *groupmembership*.

This field is required for all LDAP group implementations.

- 4 After specifying the required group information, click OK to create the LDAP authentication profile.
- 5 Assign the profile to one or more proxy services as described in each service tab section in ["Using the Cache Panel" on page 333](#).

Enabling NDS Single Sign-On for an LDAP Authentication Profile

If your LDAP-compatible directory is NDS e-Directory, you can enable NDS single sign-on by completing the following steps:

- 1 Complete the instructions in ["Setting Up and Enabling NDS \(eDirectory\) Single Sign-On" on page 154](#), then return to this procedure.
- 2 At the cache device's System prompt, enter the following commands:

```
set authentication name ldap tryndssinglesignon=yes
set authentication name ldap ndssinglesignonreplytime=seconds
set authentication name ldap ndssinglesignonnoresponsetime=seconds
```

where *name* is the name of the LDAP profile and *seconds* represents the time the service will wait for responses from the Novell client and NDS server, respectively.

- 3
- At the command line, enter
- apply

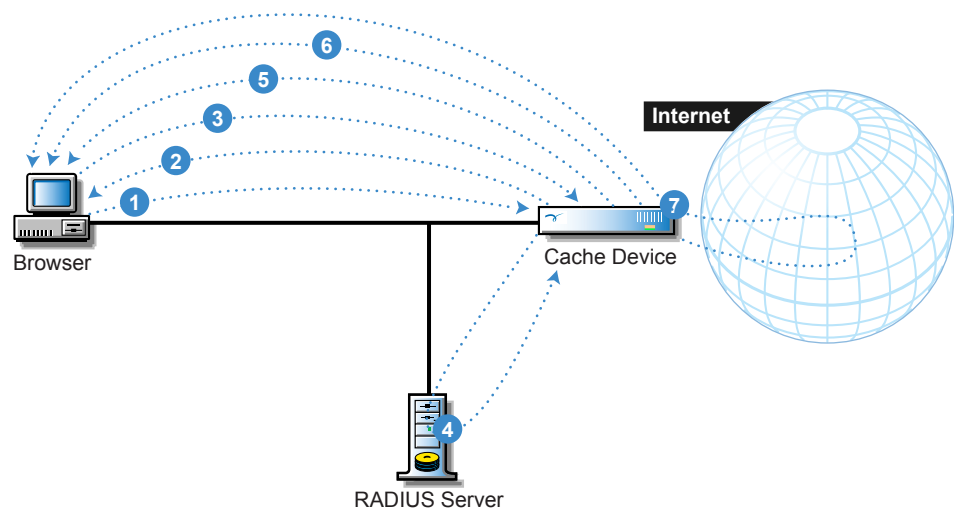
Using RADIUS Authentication

Use the information in this section to understand, create, and use RADIUS authentication profiles.

How RADIUS Authentication Works

Figure 53 illustrates how RADIUS authentication can be used to control access to proxy services

Figure 53



- 1
- A browser requests access to a Web page through service on a cache device.
- 2
- The cache device sends the RADIUS authentication form to the browser over a secure channel.
- 3
- The user provides the username and password information.
- 4
- The cache device verifies the username and password with the RADIUS server.
- 5
- The cache device sends a session cookie to the browser with an inactivity timeout limitation.
- 6
- The cache device returns the requested Web page to the browser.
- 7
- Subsequent browser requests contain the session cookie, and reauthentication is not required again unless an inactivity timeout occurs.
- 8
- If the information is not verified, the cache device resends the authentication form with a message that the login attempt failed.

Platforms Supported

The following table summarizes the platforms supported for RADIUS authentication profiles:

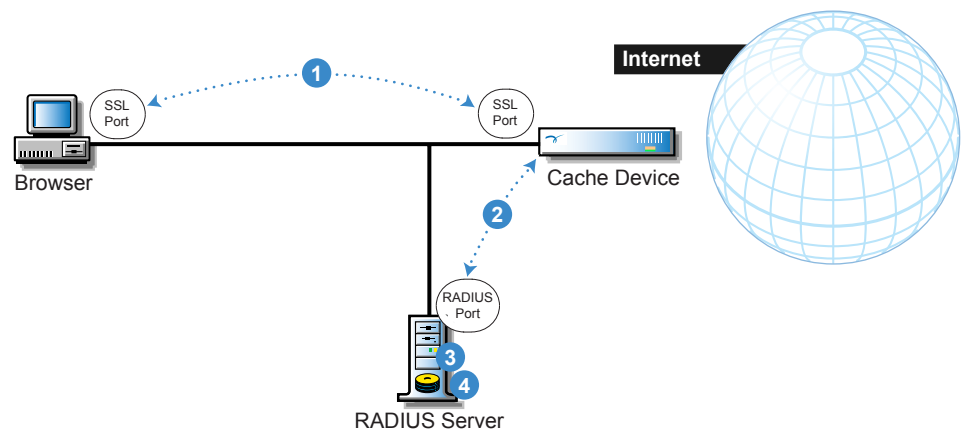
Table 6

Network Component	Software Requirements
Workstation	Any SSL-capable Internet browser
Cache Device	Excelerator 2.x
RADIUS Server	RADIUS database as specified in the authentication profile

Preparing Your Network for RADIUS Authentication

Figure 54 summarizes the configuration requirements for RADIUS authentication:

Figure 54



- 1 The browser and the cache device must be able to communicate using an SSL connection.
- 2 The cache device must be able to communicate with the RADIUS server on the port it is using for RADIUS communications. For RADIUS version 1 this is normally port 1645, for version 2 it is normally 1812.
- 3 The user list on the RADIUS server must have an entry for the user requesting access.
- 4 The client list on the RADIUS server must have an entry for the cache device's IP address and Shared Secret.

Setting Up RADIUS Authentication

After you have completed the steps in “Preparing Your Network for RADIUS Authentication” on page 149, you can set up a RADIUS authentication profile by completing the following procedure.

- 1 In the browser-based management tool, click Cache > Authentication > Insert.
- 2 Type a name for the profile in the Authentication Profile Name field.
IMPORTANT: Each profile name created on a cache device must be unique. Excelerator doesn't recognize case differences (MyProfile and myprofile are the same name to Excelerator) and it will overwrite and concatenate previously created profiles without warning if a duplicate name is used. For more information, see “Authentication Dialog Box” on page 378.
- 3 Check RADIUS Authentication > click Options.
- 4 In the RADIUS Server Address field, type the IP address of the RADIUS server.

- 5** In the RADIUS Server Listening Port field, type the port number on which the server will listen for incoming authentication requests.
- 6** In the RADIUS Shared Secret field, type the string the RADIUS server uses to verify that the appliance can request authentication of users.
- 7** Click OK > OK.
- 8** Assign the profile to one or more proxy services as described in each service tab section in [“Using the Cache Panel” on page 333](#).

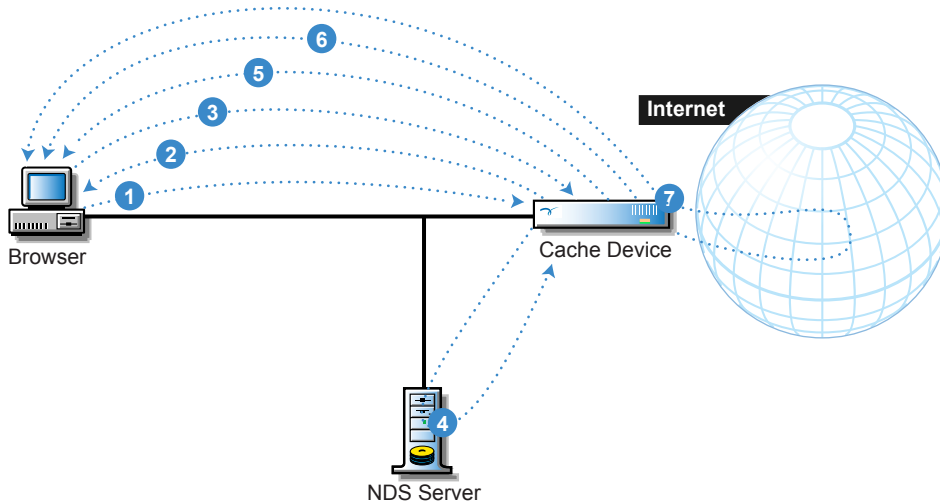
Using NDS (eDirectory) Authentication

Use the information in this section to understand, create, and use NDS authentication profiles.

How NDS (eDirectory) Authentication Works

Figure 55 illustrates how NDS authentication can be used to control access to proxy services

Figure 55



- | | |
|---|---|
| <p>1 A browser requests access to a Web page through service on a cache device.</p> <p>2 The cache device sends the NDS authentication form to the browser over a secure channel.</p> <p>3 The user provides the username and password information.</p> <p>4 The cache device verifies the username and password with the NDS server.</p> <p>If the information is verified, the process continues with Step 5.</p> | <p>If the information is not verified, the cache device resends the authentication form with a message that the login attempt failed.</p> <p>5 The cache device sends a session cookie to the browser with an inactivity timeout limitation.</p> <p>6 The cache device returns the requested Web page to the browser.</p> <p>8 Subsequent browser requests contain the session cookie, and reauthentication is not required again unless an inactivity timeout occurs.</p> |
|---|---|

Platforms Supported

The following table summarizes the platforms supported for NDS authentication:

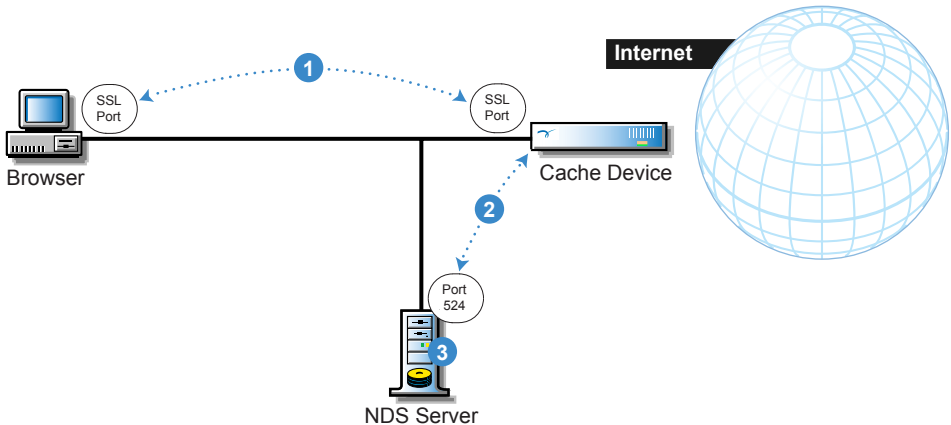
Table 7

Network Component	Software Requirements
Workstation	An SSL-capable Internet browser (For NDS single sign-on, Windows 95, 98, NT, or 2000)
Cache Device	Excelerator 2.x
NDS Database	NetWare 5 or later

Preparing Your Network for NDS (eDirectory) Authentication

Figure 56 summarizes the configuration requirements for NDS authentication:

Figure 56



- 1□ The browser and the cache device must be able to communicate using an SSL connection.
- 2□ The cache device must be able to communicate with the NDS server using port 524.
- 3□ The DNS directory must contain an entry for the user requesting access.

Setting Up NDS (eDirectory) Authentication

After you have completed the steps in “Preparing Your Network for NDS (eDirectory) Authentication” on page 151, you can set up an NDS authentication profile by completing the following procedure.

- 1 In the browser-based management tool, click Cache > Authentication > Insert.
- 2 Type a name for the profile in the Authentication Profile Name field.

IMPORTANT: Each profile name created on a cache device must be unique. Excelerator doesn't recognize case differences (MyProfile and myprofile are the same name to Excelerator) and it will overwrite and concatenate previously created profiles without warning if a duplicate name is used. For more information, see “Authentication Dialog Box” on page 378.

- 3 Check NDS Authentication > click Options.

- 4** In the NDS Server Address field, type the IP address of the eDirectory server (can be NetWare, Microsoft Windows 2000/NT, Linux, or Solaris) with the read/write or master NDS partition to which users will authenticate.
- 5** In the Users' Default Context List, include the contexts for all users who will be authenticating by clicking Insert and typing the NDS Context and NDS tree for each context.
- 6** When you have inserted all the contexts, click OK.
- 7** Click OK > OK.
- 8** Assign the profile to one or more proxy services as described in each service tab section in [“Using the Cache Panel” on page 333](#).

Enabling NDS Single Sign-On for an NDS Authentication Profile

You can enable NDS single sign-on by completing the following steps:

- 1** Complete the instructions in [“Setting Up and Enabling NDS \(eDirectory\) Single Sign-On” on page 154](#), then return to this procedure.
- 2** At the cache device's System prompt, enter the following commands:

```
set authentication name nds tryndssinglesignon=yes  
set authentication name nds ndssinglesignonreplytime=seconds  
set authentication name nds ndssinglesignonnoresponsesettime=seconds
```

where *name* is the name of the NDS profile and *seconds* represents the time the service will wait for responses from the Novell client and NDS server, respectively.
- 3** At the command line, enter

```
apply
```


Using NDS (eDirectory) Single Sign-On Functionality

NDS single sign-on authentication is not a separate profile type. Rather it is additional functionality enabled as part of either an LDAP authentication profile that uses an NDS database or an NDS authentication profile.

How NDS (eDirectory) Single Sign-On Works

[Figure 55](#) illustrates how NDS single sign-on functionality works

Figure 57



- 1 An NDS user on a Novell Client 32 workstation requests access to a proxy service that has an NDS single sign-on authentication profile associated with it.

2 The cache device requests confirmation from the NetWare server that the user is currently authenticated to NDS.

 - If the user is authenticated to the NDS partition on the profile-specified NetWare server, the process skips to Step 5.
 - If the user is authenticated to an NDS partition in the same tree on a different NetWare server, the process skips to Step 5.

3 If the user is not currently authenticated to the tree, the NetWare server notifies the cache device, which sends a 403 error to the browser indicating that access is restricted to authenticated NDS users.
- You can prevent users' receiving 403 errors by creating additional authentication profile types and configuring the service to use them in an "or" relationship with the NDS single sign-on profile.

4 The client workstation confirms its authentication status with the NetWare server using background authentication.

5 The NetWare server confirms the user's authentication status with the cache device.

6 The cache device processes the browser request as usual.

7 Subsequent user requests don't require reauthentication.

Platforms Supported

The following table summarizes the platforms supported for NDS authentication:

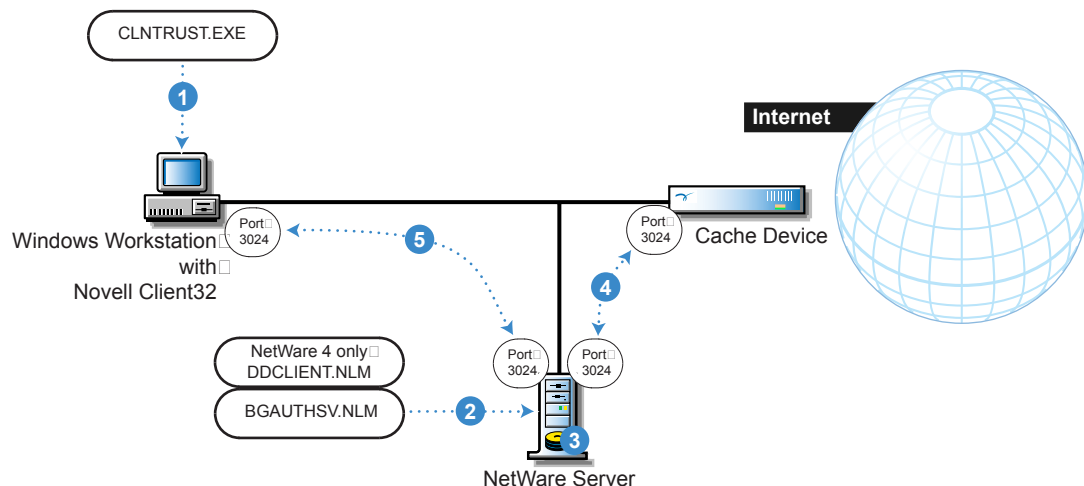
Table 8

Network Component	Software Requirements
Workstation	An SSL-capable Internet browser (For NDS single sign-on, Windows 95, 98, NT, or 2000)
Cache Device	Excelerator 2.x
NDS Database	NetWare 4 or later

Preparing Your Network for NDS (eDirectory) Authentication

Figure 56 summarizes the configuration requirements for NDS authentication:

Figure 58



1 CLNTRUST.EXE must be running on every Novell Client32 workstation using NDS single sign-on.

2 DDCLIENT.NLM must be loaded on all NetWare 4 servers and BGAUTHSV.NLM must be loaded on all NetWare 4, 5, and 6 servers providing NDS authentication.

3 The NDS partition on the NetWare server providing NDS authentication must be either a read/write or a master partition.

4 The cache device and the NetWare server must be able to communicate using UDP through port 3024.

5 The NetWare server and the Client32 workstation must be able to communicate using UDP through port 3024.

Setting Up and Enabling NDS (eDirectory) Single Sign-On

Complete the following steps to set up and enable NDS Single Sign-on:

1 Create one of the following:

- ♦ An NDS authentication profile
- ♦ An LDAP profile that references an NDS database

2 Using an FTP client, access the Excelerator 2.3 device's default FTP directory (/etc/proxy/appliance/config/user) and retrieve copies of the following files:

- ♦ CLNTRUST.EXE
- ♦ BGAUTHSV.NLM
- ♦ DDCLIENT.NLM

3 Install CLNTRUST.EXE on every Novell Client32 workstation that you want to enable for single sign-on authentication.

4 Copy BGAUTHSV.NLM to the SYS:\SYSTEM directory on every NetWare 4, 5, or 6 server running the NDS database to which users will authenticate.

IMPORTANT: BGAUTHSV cannot be run on a NetWare server that is running any version of Novell Border Manager.

5 Copy DDCLIENT.NLM to the SYS:\SYSTEM directory on every NetWare 4 server running the NDS database to which users will authenticate.

- 6 Load the NLMs you copied to the NetWare servers and include references to them in the servers' AUTOEXEC.NCF files.
- 7 Complete the relevant instructions for enabling NDS single sign-on:
 - ♦ “Enabling NDS Single Sign-On for an LDAP Authentication Profile” on page 147.
 - Or
 - ♦ “Enabling NDS Single Sign-On for an NDS Authentication Profile” on page 152.

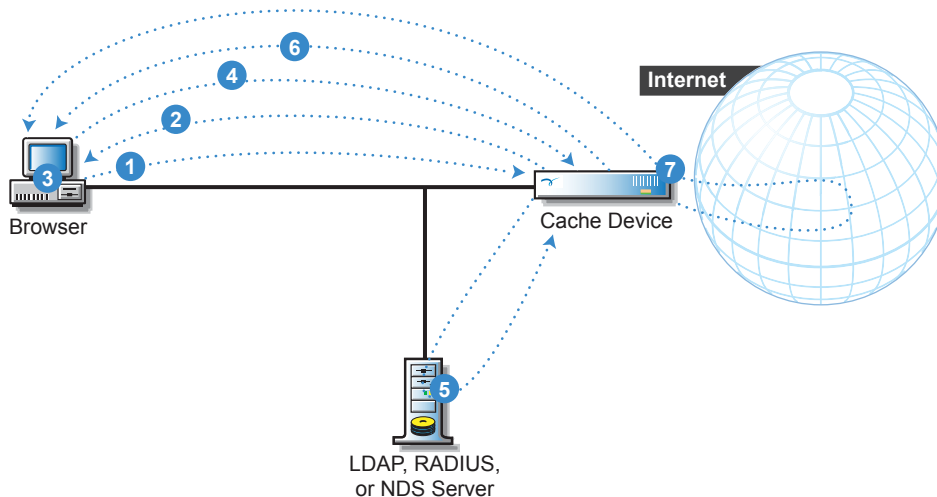
Using Basic Authentication

Use the information in this section to understand, create, and use Basic authentication profiles.

How Basic Authentication Works

Figure 59 illustrates how basic authentication can be used to control access to proxy services:

Figure 59



- 1 A browser requests access to a Web page through forward, transparent, or Web acceleration service on a cache device.
- 2 The cache device returns an Authorization-Required header.
- 3 The browser prompts the user for a username and password.
- 4 The browser sends the username and password to the cache device as part of the Authorization header of the request.
- 5 The cache device verifies the username and password with the LDAP, RADIUS, or NDS server specified in the Basic authentication profile.

If the information is verified, the process continues with Step 6.

If the information is not verified, the cache device allows two retries by the browser and user.

If the retries are not successful, the cache device sends a 407 error to the browser indicating that authentication is required.

- 6 The cache device returns the requested Web page to the browser.
If the profile is assigned to a transparent service, the device also sends a session-based cookie to the browser.
- 7 Because subsequent browser requests contain either a Basic authentication header (for forward and Web acceleration services) or a session-based cookie (for transparent services), reauthentication is not required again during the session.

Platforms Supported

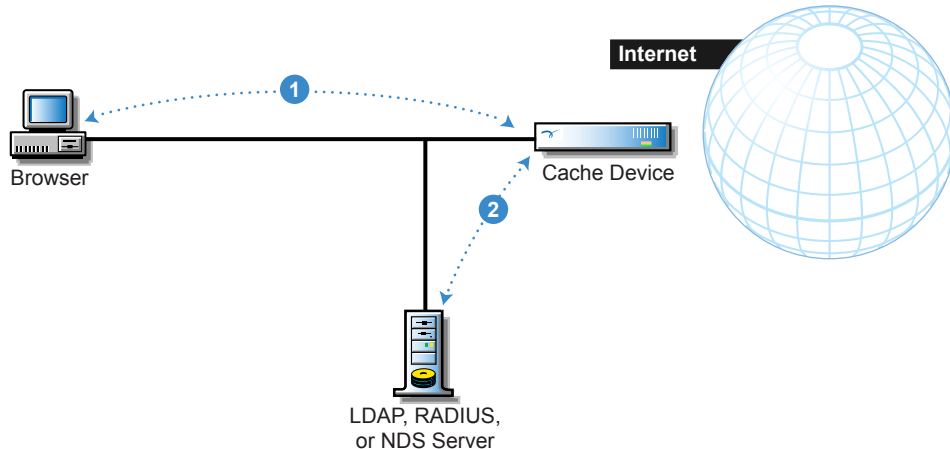
Platform requirements are dictated by the profile on which the Basic profile is based (LDAP, RADIUS, or NDS).

Preparing for Basic Authentication

Setting Up the Network

Figure 60 summarizes the configuration requirements for basic authentication:

Figure 60



1 ☐ The browser and the cache device must be able to communicate using HTTP.

2 ☐ The cache device must be able to communicate with the server specified in the authentication profile.

Creating the Underlying Profiles

Basic authentication profiles are based on previously created LDAP, RADIUS, or NDS profiles. You must, therefore, create the underlying profile before you create the basic profile that uses it.

Setting Up Basic Authentication

IMPORTANT: You cannot assign a basic authentication profile to a Web Server Accelerator (reverse proxy) for an origin Web server that is also using basic authentication. Authentication to the proxy service will succeed, but authentication to the origin Web server will fail.

After you have completed the steps in “[Preparing for Basic Authentication](#)” on page 156, you can set up a basic authentication profile by completing the following procedure.

1 In the browser-based management tool, click Cache > Authentication > Insert.

2 Type a name for the profile in the Authentication Profile Name field.

IMPORTANT: Each profile name created on a cache device must be unique. Excelerator doesn't recognize case differences (MyProfile and myprofile are the same name to Excelerator) and it will overwrite and concatenate previously created profiles without warning if a duplicate name is used. For more information, see “[Authentication Dialog Box](#)” on page 378.

3 Check Basic Authentication > click Options.

- 4 In the Authentication profile drop-down list, select the profile the Basic authentication profile will use.
- 5 Click OK > OK.
- 6 Assign the profile to one or more proxy services as described in each service tab section in [“Using the Cache Panel” on page 333](#).

Using NTLM Authentication

Use the information in this section to understand, create, and use NTLM authentication profiles.

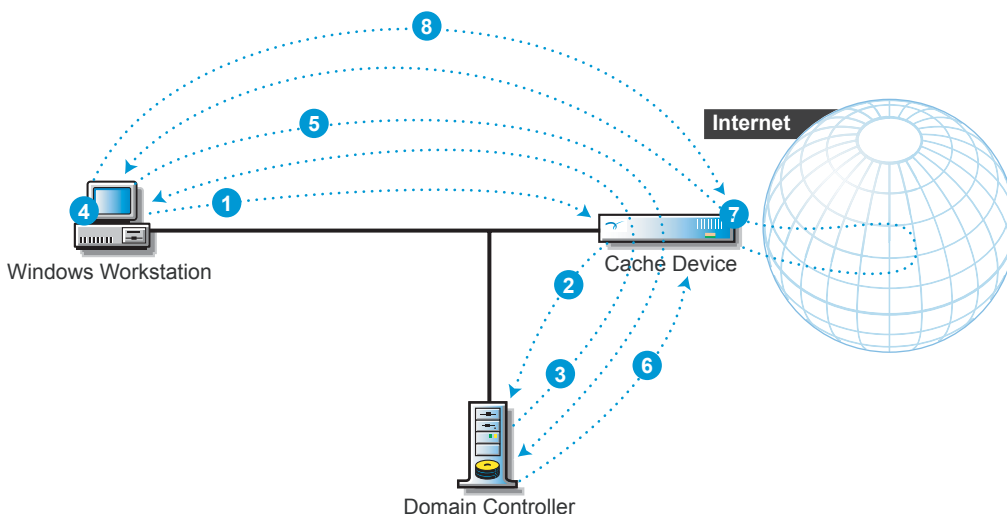
IMPORTANT: NTLM authentication profiles can only be assigned to forward proxy services.

Exceleator supports pass-through NTLM authentication for Web servers that require NTLM authentication.

How NTLM Authentication Works

Figure 61 illustrates how NTLM authentication can be used to control access to a forward proxy service.

Figure 61



- 1 A Windows user requests access to a proxy service that has an NTLM authentication profile associated with it, using a browser that supports NTLM authentication, such as Internet Explorer.
- 2 The cache device requests confirmation from the Domain Controllers that the user is currently authenticated to the Domain.

If the user is authenticated to the Domain on the profile-specified Domain Controller, the process skips to Step 6.
- 3 If the user is not currently authenticated to the Domain or the browser doesn't support NTLM authentication, the Domain Controller notifies the cache device, which sends a request for NTLM user credentials to the browser.
- 4 The user provides the username and password information to the browser.
- 5 The workstation encrypts the user information, the browser sends the credentials to the cache device, and the cache device forwards the credentials to the Domain Controller.
- 6 The Domain Controller confirms the user's authentication status with the cache device.
- 7 The cache device processes the browser request as usual.
- 8 Subsequent requests from this user don't require reauthentication.

Platforms Supported

The following table summarizes the platforms supported for NTLM authentication:

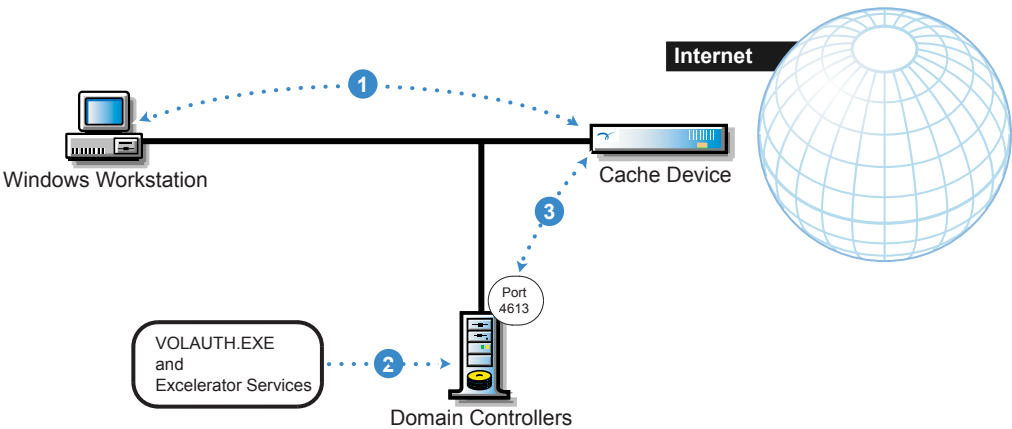
Table 9

Network Component	Software Requirements
Workstation	<ul style="list-style-type: none">Windows NT, 2000Internet Explorer 5.x or later orNetscape, Mozilla, and Opera browsers if failover support is sufficient (see “How NTLM Authentication Works” on page 157).
Cache Device	<ul style="list-style-type: none">Excelerator 2.3
Domain Controller	NT Domain database as specified in the profile

Preparing Your Network for NTLM Authentication

Figure 62 summarizes the configuration requirements for NTLM authentication:

Figure 62



- 1□The browser and the cache device must be able to communicate using HTTP.

2□The Domain Controller must have VOLAUTH.EXE installed and be running Excelerator Services
- 3□The cache device must be able to communicate with the Domain Controllers using port 4613.

Complete the following Steps

- 1 Using an FTP client, access the Excelerator 2.3 device’s default FTP directory (/etc/proxy/appliance/config/user) and retrieve the VOLAUTH.EXE file.

2 Copy VOLAUTH.EXE to the WINNT directory on each Domain Controller.

The VOLAUTH.EXE file is also located on the [Volera.product Web pages on Novell.com](http://support.novell.com/volera) (<http://support.novell.com/volera>).

- 3** Open a DOS session window and enter the following command:

```
run volauth -install
```

- 4** Start the Excelerator Services on the Domain Controller by doing one of the following:

- 4a** Restart the Domain Controller machine

Or

- 4b** On Windows NT, click Start > Settings > Control Panel > Services > Excelerator Services > Start.

Or

- 4c** On Windows 2000, click Start > Programs > Administration Tools > Services > Excelerator Services > Start.

Setting Up NTLM Authentication

After you have completed the steps in “[Preparing Your Network for NTLM Authentication](#)” on [page 158](#), you can set up an NTLM authentication profile by completing the following procedure.

- 1** In the browser-based management tool, click Cache > Authentication > Insert.

- 2** Type a name for the profile in the Authentication Profile Name field.

IMPORTANT: Each profile name created on a cache device must be unique. Excelerator doesn't recognize case differences (MyProfile and myprofile are the same name to Excelerator) and it will overwrite and concatenate previously created profiles without warning if a duplicate name is used. For more information, see “[Authentication Dialog Box](#)” on [page 378](#).

- 3** Check NTLM Authentication > click Options.

- 4** In the Addresses list, insert the IP addresses of the Domain Controllers to which users will authenticate in the order you want the controllers accessed.

- 5** Click OK > OK.

- 6** Assign the profile to one or more proxy services as described in each service tab section in “[Using the Cache Panel](#)” on [page 333](#).

- 7** Click Apply.

NTLM Authentication Multiple Domain Support

In Excelerator 2.3, NTLM authentication profiles now support multiple domains.

The username and groupname strings used by Excelerator always include the domain name followed by a back slash (\). Therefore, log file entries will contain the domain combined with the other names.

Access control rules that refer to NTLM authentication profiles must now contain the domain name followed by the username or groupname.

The procedure for creating profiles has not changed. Multiple domain support is handled by the trust relationships between domains.

To use NTLM multiple domain support, you will need to do the following:

- 1** On each Domain Controller used by an existing NTLM authentication profile, install the VOLAUTH.EXE file contained in the Excelerator 2.3 device's default FTP directory (/etc/proxy/appliance/config/user).
- 2** Use the VOLAUTH.EXE file when configuring additional Domain Controllers for NTLM authentication profiles.
- 3** Establish trust relationships between the Domain referenced in a profile and any other domains being used for authentication.
- 4** If you have previously created access controls that refer to NTLM-based authentication profiles, edit the controls and insert a domain name and back slash (\) before any usernames or groupnames included in the controls.
- 5** Include the domain name and back slash (\) with all usernames and groupnames included in new access controls you create.

23 Access Control

This section contains information about access control.

Overview

The access control functionality in Excelerator 2.3 allows you to set access rules for the source or destination of a request through the proxy server.

Access control is a per-service feature, not a per-system feature. It operates on a granular, not a global, level. It can be controlled on a per-user basis as well.

Access control lists (ACLs) are created using the Excelerator 2.3 HTML Annotator program. These lists are then implemented using the browser-based management tool.

IMPORTANT: By default, access control is not enabled and no rules or policies have been created or are active.

Process

NOTE: These are just general guidelines. The access control policies you set up will all depend on your situation.

The general process to set up and implement an ACL is as follows:

1. Determine your ACL strategy.
2. Using the HTML Annotator, create an ACL.
3. Using the Excelerator 2.3 browser-based management tool, implement the ACL.

Determining Your ACL Strategy

When you are determining your access control strategy, some questions to consider are:

- ♦ Should everything be blocked at a high level and access provided only to certain users?
- ♦ Should the rules be applied to all users or to specific users?
- ♦ Are the connections based on a single IP address, an IP address range, or an IP subnet?
- ♦ Are the destination addresses mixed?
- ♦ Should destinations be blocked by URL, a single IP address, an IP address range, or an IP subnet?

Wildcarding

Wildcards are used to define broad rules. For example, if you don't want company-wide access to the accounting web pages, you can define a rule to block access to `http://www.company.com/groups/accounting/*`.

Further, you can add a rule that allows access to the same URL wildcard for people coming from a set of IP addresses or logged in users.

IMPORTANT: While wildcarding is supported, the product does not support top-level domain blocking with wild carding.

Authentication

Anytime there is a user-specified rule, authentication needs to be configured and enabled on the service on which you're applying the access control policy. If the authentication is not set up properly, the user will not be prompted to authenticate.

Also, if an administrator enables authentication, applies an authentication profile to a proxy and enables "Authenticate only when user attempts to access a restricted page", but does not enable Access Control, the users will never be prompted to authenticate.

Examples

Here are two examples of access control policies.

Example 1: AuthOnFailure (a.k.a. Authentication Bypass). This example describes authenticating when the user runs into an ACL. This policy would have three rules. Rule 1 blocks access. Rule 2 allows access to a specified user. Rule 3 allows access to all users to `innerweb.company.com`. If a client navigates to `www.company.com`, access is blocked, they are prompted to authenticate. If authentication is successful, universal access is allowed. On the other hand, if authentication fails, the client is blocked access to the specified URL and all other sites. If a client attempts navigation to `innerweb.company.com`, access is allowed and no authentication is required.

Example 2: Blocking by Subnet. In this example, the organization is divided into two sections: R&D and Temp. R&D needs universal access; Temp needs access to `innerweb` only. Rule 1: block access to all users connecting from anywhere destined to anywhere. Rule 2: allow access to all users connecting from R&D subnet destined to anywhere. Rule 3: allows access to all users connecting from Temp subnet to destined to `innerweb.company.com`.

Creating an ACL

Access control policies are created using the HTML Annotator program.

To launch the HTML Annotator, do the following:

- 1 Launch the browser-based management tool.
- 2 Click Home > Add Ons.
- 3 Select Access Control Policies.
- 4 Click Launch.

Implementing the ACL

In order to use the access control policies you've created with the HTML Annotator, you must enable access control using the browser-based management tool.

To enable access control, do the following:

- 1** Launch the browser-based management tool.
- 2** Click Cache > Client Accelerator or Transparent Handling (depending on the type of service you are running).
- 3** Check the Enable Access Control checkbox.
- 4** Click Access Control Options.
- 5** Configure your setup, as needed.

IMPORTANT: If Excelerator is hosting the Web content, the port number must be included in the URL (i.e., <http://www.company.com:80>).

Other Guidelines

Some other guidelines for using access control as follows:

- ♦ An over-the-wire upgrade (a.k.a. "OTWUG") will fail if an access control policy exists that blocks all and the OTWUG box is not included in the exceptions list.
- ♦ A destination rule cannot have multiple source types. Due to the architecture, the Source is specified as part of the policy (i.e., if specifying an IP Address Range, you would not want to specify a URL within the destination field of the same rule). A possible solution is to create another rule with the URL destination and so on with Single IP Address, Subnet, etc.
- ♦ Administrators will want to note case sensitivity with user IDs for access control policies, especially with NDS and RADIUS users since there is no method to copy the user object into the User Specific window.
- ♦ A subnet entry must be followed with a hyphen. For example, if the Source or Destination is a Subnet Type, then the Subnet IP Address must be in the form: 63.104.0.0-. This distinguishes the Subnet Type from a Single IP Address. If the hyphen is not present then the rule will not take effect unless a requesting client is connecting from an IP address of 63.104.0.0, which is not very likely.
- ♦ The product does not perform reverse DNS lookups. If a URL maps to a six unique IP addresses, an administrator would have to block the entire range or block each single IP address specifically.
- ♦ For help with syntax for different authentication profiles, see [this TID \(http://support.novell.com/cgi-bin/search/searchtid.cgi?/10077674.htm\)](http://support.novell.com/cgi-bin/search/searchtid.cgi?/10077674.htm).

24

Managing Appliance Certificates

The Excelerator appliance has public key infrastructure mechanisms for generating, importing, using, and maintaining public key certificates. These include:

- ♦ An appliance-specific certificate authority (CA) which automatically generates certificates for each assigned IP address and other appliance resources.

The appliance uses these auto-generated certificates for certain appliance-specific secure communications, such as obtaining filtering lists.

These can also be used for secure connections with browsers using appliance caching services. However, browsers won't recognize the appliance CA unless they are specifically configured to do so. This causes confirmation messages to be generated that can confuse users and cause them to not use the appliance's caching services.

To create appliance-specific certificates, see the instructions in [“Creating Certificates Using the Appliance CA” on page 166](#).

- ♦ Mechanisms for generating a certificate signing request (CSR) and storing issued certificates on the appliance.

Generating a CSR is the first step to obtaining a certificate from an external CA.

After you obtain certificates from one or more external CAs, you can use the appliance certificate maintenance features to monitor certificate status, back up certificates in case the appliance fails, and replace certificates when they expire.

To generate a CSR and store the issued certificate, complete all the instructions in [“Obtaining a Certificate from an External CA” on page 166](#)

Because the creation process is different for internal and external certificates, they are described separately. See the instructions in [“Creating Certificates Using the Appliance CA” on page 166](#) and [“Obtaining a Certificate from an External CA” on page 166](#).

Naming Certificates

As you create certificates on the appliance, you should observe the following guidelines:

1. Identify the caching service for which the certificate will be used.
2. Pick a name for the certificate that you will easily associate with its corresponding caching service. The name must contain only alphanumeric characters and no spaces.

For example, you might pick Foo for the name of the foo.gov Web server accelerator or Marketing for the transparent service in the marketing department.

3. Identify the DNS hostname that the browser expects to find in the certificate.

Creating Certificates Using the Appliance CA

Use the instructions in this section if you plan to configure the browsers that will access the appliance's caching services. Browsers will need to import the appliance's CA in order to accept its certificates as legitimate.

If this is not done, users will get certificate confirmation messages that might confuse them.

To create an appliance CA certificate

- 1** In the browser-based management tool, click Home > Certificate Maintenance > Create.
- 2** Type an appropriate name for the certificate, as explained in [“Naming Certificates” on page 165](#).
- 3** Type an appropriate subject name, as explained in [“Naming Certificates” on page 165](#).
- 4** Click the Signature Algorithm drop-down list > select the algorithm you want to use (SHA-1 or MD-5).
- 5** Click the RSA Key Size drop-down list > select the RSA key size that you want to use.
You cannot select a key size larger than the maximum key size on the Appliance.
- 6** Check Use Local Certificate Authority.
- 7** Click the Validity Period drop-down list > select the length of time that you want the certificate to be valid.
- 8** Click OK.
- 9** Look at the Action and Status fields.

The Action field should have red arrows on the left and the word Create displayed on a green background. The Status should be Building.

The red arrows and green background indicate that you need to click Apply.

- 10** Click Apply.
If any errors occur during the certificate creation process, they are displayed in the Error field on a red background.
- 11** If an error occurs, click Modify.
- 12** In the Modify Certificate dialog box, make the changes necessary to resolve the errors > click OK.
- 13** Click Apply and repeat the modification process until the Status field displays the word Active.

Obtaining a Certificate from an External CA

Requesting the CSR

- 1** In the browser-based management tool, click Home > Certificate Maintenance > Create.
- 2** Type an appropriate name for the certificate, as explained in [“Naming Certificates” on page 165](#).
- 3** Type an appropriate subject name, as explained in [“Naming Certificates” on page 165](#).

- 4** Click the Signature Algorithm drop-down list > select the algorithm you want to use (SHA-1 or MD-5).
- 5** Click the RSA Key Size drop-down list > select the RSA key size that you want to use.
You cannot select a key size larger than the maximum key size on the Appliance.
- 6** Click Use External Certificate Authority.
- 7** If you are requesting a VeriSign certificate, check the VeriSign CA checkbox. Otherwise, leave the box unchecked.
- 8** If desired, type a name for your organization or division.
This is commonly referred to as the Organizational Unit and is used to differentiate organizational divisions or to describe departments or divisions.
- 9** Type the city or town where your organization does business.
This is commonly referred to as the Locality.
- 10** Type the unabbreviated name of the state or province where the organization does business.
This is commonly referred to as the State.
- 11** Type the ISO country code for the country where the organization does business.
This is commonly referred to as the Country and must be a valid, two-character ISO country code.
- 12** Click OK.
- 13** Look at the Action and Status fields.
The Action field should have red arrows on the left and the word Request displayed on a green background. The Status should be Building.
The red arrows and green background indicate that you need to click Apply.
- 14** Click Apply.
If any errors occur during the certificate creation process, they are displayed in the Error field on a red background.
- 15** If an error occurs, click Modify.
- 16** In the Modify Certificate dialog box, make the changes necessary to resolve the errors > click OK.
- 17** Click Apply and repeat the modification process until the Status field displays the words CSR in Progress on a yellow background.

Sending the CSR

- 1** To open a new browser window that displays the CSR contents, click View CSR.
- 2** Select and copy the complete CSR text into your computer's clipboard. Internet Explorer and other browsers sometimes combine them with the CSR text that is in between. Clicking the browser refresh/reload button will often fix the problem. If it doesn't, simply insert appropriate carriage returns during the next step. After you have copied the text, you can close that browser window.
- 3** Paste the CSR text from the clipboard to the e-mail message or HTML form as required by your CA.

The method for sending the CSR will vary, depending on the authority. VeriSign, for example, uses a web page interface.

IMPORTANT: The header and trailer must be on lines separate from the body of the CSR.

The header line will be similar to the following:

```
----- BEGIN NEW CERTIFICATE REQUEST-----
```

The trailer line will be similar to the following:

```
-----END NEW CERTIFICATE REQUEST-----
```

If required, you must use hard returns to separate these two lines from the body of the CSR.

- 4** Wait for the certificate to be returned from the external CA.

Storing the Certificate

After the external CA responds with the certificate, do the following:

- 1** In the browser-based tool, click Home > Certificate Maintenance > *the name of the certificate you want to store* > Store Certificate.

- 2** In the Store Certificates dialog box, paste the CA certificate in the CA Certificate Contents box.

NOTE: If you requested a VeriSign certificate and you checked the VeriSign box in [Step 7 on page 167](#), the CA Certificate Contents box is grayed out. You will not need to paste the VeriSign CA certificate because VeriSign certificates are already stored on the appliance.

- 3** Paste your newly issued certificate in the Server Certificate Contents box.

- 4** Click Create.

- 5** Look at the Action and Status fields.

The Action field should have red arrows on the left and the word Create displayed on a green background. The Status should be CSR in Process.

The red arrows and green background indicate that you need to click Apply.

- 6** Click Apply.

If any errors occur during the certificate creation process, they are displayed in the Error field on a red background.

- 7** If an error occurs, click Store Certificate.

- 8** In the Store Certificate dialog box, make sure the correct certificates are pasted in the boxes > click OK.

- 9** Click Apply and repeat the modification process until the Status field displays the words Active on a green background.

Viewing (Exporting) a Certificate's CA

To view (export) a certificate's Certificate of Authority (CA), do the following:

- 1** In the browser-based management tool, click Home > Certificate Maintenance > *the certificate you want to export* > Export CA Certificate.

The contents of the CA certificate is displayed in a new browser window.

Modifying a Certificate

Only certificates that have an error or the status “Building” can be modified.

- 1 In the browser-based management tool, click Home > Certificate Maintenance > *the certificate you want to modify* > Modify.
- 2 After making the necessary changes, click OK to accept the changed values.
- 3 In the Modify Certificate dialog box, make the desired changes.
- 4 If the Action field displays the word Request or Create on a red background, you must click Apply to make the changes.

Importing a Trusted Root to a Cache Device

Mutual authentication profiles and LDAP authentication profiles that rely on a secure LDAP server both require that the trusted root of their associated CAs be imported to the appliance. For more information, see “[Using Mutual \(Certificate-Based\) Authentication](#)” on page 141 and “[Using LDAP Authentication](#)” on page 143.

When creating these profiles, you will be required to access the Import Trusted Root dialog box and copy in the appropriate trusted root file.

To create a trusted root file, do the following:

- 1 In the Imported Filename field, type a path and filename for the trusted root file.

The filename can contain up to eight alphanumeric characters and a .DER extension. The appliance automatically appends the .DER extension if you don’t include it.

IMPORTANT: Be sure you use a unique filename for each .DER file. The appliance overwrites files without warning if you use duplicate filenames.

Remember that Excelerator is not case-sensitive, so MyCert.DER and mycert.der are, effectively, the same filename.

The path must be a directory path that already exists. You cannot create directories on the appliance.

If you want to list your trusted root files later, use an FTP-accessible directory, such as SYS:\ETC\PROXY\DATA, as the path. Otherwise, you won’t be able to list the files. For a list of FTP-accessible directories, see “[Functionality Limitations of the Appliance’s Mini FTP Server](#)” on page 262.

If you don’t include a path with the filename, Excelerator creates the file at the root of the SYS: volume. You cannot see the root of the SYS: volume using FTP.

- 2 Using a text editor on your configuration workstation, open the .DER file for the Certificate Authority > select the file contents > paste the contents to the clipboard.

To obtain .DER files, contact a Certificate Authority vendor.

- 3 Return to the Import Trusted Root dialog box > paste the clipboard contents into the text box above the OK and Cancel buttons.
- 4 Click OK.

Deleting a Certificate

If a certificate has expired or you are unable to resolve an error, you might want to delete a certificate.

IMPORTANT: Use caution when deleting certificates. You should never delete system-generated certificates.

1 In the browser-based management tool, click Home > Certificate Maintenance > *a certificate you have generated that has expired or has an unresolvable error*.

2 Click Delete.

3 In the Delete Certificate dialog box, click Yes.

The certificate is removed from the Certificates list. If you have deleted the certificate in error, click Cancel.

4 Click Apply to remove the certificate from the appliance.

After clicking Apply, the certificate cannot be restored unless you have created a backup copy.

Backing Up a Certificate

Only active certificates can be backed up.

1 In the browser-based management tool click Home > Certificate Maintenance > *the certificate you want to back up*.

2 Click Backup.

3 In the Backup Certificate dialog box, type a password to use when restoring the certificate.

4 In the Confirm Password field, retype the same password.

IMPORTANT: Although the password is optional, we strongly suggest you use one. If you don't enter a password, the backed up certificate can be used by anyone who has access to the file.

5 Check either Disk or Floppy to indicate where the backup file should be saved.

6 Click OK.

The Action field should display red arrows and either Backup (Disk) or Backup (Floppy) on a green background.

If you want to cancel the backup action, click Cancel Backup (by the Action field).

7 If the Action field is green, click Apply.

The Backed Up status field for each certificate indicates whether a certificate has been backed up and where the backup file was placed (disk, floppy, or both).

If any errors occur during the backup process, they are displayed on the Error line and the background turns red.

You can then click Backup and repeat the process, taking care to avoid the errors indicated.

Backed up certificates are stored in a file named *certificate*.PFX, where *certificate* is the name of the certificate that was backed up.

IMPORTANT: If the certificate was backed up to the appliance hard disk, you should transfer the file from the appliance to another secure location. Otherwise, the backup copy will be lost if the appliance fails and has to be re-imaged.

Certificate backup files are stored in ETC/PROXY/APPLIANCE/CONFIG/USER/CERT/BACKUP. See [Chapter 35, "FTP Services," on page 261](#) for help using appliance FTP services.

If the certificate was backed up to a floppy disk, the file is in the root directory of the disk and the floppy should be stored in a safe place in case the certificate must be restored.

Restoring a Certificate

Only certificates that were previously backed up can be restored.

Prior to completing the following steps, make sure the backup file is in one of the following locations:

- ♦ On a floppy disk in the appliance's floppy drive
- ♦ In ETC/PROXY/APPLIANCE/CONFIG/USER/CERT/BACKUP on the appliance's hard disk

Unless the appliance has been damaged or reimaged, the backup file will be in the expected location.

If the file is not on the appliance, you must retrieve a copy from your secure location and either copy it to a floppy disk or to the appliance using FTP.

- 1** In the browser-based management tool, click Home > Certificate Maintenance > Restore.
- 2** In the Restore Certificate dialog box, type the certificate name, which is the PFX filename.
- 3** Type the same password you used when creating the backup file.
- 4** Retype the password.

The passwords must match exactly, or you won't be able to exit the dialog box.

- 5** Click OK.
- 6** Click Disk or Floppy to indicate where the backup file is.
- 7** Click OK.

The Action field should display red arrows and either Restore (Disk) or Restore (Floppy) on a green background. The Status field should display Building.

If you want to cancel the restore action, click Cancel Restore (located next to the Action field).

- 8** Click Apply.

If any errors occur during the restore process, they are displayed on the Error line. The background for the text will turn red.

The only way to fix a restore error is to delete the certificate and try the restore process again.

A restoration failure might mean that the backup file didn't exist or you had the wrong password.

25 Transforming Content for Internet Delivery

Today's Web economy often dictates that organizations expose information that resides inside a firewall to users on the Internet.

This chapter explains

- ♦ Some of the issues that Web administrators face as they seek to accomplish this goal
- ♦ How Excelerator can help to solve the issues that administrators face

Identifying the Issues

It is fairly simple to create a reverse proxy service with an external DNS name to serve as a portal through the firewall to internal Web servers. However, further planning reveals that the issues involved in exposing internal content on the Internet are much more complex than simply providing a reverse proxy.

- ♦ Sensitive data that can flow freely within the safety of the firewall requires SSL encryption when transmitted over the Internet.
- ♦ Concerns regarding hackers dictate that internal DNS names, ports, etc. not be exposed on the Internet.
- ♦ Embedded object references that use internal DNS names, ports, etc. don't work when browsers are requesting them from the Internet.

The following section summarizes the Excelerator solutions to each of these issues.

Understanding URL Overrides

The key mechanism that Excelerator uses to prepare internal content for exposure on the Internet is called the URL override.

URL overrides can be

- ♦ Automatically created by Web Server Accelerator services and/or Secure Excelerator.
- ♦ Manually specified by administrators using the [URL Override Dialog Box](#).

For more information regarding automatically and manually created URL overrides, see [“Automatic Vs. Manual URL Overrides” on page 174](#).

Secure Excelerator: The All-in-One Solution

Secure Excelerator solves each of the issues identified in [“Identifying the Issues” on page 173](#) using SSL version 3 technology to protect your sensitive data and transform embedded object

links. It does not require any additional software or special configuration of either your Web servers or the browsers accessing your Web site.

The product was specifically designed for e-businesses, enterprises, and other organizations that want to conduct business using the Web.

For more information on obtaining and installing Secure Excelerator, see the *Volera Secure Excelerator 1.1 Administration Guide*.

URL Overrides Transform Object Reference URLs in Cached HTML

The base Excelerator product and Secure Excelerator both use URL overrides to address security and object embedding issues by transforming content before it is cached. Specifically, URL overrides affect one or more of the following fields in object reference URLs in the HTML:

- ♦ **DNS Name Overrides:** These change each instance of the origin Web server's DNS hostname when it appears in an object reference URL within the HTML. The override changes the origin Web server's DNS name to the DNS name of the accelerator service to which the override belongs (if the two DNS names are different).
- ♦ **Port Overrides:** When the object reference URL includes a port number with the origin server's DNS name, port overrides change the port number of the origin Web server to the port numbers of the accelerator service (assuming the two ports are different).
- ♦ **Pathrule:** Excelerator's path-based multihoming option uses pathrule overrides to let parent accelerator services know which child accelerator service has file in the object reference URL.
- ♦ **Protocol Scheme:** Secure Excelerator uses scheme overrides to change the HTML protocol scheme from HTTP to HTTPS when vending the objects on the Internet.

Automatic Vs. Manual URL Overrides

Understanding when to create manual URL overrides requires that you clearly understand when the system creates automatic overrides.

Understanding When Automatic URL Overrides Are Created

URL overrides are automatically created by Web Server Accelerators and Secure Excelerator as indicated in the following table:

Table 10

Object Reference URL Component	Trigger for Automatic Override Creation	How the Override Works
DNS Name	<p>If the accelerator service's DNS Name is different from the origin Web server's DNS name listed in the Web Server Addresses list, the system automatically creates a <i>DNS name override</i> for the service that covers references to objects cached on the service.</p> <p>Usually the origin Web server's DNS name is used in internal browser requests, and the accelerator service's DNS name is used in requests coming from the Internet.</p> <p>IMPORTANT: If the origin Web server's name entry is the server's IP address, an automatic DNS Name override is not created.</p>	<p>Prior to caching an HTML object for the service, the DNS name override changes all object references that contain the origin Web server's DNS name, replacing it with the DNS Name specified for the accelerator service.</p> <p>Since DNS servers on the Internet have no knowledge of internal DNS names, this action ensures that a DNS server on the Internet can properly resolve browser requests for objects whose URLs have been changed to contain the correct, Internet DNS name.</p>
Port	<p>If the Web Server Port and the Accelerator Proxy Port fields are different in the accelerator service definition, the system automatically creates a <i>port override</i> for the service that covers references to objects on the origin Web server.</p> <p>Sometimes, internal networks use non-standard port numbers. For example, the internal HTTP port number might be 81 instead of 80.</p> <p>IMPORTANT: If you are using Secure Excelsator, the port override is specified in the Secure Excelsator Options dialog box.</p>	<p>Prior to caching an HTML object for the service, the port override changes all references in the HTML that point to the accelerated Web server and contain the Web Server Port value specified, replacing the Web Server Port value with the Accelerator Proxy port value specified.</p> <p>This action ensures that the request is routed using the correct port.</p> <p>NOTE: In the case of port overrides, if the accelerator proxy port value is the industry standard port number, the URL override simply removes the port number from the URL.</p>
Pathrule	<p>If path-based multihoming is enabled for a child accelerator service and a <i>starts-with</i> rule is specified for the Sub-Path Match String, the system automatically creates a <i>pathrule override</i> for the service.</p>	<p>Prior to caching an HTML object, the pathrule override changes all object reference URLs that point to the accelerated Web server by inserting the Sub-Path Match String at the root of the URL path.</p> <p>When the parent accelerator service gets a browser request from the Internet with the Sub-Path Match String embedded, the parent accelerator service knows the request must be routed to the child accelerator service.</p>

Object Reference URL Component	Trigger for Automatic Override Creation	How the Override Works
Protocol Scheme	If Secure Excelerator is installed and enabled for the service, the system automatically creates a scheme override for the service.	Prior to caching an HTML object, the scheme override changes all object reference URLs in the HTML that point to the accelerated Web server by changing the protocol scheme from HTTP to HTTPS.

Understanding When to Create Manual URL Overrides

The key point to remember when deciding whether a manual override is required is that automatic overrides apply only to objects that originate on the origin Web server from which the accelerator service fills content. Object reference URLs that point to other origin Web servers are not changed.

The following table can help you decide whether you need to create any manual overrides:

Table 11

Object Reference Component	A Manual Override is Required If
DNS Name	<ul style="list-style-type: none"> The origin Web server for an accelerator service contains object reference URLs that point to objects on another origin Web server that is being accelerated by another accelerator service. and The DNS names in the other object reference URLs must be changed for the references to resolve correctly on the Internet. <p>IMPORTANT: You will also need to create a manual DNS Name override if an override is required and the origin Web server's DNS name is an IP address because the system will not create an automatic override in that case.</p>
Port	<ul style="list-style-type: none"> The origin Web server for an accelerator service contains object reference URLs that point to objects on another origin Web server that is being accelerated by another accelerator service. and The port number in the other object reference URLs must be changed for the references to resolve correctly on the Internet.
Pathrule	References to objects on a child service occur on either the parent service or on one of the other child services. (References to objects on a given child service are automatically changed only on that child service. References to the given child service from the parent service or other children are not changed unless you create a manual URL override.)
Protocol Scheme	The origin Web server for a Secure-Excelerator-enabled accelerator service contains object reference URLs that point to objects on other accelerated origin Web servers (whether or not they have Secure Excelerator enabled).

Using URL Overrides

This section provides visual summaries of how overrides work and explains how to create them. Before creating manual URL overrides, however, it is essential that you understand the

information presented in “[Understanding URL Overrides](#)” on page 173. Otherwise, you might either duplicate overrides the system creates automatically, or you might be frustrated when URLs in accelerated HTML content are not transformed as you think they should be.

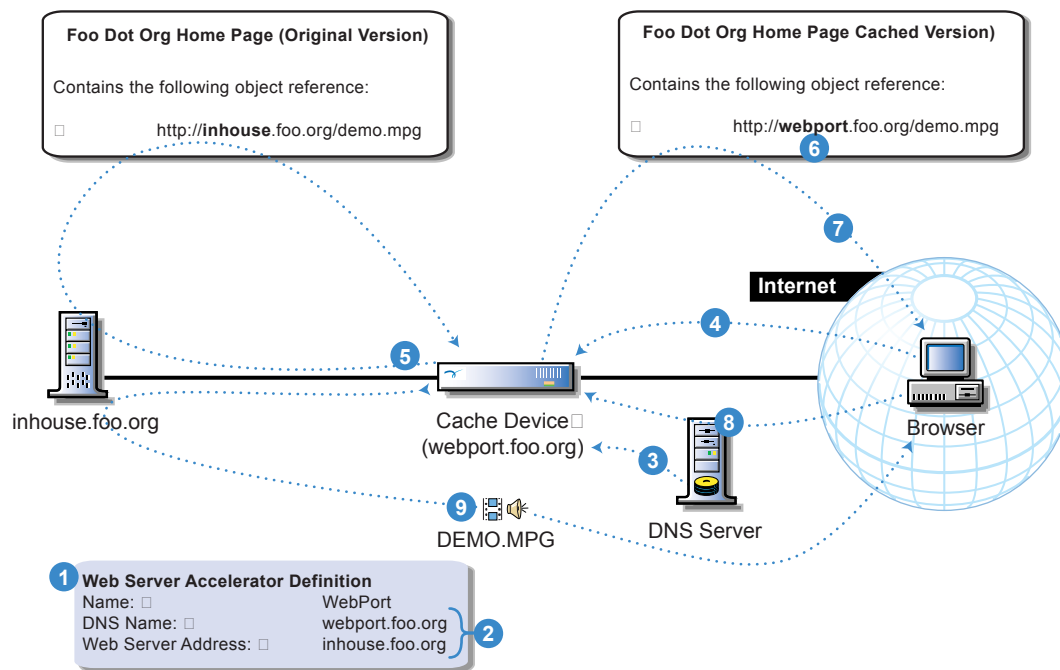
IMPORTANT: For simplicity’s sake, the following sections present each override type by itself rather than in combination with other override types. In actual implementations, however, overrides will generally be used in combination with each other because transforming object references usually involves changing more than one part of the reference URL.

DNS Name Overrides

DNS overrides replace internal DNS names in object references with external DNS names.

How DNS Name Overrides Work

Figure 63



1 □ Foo Dot Org creates a Web Server Accelerator service on the cache device to make the content available on its in-house Web server (inhouse.foo.org) available on the Internet.

2 □ Because the service's DNS Name and Web Server Address fields are different, the system creates an automatic DNS Name override.

3 □ Foo Dot Org registers an IP address assigned to the cache device with the Internet Domain Name Service (DNS) as webport.foo.org.

4 □ The cache device receives a request for the webport home page.

5 □ The cache device routes the request to the Web Port acceleration service, which retrieves the home page on inhouse.foo.org.

6 □ Prior to caching the home page, the service applies the DNS name override created in Step 2 to all object references that include the DNS name inhouse.foo.org, changing the name to webport.foo.org.

7 □ After caching the home page, the service sends a copy to the requesting browser.

8 □ The browser then requests the DEMO.MPG file, and since the DNS name in the object reference was changed to webport.foo.org, the DNS server resolves the request to the cache device. (If the name hadn't been changed, the request could not have been resolved, since inhouse.foo.org is an internal DNS name.)

9 □ The cache device routes the request to the WebPort acceleration service, which retrieves DEMO.MPG from inhouse.foo.org, caches the object, and sends a copy to the browser.

Creating DNS Name Overrides

NOTE: Prior to creating a DNS name override, you must create the Web Server Accelerator Service to which the override applies.

IMPORTANT: If an automatic override has been created for an origin Web server's DNS name, you cannot create another override for the same name.

To create manual DNS name overrides, complete the following steps:

- 1 In the browser-based tool, access the accelerator service for which you are creating the override.
- 2 Click URL Override.

- 3** In the URL Override dialog box, Insert an Origin Host Name or select one you have previously created from the Origin Host Name drop-down list.

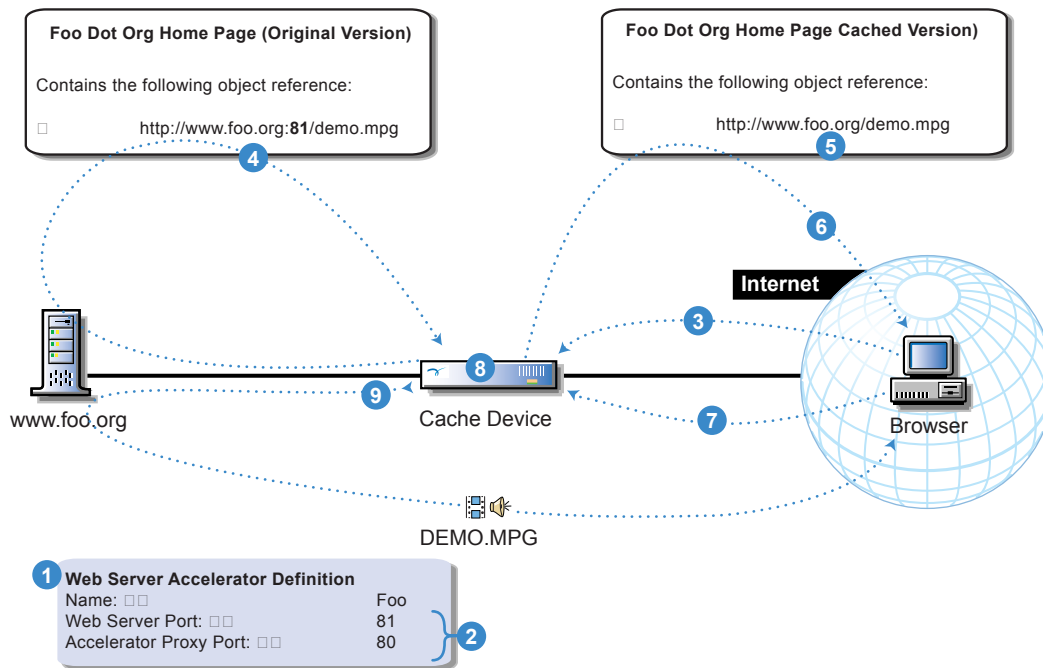
This is the internal DNS name of the origin Web server that you want changed in the object references. It corresponds to the Web Server Addresses list entry in a Web Server Accelerator definition.
- 4** In the Host drop-down list, select Override.
- 5** In the Accelerator Host field, type the DNS name that will replace the name you specified in **Step 3**.
- 6** In the left frame, click Apply.
- 7** In the Home panel, click Apply.

Port Overrides

Port overrides replace internally used port numbers with external (standard Internet) port numbers.

How Port Overrides Work

Figure 64



1 □ Foo Dot Org creates a Web Server Accelerator service on the cache device to make the content on its Web server available on the Internet.

2 □ Because the service's Web Server Port and Accelerator Proxy Port fields are different, the system creates an automatic port override.

3 □ The cache device receives a request for the www.foo.org home page on port 80.

4 □ The cache device routes the request to the Foo acceleration service which retrieves the home page using port 81.

5 □ Prior to caching the home page, the service applies the port override created in Step 2 to all www.foo.org object references that include port 81, removing the port indicator since the replacement value (80) is the default HTTP port.

6 □ After caching the home page, the service sends a copy to the requesting browser.

7 □ The browser then requests the DEMO.MPG file on port 80 rather than port 81 which would have resulted in the request not being filled.

8 □ The cache device routes the request to the Foo acceleration service.

9 □ The service retrieves DEMO.MPG from foo.org on port 81, caches the object, and sends a copy to the browser on port 80.

Creating Port Overrides

IMPORTANT: Prior to creating a port override, you must create the Web Server Accelerator Service to which the override applies.

To create manual port overrides, complete the following steps:

- 1 In the browser-based tool, access the accelerator service for which you are creating the override.
- 2 Click URL Override.
- 3 In the URL Override dialog box, Insert an Origin Host Name or select one you have previously created from the Origin Host Name drop-down list.

This is the internal DNS name of the origin Web server whose object references will have their port number changed.

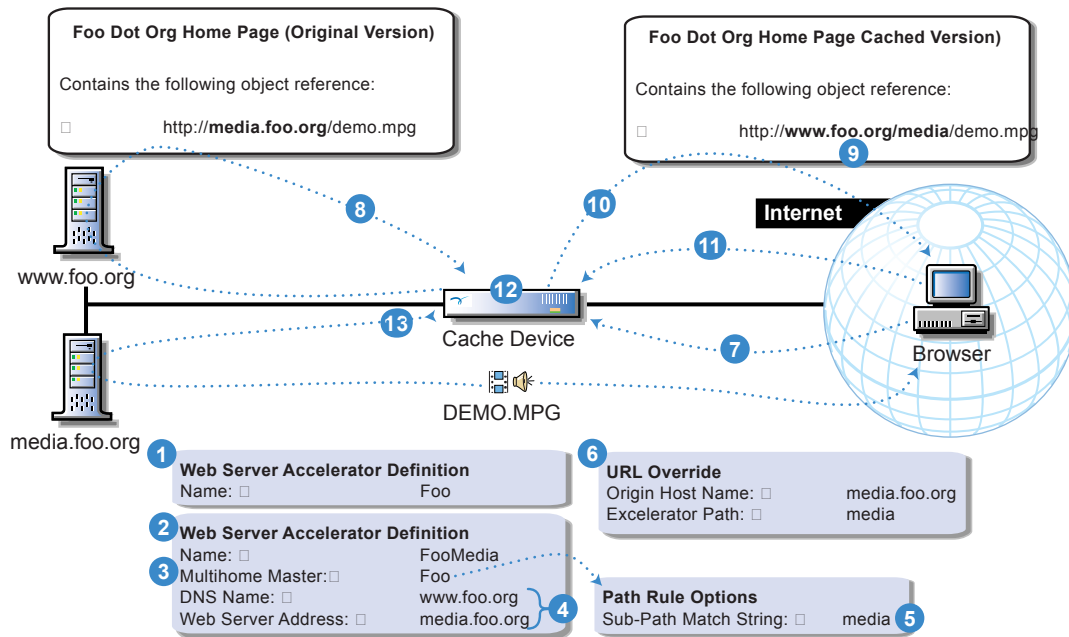
- 4** In the Origin Port field, type the internal port number you want to change
- 5** In the Origin Port drop-down list, select Override.
- 6** In the Excelerator Port field, type the port that will replace the port you specified in [Step 4](#).
- 7** In the left frame, click Apply.
- 8** In the Home panel, click Apply.

Pathrule Overrides

Pathrule overrides are the mechanism that Excelerator's path-based multihoming services use to let parent accelerator services know which child accelerator service has the referenced object.

How Pathrule Overrides Work

Figure 65



- 1 □ Foo Dot Org creates a Web Server Accelerator service named `Foo` on the cache device.
- 2 □ Foo Dot Org creates another accelerator service named `FooMedia` for the content on its `media.foo.org` Web server.
- 3 □ Foo Dot Org enables path-based multihoming in the `FooMedia` accelerator service and specifies `Foo` as the master accelerator. (`FooMedia` is now a child accelerator of the `Foo` accelerator.)
- 4 □ Specifying a multihome master for the `FooMedia` accelerator places `foo.org` in the DNS Name field, creating an automatic DNS Name override.
- 5 □ Foo Dot Org specifies the sub-path match string `media` in the Path Rule Options of the `FooMedia` accelerator, creating an automatic pathrule override for all object references on `media.foo.org` whose URLs include `media.foo.org`.
- 6 □ Because the `foo.org` server has object references to `media.foo.org`, Foo Dot Org also creates a manual pathrule override in the `Foo` accelerator for `media.foo.org` with the Excelsator Path `media`.

NOTE: Excelsator Path and Sub-Path Match String are different names for the same thing.

- 7 □ The cache device receives a request for the `foo.org` home page.
- 8 □ The cache device routes the request to the `Foo` acceleration service, which retrieves the home page from `www.foo.org`.
- 9 □ Prior to caching the home page, the service applies the overrides from Steps 4, 5, and 6.
- 10 □ After caching the home page, the service sends a copy to the requesting browser.
- 11 □ The browser then requests `DEMO.MPG`.
- 12 □ The cache device routes the request to the `FooMedia` acceleration service, which detects the `/media` trigger in the object path and routes the request to the `FooMedia` acceleration service.
- 13 □ The `FooMedia` acceleration service retrieves `DEMO.MPG` from `media.foo.org`, caches it and sends a copy to the browser.

Creating Pathrule Overrides

IMPORTANT: Prior to creating a pathrule override, you must create the Web Server Accelerator Service to which the override applies.

You must also specify a DNS name override at the same time as outlined in the following procedure. This is required because the multihome master DNS name must appear in conjunction with the pathrule for the override to work properly.

If you need more information on why a DNS name override is also required, review and consider the information in [“How Pathrule Overrides Work” on page 182](#).

To create manual pathrule overrides, complete the following steps:

- 1** In the browser-based tool, access the accelerator service for which you are creating the override.
- 2** Click URL Override.
- 3** In the URL Override dialog box, Insert an Origin Host Name or select one you have previously created from the Origin Host Name drop-down list.

This is the internal DNS name of the origin Web server whose object references will have a path string inserted.
- 4** In the Path-Based Multihoming drop-down list, select Yes.
- 5** In the Excelerator Path field, type the string that will be inserted to identify the child accelerator service.
- 6** In the Host drop-down list, select Override.
- 7** In the Excelerator Host field, type the DNS name of the multihome master (parent accelerator).

This name will replace the name you specified in [Step 3](#). The system will route the request to this service, which will then detect the path trigger and route the request on to the child accelerator service.
- 8** In the left frame, click Apply.
- 9** In the Home panel, click Apply.

Scheme Overrides

Scheme overrides are used by Secure Excelerator to change the HTML protocol scheme from HTTP to HTTPS when vending internal objects on the Internet.

How Scheme Overrides Work

Scheme overrides apply only when Secure Excelerator is installed. For a detailed explanation of the value they add and how they work, see the [Volera Secure Excelerator 1.1 Administration Guide](#).

Creating Scheme Overrides

IMPORTANT: Prior to creating a scheme override, you must create the Web Server Accelerator Service to which the override applies and have Secure Excelerator installed and active on the cache device.

To create a manual scheme override, complete the following steps:

- 1** In the browser-based too, access the accelerator service for which you are creating the override.
- 2** Click URL Override.
- 3** In the URL Override dialog box, insert an Origin Host Name or select one you have previously created from the Origin Host Name drop-down list.

This is the internal DNS name of the origin Web server whose object references will have their scheme changed.

- 4** In the Scheme drop-down list, select http to https.
- 5** In the left frame, click Apply.
- 6** In the Home panel, click Apply

Reviewing URL Overrides

You can view all the overrides you have created on a cache device by completing the following steps:

- 1** In the browser-based tool, click Home > Add Ons > Service-Based Rewriters > Launch.
- 2** In the Site Administration dialog box > Web Accelerator DNS Names drop-down list, select the DNS name of the accelerator which you want to review.
- 3** Click Properties.
- 4** In the URL Override dialog box, select an origin Web server DNS name and review its configuration.

26

Cache Freshness

This section contains information about cache freshness.

Overview

When first introduced to Web content caching, many network administrators assume that the object cache on an Excelerator appliance is basically the same as a browser's cache, which all users access when they click the Back button. The logical extension from this assumption is the fear that Excelerator will serve stale content that doesn't accurately reflect the fresh content on the origin Web server.

Actually, most time-sensitive Web content is flagged by Webmasters in such a way that it cannot become stale unless a caching system ignores the Webmaster's settings. And in fact, the Excelerator appliance honors all flags that affect cache freshness, including Time to Expire, Don't Cache, and Must Revalidate directives.

In addition, the Excelerator appliance can be fine-tuned for cache freshness in the following ways:

- ♦ Accelerated checking of objects that have longer than desirable Time to Expire headers
- ♦ Delayed checking of objects that have shorter than desirable Time to Expire headers
- ♦ Checking for freshness of objects that do not include Time to Expire headers

For more information on configuring Excelerator for cache freshness, see [“Managing Cache Freshness” on page 185](#) and [“Cache Freshness Dialog Box” on page 401](#).

Managing Cache Freshness

Cache freshness is a primary concern of most appliance administrators. The following sections briefly explain how your appliance ensures fresh content for network users and the options you have for adjusting this appliance feature.

How the Excelerator Appliance Checks for Object Freshness

Although the following explanation is an over-simplification, it lays the foundation for the specific examples that follow this section.

An Excelerator appliance has timers that it applies to every cached object.

Each time an object is cached or revalidated, the appliance starts a timer for that object. As long as the timer is running, the appliance will vend the object from cache. After the time has expired and when the appliance receives a request for the object, it will issue an IF-MODIFIED-SINCE request to the origin Web server.

If the object has changed, Excelsator retrieves the updated object into cache and serves it to the requesting browser before restarting the timer.

If the object has not changed, the appliance vends the object from cache and resets the timer, and the countdown for vending the object from cache begins again.

If a browser forces a refresh of the object, Excelsator honors the browser request, retrieves and caches the object regardless of whether it has changed, and restarts the timer.

How an Excelsator Appliance Keeps the Oldest Cached Objects Fresh

More than 80% of all Web objects either have no Time to Expire directives or are set to stay cached for as long as weeks or even months.

Because many of these objects actually change fairly frequently, the appliance has two timers for ensuring their freshness. You can configure these timers in the [Cache Freshness Dialog Box](#).

HTTP Maximum: This timer overrides an object's Time to Expire settings if it is longer than the timer's value.

The default timer value is six hours. This means that Excelsator will not vend an object that has been in cache longer than six hours without first determining if it should be refreshed.

HTTP Default: Excelsator applies this timer to objects that don't have Time to Expire settings.

The default timer value is two hours. This means that Excelsator will not vend an object that has no Time to Expire setting that has been in cache longer than two hours without determining if whether it should be refreshed.

How Excelsator Handles the Freshest Objects in Cache

Most Webmasters ensure that their time-sensitive objects have appropriate Time to Expire directives. Late-breaking news stories and photographs, for example, might stay in cache for only a few minutes before expiring.

By default, the Excelsator appliance simply honors the Webmasters' instructions and revalidates the objects in cache as directed.

However, some appliance installations, such as those connected through a modem, might need to limit how often these objects are refreshed. The appliance has a third timer for this purpose, also accessible in the [Cache Freshness Dialog Box](#).

HTTP Minimum: This timer sets the minimum number of hours or minutes Excelsator will serve HTTP data from cache before revalidating it against content on the origin Web server. No requested object will be revalidated sooner than specified by this value.

The default value for this timer is 0, meaning that Excelsator honors the Time to Expire directive for each object (assuming, of course, it is not longer than the HTTP Maximum timer).

If the timer is set to a value other than 0, it then overrides any object's Time to Expire directive that is shorter than the value set.

Fine-Tuning Cache Freshness on Your Appliance

The default timer settings explained in the previous sections are tuned for most appliance installations. However, you might have special requirements that could be met better if the default settings were adjusted.

Perhaps you are accelerating content that doesn't contain Time to Expire directives but changes frequently and needs to be refreshed more often than every two hours. You can adjust the HTTP Default timer in the **Cache Freshness Dialog Box** so that Excelsior refreshes the objects more frequently.

Perhaps you have severe Internet bandwidth restrictions and an environment with users who don't require object freshness checks every six hours. You can adjust the HTTP Maximum timer in the **Cache Freshness Dialog Box** to a different setting that meets your requirements and conserves bandwidth.

If you choose to adjust the timer values, avoid settings that result in objects being refreshed more often than is necessary. Otherwise, you could easily negate the bandwidth and response-time benefits of having the appliance on your network.

Using Custom Cache Control Headers

In addition to fine-tuning cache freshness using the system's global HTTP timers, as explained in **Fine-Tuning Cache Freshness on Your Appliance**, you can configure each proxy service to recognize custom headers in HTTP packets. Your Web server can then use these headers for transmitting caching instructions that only the configured appliance services will recognize and follow.

An Overview of How Headers Work

Only the accelerator service containing the custom header definition follows the cache policies specified in the custom headers.

All other caches, including the non-configured appliance caches, requesting browsers, and external proxy caches (transparent caches, client accelerators, etc.), do not recognize the custom headers. They follow only the cache policies specified by the standard cache control headers.

This means that you have the following options for configuring your Web server:

- ♦ You can specify that browsers and/or external caches cannot cache the objects, but the accelerator can.

This lets you offload request-processing from the origin Web server while still requiring that users return to the site each time they request an object.

- ♦ You can also specify separate cache times for browsers, external caches, and the accelerator you are defining.

Implementing Custom Cache Control Headers

To implement custom cache control headers, you must do the following:

- ♦ Enter a header string, such as MYCACHE, in the Custom Cache Control Header dialog box.
- ♦ Configure the Web server to send an HTTP header containing the defined string and the time in seconds that the object should be retained in cache (MYCACHE: 60, for example).

If the number is non-zero, Excelerator treats the reply as if it had the following headers:

```
Cache-Control: public  
Cache-Control: max-age=number
```

If the number is zero (0), Excelerator treats the reply as if it had the following header:

```
Cache-Control: no-cache
```

- ◆ Ensure that the Web server continues to send standard HTTP cache-control headers so that browsers and external caches follow the caching policies you intend them to.

For example, you could do the following:

- ◆ Use an Expires or Cache-Control: Max-Age header to specify that browsers should cache an object for two minutes.
- ◆ Use a Cache-Control: Private header to prevent external caches from caching the object at all.
- ◆ Use a custom cache control header, such as MYCACHE: 1800, to indicate that the accelerator should cache the object for 30 minutes.

Custom Cache Control Headers override the following standard HTTP cache-control headers on the appliance, but they do not affect how browsers and external caches respond to them:

```
Cache-Control: no-store  
Cache-Control: no-cache  
Cache-Control: max-age=number  
Cache-Control: private  
Cache-Control: public  
Pragma: no-cache  
Expires: date
```

An Implementation Example

For example, you might do the following:

1. While configuring a Web server accelerator service, you insert a string in the Custom Cache Control Header list with the value of FOOTTL.

The appliance will now recognize FOOTTL as a custom cache control header on objects requested through the service you are configuring.

2. You then configure the accelerated Web server to send FOOTTL: 600 in the headers of objects you want to be cached at the appliance.

The appliance will recognize this header as overriding the standard HTTP cache-control headers listed above when objects are requested through the accelerator service you are configuring.

3. Finally, you ensure that the Web server continues to send the following standard HTTP cache-control headers:
 - ◆ Cache-Control: Max-Age headers that cause browsers to cache objects for no longer than two minutes
 - ◆ Cache-Control: Private headers that cause external caches to not cache the objects

When your Web server sends an object with the FOOTTL header in response to an appliance request made through the accelerator service, your appliance recognizes the custom header and

caches the object for 10 minutes. Requesting browsers cache the object for only two minutes, and external caches do not cache the object.

Thus, the appliance off-loads a processing burden from the Web server by caching the frequently requested objects for 10 minutes (the value you specified in Step 2). Browsers, on the other hand, must always access the appliance to get the objects if their previous requests are older than two minutes. And the objects in the appliance's cache are kept fresh due to their relatively brief time-to-live value.

27

Managing Appliance Security Features

This section contains information about appliance security features.

The Console Lock Feature

The Excelsior console is locked by default to prevent unauthorized access. The password to unlock the console is the `config_user_password` you specified during the initial configuration.

To use the command line interface, you must unlock the console by entering the following command:

```
unlock  
config_user_password
```

NOTE: If a `config_user_password` is not set, the password is null.

After the console has been unlocked, it remains unlocked until you lock it using the **lock** command.

Managing HTTP CONNECT Method Support

The HTTP protocol supports a number of different access methods, such as GET, POST, and CONNECT.

The CONNECT method is normally used to establish a tunneled connection through which encrypted SSL traffic can be sent.

How the CONNECT Method Works

When proxy servers receive CONNECT requests, they are expected to establish a tunneled connection to a specified DNS host or IP address on a specified port. This allows the tunneled connection to be set up to any address and port.

An Unverified CONNECT Connection Is a Security Risk

It is not safe to assume that all CONNECT requests received by proxy servers are actually for SSL traffic. Outsiders frequently scan the Internet for proxies on port 8080 and other commonly used proxy ports to discover proxy servers that are accessible to them.

Outsiders Can Use Proxies to Attack Other Machines Inside a Firewall

If a proxy server that supports the CONNECT method is inside a firewall and the proxy server is accessible from outside the firewall, outsiders can request a tunneled connection to any address and port inside the firewall.

The proxy will set up the connection, thus allowing the outsider to have access to machines inside the firewall which would normally be inaccessible.

To the machines inside the firewall, it appears that the connection is originating from the proxy server address inside the firewall rather than from outside the firewall.

Outsiders are known to have used this capability to break through the protection normally provided by firewalls.

Attackers Can Hide Their Location

Attackers can use the CONNECT method to request that a publicly accessible proxy server set up a tunneled connection which passes through the proxy. This hides the real address of the attacker.

Attackers can then chain through several proxies by having the first proxy connect to the second, the second connect to the third, and so on, making it very difficult for law enforcement to discover where the attack is actually originating.

How Excelerator Protects Your Network

By default, Excelerator enables the CONNECT method for forward proxy services and disables it for transparent proxy services because:

- ◆ Users accessing SSL sites through a forward proxy must be able to use the CONNECT method to establish a tunneled connection with the origin Web server. Otherwise, they will not be able to access any SSL sites through the forward proxy service.
- ◆ Transparent services do not ordinarily need to use the CONNECT method.

For proxy services that require CONNECT method support, Excelerator monitors the data flowing through the tunneled connection. If the data is not SSL-related, Excelerator immediately tears down the tunneled connection and doesn't forward the data requests. This prevents outsiders from gaining access through firewalls and attackers from establishing chains to hide their identity.

Configuring Excelerator to Meet Your CONNECT Method Requirements

As explained in [“How Excelerator Protects Your Network” on page 192](#), Excelerator is configured by default to use the CONNECT method only when necessary and to verify CONNECT method requests to protect your network against security risks.

We recommend you use the default CONNECT method settings whenever possible.

If you have special configuration requirements and are considering a non-default CONNECT method configuration, consider the following points:

- ◆ If you disable the CONNECT method for client acceleration (forward proxy) services, you are blocking access to all SSL sites for browsers using the service.
- ◆ You do not normally need to enable the CONNECT method for transparent proxy services.

For most installations, only forward proxies need to support the CONNECT method.

However, if you configure your transparent proxy to intercept traffic normally intended for a forward proxy (such as port 8080), then the transparent proxy might also need to allow the CONNECT method.

- ◆ The Allow HTTP CONNECT Method and Allow only SSL CONNECT Traffic options should always be used together.

This protects your appliance and network against the establishment of tunneled connections that are not SSL-related. Not using this protection makes your installation vulnerable to the risks described in “**An Unverified CONNECT Connection Is a Security Risk**” on page 191.

28 Automatic Configuration Mechanisms

The following table summarizes the tasks discussed in this chapter.

To	See
Learn about appliance configuration files	“About Appliance Configuration Files” on page 195
Learn about the three methods of managing configuration files	“Managing Configuration Files” on page 197
Save appliance configurations	“Using Customized Configuration Files to Change the System Configuration” on page 197
Change the current appliance configuration	“Using Customized Configuration Files to Change the System Configuration” on page 197
Back up the appliance configuration	“Backing Up the Appliance Configuration” on page 198
Troubleshoot appliance configurations	“Verifying Appliance Configurations” on page 199
Configure multiple appliances	“Creating Appliance Configuration Shortcuts” on page 199
Restore the original factory settings	“Restoring Factory Settings” on page 199
Change the clone image	“Restoring the Appliance to the Clone Image” on page 200
The appliance uses this to restore the system if it senses the system has become unstable.	
Reimage the appliance	“Reimaging and Restoring the Appliance System” on page 201

About Appliance Configuration Files

Configuration files are ASCII text files that store the command line syntax used to configure the appliance. Each line in the file represents a single configuration command. When you use the browser-based management tool, the system generates multiple commands in the correct order to cause the configuration changes you specify. These commands are then recorded, in the correct sequence, in configuration files on the appliance.

The following is a clip from a configuration file with the missing portion indicated by the ellipsis (...).

```
set eth1 name=eth1
...
set eth1 speed=default
```

```
set eth1 duplex=default

clear eth1 address

add eth1 address=10.1.1.2,mask=255.255.255.0

set eth0 name=eth0

set eth0 speed=default

set eth0 duplex=default

clear eth0 address

add eth0 address=10.1.1.1,mask=255.255.255.0

set floppy poll=no

set floppy interval=120

set floppy saveonapply=no

. . .

apply
```

System-Generated Configuration Files

Excelsior employs three configuration files, as explained in the following three sections.

The FACTORY.NAS File

This file contains the appliance configuration as it came from the factory. This is a system file that is never modified.

The CURRENT.NAS File

This file contains the appliance's current configuration settings since the last apply command was issued.

You can view this file in the browser-based management tool if you are interested in seeing all the commands used to create the current appliance configuration. To view the file, click System > click Import/Export > select Current under Configuration Files on Appliance > click Download.

The AUTOLOAD.NAS File

This file is saved by Excelsior whenever a floppy disk is in the appliance's floppy disk drive and automatic polling is enabled.

NOTE: System monitoring of the AUTOLOAD.NAS file is enabled by default with a polling interval of 30 seconds. You can change these settings on the Import/Export tab in the browser-based management tool. See ["Import/Export Tab" on page 308](#).

By default, the AUTOLOAD.NAS file contains the appliance's configuration settings since the last Apply command was issued.

After the system is re-imaged and after the clone image is applied, Excelsior checks the floppy disk for an AUTOLOAD.NAS file. If the file is found, Excelsior immediately applies the commands it contains.

AUTOLOAD.NAS is useful in two situations:

- ♦ If you have reimaged the appliance, you can quickly configure it by inserting a floppy disk containing an AUTOLOAD.NAS file you have previously saved. For more information, see [“Reimaging and Restoring the Appliance System” on page 201](#).
- ♦ If you want to apply specific configuration settings (a filtering configuration, for example), you can save only these settings in an AUTOLOAD.NAS file and have the system automatically apply the settings when you insert the diskette.

IMPORTANT: Remember that the last command in AUTOLOAD.NAS is always the Apply command, which causes the system to immediately update AUTOLOAD.NAS.

You can prevent automatic updating of AUTOLOAD.NAS by opening the write-protect tab on the floppy disk after the desired AUTOLOAD.NAS file is in place or by setting the file properties of AUTOLOAD.NAS to read-only using a separate workstation.

Using Customized Configuration Files to Change the System Configuration

In addition to using system-level configuration files, Excelerator lets you save the appliance’s current configuration to arbitrarily named .NAS files and apply the configuration files you have created back to the system.

The Import feature lets you save backup copies of the appliance configurations you have created, and the Export feature lets you quickly apply any previously backed-up configuration to the appliance.

For more information about importing and exporting configuration files, see [“Managing Configuration Files” on page 197](#), [“Backing Up the Appliance Configuration” on page 198](#), and [“Import/Export Tab” on page 308](#).

Configuration files have an 8.3 DOS-style filename, the last three characters of which must be NAS.

You can save the configuration settings on the appliance or to a floppy disk on the appliance through the browser-based management tool, Telnet, and the command line interface. You can then quickly reconfigure the appliance using the configuration files.

IMPORTANT: We recommend storing copies of your customized configuration files on a floppy disk. This ensures you have the files if the clone image is ever applied or the appliance is ever reimaged.

[“Backing Up the Appliance Configuration” on page 198](#) and [“Creating Appliance Configuration Shortcuts” on page 199](#), describe two situations in which having customized configuration files is an advantage.

Managing Configuration Files

You can manage appliance configuration files using the browser-based management tool, Telnet or the command line interface, or using the appliance’s FTP functionality. The next three sections briefly explain how to use each of these management options.

Using the Browser-Based Management Tool

You can export and import configurations and manage the creation of the autoloading configuration from the browser-based management tool. For more information, see [“Import/Export Tab” on page 308](#).

Using Telnet or the Command Line

From Telnet or a command line, you can import and export configuration files. Do *not* specify the three-digit NAS extension when using either of these methods.

If You Want To	Then Enter	Notes
Apply an autoloading file from floppy	<code>import floppy</code>	First verify that the disk containing AUTOLOAD.NAS is inserted into the appliance.
Export a named configuration file to the appliance's hard drive	<code>export filename</code>	<i>Filename</i> is the name of the configuration file without the .NAS extension specified.
Export a named configuration file to a floppy	<code>export filename, floppy</code>	The file will be saved on the DOS-formatted floppy disk inserted into the appliance.
Apply a named configuration file on the appliance's hard drive	<code>import filename</code>	<i>Filename</i> is the name of the configuration file without the .NAS extension.
Apply a named configuration file on a floppy	<code>import filename, floppy</code>	The file will be loaded from the DOS-formatted floppy disk inserted into the appliance.

Using FTP

You can use FTP to move the configuration files to and from the appliance using the get and put commands. You can also apply a configuration file you are moving by using the execute option specified after a comma on the command line.

After starting the FTP client and pointing it to an IP address for the appliance (see [“Starting an FTP Session with the Appliance” on page 263](#)), use one of the following commands, where *filename* is the name of your configuration file.

Command	Description
<code>get filename.nas</code>	Downloads the specified configuration file to your FTP local directory on your client workstation
<code>put filename.nas</code>	Uploads the specified configuration file from the FTP local directory to the appliance
<code>put filename.nas,execute</code>	Uploads the specified configuration file and applies it to the appliance

Backing Up the Appliance Configuration

Because the automatic backup configuration file, AUTOLOAD.NAS, is automatically updated by the caching system each time any configuration change is applied (see [“The AUTOLOAD.NAS File” on page 196](#)), we recommend saving your appliance configuration to a .NAS file with another name.

If you ever need to reimage your appliance for some reason, having an alternatively named .NAS file will provide a configuration backup in case the AUTOLOAD.NAS file is overwritten with factory settings before your configuration is restored.

If you want to automatically load a configuration using the settings in an alternatively named file, you can use a workstation and replace the system-created AUTOLOAD.NAS file with a copy of your file on the floppy disk before inserting it into the appliance's floppy disk drive.

Verifying Appliance Configurations

An IMPORT.LOG file is created on the appliance whenever a .NAS file is imported. This file reports a success or failure for every parameter executed in the .NAS file. You can view the file via Netscape or FTP.

You can use the IMPORT.LOG file to verify if all of the commands in your .NAS file properly executed without errors. If you suspect that an error may have occurred after the execution of a .NAS file, use one of the methods explained below to access the IMPORT.LOG file.

Viewing the IMPORT.LOG File via Netscape

NOTE: The IMPORT.LOG file does not successfully display in Internet Explorer at this time.

To view the IMPORT.LOG file via Netscape, enter the following URL (the IP address and the port number are those of the desired appliance):

`http://ip_address:port_number/log/logs/import/import.log`

Viewing the IMPORT.LOG via FTP

To view the IMPORT.LOG file via FTP, use the following path:

LOG:ETC/PROXY/DATA/LOGS/IMPORT/IMPORT.LOG

For more information on using FTP with log files, see [“About the FTP Log Push Feature” on page 245](#).

Creating Appliance Configuration Shortcuts

You might want to have more than one configuration for an appliance, depending on business or other conditions. An alternate method to manually reconfiguring the appliance is to save various configurations in separate configuration files and use these to turn services on and off through FTP services.

For example, you could use two files named FORWARD.NAS and REVERSE.NAS to quickly configure the appliance to provide the services indicated by the filenames.

Restoring Factory Settings

You can quickly return the appliance to its original factory settings from the browser-based management tool, a Telnet session, or the command line. After restoring factory settings, you must either reinitialize the appliance as described in the *Getting Started* guide or use a previously created AUTOLOAD.NAS file on a floppy diskette to restore the appliance's configuration settings.

An appliance's original factory settings include the following:

- ♦ The eth0 network adapter is bound to IP address 10.1.1.1 on subnet 10.1.1.0 with a subnet mask of 255.255.255.0.
- ♦ Other network adapters have no addresses bound.
- ♦ No caching, proxy cache, appliance cluster, caching hierarchy, filtering, or other appliance services are configured.

WARNING: Restoring factory settings removes all the settings you have configured except passwords. This includes network addresses and all appliance cache services.

In most cases, you can automatically restore the settings if you have prepared an AUTOLOAD file on a floppy disk. See [“The AUTOLOAD.NAS File” on page 196](#) and [“Import/Export Tab” on page 308](#).

You should also prepare an alternatively named backup configuration file as a precaution. For further details, see [“Backing Up the Appliance Configuration” on page 198](#).

From the Browser-Based Management Tool

- 1** Click System > Actions > Factory Settings.
- 2** Restore factory settings by clicking Restore.
or
Cancel the action by clicking Do Not Restore.

From a Telnet Session or the Command Line

- 1** At the system prompt, enter
factorysettings
- 2** Do one of the following:
Restore factory settings by entering **apply**.
or
Cancel the action by entering **cancel**.

After restoring factory settings, you must either reinitialize the appliance as described in the *Getting Started* guide or use a previously created AUTOLOAD.NAS file on a floppy diskette to restore the appliance's configuration settings. See [“The AUTOLOAD.NAS File” on page 196](#) and [“Import/Export Tab” on page 308](#).

Restoring the Appliance to the Clone Image

Each appliance stores a clone image that, initially, is the same as the factory image. If the appliance experiences an abnormal shutdown four times within a half hour period, or if it is restarted six times within a half hour period, the appliance assumes the current configuration is faulty and automatically replaces it with the clone image.

If the default factory image is restored, you must either reinitialize the appliance using the instructions in the *Getting Started* guide, or, if you have saved the appliance configuration, use an AUTOLOAD.NAS file to restore the configuration. See [“The AUTOLOAD.NAS File” on page 196](#).

To prevent automatic restoration to the default factory settings in the event of system problems, you can overwrite the default clone image after you have applied an alternate configuration to the

appliance. You can also apply the clone image as an alternate method for reconfiguring the appliance. For more information, see [“Actions Tab” on page 303](#).

IMPORTANT: You should update the clone image whenever you perform an upgrade. Be aware, however, that this process causes the appliance to reboot, resulting in a temporary interruption of services. See [“Upgrade Tab” on page 310](#) and [“Actions Tab” on page 303](#).

Reimaging and Restoring the Appliance System

The appliance comes with a CD that can be used to reimage the system. This reformats the hard disks and reinstalls the Excelsior system. After reimaging an appliance, you must either reinitialize it as described in the *Getting Started* guide or use a previously created AUTOLOAD.NAS file to restore your configuration settings.

WARNING: Reimaging the system removes all the settings you have configured, including passwords, network addresses, and all cache services.

In most cases, you can automatically restore the settings if you have prepared an AUTOLOAD file on a floppy disk. See [“System-Generated Configuration Files” on page 196](#) and [“Import/Export Tab” on page 308](#).

You should also prepare an alternatively named backup configuration file as a precaution. For further details, see [“Backing Up the Appliance Configuration” on page 198](#).

Complete the following steps to reimage and restore an appliance using an AUTOLOAD.NAS file:

- 1 Locate the appliance system CD.

IMPORTANT: If the system CD is in the appliance, remove the CD, shut down the appliance, recycle the appliance's power switch, and allow the appliance to restart.

- 2 If the appliance configuration has not been previously saved, insert a formatted, blank floppy disk into the appliance's floppy diskette drive.

If you have previously saved the appliance configuration, skip to [Step 6](#).

- 3 If you have access to the appliance through the browser-based management tool, click System > Import/Export; otherwise, skip to [Step 4](#).

If the floppy disk contains an AUTOLOAD file, skip to [Step 5](#).

If you need to create an AUTOLOAD file on the floppy disk, type **autoload** in the Export Configuration File to Floppy field > click Export To > skip to [Step 5](#).

- 4 If you do not have appliance access through the browser-based management tool, establish a Telnet or null-modem session with the appliance. (You can also use an attached keyboard and monitor if your appliance has the required connections.)

At the appliance command line, enter the following:

```
export autoload floppy
```

An AUTOLOAD.NAS file is created on the floppy disk.

- 5 Remove the floppy disk from the appliance.
- 6 Open the appliance CD tray and insert the appliance system CD.
- 7 Turn the appliance's power switch off, wait a few seconds, and then turn the appliance back on.

The CD automatically launches and the appliance reinitializes.

- 8 After the initialization process starts, insert the configuration diskette with the AUTOLOAD.NAS file into the appliance.

- 9** After all disk activity ceases and the system prompt appears, remove the appliance system CD and the floppy disk.
- 10** Shut down the appliance, recycle the appliance power switch, and wait for the system prompt to appear or for the start-up beep sequence to sound.

The appliance should now be restored to its previous operating configuration.

29

Content Filtering

The following table summarizes the tasks discussed in this chapter.

To	See
Learn about appliance-based content filtering	“Overview of Filtering” on page 203
Understand when filtering starts	“Understanding When Filtering Actually Starts” on page 204
Learn the order in which Excelerator processes the various content filtering components	“The Filter Processing Sequence” on page 204
Configure a filtering service	“Configuring a Filtering Service” on page 205
Understand and use filtering’s virtual memory management feature	“Configuring RAM Usage on a Low-Memory Cache Device” on page 205
Change the filter list download schedule	“Changing the Default Download Schedule for a Filtering Service” on page 206
Learn about the bypass list	“The Bypass List” on page 206
Learn about and use the override list	“The Override List” on page 206
Use filtering in CERN and ICP hierarchies	“Critical Information about Filtering in CERN and ICP Hierarchies” on page 207
Bypassing an N2H2 Filtering Service	“Bypassing an N2H2 Filtering Service” on page 207

Overview of Filtering

Appliance-based filtering uses lists of URLs as filters for browser requests. If a URL matches a filter setting, access to its objects is blocked, monitored, or immediately vended, depending on the setting.

You can have the appliance perform URL filtering to block specific URLs or always allow them to be vended. You can create your own Never Allow and Always Allow lists. You can also subscribe to one or more cooperating filter services and use their lists.

Filter service providers create category names that identify some aspect of Web data. Typical category names might include Violence, Language, or Adult Content. Service providers rate Web pages as true or false for each category name. A page rated true for Violence has violent content as defined by the service provider.

The appliance uses filter rules you create in combination with category ratings from your service provider to determine whether to vend, monitor, or block a page.

Understanding When Filtering Actually Starts

After you subscribe to a filtering service, appliance-based filtering doesn't start until a rating list is successfully downloaded from the filtering service. The appliance will not download the rating list until you check Enable in the Service List. (See [“Filtering Tab” on page 390.](#))

The PICS description file, which normally downloads immediately after the service is defined and which you use to configure your service, contains only category names. It does *not* supply the data required for content filtering.

No Filtering During the Initial Download

The initial download process can take several minutes to complete. During this time, content filtering might be enabled, but it is not effective. Filtering begins immediately after the initial rating list is downloaded.

Subsequent downloads are made directly into memory so that the changes to the filtering list are effective immediately.

Filtering During Appliance Startup

Each time an appliance starts up, content filtering doesn't begin until the rating list is loaded into appliance memory. Because rating lists can be quite large, the memory-loading process can require up to one minute to complete.

Also, new rating lists are downloaded each time the appliance is restarted. The appliance will repeat a download request on the quarter hour (at 12:00, 12:15, 12:30, and so on) until the download is successful.

If You Remove a Filtering Service

If you remove a defined service for whatever reason and then reinstall it, you should verify the next rating list download time. It might be set for the next day, leaving your network without appliance-based filtering until the download is completed and the filter file is loaded into memory.

The Filter Processing Sequence

To understand appliance filtering behavior, you need to know the order in which filtering processes are applied.

1. First, the appliance checks the Bypass list. Requests from IP addresses in the list are always serviced without filtering.
2. Next, the appliance processes the Override list. URLs covered by the IP address masks in the Override list are either always vended or never vended, as specified. They are not affected by other filtering.
3. Finally, the appliance processes the enabled filter services.

For details on how the appliance processes filter categories, see [“The Categories List” on page 394.](#)

For more information about appliance filtering, see [“Filtering Tab” on page 390.](#)

Configuring a Filtering Service

To set up a filtering service on the appliance, do the following:

- 1 Subscribe to a cooperating filter service.

After subscribing, you will receive a download URL, an account name, and an account password.

- 2 In the browser-based tool, click Cache > Filtering > Insert under Service List.

- 3 In the Insert Filter Service dialog box, fill in the following fields:

Service Name: A name, with no spaces or special characters, that you use to identify the filter service.

Configuration File URL: The URL you received when you subscribed to the filter service.

Account Name: The account name you received when you subscribed to the filter service.

Account Password: The password you received when you subscribed to the filter service. This appears as clear text in the box, but it will be transmitted over the wire using HTTPS.

- 4 Click OK.

- 5 Ensure that Enable is checked > click Apply.

- 6 After the PICS file downloads successfully, click Modify and configure appliance behavior for each category name.

See the help file provided by your service provider or the information in [“Modify Filter Service Dialog Box” on page 393](#) for more information.

Configuring RAM Usage on a Low-Memory Cache Device

Excelsior 2.3 has virtual memory management of filtering rating lists. This means that filtering can be installed on low-memory cache devices. (The minimum Excelsior RAM requirement is 256 MB.)

By default, Excelsior allocates ten percent of a cache device’s RAM for use by filtering rating lists. On low-memory devices, portions of the rating lists often won’t fit in RAM and are therefore swapped between RAM and the device’s hard disks as required.

For most low-memory devices, the ten percent allocation is optimal. However, you can set the percentage of RAM used for rating lists to any value from 10 to 25 percent by entering the following command at the System Console prompt:

```
set filter memorypercent=percent
apply
```

where percent is a number from 10 to 25.

IMPORTANT: Memory that you allocate to filtering services can no longer be used for proxy services. Therefore, you should closely monitor proxy service performance to ensure your changes are properly balanced.

One scenario that might justify a memorypercent setting higher than 10 would be if you add RAM to your low-memory device without installing additional disk space. In this case, chances are good that the additional RAM normally allocated to proxy services would not be completely utilized by the services.

Changing the Default Download Schedule for a Filtering Service

A filtering service is scheduled for its first download when you check Enable and click Apply. If you want to view or modify the schedule, or any of the attributes of the filtering service, use the browser-based management tool. Click Cache > click Filtering > select the intended list > click Modify.

You can change any of the original values you set when you created the service list item, and you can specify the frequency of downloads. You can also view or change when the next download will occur.

It is a good idea to schedule filter downloads for off-peak times on your network.

The Bypass List

You can create a bypass list of IP addresses. Excelsator automatically services all requests from these addresses without filtering and without checking the override list. For information regarding how Excelsator processes requests, see [“The Filter Processing Sequence” on page 204](#).

If people on your network require filter-free access, access the [Filtering Tab](#) and insert their workstations' IP addresses into the bypass list.

The Override List

You can create an override list of URLs that should be either always vended or never vended. The override list always takes precedence over any conditions specified in your filtering services, but it does not affect requests from IP addresses in the bypass list.

Creating an Override List

To add entries to the override list, do the following:

- 1** In the browser-based management tool, click Cache > Filtering > Insert under Override List.
- 2** Enter the URL or IP address mask you want to add to the list.
Use the asterisk (*) wildcard character to limit or broaden the effect. See [Critical Information about Wildcards in the Override List](#) below for help.
- 3** Click the Allow drop-down list > select Always or Never.

Always will always vend the page; Never will never vend the page.

Critical Information about Wildcards in the Override List

An asterisk (*) wildcard in the override URL mask causes the appliance to interpret everything between the asterisk and the following delimiter as a wildcard.

Delimiters include the forward slash (/), the period (.), and the colon (:).

The appliance automatically appends http:// at the beginning of the URL mask. For example, the mask *.*.edu would instruct the appliance to allow URLs such as the following to bypass filtering.

http://www.*.edu

http://ww1.*.edu

`http://*.*.edu/*`

Asterisks must be used with caution. Otherwise, you could allow undesirable content to bypass filtering. For example, some education sites contain personal pages which would be undesirable in certain environments.

When defining override URL masks, be specific and include as much of the complete path as possible.

IMPORTANT: Do not include asterisks in the protocol indicator (scheme) of override list URLs.

For example, the protocol indicator `http*://` is invalid.

If you include an asterisk in the protocol indicator, you can only remove the entry by clearing the entire override list.

Critical Information about Filtering in CERN and ICP Hierarchies

Filtering services rely on the appliance performing DNS resolution.

If an appliance has filtering configured, is configured as a client in a CERN or ICP hierarchy, and has the Must Only Forward through Hierarchy option checked in Hierarchy > ICP/CERN Configuration, it must be able to access a DNS server.

For more information on the Must Only Forward Through Hierarchy option, see [“ICP/CERN Configuration Tab” on page 405](#).

Bypassing an N2H2 Filtering Service

Filtering functionality is enhanced in Excelsior 2.3 so that an N2H2 filtering service or an override list entry can be temporarily bypassed, allowing objects to be cached that would otherwise be blocked.

NOTE: The filter service bypass functionality does not apply to installed X-Stop filtering services.

This feature can be useful for school situations wherein a teacher can either cache content or permit content to be cached that would normally be blocked. For example, sites dealing with breast cancer research might normally be blocked, in which case a teacher could cause the sites to be cached and make the information available to students writing research papers.

Critical Information Regarding this Feature

Before deploying this feature, it is critical that you understand exactly how it works as outlined in the following sections.

Understanding How the Feature Works

Consider the following points:

- ♦ **Blocking Prevents Objects from Being Cached:** Excelsior filtering services and override lists block specified Web content from coming through the proxy service. Content that can't get through the proxy service can't be cached.

Conversely, if content is not blocked and if it is cachable, it is cached automatically as part of the retrieval process.

- ♦ **Filtering Services Prevent Access Only Where Installed:** An installed filtering service or an override list will prevent general access to cached content on the device where they are installed or active.

However, they do not prevent general access to cached content on other devices. For example, a filtering service on a hierarchical parent will not prevent access to cached content on a hierarchical child that doesn't have a filtering service installed.

- ♦ **Unblocked Cached Content Is Openly Available:** Once content is cached, anyone accessing the Web through the cache device has access to the cached content if a filtering service or override list on the box doesn't block the access.
- ♦ **Hierarchical Caching Broadens the Distribution of Cached Objects:** When a cachable object is requested and retrieved through a hierarchy, any hierarchical parents or peers involved in retrieving the objects will also have cached copies of the objects.

Therefore, anyone accessing the Web through these cache devices will also have access to the objects unless they are blocked by a filtering service or override list on the device.

- ♦ **Unblocked Access Is Unrestricted Access:** Once filtering has been bypassed for a workstation, anyone using the workstation has unrestricted Web access for the period of time specified when access is granted.

It is imperative, therefore, that unblocked access is carefully monitored and controlled.

What the Feature Doesn't Provide

- ♦ **Access Control to the Web:** After a workstation is authorized to bypass filtering, Web access is unrestricted.
- ♦ **Object Caching Controls:** Once filtering is bypassed, all cachable objects accessed through the cache device will be cached.
- ♦ **Access to Cached Objects:** Once objects are cached, they become generally available unless one of the following happens:
 - ♦ A filtering service or override list *on the box* blocks access
 - ♦ The objects expire
 - ♦ The objects are purged from cache

Hierarchical Considerations

Organizations often set up hierarchies wherein all content requests are funneled through one or two parent cache devices. Such organizations usually install filtering services only on the parent devices since these can then act as gatekeepers for all content entering the hierarchy.

Obviously, when filtering is bypassed and content that would normally be blocked is instead brought into the hierarchy, the scenario changes substantially. Requests made directly to the parent devices with filtering installed will still be blocked. However, requests to devices that don't have filtering installed will be filled directly from cache.

Managing the Impacts of Using this Feature

Organizations generally deploy filtering services and override lists to ensure that only appropriate content is cached and made generally available. For this reason, any decisions to bypass these safeguards must be carefully considered.

We recommend you consider the following usage suggestions:

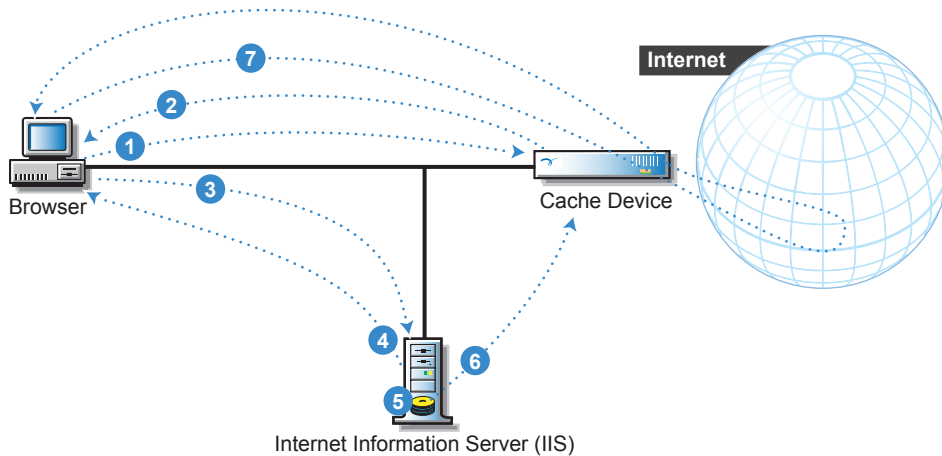
1. Use the feature only as a means for caching content needed for specific purposes. For example, a teacher might bypass filtering to access sites on breast cancer research. If site objects are cachable, the information will be cached and become generally available to everyone accessing the Web through the cache device, as long as a filtering service or override list on the device doesn't block access.
2. In caching hierarchies, consider that content will be cached on all parents and peers involved in retrieving objects. Evaluate whether the content is appropriate to other groups who will access the cache devices that don't have filtering services or override lists. If the content is not appropriate, perform a purge on the devices that the other groups will access.
3. When the cached content is no longer needed, use a selective purge to remove it from cache.

IMPORTANT: If undesirable content is accessed and cached for whatever reason, you will need to use the selective purge to remove it from the cache device and any hierarchical parents or peers involved in retrieving the content.

Overview

Filter bypass functionality is summarized in **Figure 66**.

Figure 66



- 1 A user requests access to a Web page through a service on a cache device with N2H2 installed or an override list defined.
- 2 N2H2 or the override list indicates that the Web page is blocked. The device then sends the browser an Access Denied page containing a form for requesting permission to bypass the block.
- 3 The user provides a user id, password, and bypass duration, and the browser sends this to the Internet Information Server (IIS).
- 4 If the information is verified, the IIS notifies the browser that the request has been granted and the process continues with Step 5.
If the information is not verified, the process repeats from Step 2 until an unblocked URL is requested.
- 5 The IIS adds a line to its bypass file that contains all IP addresses currently authorized to bypass filtering and override blocks.
- 6 Within 60 seconds, IIS transmits an updated bypass list to the device through FTP.
- 7 After waiting up to 60 additional seconds, the user has unrestricted Web access. Any content accessed is cached.
After the specified bypass duration expires, the user no longer has unrestricted access. However, on any devices that are members of hierarchies and that don't have filtering services or override lists, content that has been cached is generally available until it either expires or is purged from cache.

Preparing to Install the Filter Bypass Software

To deploy the filter bypass solution, you need the following:

- ♦ At least an intermediate understanding of the following:
 - ♦ Excelerator proxy services
 - ♦ N2H2 filtering on Excelerator
 - ♦ General networking
 - ♦ Microsoft Internet Information Server Web services
- ♦ N2H2 filtering software or override lists configured and active on the cache devices
- ♦ A Windows 2000 computer with at least a Pentium III 500 Mhz processor and a minimum of 256 MB RAM installed
- ♦ A static (non-DHCP), publicly accessible IP address assigned to the Windows 2000 computer.
- ♦ Microsoft Internet Information Server (IIS) version 5.0 or later installed and running as a Web server on the Windows 2000 machine
- ♦ Active State's ActivePerl 5.6.x or later installed and running on the Windows 2000 IIS computer

Complete the following steps:

- 1** Install and configure an N2H2 filtering service or create an override list on each Excelerator 2.3 cache device.

For help installing N2H2 filtering services, refer to the N2H2 product documentation. For more information on creating filter override lists, see [“Creating an Override List” on page 206](#).

- 2** Enable the Mini-FTP server on each Excelerator cache device.

For more information, see [Chapter 35, “FTP Services,” on page 261](#).

- 3** Install the Microsoft Internet Information Server (IIS) on the Windows 2000 computer.

This software is available on the Windows 2000 CD if one came with the computer or on [Microsoft's Web site \(http://www.microsoft.com/\)](http://www.microsoft.com/).

- 4** Download and complete a default installation of ActiveState's ActivePerl 5.6.x or later from [ActiveState's Web site \(http://www.activestate.com/\)](http://www.activestate.com/).

NOTE: The download is free and should be installed using the default options for use with Microsoft IIS as outlined in the accompanying documentation.

If the default options are not used, for example if you install in another location, be sure to follow the tips presented in [“Tips for Non-Default Installations” on page 214](#).

- 5** Ensure that the Perl executable appears in the path on the Windows 2000 computer by running CMD.EXE and entering PATH at the command prompt.

Installing the Filter Bypass Software

Complete the following steps:

- 1** Using an FTP client, access the Excelerator 2.3 device's default FTP directory (/etc/proxy/appliance/config/user) and retrieve the N2H2FBYP.EXE file.

- 2** On the Windows 2000 computer, run the N2H2FBYP.EXE from the Start menu and unzip the files to the default location.
- 3** Run the command processor (CMD) from the Start menu, and using the CD command, change to Documents and Settings*username*\Local Settings\Temp, where *username* is your Windows 2000 user name.

NOTE: Local Settings is a hidden directory and won't appear in a directory listing (DIR) at the command line. However, the directory is accessible and the CD command should work.

- 4** In the command processor window, enter INSTALL.BAT.

This batch file copies the following files to the directories indicated:

- ♦ The root of the IIS Web server (default is C:\inetpub\wwwroot)
Adduser.html
- ♦ The IIS Web server's cgi-bin directory (default is C:\inetpub\wwwroot\cgi-bin)
adduser.pl
login.pl
login.htx
logout.pl
logout.html
- ♦ The filter work directory (default is C:\filter)
resetAdminPass.pl
configex.pl
sysAdminPassword.dat

Configuring the Main Perl Script on the IIS Server

After installing the filter bypass software on the IIS server, you must modify its Perl scripts to communicate with your Excelerator cache devices.

Complete the following steps:

- 1** Using an ASCII text editor, open the configex.pl file in the C:\filter directory.

NOTE: This is the main Perl script that communicates with the Excelerator cache device.

- 2** Locate the following line in the file:

```
$uploadServers = "<Mini-FTP_IP_Address>:config:<Pwd>";
```

This line specifies the Excelerator Mini-FTP servers to which the IIS server will send bypass list files.

IMPORTANT: In the steps that follow, modify only the text between the double quotation marks ("...").

- 3** Replace <Mini-FTP_IP_Address> with the IP address of one of the Excelerators' mini-FTP servers to which you want the bypass lists sent.
- 4** Do not change the word *config* in the string. This references the Excelerator config user - the user with rights to copy files using FTP.
- 5** Replace the word <Pwd> with the password of the config user on the Excelerator.
- 6** If you want the bypass lists sent to multiple Excelerators, create new entries (<Mini-FTP_IP_Address>:config:<Pwd>) within the double quotation marks for each Excelerator and separate each entry using commas (,).

For example, "1.1.1.1:config:pwd1,2.2.2.2:config:pwd2, ...".

- 7 If you have not installed all components to their default locations, complete the modifications outlined in [“Tips for Non-Default Installations” on page 214](#).

Configuring the IIS Server To Run the Perl Scripts

Complete the following steps to enable the IIS server to run Perl scripts located in the cgi-bin directory:

- 1 On the IIS server, click Start > Settings > Control Panel > Administrative Tools.
- 2 Run the Internet Services Manager.
- 3 Browse in the tree to the Default Web Site > cgi-bin folder.
- 4 Right-click on cgi-bin > select Properties.
- 5 Under Execute Permissions, select Scripts and Executables.

Changing Each Excelerator’s Default Blocked Content Page

You must change the default HTML page that Excelerator serves when access to blocked content is attempted so that the completed form can be sent to the IIS. You must then FTP this page to each Excelerator.

Complete the following steps:

- 1 Browse to C:\Documents and Settings\username\Local Settings\Temp\misc where username is the Windows 2000 user name you used when installing the filtering bypass software, and make a copy of the picsblok.htm file.
- 2 Using an ASCII text editor, open the picsblok.htm file.
- 3 Locate the string <IP_OF_WEB_SERVER> in the file and replace only this string with the IP address of the IIS server.

IMPORTANT: If you choose to modify other portions of the picsblok.htm file, *do not* modify the following lines:

```
<input type="hidden" name="REFERER" value="#REFERER#">
```

```
<input type="hidden" name="FORWARDED_FOR_IP" value="<X-FORWARDED-FOR>"
```

These two lines are critical elements of the form that cannot be changed or deleted.

- 4 Using an FTP application, copy the modified picsblok.htm file to the etc/proxy/data directory on the Excelerator cache device where filtering is installed.

Setting Up Administrative Access to the List of Bypass Users

Access to the list of filter bypass users and their passwords is controlled by a single administrative password. To set or change this password, complete the following steps:

- 1 On the IIS server, run the command processor (CMD) from the Start menu.
- 2 Change to the C:\filter directory, or to the directory where the setadminpass.pl script file is located.
- 3 Enter the following command:

```
perl setadminpass.pl xxxxxx
```


where *xxxxx* is the administrative password you want to use when modifying the filter bypass user database.

Adding Authorized Filter Bypass Users

IMPORTANT: This section lets you designate the people who will be able to bypass filtering. In most situations, this will be a supervisor (or teacher in a school environment) who has the responsibility to see that only appropriate content is accessed and cached.

To add authorized users to the filter bypass database, complete the following steps:

- 1** In a browser, access the Add User page on the IIS by entering the following URL:
`http://iii.iii.iii.iii/AddUser.html`
where *i* is the IP address of the IIS Web server.
- 2** In the SysAdmin Password field, type the administrative password created in [“Setting Up Administrative Access to the List of Bypass Users”](#) on page 212.
- 3** In the UserName field, type the user id for one of the filter bypass users.
- 4** In the Password field, type the password this user will use to bypass filtering.
- 5** In the Verify Password field, re-type the password for this user.
- 6** Click Add User.
- 7** To create additional users, repeat the process starting with [Step 2](#).

Starting the Filtering Bypass Service

NOTE: All maintenance of bypass lists, expiration checking, etc. occurs on the IIS server. It is therefore critical that the server communicate regularly with the Excelsator cache devices. Otherwise, the devices won't receive timely bypass list updates, and the expiration of IP addresses will also not be communicated to the cache devices.

To start the IIS server sending bypass lists to your Excelsators, you need to run a Perl script on the server.

Complete the following Steps:

- 1** On the IIS server, run the command processor (CMD) from the Start menu.
- 2** Change to the C:\filter directory, or the directory containing the configex.pl Perl script.
- 3** Enter the following command:

```
perl configex.pl
```

By default, this script runs continually, processing new and expired bypass authorizations every 60 seconds. It also provides a status update every 60 seconds indicating that the script is still running.

Running in a Hierarchy with Filtering on the Parent

If your users are accessing the Web through a proxy that is a member of a hierarchy in which filtering is installed on the parent cache device, you must do the following on the child device:

- ♦ Check the X-Forwarded-For option in the proxy service definition
- ♦ Check the Must Forward Through Hierarchy option in the hierarchy definition

IMPORTANT: Be aware that any objects cached on a device that fills through a hierarchy will also be cached on all other devices involved in retrieving the content. For a review of the implications of caching through a hierarchy, see “[Managing the Impacts of Using this Feature](#)” on page 208.

Starting the Filtering Bypass Service Automatically

IMPORTANT: You should start the bypass service from the command line at first and watch for any error messages or other issues requiring troubleshooting.

If you want to have the script start automatically each time the server reboots, you can use `srvany.exe` and `instsrv.exe`, available with the W2K Resource Kit, to convert `configex.pl` to an NT service. More information is available on the [Web \(http://www.perlguy.com/articles/nt_service.html\)](http://www.perlguy.com/articles/nt_service.html).

Alternatively, you can also use a third party CRON utility to run the script every 45 to 60 seconds.

Tips for Non-Default Installations

If you have installed any components (ActivePerl, the bypass software, Internet Information Server, etc.) to non-default locations, you must check the `.html`, `.htx`, and `.pl` files in the following locations and modify any affected paths.

- ♦ `non_default_directory_path/filter`
- ♦ `non_default_directory_path/inetpub\wwwroot`
- ♦ `non_default_directory_path/inetpub\wwwroot\cgi-bin`

where `non_default_directory_path` represents the directory where you’ve installed the bypass software.

Notes About Product Security

The filtering bypass software manages a simple database of userids and passwords for the users authorized to grant temporary unrestricted Web access, such as teachers who need to cache specific content in an educational environment.

Passwords in this database are stored in their hashed format, and the database is not stored in a path that is accessible to the HTTP server. A would-be hacker would have to compromise Windows 2000 security to crack the passwords in the database.

Creating the userids and passwords in the database is handled through the HTML input password mechanism. Anyone sniffing the wire would be able to see the password in plain text.

Access to the database requires an administrative password. For more information on setting this password, see “[Setting Up Administrative Access to the List of Bypass Users](#)” on page 212. Since this password is created on the box, it cannot be sniffed on the wire. One would have to look over the administrator’s shoulder to discover this password.

The IIS server sends the bypass list to the Excelerator using Excelerator’s mini-FTP functionality. All FTP security liabilities apply to this transaction.

30 DNS Name Resolution

As Excelerator processes browser requests, it uses the DNS system to obtain the IP addresses of origin Web servers.

Because the DNS names in browser requests are not always straight-forward, Excelerator tries various permutations to try and locate the Web server. As a result, DNS names ending with domain extensions other than .com, .org, and so on, are sometimes resolved in unexpected ways.

If users of your appliance are experiencing this problem, you can customize how the appliance resolves DNS names.

How the Appliance Resolves DNS Names

When the appliance receives a browser request, it creates a DNS query based on the URL in the request and sends the query to one of the DNS name servers defined for the appliance.

How the Appliance Formulates Subsequent DNS Queries

If the DNS name server can't resolve the query, the appliance formulates subsequent DNS queries based on the information in its DNSINFO.CFG file:

The DNSINFO.CFG File

The DNSINFO.CFG file lets you control the processing of partial host names. It contains three sections:

- ♦ **Domain:** The <DOMAIN> keyword instructs the DNS name resolver to try appending the appliance's domain and subdomain names first. You can prevent this initial lookup sequence by removing the <DOMAIN> keyword from the file.
- ♦ **DNSPingHost:** The DNSPingHost keyword is an optional entry that specifies the domain that Excelerator uses to validate communications with its DNS servers.

If multiple entries exist, Excelerator uses only the first. If no entry exists, the system uses www.volera.com by default. This entry is only used to verify the DNS server's availability. It is skipped by the name resolution process.

- ♦ **Format strings:** These strings give the DNS name resolver alternate names to try.

An Example

For example, assume the following:

- ♦ The browser request URL is webserver.
- ♦ The appliance's domain name is support.acme_ex2.com.

- ♦ The appliance's DNSINFO.CFG file has the following content:

<DOMAIN>

DNSPingHost=www.volera.com

www.%s.com

www.%s.ed

www.%s.org

www.%s.gov

www.%s.net

%s.com

%s.edu

%s.org

%s.gov

%s.net

www.%s

After the initial request fails, the appliance formulates subsequent requests as follows:

1. If the <DOMAIN> keyword is not found, the process skips to Step 3. Otherwise, the appliance formulates a second query by appending the appliance's domain name to the URL as follows:

webserver.support.acme_ex2.com

2. If this query fails, the appliance appends the appliance's subdomain names to the URL in order as follows:

webserver.acme_ex2.com

and then

webserver.com

If both these queries fail to return an IP address, the appliance continues with Step3.

3. The appliance skips the DNSPingHost entry and then appends each format string in the DNSINFO.CFG file in the order listed until one of the following occurs:
 - ♦ The DNS server returns an IP address for the name.
 - ♦ The appliance's query options are exhausted and it returns a DNS error to the browser.
4. If a DNS name has already been tried, the appliance skips the query and moves to the next item in the list.

Continuing with the example, the appliance would submit the following queries, substituting webserver for the %s variable in the format strings of the DNSINFO.CFG file.

www.webserver.com

www.webserver.edu

www.webserver.org

www.webserver.gov

www.webserver.net

webserver.edu

webserver.org

webserver.gov

webserver.net

www.webserver

Because webserver.com was tried previously, the appliance skips the sixth line (%s.com) in the DNSINFO.CFG file.

Modifying the DNS Lookup Sequence

To modify the DNSINFO.CFG file, complete the following steps:

- 1 Start an FTP client on a workstation with access to the appliance.

For help, see [“Setting Up Appliance FTP Services” on page 262](#) and [“Starting an FTP Session with the Appliance” on page 263](#).

- 2 Point the FTP client to one of the appliance’s IP addresses.

- 3 Enter the following command:

```
get dnsinfo.cfg
```

The file is transferred to the FTP client's default directory, which is /etc/proxy/appliance/config/user.

- 4 Referring to the example in the previous section, modify the DNSINFO.CFG file using an ASCII editor.

Ensure that the lines in your file reflect the query order and content you want the appliance to use when attempting DNS name resolution. For example, you might want to reorder the domains listed or include two-letter country codes in the list.

- 5 Use the **put** command to place the modified DNSINFO.CFG file back in the appliance’s default FTP directory.

The system polls the DNSINFO.CFG file for changes every minute. If a change occurs, the system loads the revised DNSINFO.CFG file and displays the revised name resolution scheme on the system monitor.

Managing the HOSTS File

The system uses the HOSTS file to specify entries for the DNS cache. Entries in this file override responses from a DNS server. In addition, DNS entries in DNS cache that result from entries in the HOSTS file cannot be purged. (See [“Purging the Appliance’s DNS Cache” on page 218](#).)

You can manage the contents of a cache device’s HOSTS file using the following procedure:

- 1 Start an FTP client on a workstation with access to the appliance.

For help, see [“Setting Up Appliance FTP Services” on page 262](#) and [“Starting an FTP Session with the Appliance” on page 263](#).

- 2 Point the FTP client to one of the appliance’s IP addresses.

- 3 Enter the following command:

```
get hosts
```

The file is transferred to the FTP client's default directory, which is /etc/proxy/appliance/config/user.

- 4 Modify file contents as required.

- 5 Use the **put** command to place the modified HOSTS file back in the appliance default FTP directory.

The system polls the HOSTS file for changes every minute. If a change occurs, the system loads the revised HOSTS file and displays its contents on the system monitor.

Resolving DNS Names in Hierarchies

If your cache device is a member of an ICP or CERN hierarchy, you can have it resolve DNS requests through its hierarchical parents.

To configure the device to use the hierarchy for resolving DNS requests, enter the following at the System Console prompt:

```
set icpclient forwarddns=yes  
apply
```

Purging the Appliance's DNS Cache

You can remove incorrect DNS names and IP address combinations from a cache device's DNS cache.

To purge a cache device's entire DNS cache, enter the following at the System Console prompt:

```
purgednscache
```

To purge a cache device's DNS cache for a specific host name, enter the following at the System Console prompt:

```
purgednscache hostname
```

where *hostname* is the DNS name in the DNS cache.

IMPORTANT: DNS cache entries that were specified in the HOSTS file cannot be purged as long as their corresponding entry remains in the HOSTS file. For information on managing the HOSTS file, see [“Managing the HOSTS File” on page 217](#).

31

Controlling Referred Access to Content

One of the most common practices among Web site owners is providing links to other sites. Some Web pages consist almost entirely of links to other sites.

While many Web site owners welcome all exposure to their information, others have ownership and cost issues that require them to limit which other sites can provide links (referred access) to their content.

The most common mechanism for limiting referred access utilizes the *referer* (note the misspelling) header which all browser requests contain.

When a user enters a target URL in a browser, the request doesn't contain a referer header. When the user clicks a link to this same target URL, the requesting browser includes the referring URL (for the page containing the link) in a referer header in the resulting request.

By requiring that referring URLs be validated against a list of approved URLs, Web site administrators can limit referred access to their content.

This chapter describes how to use Excelsior's referer header validation feature to control access to cached content.

How Referer Header Validation Works

Excelsior's referer header validation mechanism uses an ASCII configuration file that you create. This file tells Excelsior the following:

- ♦ Which URLs are valid referrers
- ♦ Which directories and/or objects each of these URLs can refer (link) to
- ♦ The error page that is sent to requesting browsers when objects can't be vended because the referring URL is not valid.

Understanding the Referer Header Configuration File

To set up referer header validation for a Web accelerator service, you create a referer header configuration file for the service.

Configuration File Specifications

Table 12 explains the entries that can appear in a referer header configuration file.

IMPORTANT: Each entry described in the table must appear on a separate line in the referer header configuration file.

Table 12

File Line Entry	What the Entry Does	Rules for Including in the Configuration File
<code>[errorpage: url]</code>	<ol style="list-style-type: none"> 1. This line specifies the error message file that is copied to the cache device when referer header validation is enabled on the device. 2. Excelerator responds with the contents of this file when a browser request contains a referer header that is not in the configuration file. 	<ul style="list-style-type: none"> ♦ This line is optional. If it is not included in the file, users receive the standard 403 Forbidden error page. ♦ If included, this must be the first line in the file. ♦ The square brackets are required. ♦ <i>url</i> is the URL path to a customized HTML error message file that Excelerator copies using HTTP, for example, <i>http://www.foo.org/messages/not_allowed.html</i>. <p>For more information, see “Customizing Referer Header Error Messages” on page 227.</p>
<code>[referer: url]</code>	<ol style="list-style-type: none"> 1. This line specifies a URL that Excelerator will recognize as a valid referrer for the directories and objects that are listed below it. <p>NOTE: For clarity, we refer to this entry and the directories and objects listed below it as a <code>[referer:]</code> section in the instructions and information that follows.</p>	<ul style="list-style-type: none"> ♦ The square brackets are required. ♦ There can be any number of <code>[referer:]</code> sections in the referer header configuration file. ♦ <i>url</i> is the URL as it will appear in the referer headers of browser requests you want to service (i.e. <i>http://www.foo.org/</i>)
<code>[referer: no-referer]</code>	<ol style="list-style-type: none"> 1. This line sets up a specialized <code>[referer:]</code> section that tells Excelerator to service requests that do not contain a referer header for the directories and objects that are listed below it. For example, you must use this if you want to provide direct browser access to any of the directories or objects that are listed in other <code>[referer:]</code> sections. 	<ul style="list-style-type: none"> ♦ The square brackets are required. ♦ There can be any number of <code>[referer:]</code> sections in the referer header configuration file. See the note in the column to the left.

File Line Entry	What the Entry Does	Rules for Including in the Configuration File
<i>target_url</i>	<ol style="list-style-type: none"> 1. Including a directory or object in a [referer:] section sets up a referer control for that directory or object. 2. Once you include a directory or object in a [referer:] section, you must ensure that you include it in other [referer:] sections for all other URLs that could legitimately contain a link to it. For more information, see “Two Key Points About Target_URL Entries” on page 221. 	<ul style="list-style-type: none"> ♦ If the <i>target_url</i> is on the hostname shown in the [referer:] entry, only a relative path from the root of the server is required. Otherwise, the full URL path, including the scheme and DNS name, is required. ♦ There can be any number of object URL entries in each [referer:] section. <p>IMPORTANT: Be sure you understand the information in About URL Line Entries in the Configuration File: Conventions, Wildcards, Etc. before creating your referer header configuration file.</p>

Access Is Either Totally Limited or Unrestricted

Access to directories and objects that are not covered by *target_url* entries in the configuration file is unrestricted.

Access to all directories and/or objects that match any *target_url* entry in the file is limited as follows:

- ♦ The referring URL (containing the link) must have a [referer:] section in the configuration file
- ♦ The [referer:] section must contain a *target_url* entry that matches the directories and/or objects being requested. If the [referer:] section contains a matching entry, the request is processed. If the [referer:] section contains no matching entry, the request is not filled and an error message (customized or default) is sent.

IMPORTANT: The [referer:] section requirement also applies to links on a Web server’s pages to its own directories or objects.

Two Key Points About Target_URL Entries

There are two key points to remember when setting up referer header configuration files:

- ♦ The only objects affected by referer header validation are those covered by at least one *target_url* entry in the configuration file.
An object might be covered by an object-specific entry, an asterisk (*) wildcard at the end of a URL directory path, or by multiple entries using either method.
- ♦ Once an object has been covered by one entry in the file, each [referer:] header section that you want to be able to link to objects in the directory must also have a *target_url* entry that covers (and therefore allows access to) the object.

About URL Line Entries in the Configuration File: Conventions, Wildcards, Etc.

Except for the [referer: no-referer] entry, each line in the referer header configuration file contains a URL or URL pattern.

These URLs and patterns can reference any of the following:

Table 3

Description	Useful in Entry Types	Entry Examples	Notes
A Web server and all of its subdirectories and objects	[referer:] target_url	http://www.foo.org/*	
A directory and all of its subdirectories and objects	[referer:] target_url	http://www.foo.org/flowers/*	
A range of directories	[referer:] target_url	http://www.foo.org/flower*/	This makes sense as a target entry only if the directories each contain a file that opens automatically, such as index.html.
A single directory	[referer:] target_url	http://www.foo.org/flowers/	This makes sense as a target entry only if the directory contains a file that opens automatically, such as index.html.
A range of objects	[referer:] target_url	http://www.foo.org/flowers/*.gif	
A single object	[referer:] target_url [errorpage:]	http://www.foo.org/flowers/rose.gif http://www.foo.org/error.html	

The following sections provide more information about the entries in [Table 3](#).

Using Wildcards in URL Entries

Each URL entry in a referer header configuration file can contain a single asterisk (*) wildcard. Wildcard matching works as follows:

- ◆ The wildcard matches zero or more characters up to the first occurrence of the string that follows the wildcard in the URL.
- ◆ Wildcard matching doesn't cross directory boundaries unless it occurs immediately following the final forward slash (/) in the entry (see [“Including All Subdirectories and Objects in a Single Entry” on page 223](#)).

For example, the following entry:

```
http://www.foo.org/flower*Red.gif
```

would match

```
http://www.foo.org/flowerRed.gif
```

```
http://www.foo.org/flowerRoseRed.gif
```

```
http://www.foo.org/flowerRedRed.gif
```

but not

`http://www.foo.org/flower/rose/Red.gif`

`http://www.foo.org/flowerRed.txt`

NOTE: At first glance, some assume that the third example of the matching entries is incorrect. Remember that the string that follows the asterisk is *Red.gif* and not just Red.

Matches Are Case Sensitive

The matching of directory or object requests against the configuration file is case sensitive. Continuing the example above,

`http://www.foo.org/flower*Red.gif`

would not match

`http://www.foo.org/flowerRed.GIF`

or

`http://www.foo.org/FlowerRed.gif`

because of case differences.

Including Special Characters in URL Entries

Any URLs that contain special characters, such as characters above ASCII 0x7 for multi-byte characters, must be specified using the standard URI escaping mechanism `%HH`, where HH is the hexadecimal notation of the byte value of the special character.

Including All Subdirectories and Objects in a Single Entry

If the URL entry ends with a forward slash and an asterisk (`/*`), it will match all items in that directory and in all its subdirectories. For example,

`http://www.foo.org/*`

would match all of the following entries.

`http://www.foo.org/flowerRed.gif`

`http://www.foo.org/flowerRoseRed.gif`

`http://www.foo.org/flowerRedRed.gif`

`http://www.foo.org/flower/rose/Red.gif`

`http://www.foo.org/flowerRed.txt`

Allowing a Web Server to Link to Its Own Objects

As mentioned in “[Access Is Either Totally Limited or Unrestricted](#)” on page 221, all links to any objects covered by a `target_url` entry in the configuration file must be explicitly allowed for each referring URL. This includes links to objects from the Web server containing the objects.

If all links on your Web server to its own objects are valid, you would only need to include two lines in the file:

- ♦ A `[referer:]` section that covers the Web server and all its subdirectories
- ♦ A `target_url` entry that covers all potential link targets

NOTE: In most cases you would also want to create a `[referer: no referer]` section as well to provide direct browser access to the same link targets.

For example, if objects on foo.org are included as targets in a configuration file, and you want all links on foo.org to its own objects to be valid, you would include the following entries in the configuration file.

```
[referrer: http://www.foo.org/*]  
http://www.foo.org/*
```

If the links on foo.org to its own objects need to be restricted, you would need to create [referrer:] sections for all of the server's valid referring URLs.

How Requests Are Processed

The following is a simplified explanation of the process Excelerator follows to enforce the rules established by a referer header configuration file.

1. A Web Server Acceleration service receives an object request and checks to see whether the object's URL is covered by a target_url entry in the configuration file.
2. If the object is not covered, the service vends it.
3. If the object is covered, the service checks the request for a referer header.
4. If the request has a referer header, the service checks for a [referrer:] section that covers the header and whether the section has a target_url entry that covers the requested object.
5. If a [referrer:] section exists and the object is covered in it, the service vends the object. If a section doesn't exist or the object isn't covered, the service issues a 403 error (default or customized).
6. If the request has no referer header, the service checks for a [referrer: no referer] section and whether the section has a target_url entry that covers the requested object.
7. If a [referrer: no referer] section exists and the object is covered in it, the service vends the object. If the section doesn't exist or the object isn't covered, the service issues a 403 error (default or customized).

A Hypothetical Example

Consider the following hypothetical scenario:

1. The Web site InfoAndStreams.com has both general information and streaming content for distribution on the Web.

Information and streaming content are available at *www.InfoAndStreams.com/Information* and *www.InfoAndStreams.com/Streams*, respectively. In addition, there are short clips of the streams available at *www.InfoAndStreams.com/SClips/*.
2. The Web site administrator installs an Excelerator cache device, creates a Web Server Accelerator service for InfoAndStreams.com on the device, and assigns the InfoAndStreams.com Web server's original IP address to the cache device.

DNS requests to InfoAndStreams.com now resolve to the cache device, which is a reverse proxy for the origin Web server.
3. A second Web site administrator oversees another Web site named MoreInfo.com and creates links to various information and streaming objects on InfoAndStreams.com.

Some of the links on MoreInfo.com point to information and streaming content on InfoAndStreams.com.

4. The InfoAndStreams.com administrator has no problem with MoreInfo.com providing links to the information and the streaming clips on InfoAndStreams.com. The administrator also wants users to be able to enter URLs to either of these areas directly in their browsers. However, requests for streaming content that don't originate from InfoAndStreams are resulting in site bandwidth charges for which InfoAndStreams is not receiving corresponding revenue.
5. To stop the revenue drain, the InfoAndStreams site administrator creates the referer header configuration file shown in **InfoAndStreams Referer Header Configuration File** for the InfoAndStreams accelerator service on the cache device. This configuration file provides for the following referer header validation:
 - ♦ Access to the */Information*, */Streaming*, and */SClips* subdirectories for requests originating from InfoAndStreams.com
 - ♦ Access to only the */Information* and */SClips* subdirectories (not to the */Streaming* subdirectory) for requests originating from MoreInfo.com
 - ♦ Access to only the default pages in only the */Information* and */SClips* subdirectories for users entering direct browser requests.
6. The InfoAndStreams site administrator also creates an error page for incoming requests that don't pass referer header validation. The error page indicates that the content can't be accessed through the URL that was tried, but it is available to users of the InfoAndStreams.com Web site.
7. As a result, users experience the following:
 - ♦ Users accessing content through the InfoAndStreams Web site can access both information and streaming content.
 - ♦ Users accessing content through the MoreInfo.com Web site can access InfoAndStreams information and streaming clips. However, when they click links to any of the full streams on InfoAndStreams.com, they receive the error message with instructions to try the InfoAndStreams Web site.
 - ♦ Users entering InfoAndStreams URLs into their browsers directly can access information and streaming clips, but not the full streams.
 - ♦ Users attempting to access InfoAndStreams content from other Web sites receive the error message with instructions to try the InfoAndStreams Web site directly.

InfoAndStreams Referer Header Configuration File

The referer header configuration file for the hypothetical scenario given in “**How Referer Header Validation Works**” on page 219 might look something like this:

```
[errorpage: http://10.1.1.200/cfg_files/InfoError.html]

[referer: http://www.InfoAndStreams.com/*]

/Information/*
/Streams/*
/SClips/*

[referer: http://www.MoreInfo.com/*]

http://www.InfoAndStreams.com/Information/*
http://www.InfoAndStreams.com/SClips/*
```

```
[referrer: no-referer]
```

```
http://www.InfoAndStreams.com/Information/
```

```
http://www.InfoAndStreams.com/SClips/
```

Table 4 contains an analysis of the file.

Table 4

File Line Entry	Effect
<code>[errorpage: http://10.1.1.200/ cfg_files/InfoError.html]</code>	<p>When the referer header configuration file is loaded, the cache device retrieves the InfoError.html file using HTTP and associates it with the Web Acceleration service being configured.</p> <p>Thereafter, when an object cannot be vended because the referring URL isn't valid, the cache device sends this page to the requesting browser.</p>
<code>[referrer: http:// www.InfoAndStreams.com/*]</code>	<p>http://www.InfoAndStreams.com (and any of its subdirectories) is a valid referer header for incoming requests.</p>
<code>/Information/*</code>	<p>Because this directory is at the root of the referer Web site, only the directory path is required.</p> <p>The asterisk (*) specifies that all subdirectories and files below the Information directory must pass the referer header validation requirement.</p>
<code>/Streams/*</code>	<p>All subdirectories and objects must pass referer header validation.</p>
<code>/SClips/*</code>	<p>All subdirectories and objects must pass referer header validation.</p>
<code>[referrer: http://www.MoreInfo.com/*]</code>	<p>http://www.MoreInfo.com is set up as a valid referer header for incoming requests.</p>
<code>http://www.InfoAndStreams.com/ Information/*</code>	<p>Because this directory is not on the MoreInfo.com Web site, the full URL directory path is required.</p> <p>The asterisk (*) after the forward slash makes all links to any objects in the Information subdirectory valid.</p>
<code>http://www.InfoAndStreams.com/ SClips/*</code>	<p>Because this directory is not on the MoreInfo.com Web site, the full URL directory path is required.</p> <p>The asterisk (*) after the forward slash makes all links to any objects in the SClips subdirectory valid.</p>
<code>[referrer: no-referer]</code>	<p>This entry sets up access for requests that do not contain a referer header, such as initial browser requests or requests that users enter directly.</p>

File Line Entry	Effect
<code>http://www.InfoAndStreams.com/ Information/</code>	<p>Because this applies to requests without a referer header, the full URL directory path is required.</p> <p>Only requests that end with <code>Information/</code> are covered by this <code>target_url</code> entry. (The directory is assumed to contain an <code>index.html</code> file that will automatically load in the browser.)</p>
<code>http://www.InfoAndStreams.com/ SClips/</code>	<p>Because this entry applies to requests without a referer header, the full URL directory path is required.</p> <p>Only requests that end with <code>SClips/</code> are covered by this <code>target_url</code> entry. (The directory is assumed to contain an <code>index.html</code> file that will automatically load in the browser.)</p>

The Net Result of the Sample Configuration File

The aggregate result of the above referer header configuration file is as follows:

- ♦ All content in the Streams directory can be accessed only through InfoAndStreams.com.
- ♦ All content in the Information directory can be accessed through InfoAndStreams.com or MoreInfo.com. A request with a blank referer header will be able to access only the default page in the Information directory. Other referring sites will be denied access to the Information directory.
- ♦ InfoAndStreams.com, MoreInfo.com, and requests without a referer header will be able to access the default page in the SClips directory as well as any subdirectories or objects it contains. Other referring sites will be able to access subdirectories or objects, but not the SClips directory.

NOTE: The InfoAndStreams administrator could have left access to the `/Information` and `/SClips` directories in an unrestricted state by not including them in the referer header configuration file. In that case, the only entries required in the file would be a `[referer:]` section for InfoAndStreams with a single `/Streams/*` URL entry below it.

Setting Up a Referer Header Control

To set up a referer control on an Exceleator cache device, you must do the following:

- ♦ Create a referer header configuration file that explicitly grants access to objects you want controlled and post the file to a network location the cache device can access through HTTP.
- ♦ Optionally, create an error page file if you want to send a customized message to those with a non-valid referer header in their browser requests.
- ♦ Use the command line interface to load the referer file and manage referer header control on the appliance. For the specific commands, see [“Managing Referer Controls” on page 229](#).

Customizing Referer Header Error Messages

By default, users receive the standard 403 Forbidden error page when their request contains an invalid referer header.

You can provide customized error messages which the accelerator service will send to browsers. You might want to provide a customized error message file to provide instructions to users for

accessing the origin Web site directly or to provide other information helpful to those seeking access to your information.

To create a customized referer header error message file, complete the following steps:

- 1** Create an HTML file that provides the information you want to pass along to users who request information access through an invalid referer.
- 2** Include any of the following customized tags in the HTML file to have Excelerator substitute the indicated values at display time as indicated.
 - ♦ **<SERVICE_NAME>**: The name of the Web Server Accelerator that rejected the request and issued the error message.
 - ♦ **<REFERER_URI>**: The URL of the Web site that didn't pass referer header validation.
 - ♦ **<REQUEST_URI>**: The target object that couldn't be vended because the referring Web site didn't pass referer header validation.
- 3** Save the file to a location accessible by the cache device through an HTTP connection.

Creating the Referer Configuration Files

To create a referer configuration file, complete the following steps:

- 1** Study the information in [“Understanding the Referer Header Configuration File” on page 219](#).
- 2** Using the information you have learned, identify the Web acceleration services for which you need to provide referred access controls.
- 3** If desired, create a customized error message files as explained in [“Customizing Referer Header Error Messages” on page 227](#)
- 4** For each acceleration service that requires referred access controls, plan which objects you want to cover.
- 5** Identify all the URLs that might appear in valid referer headers in browser requests to the service.
- 6** Using an ASCII editor, create a DOS or UNIX text file for each acceleration service that includes the entries explained in [“Configuration File Specifications” on page 219](#), specifically in [Table 12 on page 220](#).
- 7** If applicable, include an [errorpage:] entry that refers to the customized error file created in [Step 3](#).
- 8** Create a [referer:] section for each URL identified in [Step 5](#).
- 9** In each [referer:] section, create target_url entries for all controlled directories and objects to which the [referer:] can link.

NOTE: Remember to create a [referer:] section for the accelerated Web server. For more information, see [“Allowing a Web Server to Link to Its Own Objects” on page 223](#).
- 10** If desired, allow for access through requests that do not contain a referer header, such as direct browser requests. For more information, see [Table 12 on page 220](#).
- 11** Save the file to a location accessible by the cache device through an HTTP connection.

Loading the Referer Configuration File

To load a referer header configuration file, use the following commands from the System Console prompt:

```
set accelerator name referercontrol url=url  
apply
```

where *name* is the name of the Web Server Acceleration service on the cache device and *url* is the path to the referer header configuration file you created for the service in [“Creating the Referer Configuration Files” on page 228](#).

For example: if the Web Server Acceleration service is named InfoAndStreams and the URL path to the referer configuration file is `http://10.1.1.200/cfg_files/InfoAndStreams.txt`, you would enter the following:

```
set accelerator InfoAndStreams referercontrol url=http://10.1.1.200/  
cfg_files/InfoAndStreams.txt  
apply
```

Exceclerator would then use HTTP to access the configuration file and copy it to the cache device.

NOTE: You must activate a referer header control to have it affect incoming requests as explained in the following section.

Activating the Referer Control

To activate the referer header control, use the following command from the System Console prompt:

```
set accelerator name referercontrol enable=yes  
apply
```

where *name* is the name of the Web Server Acceleration service on the cache device.

For example: to activate the referer header control for the configuration file loaded in [“Loading the Referer Configuration File” on page 229](#), you would enter the following:

```
set accelerator InfoAndStreams referercontrol enable=yes  
apply
```

Managing Referer Controls

After you have created a referer header configuration file and put it in an HTTP-accessible location on the network, you can do the following from the cache device’s System Console prompt.

De-activating a Referer Control

To de-activate a referer header control, use the following command from the System Console prompt:

```
set accelerator name referercontrol enable=no  
apply
```

where *name* is the name of the Web Server Acceleration service on the cache device.

IMPORTANT: When a referer header control is deactivated, the referer header configuration file is deleted from the cache device. If you subsequently activate the control, the configuration file is again retrieved from the network using HTTP.

If you delete an Web acceleration service, the referer header configuration file is also deleted.

Refreshing the Referer Header Configuration File

After you load a referer header configuration file, you can change the contents of the file on the network and have Excelerator reload the file by entering the following command:

```
set accelerator name referercontrol refresh=yes
```

where *name* is the name of the Web Server Acceleration service on the cache device.

Getting Information for a Referer Header Control

To get information about a referer header control you have set up, enter the following command:

```
get accelerator name referercontrol
```

where *name* is the name of the Web Server Acceleration service on the cache device.

32

Dynamic Bypass

For help configuring the Dynamic Bypass feature, see [“Enable Dynamic Bypass” on page 398](#).

The Dynamic Bypass feature lets you configure the appliance so that specific errors from Web sites are explicitly not cached and subsequent requests to the Web sites are simply passed through for a specific time period.

Why Dynamic Bypass Is Needed

The appliance follows the HTTP 1.1 specification, which stipulates that no 400- or 500-series errors are cached (except HTTP 410: Data Moved Permanently). If these errors occur, Excelerator passes them through to the requesting browsers.

Some Web servers erroneously pass cache control public headers along with the errors they generate. These headers override the default error handling by the caching system and cause the error to be cached on the appliance.

For example, if a Web server receives an unauthorized request for data and responds with an HTTP 403: Forbidden error accompanied by a cache control public header, the error will be cached. The appliance will then respond to all subsequent requests, including those from authorized users, with the 403 error.

What Excelerator Does

If dynamic bypass is enabled for a given error and Excelerator receives the error from an origin Web site, the site’s URL is dynamically added to the bypass list and retained in the list for the period of time specified. All subsequent requests to that URL during the dynamic bypass time period are passed through to the Web site and do not go through the cache.

To continue with the above example, if the appliance has dynamic bypass enabled and the 403 error checked, it would not cache the error but pass it directly back to the browser. It would also add the site’s URL to its dynamic bypass list. On subsequent requests to the same URL, Excelerator would bypass cache and pass through requests to the Web site for the period of time specified. This lets the Web server determine whether to accept or reject access requests on a case-by-case basis.

33 Appliance Error Messages

The appliance lets you specify a language for the error messages it sends to browsers. This section explains appliance error message support and provides instructions to modify error message text and create support for additional languages.

What You Need to Know about Appliance Error Messages

The appliance provides error messages to browsers through a set of language-specific directories, each of which contains three files.

When you select a language in the browser-based management tool, you are actually selecting one of these language-specific directories and the files it contains.

The appliance uses the following files:

- ♦ `ERRPAGE.CFG` contains the text of all appliance error messages
- ♦ `PXYERR.HTM` the first of two HTML template files. It applies a format to the applicable error text in `ERRPAGE.CFG` and other error information and is then sent to the receiving browsers. It is used in all cases except when a Transparent Handling service has transparent error handling enabled.
- ♦ `TRANSERR.HTM` the second of two HTML template files. It applies a format to the applicable error text in `ERRPAGE.CFG` and other error information and is sent to browsers accessing a Transparent Handling service that has transparent error handling enabled.

The language-specific directories that contain these three files are located in `\ETC\PROXY\DATA\ERRPAGE\NLS\LANGUAGE`, where *LANGUAGE* is the English name of the language.

For example, the English files are stored in `\ETC\PROXY\DATA\ERRPAGE\NLS\ENGLISH`.

Other common appliance error message directories include:

```
\ETC\PROXY\DATA\ERRPAGE\NLS\GERMAN
\ETC\PROXY\DATA\ERRPAGE\NLS\SPANISH
\ETC\PROXY\DATA\ERRPAGE\NLS\PORTUGUESE
\ETC\PROXY\DATA\ERRPAGE\NLS\FRENCH
\ETC\PROXY\DATA\ERRPAGE\NLS\JAPANESE
```

You can specify the language for error page vending in `Cache > Mini Web`. For more information, see [“Mini Web Tab” on page 402](#).

Checking the Language Directories on Your Appliance

To see a list of the directories on your appliance, do the following:

- 1 In the browser-based management tool, click Cache > Mini Web > the drop-down list.

Customizing the Appliance Error Template and Message Files

You can create error message support for additional languages and customize existing error message text and format.

Creating a New Language

- 1 Start FTP and log in as explained in “Starting an FTP Session with the Appliance” on page 263.
- 2 Enter the following:

```
get /etc/proxy/data/errpage/nls/english/pxyerr.htm
get /etc/proxy/data/errpage/nls/english/transerr.htm
get /etc/proxy/data/errpage/nls/english/errpage.cfg
```
- 3 Modify the ERRPAGE.CFG file.

Explicit instructions in the file clearly indicate which parts can be translated.

You cannot delete or add messages, nor can you change the number or order of the messages.
- 4 Modify the PXYERR.HTM and TRANSERR.HTM files.

Because these are HTML files, you can customize them in a variety of ways. To interface with the appliance's message delivery mechanisms, you must ensure the following:

 - ♦ The keywords <PROXY_ADDRESS>, <ERROR_STATUS>, and <ERROR_DESCRIPTION> must be retained because they are dynamically replaced with the information that their names imply.
 - ♦ Graphics that you add should be put in the SYS:\ETC\PROXY\DATA directory. References to these graphics must use the <PROXY_ADDRESS> keyword as a starting reference point. See the usage of ALERTBAR.GIF in the PXYERR.HTM file.
- 5 Using FTP, create a new language directory in the path given in “What You Need to Know about Appliance Error Messages” on page 233.
- 6 Enter the following, where *language* is the new language directory you created:

```
put errpage.cfg /etc/proxy/data/errpage/nls/language/errpage.cfg
put transerr.htm /etc/proxy/data/errpage/nls/language/transerr.htm
put pxyerr.htm /etc/proxy/data/errpage/nls/language/pxyerr.htm
```

The new language is dynamically available in the browser-based management tool and from the command line. You do not need to restart the appliance.

Customizing the Error Message Text of an Existing Language

Referring to the procedure in “Creating a New Language” on page 234 for details, complete the following basic steps:

- 1 Get the ERRPAGE.CFG file from an existing language-specific directory.
- 2 Modify the file.

IMPORTANT: There are limitations on what you can change in this file. See “Creating a New Language” on page 234 for details.

- 3 Replace the file when modifications are completed.

Customizing the Error Message Format of an Existing Language

Referring to the procedure in “[Creating a New Language](#)” on page 234 for details, complete the following basic steps:

- 1 Get the PXYERR.HTM and TRANSERR.HTM files from an existing language-specific directory.

- 2 Modify the files.

IMPORTANT: There are limitations on what you can change in this file. See “[Creating a New Language](#)” on page 234 for details.

- 3 Replace the files when modifications are completed.

34 Logging

Logging appliance caching activity can be useful for a number of reasons, such as billing for services rendered. Excelerator lets you specify how often a new log file will be started (rolled over), how long old log files will be retained, and the format of the log files.

Using Appliance Logging Services

Your appliance offers the following logging services:

- ◆ You can turn on logging for forward, transparent, and reverse proxy as well as for URL filtering.
- ◆ You can have the appliance automatically download files to an FTP server and automatically delete downloaded files.
- ◆ You can control the deleting of old log files based on an older-than-x time period or the number of log files in the system.

The appliance can create logs using both the common and extended log formats. A wide variety of tools exist for manipulating and processing these files.

Overview of Appliance Logging

The Volera Excelerator appliance provides a high-performance proxy cache system capable of handling thousands of transactions per second. Even though Excelerator can log extensive details of each transaction and the disk space reserved for log files is quite generous on most appliances, Excelerator can fill up the available disk space in a matter of minutes if transaction volume is high and log entries consume a few hundred bytes each.

This section explains how appliance logging works and presents management options you can use to ensure optimal use of the available log file disk space and timely migration (downloading) of log files to other storage devices.

What the Appliance Can Log

The following table shows the transactions the appliance can log and the formats available for each service type.

Service	Common	Extended
Transparent/Forward Proxy	Yes	Yes
Web Server Accelerator	Yes	Yes
Clustered Services (all types)	Yes	Yes

Service	Common	Extended
Content Filtering	Yes*	No
Dynamic Bypass	No*	Yes
* These common log formats differ from the industry standard proxy cache common log format.		

The Costs of Logging

Performance

Turning on logging for a given service increases system overhead and causes some performance degradation. Therefore, logging should be used only when service transactions must be tracked for customer billing purposes or other compelling reasons.

Disk Space

Transaction volume and log entry size can cause available log disk space to fill up quickly. Proxy cache disk space is unaffected by log files.

See [“Planning Step 3: Calculating Log Rollover Requirements” on page 241](#) for formulas you can use to estimate how quickly your logging disk will fill.

System Constraints

To plan a logging strategy, you must know the capacity and limitations of your appliance.

Disk Space Is Preset

It is essential that you know how much logging disk space is available on your appliance.

Logging disk space is not user-configurable; it is preset by the vendor who produced your appliance. Most vendors allocate different amounts of disk space for log files on each level or tier of appliance they produce.

If you don’t have this information, you must get it from your appliance vendor before you can plan a logging strategy. After you determine how much disk space is available, you can plan when to download and delete log files so the disk does not become full.

Log Files Must Be Rolled Over Before Deletion

Excelsior will not allow the deletion of active log files—files that are currently in use by the caching system. Only log files that have been rolled over and closed can be deleted.

You can ensure you have closed files on the system by scheduling regular rolling over of log files. During each rollover, the current log file is closed and a new log file is opened.

You must plan your log file rollover schedule to coincide with your download and deletion schedule so that you have at least one closed log file per service when the download and delete cycle starts.

NOTE: Although you can download active log files, this is normally useful only for periodic administrative checks.

Active files contain only the transaction data up to the moment of the download and are incomplete from customer billing and other business standpoints.

Logging Ceases When the Logging Disk Is Full

When the appliance encounters a log disk full condition, it stops logging and closes all active log files. Information that would have been logged after that point is lost. Other appliance functions continue without interruption.

Log Filenames Are Limited

The appliance automatically generates log filenames as follows:

- ♦ Six numbers representing the year, month, and day the file was created
- ♦ A dash separating the date from a single-letter identifier

NOTE: The dash is not included after the letters double.

- ♦ Letter identifiers running from A through ZZ

This naming convention accommodates up to 702 log files per service per day. If the rollover options are set so that all the possible filenames are used in one day, the log file with the ZZ letter identifier is not closed until the start of the next day (unless the logging disk becomes full).

The Appliance Offers Two Log Rollover Options

Appliance log rollover options let you specify when the appliance closes active log files and opens new files so that the closed files can be downloaded and deleted.

Because of the limitations explained in this section, it is essential that you develop a solid log file rollover plan. This will ensure that your appliance doesn't run out of logging disk space or overwrite log files before they are downloaded and deleted.

You can have the appliance roll log files over according to time or file size, as explained in the following table:

Option	Considerations
Roll Over by Time	<p>If you plan to download and delete older log files at a set time, you must configure the appliance so that at least two log files exist per service at the time you've scheduled for downloading and deleting. One file will be active, the other closed and ready for download and deletion.</p> <p>For example, if you determine that your log disk space will fill every 12 hours, then you must configure the appliance to roll the log files over in intervals less than 6 hours, so that at least one log file per service is closed and ready to be downloaded and deleted.</p>
Roll Over by Size	<p>If you aren't certain how long it will take to fill your appliance's logging disk space, you can roll the log files over by size.</p> <p>For example, you might be logging transactions for three services and have a log volume size of 6 GB. Because you must have at least two log files per service before the disk space fills, each log file must be smaller than 1 GB when the appliance rolls it over.</p>

Planning Your Logging Strategy

As explained in “[The Costs of Logging](#)” on page 238 and “[System Constraints](#)” on page 238, logging of caching transactions involves system and maintenance overhead. If your situation requires logging, you should plan carefully so that the information you are tracking aligns with specific requirements. This will ensure optimal use of appliance resources.

Because logging requirements and transaction volume vary widely, it is impossible to make recommendations regarding specific logging strategies.

The following sections step you through the logging strategy planning process. We recommend you record the information you gather on a planning sheet of some kind.

Planning Step 1: Determining Your Logging Requirements

To plan a logging strategy, you should first determine the requirements driving the need for logging. We recommend you complete the following steps:

- 1** Identify the business and/or other reasons for tracking service transactions.
Examples include customer billing requirements, statistical analysis, or growth planning.
- 2** Determine which services you need to track.
- 3** Record this information for later reference.

Planning Step 2: Selecting a Log File Format and Optimizing the Log Entry Size

If you use the common log format, log entry size is fixed. If you use the extended log format, log entry size depends on the number of log fields selected.

Complete the following steps for each service you need to track:

- 1** Referring to the “[Log Options Dialog Box](#)” on page 335, record which log fields must be tracked for each service to be logged.
- 2** Carefully scrutinize the information you plan to track to ensure that the log data collected is essential. Consider the following points:
 - ♦ Logging cookie information can consume a lot of space and might provide little, if any, critical information.
 - ♦ If you plan to import Excelerator 3 log files to a third-party reporting tool, you will need to use the extended format and ensure that the URI, URI-STEM and URI-QUERY fields are all selected. If not, the reporting agents won’t be able to import the logs and compile the information.
 - ♦ If you are not using a third-party tool, however, and you select URI, also selecting URI-STEM and URI-QUERY would be redundant because $URI = URI-STEM + URI-QUERY$.

The main point is to log only the essential data because a few bytes can add up quickly when the cache device is tracking thousands of hits every second.

Planning Step 3: Calculating Log Rollover Requirements

You can have the appliance roll over log files based on time or on size, but not both.

If you already know which option you want to use, scan this section and then complete only the calculations pertinent to your choice.

If you don't know which option best matches your situation, completing the calculations in this section should help you decide.

Variable Definitions

The following variables are used in the formulas:

logvol_size: The total disk capacity reserved for log files on your appliance.

You must get this information from your appliance vendor.

logentry_size: The average log entry size.

You can determine this by configuring your appliance to track the required information, generating traffic to the appliance, downloading the log files, determining how large each entry is, and calculating the average.

request_rate: The peak rate of requests per second.

You can estimate this rate or place your appliance in service and get more accurate data by accessing the browser-based management tool's Monitoring tabs.

num_services: The number of services for which you plan to enable logging.

logs_per_service: The number of log files, both active and closed, that you want the appliance to generate for each service before the disk fills.

You must plan to have at least two logs per service before the disk is filled. See [“Log Files Must Be Rolled Over Before Deletion” on page 238](#).

Calculating DISKFULL_TIME

Using the following formula, you can calculate how long it will take the appliance to fill your logging disk space:

```
diskfull_time seconds = logvol_size / (request_rate * logentry_size *  
num_services)
```

For example, if you assume the following:

- ♦ *logvol_size* = 1 GB
- ♦ *request_rate* = 1000 requests per second
- ♦ *logentry_size* = 1KB
- ♦ *num_services* = 1

Then *diskfull_time* = (1 GB) / (1000 * 1KB * 1) = 1048 seconds (17.47 minutes).

The logging disk space will fill up every 17.47 minutes.

If this time is too short, you must reduce the log entry size by configuring the appliance to log less information per transaction. This is because you can't increase the disk space nor limit the requests being logged.

To calculate the *diskfull_time* for your appliance, complete the following steps:

- 1 Determine the values of the four variables listed above.

For more information, refer to “[Variable Definitions](#)” on page 241.

- 2 Using the *diskfull_time* formula, calculate how often you can expect your logging disk to fill; then use the result in [Calculating MAX_ROLL_TIME](#).

Calculating *MAX_ROLL_TIME*

Using the following formula, you can calculate the maximum roll-over time value you should specify in the Rollover Every field of the Log Options dialog box.

$$\text{max_roll_time} = \text{diskfull_time} / \text{logs_per_service}$$

For example, if you assume the following:

- ♦ *diskfull_time* = 12 hours
- ♦ *logs_per_service* = 2

Then *max_roll_time* = 12 / 2 = 6 hours.

If you roll your logs over by time intervals, the maximum time should be less than six hours. Otherwise, scheduling the download and deletion of log files is much more complicated and the window in which this can be done is narrower.

To calculate the *max_roll_time* for your appliance, complete the following steps:

- 1 Determine how many log files you want the appliance to generate per service before log space fills.
The minimum number is two.
- 2 Using the *max_roll_time* formula and the *diskfull_time* value obtained in “[Calculating DISKFULL_TIME](#)” on page 241, calculate how often you should have the appliance roll over the log files.
- 3 Record the *max_roll_time* result on your planning sheet.

Calculating *MAX_LOG_ROLL_SIZE*

Using the following formula, you can calculate the maximum log file size you should specify in the Rollover When File Size Reaches field of the Log Options dialog box.

$$\text{max_log_roll_size} = \text{logvol_size} / (\text{num_services} * \text{logs_per_service})$$

For example, if you assume the following:

- ♦ *logvol_size* = 600 MB
- ♦ *num_services* = 2
- ♦ *logs_per_service* = 3

Then *max_log_roll_size* = 600 MB / (2 * 3) = 100 MB.

If you roll your logs over when they reach a specific size, the file size must be no more than 100 MB. Otherwise, the system will run out of disk space before you have three complete log files and scheduling the download and deletion of log files will be much more complex.

To calculate the *max_log_roll_size* for your appliance, complete the following steps:

- 1 Determine the values of the three variables listed above.
- 2 Using the *max_log_roll_size* formula, calculate the maximum size a log file should reach before the appliance rolls it over.

Configuring Logging Options

Based on the planning you have completed in “[Planning Your Logging Strategy](#)” on page 240, you must now configure the log options for each affected service.

Configuration Step 1: Opening the Appropriate Log Options Dialog Box

- 1 For each service you are logging, open the Log Options dialog box in the browser-based management tool.

Refer to the services you selected in “[Planning Step 1: Determining Your Logging Requirements](#)” on page 240.

The following table gives the path for each service:

Service	Path
Transparent/Forward	Cache > Client Accelerator > Enable Logging for Client Acceleration > Log Options
Web Server Accelerator	Cache > Web Server Accelerator > Insert > Enable Logging for This Accelerator > Log Options
Clustered Services	Cache > Cluster > Forward or Transparent or Accelerator > Enable Logging > Log Options
Content Filtering	Cache > Filtering > Enable Filter Logging > Log Options

The following sections discuss each of the areas within the Log Options dialog box. For further information, see “[Log Options Dialog Box](#)” on page 335.

Configuration Step 2: Selecting a Log Format

- 1 In the Log Options dialog box, specify the log format for the service based on the planning you did in “[Planning Step 2: Selecting a Log File Format and Optimizing the Log Entry Size](#)” on page 240.

Remember that each bit of information you log increases the size of each log entry, thus affecting the rate at which logging disk space is used.

Configuration Step 3: Specifying Rollover Options

- 1 In the Log Options dialog box, specify how the appliance rolls over the log files for the service based on the planning you did in “[Planning Step 3: Calculating Log Rollover Requirements](#)” on page 241.

Configuration Step 4: Specifying Handling of Older Files

You must schedule the regular download and deletion of log files to avoid running out of log disk space.

Whenever possible, we recommend you use the FTP log push feature for this task. However, you can also manage log files manually using the browser-based management tool or the appliance's Mini FTP Server. See [“Manually Downloading and Deleting Log Files” on page 246](#).

The appliance also provides three options for dealing with old files as a failover precaution.

Ideally Excelsior will never actually use the old file option you select because you will schedule the downloading and deleting of log files so that the system never becomes full.

Two of these options automatically dispose of older files to avoid the disk full condition. The third option is not recommended for most situations.

- ♦ **Limit Number of Files To:** This option lets you limit the total number of log files retained for each service. After the limit for each is reached, the oldest file for the service is deleted each time a new file is created. All logging data in deleted files is lost.
- ♦ **Delete Files Older Than:** This option lets you configure the appliance to delete files when they are older than the time you specify. All logging data in deleted files is lost.
- ♦ **Do Not Delete:** This option is not recommended because it can lead to a disk full condition if files are not manually downloaded and deleted. If, however, the older logging data has more value for some reason, this option will preserve the oldest log files unless you manually delete them or specify their deletion in the FTP Log Push Configuration dialog box.

To specify how the appliance handles older files, complete the following:

- 1** In the Log Options dialog box, select an old file option that matches your requirements. (To review option specifics, see the bullet list above.)

As with log format and rollover options, you can specify different old file options for each service. We recommend, however, that you avoid potential confusion by using the same old file settings for each service.

- 2** (One time only) Click FTP Log Push > configure the FTP log push options.

For help with setting options in the FTP Log Push Configuration dialog box, refer to [“Using FTP Push to Automatically Download and Delete Log Files” on page 245](#), then return to this procedure.

- 3** In the Log Options dialog box, double-check the Old File Options settings against either your FTP log push configuration or your schedule for manual download and deletion to ensure that log files won't reach the deletion threshold (number or age) prior to a scheduled download and deletion.
- 4** If you need to configure log options for other services, return to [“Configuration Step 1: Opening the Appropriate Log Options Dialog Box” on page 243](#); otherwise, continue with the next section.

Configuration Step 5: Monitoring and Refining Your Logging Strategy

As with all appliance operations, you should monitor what is happening with your logging strategy over time and make adjustments and refinements if necessary.

When you monitor your logging strategy, you should ensure that:

- ♦ All the logging information you are gathering is being used. If not, you might be able to further reduce your logging record size.
- ♦ Your log file sizes match the estimated averages you used to plan your log file roll-over strategy. If not, you might need to adjust the frequency or even the method used to trigger log file rollover.
- ♦ Your logging strategy is leaving a buffer of free log disk space adequate for possible surges in appliance traffic.
- ♦ The external storage capacity (FTP server or other storage) is adequate.
- ♦ All aspects of your logging strategy are keeping pace with increases in traffic through the appliance.

About the FTP Log Push Feature

The FTP Log Push Feature lets you configure the appliance to push log files to an FTP server at specified intervals: on the first day of the month, on specified days of the week, or when log files roll over.

The feature operates within the following parameters:

- ♦ Excelsator will try as many times as necessary to establish one connection with the FTP server during the hour of the scheduled push. When the hour changes, Excelsator stops trying until the next interval you have specified.
- ♦ When a connection with the FTP server has been established, Excelsator assumes that the pushing of log files is going to be successful. Any FTP connection errors beyond that point are not detected by Excelsator.

For example, you specify that log files are to be pushed on every day of the week at 12 midnight. When the system clock reaches the target hour, Excelsator begins trying to establish a connection with the FTP server.

If a connection cannot be established before the hour changes to 1 a.m., Excelsator stops trying to connect and doesn't try again until 12 midnight the next day.

If a connection is established but an error occurs that prevents a successful push, the error is not detected, and Excelsator doesn't try to connect again until 12 midnight the next day.

IMPORTANT: Although FTP connection errors are not detected, log files are not deleted until they have been successfully pushed. Therefore, if files for a given period are not successfully pushed, the system will attempt to push them during the next scheduled period and will not delete them until the push succeeds.

Using FTP Push to Automatically Download and Delete Log Files

To configure your appliance to use the FTP Log Push feature, complete the following steps:

- 1 In the browser-based management tool, access the FTP Log Push Configuration dialog box by clicking FTP Log Push on any of the Log Options dialog boxes.

Paths to the dialog boxes are summarized under **“Configuring Logging Options” on page 243**.

IMPORTANT: Although the FTP Log Push Configuration dialog box is accessed through one of the service-specific Log Options dialog boxes, it is unaffected by the path used to reach it. The settings you specify affect all the log types you check in the box.

This lets you set the FTP push options for all log types in a single place.

2 In conformance with your logging strategy, specify the following information:

- ◆ Which log file types to push (all the log types to be managed through FTP push must be checked)
- ◆ Your FTP server information
- ◆ The method the appliance uses for determining when to push log files

If your FTP server is always available, we recommend using the Push Logs When the Logs Roll Over option rather than setting specific download times. This will protect your appliance from sudden surges in traffic, which can fill the disk sooner than expected.

- ◆ Whether the appliance should delete the files from the log disk once they have been pushed

We recommend deleting log files after they have been successfully pushed unless you have a compelling reason for manually deleting them. Automatic deletion also protects your appliance from sudden surges in traffic.

For more information on the FTP Log Push Configuration dialog box, refer to [“FTP Log Push Configuration Dialog Box” on page 336](#).

3 When you have configured your FTP log push options, click OK to return to the Log Options dialog box of the service you are configuring.

Manually Downloading and Deleting Log Files

Whenever possible, we recommend you use the FTP Log Push feature for downloading and deleting log files. See [“Using FTP Push to Automatically Download and Delete Log Files” on page 245](#).

If you need to manage your log files manually, we recommend that you establish a regular schedule and ensure that all those responsible for downloading and deleting log files know the following things:

- ◆ When log files are to be downloaded and deleted
- ◆ How to determine the name of each log file to be downloaded and deleted
- ◆ Where to save the log files

You will want to develop specific procedures for your situation. The following sections contain general ideas for accomplishing these tasks.

When to Download and Delete Log Files

The primary consideration is that log files must be downloaded and deleted before the logging disk space fills up.

Getting Log Filenames

Before you can download or delete a log file, you must know its exact name.

Appliance log filenames can be listed in the browser-based management tool in Monitoring > Cache Logs. They can also be listed from the command line, or through a Telnet session using the get command.

The appliance automatically generates log filenames as follows:

- ♦ Six numbers representing the year, month, and day the file was created
- ♦ A dash separating the date from a letter identifier
 - NOTE:** The dash is not included after the letters double.
- ♦ Letter identifiers running from A through ZZ

This naming convention accommodates up to 702 log files per day. If the rollover options are set so that all the possible filenames are used in one day, the log file with the ZZ letter identifier should not be closed manually until the start of the next day (unless the logging disk becomes full).

To list log files using FTP, you must know the path to the files. Use the following table to determine the paths to various log files.

File	Location
All log files	LOG:ETC/PROXY/DATA/LOGS/
Transparent and forward proxy log files in common format	LOG:ETC/PROXY/DATA/LOGS/FORWARD/COMMON/
Transparent and forward proxy log files in extended format	LOG:ETC/PROXY/DATA/LOGS/FORWARD/EXTENDED/
Filter log files in appliance filtering common format	LOG:ETC/PROXY/DATA/LOGS/FILTER/COMMON/
Web server accelerator log files in common format	LOG:ETC/PROXY/DATA/LOGS/REVERSE/COMMON/ <i>name</i> The variable <i>name</i> is the name of the Web server accelerator.
Web server accelerator log files in extended format	LOG:ETC/PROXY/DATA/LOGS/REVERSE/EXTENDED/ <i>name</i> The variable <i>name</i> is the name of the Web server accelerator.
Dynamic Bypass log files in extended format	LOG:ETC/PROXY/DATA/LOGS/DBYPASS/EXTENDED/
Clustered service log files in common format	LOG:ETC/PROXY/DATA/LOGS/CLUSTER/COMMON/ <i>name</i> The variable <i>name</i> is the name of the clustered service. IMPORTANT: Clustered services often have multiple IP addresses assigned to them. Clustered services are distributed among cluster members by IP address. This means that the log files of a clustered service will probably be stored on multiple appliances if the service has more than one IP address assigned to it.

File	Location
Clustered service log files in extended format	LOG:ETC/PROXY/DATA/LOGS/CLUSTER/EXTENDED/ <i>name</i> The variable <i>name</i> is the name of the clustered service. IMPORTANT: Clustered services often have multiple IP addresses assigned to them. Clustered services are distributed among cluster members by IP address. This means that the log files of a clustered service will probably be stored on multiple appliances if the service has more than one IP address assigned to it.

Using the Browser-Based Tool to Get Filenames

You can most easily view log filenames in the browser-based management tool. To do so, click Monitoring > click Cache Logs > select a log format > select a service.

Using FTP to Get Filenames

The Mini FTP Server in version 1.3 and later supports the CWD command for changing to the target log directories. All appliance versions let you use the LS command in connection with full paths to list log files.

For example, the following command lists transparent and forward proxy log files in common format:

```
ls log:etc/proxy/data/logs/forward/common/
```

For a complete list of log file directory paths, see [“Getting Log Filenames” on page 246](#).

Using the Command Line or Telnet to Get Filenames

You can also see a list of log filenames from the command line. However, you cannot download files from the command line.

The following table presents some command line/Telnet examples.

If You Want To	Then Enter
See a list of available forward/transparent log files in common format	<code>get comlog forward</code>
See a list of available Web server accelerator log files in common format	<code>get comlog reverse:name</code> (The variable <i>name</i> is the name of the Web server accelerator.)
See a list of available filtering log files in appliance filtering common format	<code>get comlog filter</code>

If You Want To	Then Enter
See a list of available forward/transparent log files in extended format	<code>get extlog forward</code>
See a list of available Web server accelerator log files in extended format	<code>get extlog reverse:name</code> (The variable <i>name</i> is the name of the Web server accelerator.)

Downloading Log Files

Using the Browser-Based Management Tool to Download Log Files

You can download the files in the browser-based management tool as you view them. After you click Download, when the browser asks what you want to do with the file, save it to your designated log file storage location.

Using FTP to Download Log Files

You can use FTP from the storage location to retrieve the files using the GET command. You must first obtain each filename using one of the options explained in [“Getting Log Filenames” on page 246](#).

After you have the log filename, you can transfer it to your workstation. For example, to download a forward proxy common format log file, you would use the following command after starting an FTP session with the appliance:

```
get log:/etc/proxy/data/logs/forward/common/filename.log
```

The *filename* variable is the name of the log file you have previously obtained.

You can also use the MGET command, but be aware that this command also downloads active log files that are not complete.

The appliance doesn't currently support the FTP server PUT command.

Deleting Downloaded Log Files

After the log files have been downloaded and saved to another location, delete the files using one of the following options:

- ♦ The Delete button in the browser-based management tool
- ♦ The DEL command in FTP

About Extended Log Field Headers

The following information about field values in extended log files might help you interpret the content of the files:

- ♦ Fields within the file are delimited by the tab character.
- ♦ A field is of two types: string and non-string.
- ♦ String fields are enclosed in quotation marks (").

- ◆ If a string field contains a quotation mark, that character is repeated once for every occurrence to enable unambiguous file parsing.
- ◆ If a string field has no value, it is represented by two quotation marks ("").
- ◆ Non-string fields containing no value are represented by a hyphen (-).
- ◆ Field headers starting with s- are associated with the appliance.
- ◆ Field headers starting with c- are associated with the client/browser.
- ◆ Field headers starting with sc are associated with flow from the appliance to the client/browser.
- ◆ Field headers starting with cs are associated with flow from the client/browser to the appliance.

The information in the following table is supplementary to the W3C Extended Log Format Specification found on the [Extended Log File Format Web site \(http://www.w3.org/TR/WD-logfile\)](http://www.w3.org/TR/WD-logfile). You might find it useful for interpreting the content of extended log field headers.

Name	Description	Type	Selectable	Comments
date	GMT date in YYYY-MM-DD format	non-string	No	
time	GMT time in HH:MM:SS format	non-string	No	
c-ip	Client (browser) IP address	non-string	No	
cs-authname	Username if applicable	non-string	Yes	
s-ip	The appliance IP address	non-string	Yes	
s-sitename	Reverse proxy or accelerator site name	non-string	Yes	
cs-method	The HTTP method the browser sent to the appliance	non-string	Yes	
cs-uri	The HTTP URL the browser sent to the appliance	non-string	Yes	The URL must not have spaces per the HTTP specification.
cs-uri-stem	The stem portion of the HTTP URL the browser sent to the appliance	non-string	Yes	<p>The URL stem is everything up to the first question mark (?).</p> <p>If the URL has no question mark, the cs-uri-stem is the same as the cs-uri.</p> <p>This field is redundant if cs-uri is selected.</p>

Name	Description	Type	Selectable	Comments
cs-uri-query	The query portion of the HTTP URL the browser sent to the appliance	non-string	Yes	<p>The query portion is the first question mark through the end of the URL.</p> <p>If the URL has no question mark, cs-uri-query has no value.</p> <p>This field is redundant if cs-uri is selected.</p>
c-version	The HTTP version specified in the URL the browser sent to the appliance	non-string	Yes	
sc-status	The HTTP status code the appliance sent to the browser	non-string	Yes	
sc-bytes	The number of bytes of HTTP response data the appliance sent to the browser	non-string	Yes	
cs-bytes	The number of bytes of HTTP request data the appliance received from the browser	non-string	Yes	
time-taken	The time in seconds it took appliance resources to deal with the request	non-string	Yes	
cs(User-Agent)	The User-Agent HTTP request header value the browser sent to the appliance	string	Yes	
cs(Cookie)	The Cookie HTTP request header value the browser sent to the appliance	string	Yes	The appliance doesn't cache cookie information.
cs(Referer)	The Referer HTTP request header value the browser sent to the appliance	string	Yes	The appliance reads the field header as it is.

Name	Description	Type	Selectable	Comments
cs(X-Forwarded-For)	The X-Forwarded-For HTTP request header value the browser sent to the appliance	string	Yes	Do not confuse this with the X-Forwarded-For option that causes the appliance to generate or forward headers to upstream proxies or Web servers.
cached	The value indicating whether the request was filled from cache	non-string	Yes	1 = filled from cache 0 = not filled from cache
x-fill-proxy-ip	The IP address of the upstream proxy	non-string	Yes	Assumes the appliance is configured with an upstream proxy and fetched the request from that proxy
x-origin-ip	The IP address of the origin server	non-string	Yes	Assumes the appliance fetched the request directly from the origin server

Extended Logging Enhancements in Excelerator 2.3

Excelerator 2.3 includes the following extended logging enhancements.

Standard Header Field Logging

The following standard header fields are now supported.

NOTE: *Sc* refers to server-to-client response headers, and *cs* refers to client to server request headers. Header content is included in double quotation marks, with null content indicated by a pair of marks.

sc-(Content-Range)

Purpose

This field records the byte ranges vended from an Excelerator cache device to a requesting client.

Values:

The results include the word *bytes*, are enclosed in double quotes ("), and are either one or more byte ranges separated by commas (,) or a null value (") if this field is not applicable to the specific log entry (i.e. this is not a byte range request).

NOTE: A client may request one or more byte ranges in a single request and subsequently, a server may respond with one or more ranges.

In the case of a single range being vended, the log entry will contain an entry equivalent to the HTTP Content-Range header.

In the case of a multi-part response from the server, there are multiple Content-Range headers, which are found preceding each byte range. For multi-part range responses, the log entry will contain each of the ranges vended separated by commas, followed by a slash "/" and the total size of the object.

Enabling

To enable this option, enter the following at the command line:


```
set accelerator name extlogcontentrange = yes
apply
```

where *name* is the name of your web accelerator service.

Disabling

To disable this option, enter the following at the command line:

```
set accelerator name extlogcontentrange = no
apply
```

where *name* is the name of your web accelerator service.

Getting Status

To get logging status for this option, enter the following at the command line:

```
get accelerator name extlogcontentrange
```

where *name* is the name of your Web Accelerator service.

Examples

```
"bytes 100-200/2000","bytes 500-700,952-999/2000",""
```

cs(Range)

Purpose

This field records the byte ranges requested by a client to the Excelerator cache device.

Values

The results include the word *bytes*, are enclosed in double quotes ("), and are either one or more byte ranges separated by commas (,) or a null value (") if this field is not applicable to the specific log entry (i.e. this is not a byte range request).

Enabling

To enable this option, enter the following at the command line:

```
set accelerator name extlogrange = yes
apply
```

where *name* is the name of your Web Accelerator service.

Disabling

To disable this option, enter the following at the command line:

```
set accelerator name extlogrange = no
apply
```

where *name* is the name of your Web Accelerator service.

Getting Status

To get logging status for this option, enter the following at the command line:

```
get accelerator name extlogrange
```

where *name* is the name of your Web Accelerator service.

Examples

"bytes 100-200", "bytes 500-700, 952-999", ""

cs(If-Range)

Purpose

This field records whether the client request was a conditional range request.

Values

The result is enclosed in double quotation marks (") and is either the text of the range request or a null value (") if this field is not applicable to the specific log entry (i.e. this is not a conditional byte range request).

NOTE: The value is typically either an Etag or a date.

Enabling

To enable this option, enter the following at the command line:

```
set accelerator name extlogifrange = yes  
apply
```

where *name* is the name of your Web Accelerator service.

Disabling

To disable this option, enter the following at the command line:

```
set accelerator name extlogifrange = no  
apply
```

where *name* is the name of your Web Accelerator service.

Status

To get logging status for this option, enter the following at the command line:

```
get accelerator name extlogifrange
```

where *name* is the name of your Web Accelerator service.

Examples

"Etag1234", "Tue, 17 Oct 2002", ""

sc(Content-Length)

Purpose

This field records the size (in bytes) of the entire object being vended from an Exceleator cache device to a requesting client.

Values

The result is contained in double quotes (") and is either the content-length in bytes or null (") if this field is not applicable to the specific log entry (i.e. the total size of the object is unknown).

NOTE: A null value (") in the field might indicate that the object is not in cache and has to be piped from the origin server.

Enabling

To enable this option, enter the following at the command line:

```
set accelerator name extlogcontentlength = yes
apply
```

where *name* is the name of your Web Accelerator service.

Disabling

To disable this option, enter the following at the command line:

```
set accelerator name extlogcontentlength = no
apply
```

where *name* is the name of your Web Accelerator service.

Getting Status

To get logging status for this option, enter the following at the command line:

```
get accelerator name extlogcontentlength
```

where *name* is the name of your Web Accelerator service.

Examples

```
"13245", "7776", ""
```

sc(Etag)

Purpose

This field records the Etag vended from an Exceleator cache device to a requesting client.

Values

The result is enclosed in double quotes (") and is either the etag-text or null (") if this field is not applicable to the specific log entry (i.e. no Etag provided).

NOTE: Etags may contain quotes, which will appear as "" to indicate that they are part of the field rather than the end of the field. For example, a log entry of ""123abc"" indicates an Etag value of "123abc". Etags may be used in conjunction with If-Range or other conditional requests.

Enabling

To enable this option, enter the following at the command line:

```
set accelerator name extlogetag = yes
apply
```

where *name* is the name of your Web Accelerator service.

Disabling

To disable this option, enter the following at the command line:

```
set accelerator name extlogetag = no  
apply
```

where *name* is the name of your Web Accelerator service.

Getting Status

To get logging status for this option, enter the following at the command line:

```
get accelerator name extlogetag
```

where *name* is the name of your Web Accelerator service.

Examples

```
""132mydoghasflees45"", "w/"123hiabc"", ""
```

cs(Pragma)

Purpose

This field records the pragma value associated with a client request to an Excelerator cache device.

Values

The result is enclosed in double quotes (") and is either the pragma text value or null (") if this field is not applicable to the specific log entry (i.e. a pragma header is not used).

NOTE: For HTTP 1.1 no-cache is the only valid value for pragma and although pragma: no-cache is intended for client to server requests, it has been implemented in practice in server to client responses as well.

Enabling

To enable this option, enter the following at the command line:

```
set accelerator name extlogrequestpragma = yes  
apply
```

where *name* is the name of your Web Accelerator service.

Disabling

To disable this option, enter the following at the command line:

```
set accelerator name extlogrequestpragma = no  
apply
```

where *name* is the name of your Web Accelerator service.

Getting Status

To get logging status for this option, enter the following at the command line:

```
get accelerator name extlogrequestpragma
```

where *name* is the name of your Web Accelerator service.

Examples

"no-cache", ""

sc(Pragma)

Purpose

This field records the pragma value associated with a server response from an Excelerator cache device to a requesting client.

Values:

The result is enclosed in double quotation marks (") and contains either the pragma text value or is blank if the field is not applicable to the specific log entry (i.e. a pragma header is not used)

NOTE: For HTTP 1.1 no-cache is the only valid value for pragma and although pragma: no-cache is intended for client to server requests, it has been implemented in practice in server to client responses as well.

Enabling

To enable this option, enter the following at the command line:

```
set accelerator name extlogreplypragma = yes  
apply
```

where *name* is the name of your Web Accelerator service.

Disabling

To disable this option, enter the following at the command line:

```
set accelerator name extlogreplypragma = no  
apply
```

where *name* is the name of your Web Accelerator service.

Getting Status

To get logging status for this option, enter the following at the command line:

```
get accelerator name extlogreplypragma
```

where *name* is the name of your Web Accelerator service.

Examples

"no-cache", ""

Pseudo Header Field Logging

The following pseudo header fields are also supported.

sc-completed

Purpose

This field either records that a transaction was completed successfully or it indicates the reason for a failure to complete.

Values

The field contains one of the following values:

Success	Indicates either a normal connection termination or that the transaction was completed and the persistent connection remains open
Timeout	Indicates the connection timed out, perhaps due to network failure, and Excelerator discarded the connection
Reset	Indicates the connection was terminated by the client
Administrative	indicates that Excelerator terminated the connection for reasons unrelated to the connection, such as Excelerator shutting down
-	Not applicable to the specific log entry

Enabling

To enable this option, enter the following at the command line:

```
set accelerator name extlogcompleted = yes
apply
```

where *name* is the name of your Web Accelerator service.

Disabling

To disable this option, enter the following at the command line:

```
set accelerator name extlogcompleted = no
apply
```

where *name* is the name of your Web Accelerator service.

Getting Status

To get logging status for this option, enter the following at the command line:

```
get accelerator name extlogcompleted
```

where *name* is the name of your Web Accelerator service.

Examples

Success, Reset, -

sc-header-size

Purpose

This field records the size (in bytes) of the HTTP header associated with a response to a client.

Values

The field will contain either a byte-count of the record size or a dash (-) if this field is not applicable to the specific log entry.

TIP: To determine whether an entire object was downloaded, subtract this data from sc-bytes and compare the difference to the entity size.

Enabling

To enable this option, enter the following at the command line:

```
set accelerator name extlogheadersize = yes  
apply
```

where *name* is the name of your web accelerator service.

Disabling

To disable this option, enter the following at the command line:

```
set accelerator name extlogheadersize = no  
apply
```

where *name* is the name of your web accelerator service.

Getting Status

To get logging status for this option, enter the following at the command line:

```
get accelerator name extlogheadersize
```

where *name* is replaced with the name of your Web Accelerator service.

Examples

678, 679, -

x-cache-info

Purpose

This field records brief status statements for objects vended from cache or brief reasons why a requested object was not cached.

Values

The following statements or reasons are enclosed in double quotation marks ("). Additional information is enclosed in parentheses ().

- pass thru
- error
- redirect
- in cache
- fresh
- revalidate each use
- stale
- not modified
- updated

not in cache

filled

if revalidate each use

(The reasons it was revalidate each use)

not cacheable

(The reason is was not cacheable)

NOTE: This information could be very useful in troubleshooting.

Enabling

To enable this option, enter the following at the command line:

```
set accelerator name extlogcacheinfo = yes  
apply
```

where *name* is the name of your Web Accelerator service.

Disabling

To disable this option, enter the following at the command line:

```
set accelerator name extlogcacheinfo = no  
apply
```

where *name* is the name of your Web Accelerator service.

Getting Status

To get logging status for this option, enter the following at the command line:

```
get accelerator name extlogcacheinfo
```

where *name* is the name of your Web Accelerator service.

Examples

"In Cache, Browser Requested Revalidate, Not Modified", "Not In Cache, Not Cacheable (No Validator Or Expiration Time)", "In Cache, Fresh"

35

FTP Services

You can manage several aspects of the appliance using an FTP client on a workstation connected to a network where the appliance is visible.

The appliance's FTP services let you get and put configuration files, log files, and the optional splash screen (which you can customize with your own HTML).

You can also configure the appliance to provide FTP forward proxy (client acceleration) and FTP reverse proxy (server acceleration) services. For more information, see [“FTP Tab” on page 354](#).

Tips for Using Excelerator FTP Services

The following sections contain important information about using Excelerator FTP services.

Firewalls Usually Require Passive FTP

The appliance's system supports access from both active and passive FTP clients.

The fact that passive FTP is often required to traverse a firewall has certain implications for using FTP with the appliance.

FTP Access Through a DOS Window Is Limited

Because DOS uses only active FTP, you cannot access an appliance that is outside a firewall through a DOS window on a client inside the firewall.

The reverse is also generally true. If the appliance is inside a firewall, you will not usually be able to access it through a DOS window on a client outside the firewall.

FTP Access Through a Browser Outside the Firewall Is Also Limited

You cannot generally access an appliance inside a firewall through a browser that is outside the firewall.

Browser Must Be Properly Configured

Because Netscape browsers use passive FTP, you can generally access an appliance outside the firewall from a Netscape* browser inside the firewall.

To access an appliance outside a firewall using Internet Explorer 5 inside the firewall, you must configure the browser to use passive FTP. Complete the following steps:

- 1 In the browser click Tools > Internet Options > Advanced.
- 2 Under Browsing, check Use Web-Based FTP.

NOTE: This option name varies according to browser version. In Internet Explorer 5.5, for example, the option is Use Passive FTP for compatibility with some firewalls and DSL modems.

- 3 Click OK.

Directory and File Names Cannot Contain Spaces

The Mini FTP Server will not work with directory or file names that contain spaces.

Appliance Routing and Transparent Proxy Limitations on FTP

When using the built-in appliance router capabilities in a *transparent* proxy situation, users will not be able to perform browser-based FTP using ftp:// as the protocol. Browser-based FTP works normally with forward proxy.

Switching from Anonymous FTP to Username/Password FTP

If you are using browser-based FTP with forward proxy through the Excelerator appliance and you want to switch from anonymous FTP to username/password FTP, do the following:

- 1 In the browser, enter the new URL.
- 2 Click the browser's Refresh button.

IMPORTANT: You must refresh the screen contents or you won't be able to switch to the other FTP service.

Setting Up Appliance FTP Services

Before using FTP services to manage the appliance, you must ensure the appliance's Mini FTP Server is properly configured.

- 1 Start the browser-based management tool > click Cache > FTP.
- 2 Ensure that at least one of the IP addresses in the Server IP Addresses list is checked.

You will use the checked address for your FTP session. IP address 10.1.1.1 is checked by default.

Functionality Limitations of the Appliance's Mini FTP Server

The appliance's Mini FTP Server was originally designed only for uploading and executing appliance configuration (.NAS) files. This functionality has been expanded slightly and currently supports the following commands:

- ♦ CWD
- ♦ DELE
- ♦ GET
- ♦ LIST
- ♦ MPUT
- ♦ PASS
- ♦ PASV
- ♦ PORT

- ♦ PUT
- ♦ PWD
- ♦ QUIT
- ♦ RETR
- ♦ STOR
- ♦ SYST
- ♦ TYPE
- ♦ USER

The FTP server has no support for downloading with wildcard characters.

FTP server access is limited to the following directories and their subdirectories:

- ♦ SYS:ETC\PROXY\APPLIANCE
- ♦ SYS:ETC\PROXY\DATA
- ♦ SYS:ETC\APPLIANCE
- ♦ LOG:ETC\PROXY\DATA

When you log in to the FTP server, the SYS:ETC\PROXY\APPLIANCE\CONFIG\USER directory is the default.

To execute a .NAS configuration file, you must be in this default directory and use the following syntax:

```
put local_filename remote_filename,execute
```

The *local_filename* variable is the name of the .NAS file on your local machine and the *remote_filename* variable is the name after the file is uploaded to the appliance.

Starting an FTP Session with the Appliance

- 1** Set up FTP services on the cache device by completing the steps in “[Setting Up Appliance FTP Services](#)” on page 262.
- 2** Launch your FTP application and enter a valid appliance IP address, such as
`ftp 10.1.1.1`
- 3** Log in to the appliance using the Config username and password you have set.

Changing the FTP Working Directory

The default working volume and directory for FTP sessions is SYS:\ETC\PROXY\APPLIANCE\CONFIG\USER. Therefore, the SYS: volume is implied for all FTP commands unless the LOG: volume is specifically included.

You can use the cd (change directory) command to change the current volume and directory path. For example, cd log/etc/proxy/data changes the working path to the LOG: volume and the directory path to \ETC\PROXY\DATA.

FTP commands that include only a directory path and that are issued subsequent to the cd command will use the newly specified volume.

You can specify a full path (volume and directory) when using the get and put commands to copy files to or from any FTP-accessible location. However, the default volume and directory are unaffected by these commands. Only the cd command changes the FTP working path.

Managing Configuration Files with FTP

All appliance settings are contained in configuration text files with the extension .NAS. These files can be edited and sent through FTP back to the appliance where they can be executed as a means of instant configuration.

By using several configuration files, you can apply different scenarios instantly, such as turning various proxy services on or off. The appliance updates the default CURRENT.NAS file every time you apply a setting. Additionally, if you have created a file named AUTOLOAD.NAS on a floppy disk, it is updated with each change. You can export other configuration files from the browser-based management tool or the Telnet/command line interface.

Downloading a Configuration File to Your Workstation Using FTP

- 1 Start FTP and log in as explained in “Starting an FTP Session with the Appliance” on page 263.
- 2 Enter the following, where *filename* is the name of your configuration file:

```
get filename.nas
```

The file is transferred to your FTP client’s default directory.

Moving a Configuration File to the Appliance from a Workstation

- 1 Start FTP and log in as explained in “Starting an FTP Session with the Appliance” on page 263.
- 2 Enter the following, where *filename* is the name of your configuration file:

```
put filename.nas
```

The file is transferred to the appliance.

Moving a Configuration File to the Appliance and Executing It

- 1 Start FTP and log in as explained in “Starting an FTP Session with the Appliance” on page 263.
- 2 Enter the following, where *filesrc* is the name of the source configuration file and *filedst* is the name used at the destination:

```
put filesrc.nas filedst.nas,execute
```

Customizing the Appliance Splash Screen with FTP

The appliance has an optional splash screen that can be displayed periodically before pages are vended.

The splash screen has the root filename BMSPLASH.HTM. The screen is disabled by default. For more information, see “Tuning Tab” on page 399.

To edit the splash screen, do the following:

- 1** Start FTP and log in as explained in “Starting an FTP Session with the Appliance” on page 263.
- 2** Use the following command to get the file for editing:

```
get /etc/proxy/data/bmsplash.htm
```
- 3** Modify the file as desired.
- 4** Put the file back on the appliance using the following command:

```
put bmsplash.htm /etc/proxy/data/bmsplash.htm
```

Using Non-Anonymous FTP

Non-anonymous FTP can be used with the Excelerator appliance.

Client Acceleration (Forward Proxy)

For non-anonymous FTP using Client Acceleration (forward proxy), enter from a DOS window:

syntax: `ftp appliance_proxy_address`

username: `username$ftp_host`

password:

or from a browser:

`ftp://username$ftp_host:user's password@appliance_proxy_address`

NOTE: If you have the browser set to proxy through Excelerator using HTTP, then you need to make sure that the field for FTP proxy is empty. This is because most browsers try to do FTP proxy requests via HTTP by default. Also make sure that no IP address or port numbers are listed.

Reverse Proxy

Non-anonymous FTP can also be used in a browser via FTP reverse proxy. The browser must not be set up to use any proxy, or if it is, the FTP proxy field and port number must be left blank.

From a browser, enter:

`ftp://username:user's password@appliance_proxy_address/`

Accelerating FTP Requests

FTP Forward Proxy

The workstation sends a request directly to an appliance IP address configured for FTP forward proxy services. The FTP forward proxy service obtains the objects and forwards copies back to the workstation. This is similar to the explanation of HTTP forward proxy found in “Overview of Forward Proxy” on page 35, but the FTP forward proxy service caches only objects retrieved using anonymous FTP.

For tips and guidelines on setting up FTP forward proxy, see “FTP Forward Proxy Setup” on page 266.

FTP Accelerator

DNS routes requests originally targeted at the origin FTP server to the FTP accelerator service instead. The FTP accelerator handles the request, accessing the origin FTP server only when needed objects are not cached.

FTP requests meant for FTP servers can be routed to the FTP accelerator service instead:

- ◆ Without acceleration, DNS resolves the origin FTP server's DNS name to the origin server's IP address.
- ◆ With acceleration, DNS resolves the origin FTP server's name to the IP address of an appliance FTP accelerator service.

FTP acceleration is similar to Web server acceleration found in [“Overview of Web Server Acceleration” on page 51](#). However, FTP acceleration *does not support* load balancing; you can assign only a single FTP host per appliance IP address.

For tips and guidelines on setting up an FTP accelerator, see [“FTP Accelerator Setup” on page 266](#).

FTP Forward Proxy Setup

Set up FTP forward proxy services as follows:

To	Do This	Notes
Ensure your basic network configuration is complete	1. See “Basic Network Configuration Setup” on page 27 .	
Enable FTP forward proxy services on the appliance	<ol style="list-style-type: none">1. In the browser-based tool, click Cache > FTP.2. In the FTP Forward Proxy list, check the IP addresses that FTP forward proxy services will be available on.3. Click Apply.	<p>All appliance FTP services listen on port 21. Therefore, each IP address on the appliance can be configured for only one FTP service (Mini FTP, Forward Proxy, or FTP Accelerator).</p> <p>For more information, see “Client Accelerator Tab” on page 333.</p>

FTP Accelerator Setup

Set up FTP accelerator services as follows:

To	Do This	Notes
Ensure your basic network configuration is complete for each appliance	1. See “Configuring the Accelerator Appliance” on page 28 .	
Ensure that DNS resolves FTP requests to the appliance IP addresses configured for the FTP accelerator services	1. See “Working with DNS” on page 53 .	Although the DNS configuration instructions are targeted at HTTP acceleration, the same basic principles apply.

To	Do This	Notes
Set up one or more FTP accelerator services	<ol style="list-style-type: none"> 1. In the browser-based tool, click Cache > FTP. 2. Under the FTP Accelerator list, click Insert. 3. Enter a name for the FTP accelerator. 4. Enter either the IP address or DNS name of the FTP server being accelerated. 5. Check the appliance IP address for the accelerator service. 6. Click OK > Apply. 	<p>All appliance FTP services listen on port 21. Therefore, each IP address on the appliance can be configured for only one FTP service (Mini FTP, Forward Proxy, or FTP Accelerator).</p> <p>If the FTP Server address is a DNS name, make sure it is not the name that now resolves to the appliance's numeric IP address. That would create an endless loop.</p>

36

Object Pinning

This section contains information about object pinning.

The Pin List

Pinned objects remain in cache indefinitely unless it fills up. This ensures that they are available from cache and will not be bumped out by more recently requested objects.

The pin list contains URL patterns for identifying objects on the Web. You configure each URL pattern in the list with specific handling instructions, as explained in the following sections.

Configuring Pin Lists

To configure the appliance's pin lists, do the following:

- 1 In the browser-based tool click Cache > Management.
- 2 Check Enable Pin List.
- 3 From the drop down list box, select a default Refresh Frequency.

NOTE: If you selected Timed Interval, enter a default refresh time.

For more information see [“Refresh Frequency/Time” on page 271](#).

To add a URL Mask to the pin list, do the following:

- 1 Click Insert.
- 2 Enter the URL mask.
For more information, see [“URL Mask” on page 270](#).
- 3 From the drop down list, select the pin type.
For more information, see [“Pin Type” on page 270](#).
- 4 From the drop down list, select the number of pin links.
For more information, see [“Pin Links” on page 271](#).
- 5 If you want to pin image files that reside on a different host than the page requested, check the Pin Images check box.

To add more URL Masks, complete steps 1 through 5 above until you are done.

- 6 Click Apply to save the changes.

To modify an existing URL mask, do the following:

- 1 Select the desired URL.
- 2 Click Modify.

Make the desired changes.

3 Click Apply.

To exit without applying changes, do the following:

- ◆ Click Cancel.

To disable the pin list without deleting your list, do the following:

- ◆ Click Enable Pin List.

The list grays out, indicating the list is disabled.

URL Mask

The URL mask can contain complete or partial URL patterns. A single URL mask might apply to a large set of URLs, or it might be so specific that only a single file on the Web matches it. For more information, see [“Pin List Examples” on page 274](#).

The appliance processes the masks in the pin list in order of specificity. A mask containing a host name is more specific than a mask that specifies only a file type. The action taken for an object is the action specified for the first mask that the object matches. For more information, see [“How URL Masks Are Processed” on page 271](#).

If the mask contains an asterisk, only the pin type can be specified. The Pin Links, Pin Images, and Refresh Frequency/Time options are not available for URLs containing this wildcard. Objects matching a mask with an asterisk are not automatically downloaded, but are pinned in cache only as individually requested.

Pin Type

The Pin Type options specifies whether and how the appliance will cache objects that match the URL mask:

- ◆ **Normal:** Excelerator handles objects matching the mask in the same way it handles any other requested objects. In other words, objects are cached but not pinned.

Administrators often use this pin type in combination with a broad URL mask that has a bypass pin type. This allows them to insulate specific objects from the effects of the bypass rule.

For example, you could specify a URL mask of `/*.jpg` with a pin type of bypass and a second URL mask of `www.foo.gov/graphics/*.jpg` with a pin type of normal. The result would be that the `.jpg` files in the graphics directory on the `foo.gov` Web site would be cached as requested. They would not, however, be pinned in cache. Assuming there were no other URL masks in the pin list, all other `.jpg` graphics would not be cached due to the first URL mask.

- ◆ **Cache:** Excelerator keeps the pinned objects in cache as long as possible, although they might be written to the appliance’s hard disk.
- ◆ **Memory:** Excelerator keeps the pinned objects in memory as long as possible, writes them to disk when memory gets too full, and places them back in memory as soon as they are requested by a user of the cache.
- ◆ **Bypass:** Excelerator does not cache the objects. In other words, you can use this option to prevent objects from being cached.

Pin Links

The Pin Links option specifies how many link levels Excelsior will follow the pin type rule you've established. Selecting levels 1 or 2 causes all linked objects, including the images on the host, to be downloaded and cached when the pin list is applied to the appliance configuration, and then to be periodically refreshed as specified.

For example, if the requested object is an HTML page and you have specified a pin links level of 1, the HTML page along with all the items linked from the page will be downloaded and cached when the pin list is applied. These cached objects will also be refreshed at the refresh frequency and time specified.

To use levels 1 or 2, you must specify an absolute address that includes the scheme, host, and path for the URL mask, such as `http://www.foo.gov/documents/`. The tool will let you insert masks that do not meet this requirement, but the entries are removed when you click Apply.

IMPORTANT: Attempting to include an asterisk wildcard immediately hides this option.

Pin Images

The Pin Image option is used to pin image files that reside on a different host than the page requested. It works in conjunction with the Pin Links option, which specifies how many levels of links Excelsior will follow when downloading a page.

For example, if the requested HTML page uses images that reside on another host and you have checked this option, the HTML page will be cached along with all the image files associated with the page, including those on the other host. If you have also specified a pin link level, images on the linked pages that reside on another host will also be pinned.

On the other hand, if the Pin Images option is not checked, Excelsior only pins the images that reside on the same host as the requested page.

Refresh Frequency/Time

The Refresh Frequency/Time option lets you specify a refresh frequency and time for the URL that is different from the default values shown above the pin list.

How URL Masks Are Processed

You can enter four basic types of URL masks in the pin list. The following table lists each type, provides a few examples of each, and provides information on how they are processed by Excelsior.

Type	URL Mask Examples by Specificity	Notes
Hostname	http://www.foo.gov/documents/picture.gif http://www.foo.gov/documents/ http://www.foo.gov/ http://foo.gov/documents/ http://foo.gov/ *.foo.gov/	<p>Although these entries can include the protocol or scheme, the DNS name, the path, and the filename, only the DNS name or hostname must be present in the mask. All DNS label portions must be indicated, if only by an asterisk wildcard.</p> <p>Excelerator processes hostname entries before it processes other mask types. It also processes the most specific URL mask entries first.</p> <p>When an object match occurs, Excelerator applies the pin type rule, and processing of the object is finished.</p> <p>For example, if the first URL mask in the examples column has a pin type rule of bypass, picture.gif will not be cached regardless of the pin type rules for the other URL masks.</p> <p>Hostname entries can have a dramatic impact on object pinning and cache bypassing.</p> <p>For example, if the first two URL masks in the examples column were not present, a pin type of Bypass on the third URL mask would prevent caching of all objects delivered through HTTP on the www.foo.gov Web site.</p> <p>If no scheme (HTTP, FTP, etc.) is indicated, the mask applies to all schemes. The last three masks would apply to objects delivered through any Web protocol.</p> <p>Finally, Excelerator interprets hostnames literally. For example, the sixth entry would cover www.foo.gov, ww1.foo.gov, army.foo.gov, etc., but the fourth and fifth entries would not, because a scheme is assumed to immediately precede the hostname.</p>

Type	URL Mask Examples by Specificity	Notes
Path	/documents/picture.gif	<p>Excelerator processes path entries after all hostname entries have been considered. It assumes that the first forward slash immediately follows a hostname.</p> <p>For example, the first entry would apply only to a graphics file named PICTURE.GIF that is located in a DOCUMENTS directory at the root of the host.</p> <p>The forward slash in the second entry causes Excelerator to assume that PICTURE.GIF is a directory. The pin type rules associated with this entry would apply to any matched objects that have a URL directory path that starts with a documents directory followed by a subdirectory named PICTURE.GIF.</p> <p>The third entry would apply to any matched objects that contain a DOCUMENTS directory at the start of their URL paths.</p>
	/documents/picture.gif/	
	/documents/	
Filename	/picture.gif	<p>After the path entries have all been processed, Excelerator looks for specific filenames.</p> <p>For example, if requested objects named PICTURE.GIF, WIDGET.JS, and DEFAULT.HTM have not been covered by one of the hostname or path entries above, the files will have the pin type rule for their respective filename mask applied to them.</p> <p>If the first entry carries a pin type rule of Bypass, all PICTURE.GIF files that didn't match previously processed hostname or path masks would not be cached.</p>
	/widget.js	
	/default.htm	
File Extension	/*.gif	<p>File extension entries are processed last.</p> <p>These are simply filename entries with the root of the filename replaced by an asterisk, which makes them less specific than complete filenames.</p> <p>For example, if the examples shown all had pin types of Bypass, then only those .GIF, .JS, and .HTM files that had been cached and pinned because of hostname, path, or filename masks would be stored in cache. All other files with the named extensions would not be cached.</p>
	/*.js	
	/*.htm	

About Wildcards in Pin Lists

Only the asterisk (*) wildcard is allowed in pin list entries.

Excelerator interprets everything between an asterisk and the next delimiter to the right (a forward slash [/], a period[.], or a colon [:]) as a wildcard. This effectively allows only one asterisk between delimiters.

Pin List Examples

The following table provides brief examples of sample pin list entries and their effects on appliance caching.

URL Mask	Pin Type	Pin Links	Pin Images	Effect on Cache
http://www.foo.gov/documents/	cache	1	Yes	<p>As a general rule, you should always include fully qualified DNS names or hostnames in the pin list. Excelerator resolves these more quickly than other masks, and you will be able to track the effects on pinning more easily.</p> <p>For this URL mask, Excelerator downloads, caches, and pins all objects whose URL starts with the mask. In other words, all objects below the documents directory will be downloaded, cached, and pinned. Also, all objects that are linked from one of the pinned objects will be downloaded, cached, and pinned. And finally, images that reside on other hosts will be downloaded, cached, and pinned as well.</p> <p>Objects will be refreshed according to the refresh settings (default or specific) as specified in the pin list entry.</p>
www.foo.gov/groups.html	cache	1	No	<p>Excelerator downloads, caches, and pins objects (including images) in the GROUPS.HTML page and in pages linked from that page. Any images referenced from other hosts, however, are not included.</p>

URL Mask	Pin Type	Pin Links	Pin Images	Effect on Cache
www.foo.gov/groups.html/	normal	1	Yes	<p>Excelerator downloads and caches objects in the subdirectory named groups.html and in pages linked from any of those objects.</p> <p>The forward slash tells Excelerator that this is a directory rather than a file.</p> <p>Objects are cached but not pinned in cache, meaning they might be bumped by more frequently accessed objects or objects that are pinned.</p> <p>Images linked from other hosts are downloaded and cached.</p>
www.foo.*	bypass	n/a	n/a	<p>Excelerator doesn't cache objects from any URLs whose DNS names begin with www.foo.</p> <p>All domain extensions (.com, .net, .org, etc.) are covered by the asterisk wildcard.</p> <p>Link and image pinning is not available for bypass pin types.</p> <p>If this entry appeared in a pin list with either of the previous two entries, it would not prevent caching of objects covered by them because it is less specific than they are.</p>
w*.f*.com	bypass	n/a	n/a	<p>Excelerator doesn't cache objects for any URLs whose first domain label begins with w and second domain label begins with f, providing the domain extension is .com.</p> <p>This mask doesn't prevent caching of objects on other domains such as .net or .gov.</p>
w*.f*.*	bypass	n/a	n/a	<p>This mask functions like the previous entry, but the domain is not limited to .com.</p>

URL Mask	Pin Type	Pin Links	Pin Images	Effect on Cache
.foo.	cache	n/a	n/a	<p>This causes all objects on any Web server whose second domain label is foo to be pinned in cache.</p> <p>Link and image pinning are not available because the mask contains asterisks.</p> <p>This mask would not cover DNS names that don't have a domain label before foo. For example, foo.gov would not normally be covered. However, if foo.gov happened to resolve in DNS to the same IP address as www.foo.gov, Excelsator would apply the pinning rules specified for www.foo.gov to foo.gov. To understand more about IP addresses and URL masks, see “Excelsator Records IP Addresses When Resolving URL Masks” on page 276.</p>

Excelsator Records IP Addresses When Resolving URL Masks

As stated earlier, you should include fully qualified DNS names or hostnames in URL masks whenever possible.

Excelsator resolves DNS names to their respective IP addresses and uses those addresses when pinning objects.

You can use this fact when constructing your pin list entries.

For example, if you use the DNS name www.foo.gov as the URL mask and you know that the DNS name foo.gov resolves to the same IP address, you don't need to include foo.gov in the pin list.

Excelsator will treat objects for both DNS names the same because both URLs resolve to the same IP address.

On the other hand, if www.foo.gov and foo.gov resolve to different IP addresses, separate pin list entries would be required to cover both sites.

37

Router Capabilities

Having the appliance double as a router impacts appliance performance, but it is a low-cost router option that delivers acceptable performance in some low-volume networks.

Each appliance is normally configured with a default gateway. If the appliance is not acting as a router, the default gateway is the appliance's next hop.

For more information about appliance routing, see [“Gateway/Firewall Tab” on page 326](#) and [“Router Options Dialog Box” on page 347](#).

Using Appliance Routing

If the appliance is acting as a router, it routes requests to IP addresses based on the information in its routing table. If a request could be routed through multiple gateways, the appliance chooses the gateway associated with the most restrictive mask (the smallest range of destination addresses).

Routing table entries fall within three basic groups:

- ♦ Host gateways for specific destination addresses
- ♦ Network gateways for destination addresses that fall within specific subnets
- ♦ The default gateway for destination addresses that aren't covered by host or network gateways

The syntax for this gateway is often expressed in router configuration tables as 0.0.0.0 / 0.0.0.0 / *iii.iii.iii.iii*, where the *i*'s represent the IP address of the default gateway.

You define these gateways in the browser-based management tool by clicking Network > Gateway/Firewall > Additional Gateways or by clicking Cache > Client Accelerator > Router Options.

IMPORTANT: If the appliance is acting as a router and you don't specify a default gateway, the appliance routes only those requests whose destination addresses are covered by a host or network gateway. Other requests are not routed.

For more information on routing concepts, see one of the TCP/IP references available at any bookstore carrying computer reference manuals.

38 Shutting Down and Restarting

If you need to shut down or restart an appliance, you should do shut it down properly to protect the data in memory and ensure the data is written to disk.

Restarting from the Browser-Based Management Tool

- 1 Start the browser-based management tool.
- 2 Click System > Actions.
- 3 Shut down Excelerator by clicking Shut Down or shut it down and restart it by clicking Restart.

You are given a chance to verify your selection.

If you choose to shut Excelerator down, you hear a three-beep sequence that repeats until the appliance is turned off or restarted.

If you choose to restart Excelerator, you hear a two-beep, four-beep sequence that repeats for a short period and then stops, signifying that your appliance has restarted.

If you do not have access to the physical location of the appliance, you can test to see if the appliance has restarted by pinging its address on port 1959. If the ping succeeds, the appliance has restarted.

Shutting Down and Restarting from Telnet or the Command Line

You can shut down or restart an appliance down from the command line.

NOTE: Both actions break the connection. If you *restart* the appliance from a remote connection, you will be able to reconnect after the appliance restarts. If you *shut down* the appliance, however, someone will need physical access to the appliance to restart it.

To restart the appliance from the command line, enter

Restart

If you are near the appliance, you will hear a two-beep, four-beep sequence that repeats for a short period and then stops, signifying that your appliance has restarted.

To shut down the appliance from the command line, enter the following:

Shutdown

If you are near the appliance, you will hear a continuous three-beep sequence when the system has been shut down. When you hear the three-beep sequence, you can restart the system from an attached keyboard by pressing Ctrl+Alt+Del or you can shut off the power.

39

SOCKS Client Services

This section contains information about SOCKS client services.

Using the SOCKS Client Service

The appliance provides a SOCKS client service that can be used to redirect all forward proxy traffic through the SOCKS firewall. Redirecting appliance traffic to the SOCKS firewall might significantly reduce appliance performance.

The appliance currently supports both SOCKS4 and SOCKS5 protocols.

For information about setting up SOCKS client support, see [“Gateway/Firewall Tab” on page 326](#).

40 Time Synchronization

Time settings offered within the management tool are more than adequate for most system needs. For more information on the specific parameters available, see [“Synchronizing Time” on page 283](#) and [“Date/Time Tab” on page 302](#).

Additional flexibility in setting system time, including changing the GMT offset and daylight saving parameters, is available through the command line interface.

Synchronizing Time

You can either set the time manually or synchronize it using the network time protocol (NTP). The appliance uses NTP by default and comes preconfigured with two servers:

63.192.96.3

64.243.118.2

You can add or delete servers using the browser-based management tool and the command line interface. For more information regarding NTP functionality, see [“NTP Date/Time Synchronization Is Not Immediate” on page 284](#).

Using the Browser-Based Management Tool

Adding or Deleting an NTP Server

- 1** Start the browser-based management tool > click System > Date/Time.
- 2** Check Use Network Time Protocol.
- 3** Do one of the following:
 - ♦ To add a server, click Insert > type the URL or IP address of the server.
 - ♦ To delete a server, select the server > click Delete.

Changes in the Date/Time tab are effective immediately.

Setting the Time Manually

- 1** Start the browser-based management tool > click System > Date/Time.
- 2** Check Set Time Manually.
- 3** Click Set Time.
- 4** Using the drop-down lists, select the correct time and date.
- 5** Click OK.

Using the Command Line

1 Do one of the following:

- ♦ To add an NTP server address, enter
add ntp server=63.192.96.3
- ♦ To enable NTP, enter
set ntp enable=yes
- ♦ To disable NTP, enter
set ntp enable=no

2 To have the changes take effect, enter **apply**.

For more command line options, refer to the appliance's command line help.

NTP Date/Time Synchronization Is Not Immediate

When you specify an NTP server, synchronization between the NTP server clock and the appliance clock might not be immediate.

If the NTP server clock has an earlier date or time setting than the appliance clock, the system will slow down the appliance clock until the two are synchronized. This provides for proper incrementation of log files and other time-sensitive information during the synchronization process.

If the NTP server clock has a later date or time setting than the appliance clock, synchronization between the two will generally be immediate. However, in certain situations you might observe the appliance clock incrementing by 600-minute intervals. This is normal system behavior.

IMPORTANT: The Apply button changes from Wait back to Apply. This indicates only that the NTP configuration change has been made, not that the appliance clock is fully synchronized with the NTP server.

If necessary, you can set appliance time manually to the target time and then re-enable the NTP feature.

41

Web Proxy Auto-Discovery (WPAD)

When properly configured for the browsers used on your network, the Web Proxy Auto-Discovery (WPAD) feature lets network users automatically access the appliance's forward proxy services without having to individually configure their browsers.

Customizing Web Proxy Auto-Discovery

When WPAD is enabled in Cache > Client Accelerator, the appliance activates WPAD listening on all IP addresses configured for client acceleration.

The appliance answers WPAD requests from browsers by returning a .DAT file that contains standard WPAD JavaScript* configuration instructions for the browser.

Procedures for configuring browsers to request and use WPAD information vary. Internet Explorer supports both automatic detection of WPAD settings and automatic configuration files retrieved from known URLs. Netscape* supports only the latter WPAD configuration method.

Internet Explorer's automatic detection feature requires that the DNS server databases which the workstations use contain wpad.domain.com entries that resolve to the appropriate appliance forward proxy address. See [“How the Appliance Handles WPAD Requests” on page 285](#).

If your proxy services require additional browser configuration, you can create a default configuration file for all WPAD requests, or you can customize any or all of the system-created .DAT files. See [“Creating a Default WPAD.DAT Configuration File” on page 286](#) and [“Customizing System-Created WPAD DAT Files” on page 286](#) for details.

How the Appliance Handles WPAD Requests

When the appliance receives a WPAD request, it looks in \ETC\PROXY\DATA for a WPAD.DAT file. This file exists only if an appliance administrator has created it.

If the file is found, the appliance returns it to the requesting browser regardless of which client acceleration IP address the request was received on. Thus, the same WPAD configuration is used for all requesting browsers.

If the file is not found, the appliance returns a .DAT file created by the caching system for the specific IP address through which the request was received.

The appliance automatically creates and stores a WPADxx.DAT file in \ETC\PROXY\DATA for each IP address configured for client acceleration. Files are automatically named with successive numbers, as shown in the following table.

Order in Which IP Address Are Enabled for Client Acceleration	WPAD Configuration Filename
First IP address enabled	WPADX1.DAT
Second IP address enabled	WPADX2.DAT
Third IP address enabled	WPADX3.DAT

Creating a Default WPAD.DAT Configuration File

You can create a default configuration file that will be sent in response to all WPAD requests.

This file might provide backup proxy references by containing a list of all appliance addresses configured for client acceleration. Alternatively, it might redirect all WPAD clients to use only a specific IP address for proxy services. Or it might contain additional configuration instructions to browsers using WPAD.

To create a WPAD.DAT file, complete the following steps:

- 1 If you have not already done so, enable at least one appliance IP address for client acceleration in Cache > Client Accelerator and apply the settings.
- 2 Start an FTP client on a workstation that has access to the appliance.
For help, see [“Starting an FTP Session with the Appliance” on page 263](#).
- 3 Point the FTP client to one of the appliance’s IP addresses.
- 4 Enter the following command:

```
get /etc/proxy/data/wpadx1.dat
```

The .DAT file for the first client accelerator is transferred to the FTP client's default directory.

- 5 Modify the file using an ASCII editor. Be sure to conform to JavaScript programming conventions.

Information about WPAD configuration file content and functionality is available in various locations on the Web. Be aware that some conventions are browser-specific.

- 6 Use the PUT command to rename the modified file and place it back on the appliance by entering the following command:

```
put wpadx1.dat /etc/proxy/data/wpad.dat
```

The appliance now answers all WPAD requests by returning the WPAD.DAT file you created.

Customizing System-Created WPAD DAT Files

You can customize the system-created, IP address-specific .DAT files so that browsers are configured differently depending on the IP address through which the WPAD request originates.

To customize any of the system-created .DAT files on your appliance, complete the following steps:

- 1 If you have not already done so, enable at least one appliance IP address for client acceleration in Cache > Client Accelerator and apply the settings.
- 2 Start an FTP client on a workstation that has access to the appliance.

For help, see “Starting an FTP Session with the Appliance” on page 263.

- 3 Point the FTP client to one of the appliance’s IP addresses.
- 4 Enter the following command, where *xx* represents the system-assigned ID number of the file starting with x1:

```
get /etc/proxy/data/wpadxx.dat
```

The .DAT file for the first client accelerator is transferred to the FTP client's default directory.

- 5 Modify the file using an ASCII editor. Be sure to conform to JavaScript programming conventions.

Information about WPAD configuration file content and functionality is available in various locations on the Web. Be aware that some conventions are browser-specific.

- 6 Use the PUT command to copy the file back to the appliance by entering the following command, where *xx* represents the system-assigned ID number of the file starting with x1:

```
put wpadxx.dat /etc/proxy/data/wpadxx.dat
```

- 7 Repeat this procedure for each .DAT file you need to customize.

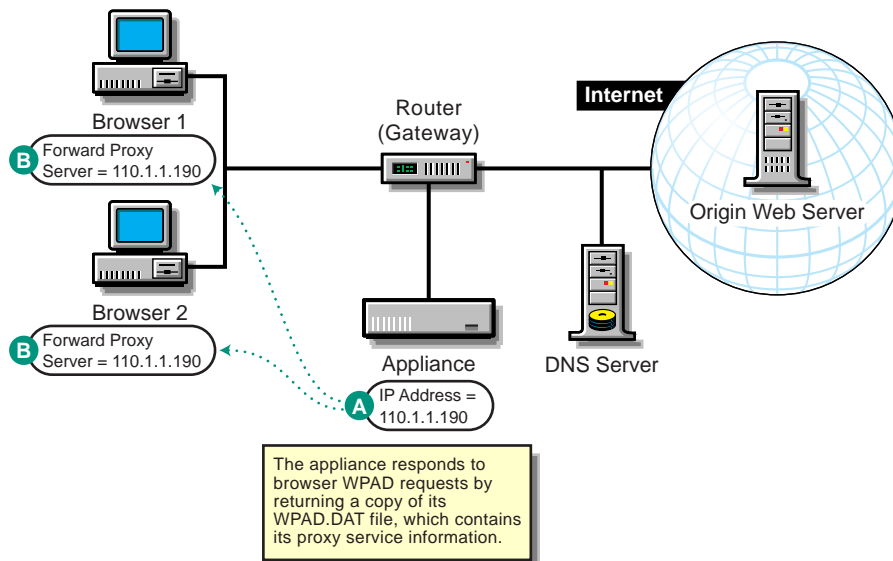
The appliance now answers WPAD requests according to the customizations you have made.

Setting Up Forward Proxy with WPAD

Figure 67 provides a visual map for the information in this section.

NOTE: The letters in Figure 67 are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 67



You can set up forward proxy and WPAD services as explained in the following table:

To	Do This	Notes
Ensure your basic network configuration is complete	1. See "Basic Network Configuration Setup" on page 27.	
Enable forward proxy and WPAD services on the appliance	<ol style="list-style-type: none"> 1. In the browser-based tool, click Cache > Client Accelerator > Enable Client Acceleration (Forward Proxy). 2. In the Proxy IP Addresses list, check the IP addresses that forward proxy services will be available on. 3. Enter the port that Excelerator will receive and process forward proxy requests on. (The default is 8080.) 4. Check Enable Automatic Proxy Configuration (WPAD). 5. Click Apply. 	<p>See A in Figure 67 on page 287.</p> <p>The WPAD listener port number is 80, which is the default port for transparent proxy and Web server accelerators. Because proxy services take precedence, you cannot have another service that uses port 80 configured on an IP address that needs to accept WPAD requests.</p> <p>After Excelerator has been enabled for WPAD, it will automatically listen for WPAD requests on all IP addresses enabled for forward proxy services as long as the addresses are not also enabled for a service that uses port 80.</p> <p>For more information, see "Client Accelerator Tab" on page 333.</p>
<p>Configure your network to provide the required information to requesting browsers</p> <p>(To work with WPAD services, you must have Internet Explorer 5 or other browsers on your network that rely on DNS, SLP, DHCP, or other protocols.)</p>	<p>Procedures for configuring Web protocols to work with WPAD vary.</p> <p>Information and instructions are found on the Web. A sample site is listed in the Notes column.</p> <p>After you understand the requirements for your network and the browsers being used, configure your network to provide WPAD services to network browsers.</p>	<p>For more information on using auto-config files like WPAD.DAT, see "Navigator Proxy Auto-Config File Format" on the Web (http://www.netscape.com/eng/mozilla/2.0/relnotes/demo/proxy-live.html).</p> <p>The WPAD RFC draft might also help you understand your network setup requirements. See a draft copy on the Web (http://www.wrec.org/drafts/draft-tomlinson-epsfw-00.txt).</p>

To	Do This	Notes
Enable the client browsers to request and use the WPAD information	<p>Read the notes to understand which of the following applies to your browsers.</p> <p>To configure Netscape Communicator 4.5 to access an appliance configuration file:</p> <ol style="list-style-type: none"> 1. Click Edit > click Preferences > check Advanced > click Proxies > check Automatic Proxy Configuration. 2. Enter the URL for the WPAD-enabled appliance's WPAD file as follows: <code>http://IP_address/wpadxx.dat</code>, where <i>IP_address</i> is the appliance's IP address and <i>xx</i> is the file numbering convention explained in "Customizing Web Proxy Auto-Discovery" on page 285. <p>To configure Internet Explorer 5 to use automatic detection:</p> <ol style="list-style-type: none"> 1. Click Tools > Internet Options > Connections > LAN Settings. 2. Check Automatically Detect Settings. 	<p>See <i>B</i> in Figure 67 on page 287.</p> <p>Excelerator sends a WPAD configuration file to requesting browsers. This file contains proxy service information for the network.</p> <p>Procedures for configuring browsers to request and use WPAD information vary. Internet Explorer supports both automatic detection of WPAD settings and automatic configuration of files retrieved from known URLs. Netscape supports only the latter WPAD configuration method.</p> <p>Internet Explorer's automatic detection feature requires that the names database on the DNS server contain a <code>wpad.domain.com</code> entry that resolves to the appliance's forward proxy address. On Windows 98 clients, you must add the same DNS entry to the HOSTS file in the WINDOWS directory. If the HOSTS file does not exist, you can either copy HOSTS.SAM to HOSTS or create a new HOSTS file.</p> <p>Excelerator automatically creates a WPADxx.DAT file for each IP address configured for client acceleration. You can customize each address' file using FTP, or you can provide a customized file named WPAD.DAT that Excelerator will use for all WPAD requests regardless of the IP address on which the request was received.</p> <p>For more information on appliance WPAD capabilities, see "Customizing Web Proxy Auto-Discovery" on page 285.</p>

VII

Browser-Based Tool Help

The following table summarizes the tasks covered in this section.

To	See
Launch and use the browser-based management tool	Appendix 42, “Using the Browser-Based Management Tool,” on page 293
Use the Home panel	Appendix 43, “Using the Home Panel,” on page 295
Use the System panel	Appendix 44, “Using the System Panel,” on page 301
Use the Network panel	Appendix 45, “Using the Network Panel,” on page 321
Use the Cache panel	Appendix 46, “Using the Cache Panel,” on page 333
Use the Hierarchy panel	Appendix 47, “Using the Hierarchy Panel,” on page 405
Use the Monitoring panel	Appendix 48, “Using the Monitoring Panel,” on page 413

42

Using the Browser-Based Management Tool

Use the information in this section to explore, understand, and use the browser-based management tool.

Prerequisites for Running the Management Tool

You need the following:

- ♦ An appliance that has been initialized and is currently running
- ♦ A Java-enabled browser, such as Netscape* Navigator* 4.07 (or higher), Netscape Communicator* 4.5 (or higher), or Internet Explorer 4.01 (or higher) running on your workstation
- ♦ SSL 2.0 and SSL 3.0 (where available) enabled on the browser
- ♦ A network or cross-over cable connection to the appliance
- ♦ The IP address of the appliance

After the appliance has been configured with an IP address and mask, a gateway server, and a DNS server, you can administer the appliance over the network via any client that can communicate with it over IP.

Until you have completed that configuration, however, you must use a crossover cable and a client with the following constraints:

- ♦ Client IP address set to 10.1.1.2 (or another available 10-net IP address) with a mask of 255.255.255.255
- ♦ Client Gateway address set to 10.1.1.1 (the management address of the appliance)
- ♦ Client DNS server address set to 10.1.1.1

Starting the Management Tool

- 1** Start the browser on your client workstation.
- 2** Point the browser to the URL of the appliance you want to manage.

The URL must contain either the 10-net management address or an IP address you have already configured on the appliance, followed by :1959/appliance/config.html, for example:

`http://10.1.1.1:1959/appliance/config.html`

- 3** Accept the SSL certificate.

IMPORTANT: You must have SSL 2.0 and SSL 3.0 (where available) enabled in your browser. Otherwise, the browser will display an error indicating that the page cannot be displayed.

- 4** Enter a password if you have previously specified one for the appliance.

The Apply and Cancel Buttons

As you make changes to appliance parameters in the management tool, these changes are tracked and accumulated in a buffer until you either apply or cancel them. You can make changes in multiple tabs and wait to apply them all at once.

This does not apply to the Actions and Date/Time tabs. Changes in these tabs are effective immediately. If you change the NTP server, the appliance time will change with the next synchronization cycle (normally about 15 minutes).

Except in the cases just mentioned, clicking Apply commits all changes made in any tab since the last time you started the appliance or clicked Cancel. Clicking Cancel cancels all changes made since the last time you started the appliance or clicked Apply. Clicking Cancel is also a quick way of requesting that the appliance reread the currently displayed settings.

After you click Apply or Cancel, the action cannot be undone.

The Help Button

You can click the Help button in the left frame to display the online documentation. A table of contents appears in the left frame. To navigate through the documentation, you can click the titles in the table of contents.

Encryption

If you have specified passwords for appliance management purposes, communications regarding the password are transmitted through HTTPS. All other communications with the appliance are not normally encrypted.

43

Using the Home Panel

The Home panel provides access to general information regarding the appliance, such as the caching system version currently running and the general health of the current configuration.

Introduction Tab

Path: Home > Introduction

Figure 68

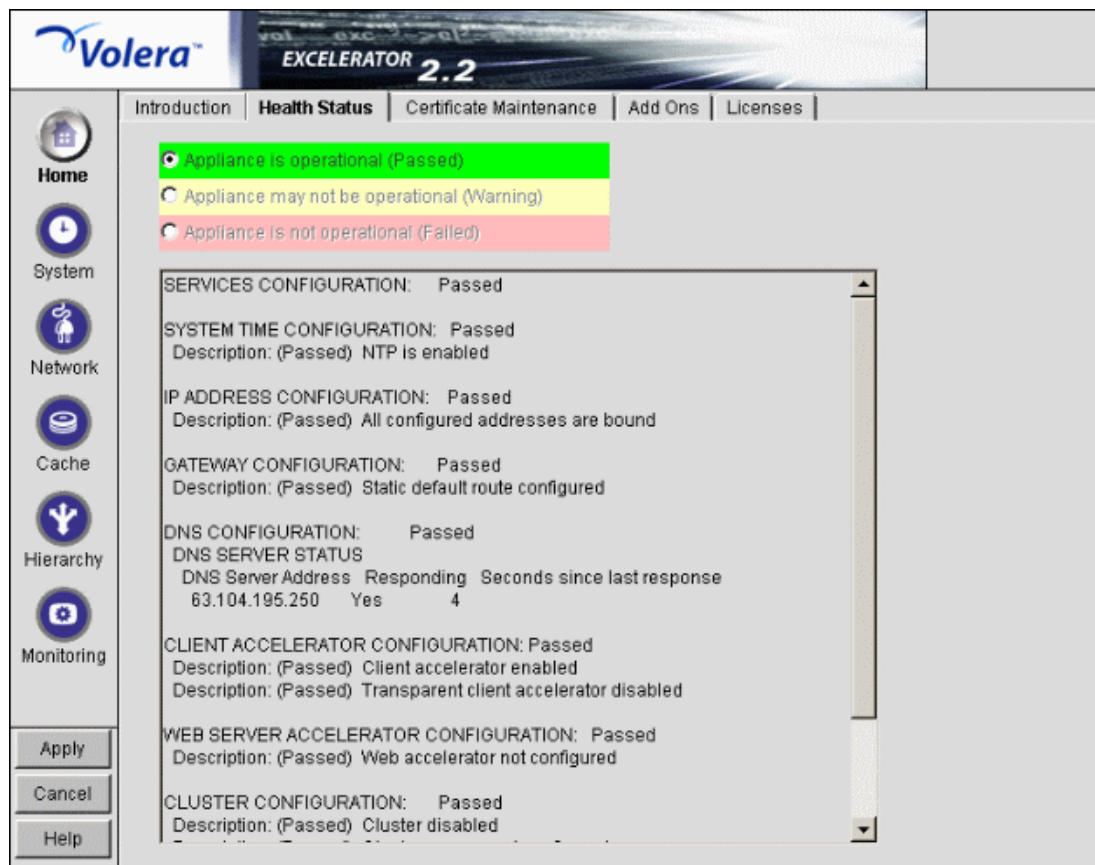


The Introduction tab displays appliance information, such as version number.

Health Status Tab

Path: Home > Health Status

Figure 69



The Health Status tab indicates general status of appliance configurations, including which services are currently configured and the operational status of selected services.

A green status indicates Excelerator has not detected any configuration discrepancies.

A yellow status indicates Excelerator might be functioning sub-optimally due to configuration discrepancies.

A red status indicates that the Excelerator configuration might be incomplete or wrong.

Services Configuration: Reports the overall configuration status of all proxy services.

System Time Configuration: Reports the current NTP status.

IP Address Configuration: Reports the status of IP address assignments to network interfaces. Excelerator requires at least one IP address assignment for proper operation.

Gateway Configuration: Reports the status of the next hop gateway configuration. Without proper gateway configuration, the appliance cannot connect to origin Web servers.

DNS Configuration: Reports the status of DNS server configuration and connectivity to configured DNS servers. Without access to a DNS server, proxy services cannot function properly.

Client Accelerator Configuration: Reports the status of the forward and transparent proxy configuration. If browser clients pointing to this appliance as their proxy server have Web browsing problems, check the status here and in [“Services Tab” on page 414](#).

Web Accelerator Configuration: Reports status of the Web server accelerator configurations. If browser clients have problems accessing a site being accelerated by this appliance, check the status of the Web server accelerator service in this section and in [“Services Tab” on page 414](#).

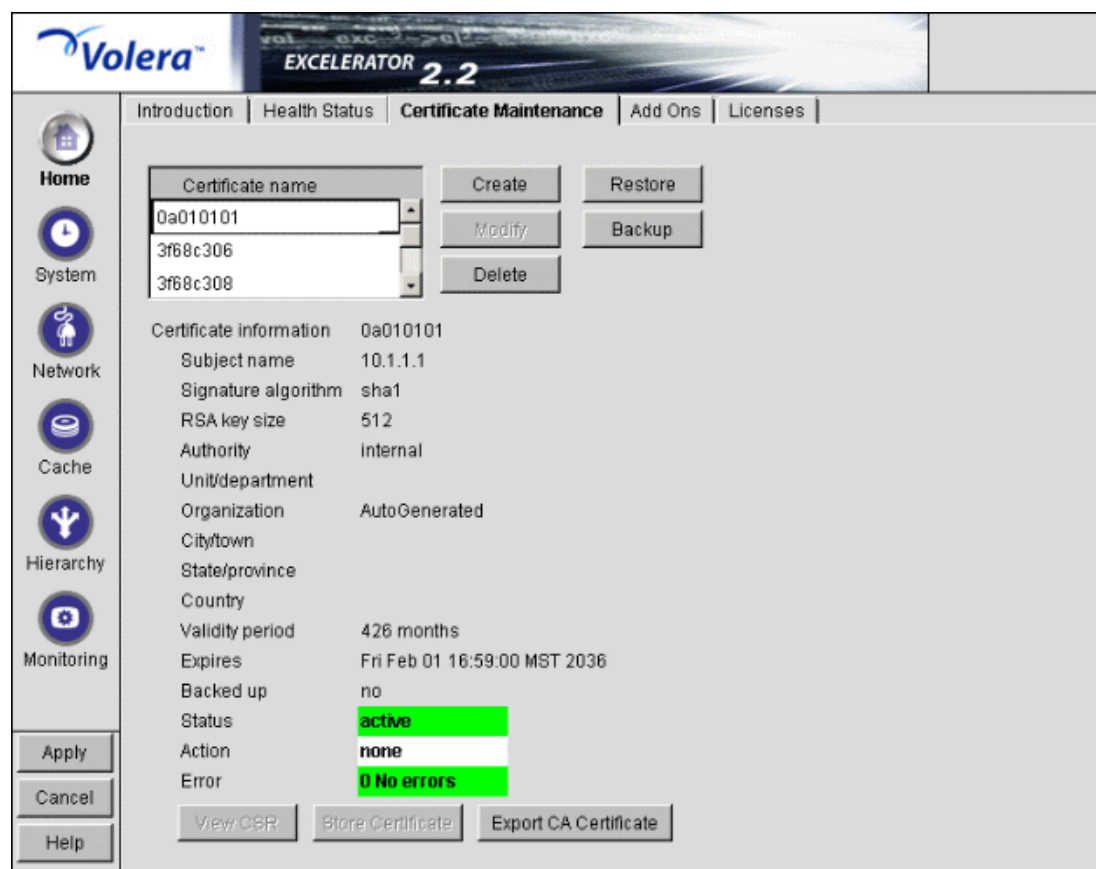
Cluster Configuration: Reports the status of cluster servers and services. Both cluster servers and clustered services need to be configured for achieving fault tolerance through clustering.

Filtering Configuration: Reports the status of filter service configurations. If filtering is enabled and waiting for a rating list to be downloaded, the status indicates that filtering is not active and the health status is yellow (warning).

Certificate Maintenance Tab

Path: Home > Certificate Maintenance

Figure 70



The Certificate Maintenance tab lets you create, delete, back up, restore, and view authentication certificates stored on the appliance. This includes internal certificates generated by the appliance's certificate authority (CA) and external certificates generated by an external CA, such as VeriSign*. For more information, see [“Managing Appliance Certificates” on page 165](#).

Certificate Name: A list of certificates created on the appliance.

Certificate Information: Information for the selected certificate.

View CSR: Use this option to display the certificate signing request of a certificate you have created.

This option is also used to request a certificate from a certificate authority. For more information, see [“Obtaining a Certificate from an External CA” on page 166](#).

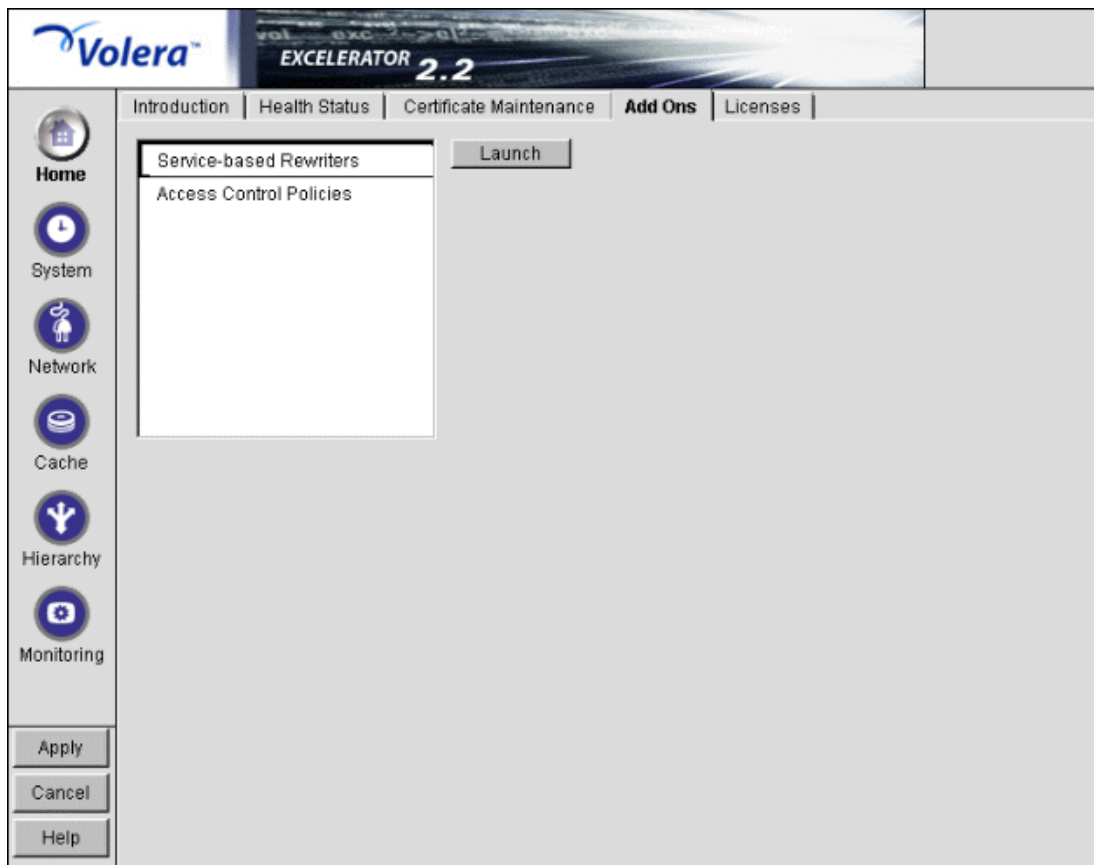
Store Certificate: Use this option to store certificate information received from a certificate authority. For more information, see [“Obtaining a Certificate from an External CA” on page 166](#).

Export CA Certificate: Use this option to display a certificate authority’s certificate. For more information, see [“Viewing \(Exporting\) a Certificate's CA” on page 168](#).

Add Ons Tab

Path: Home > Add Ons

Figure 71

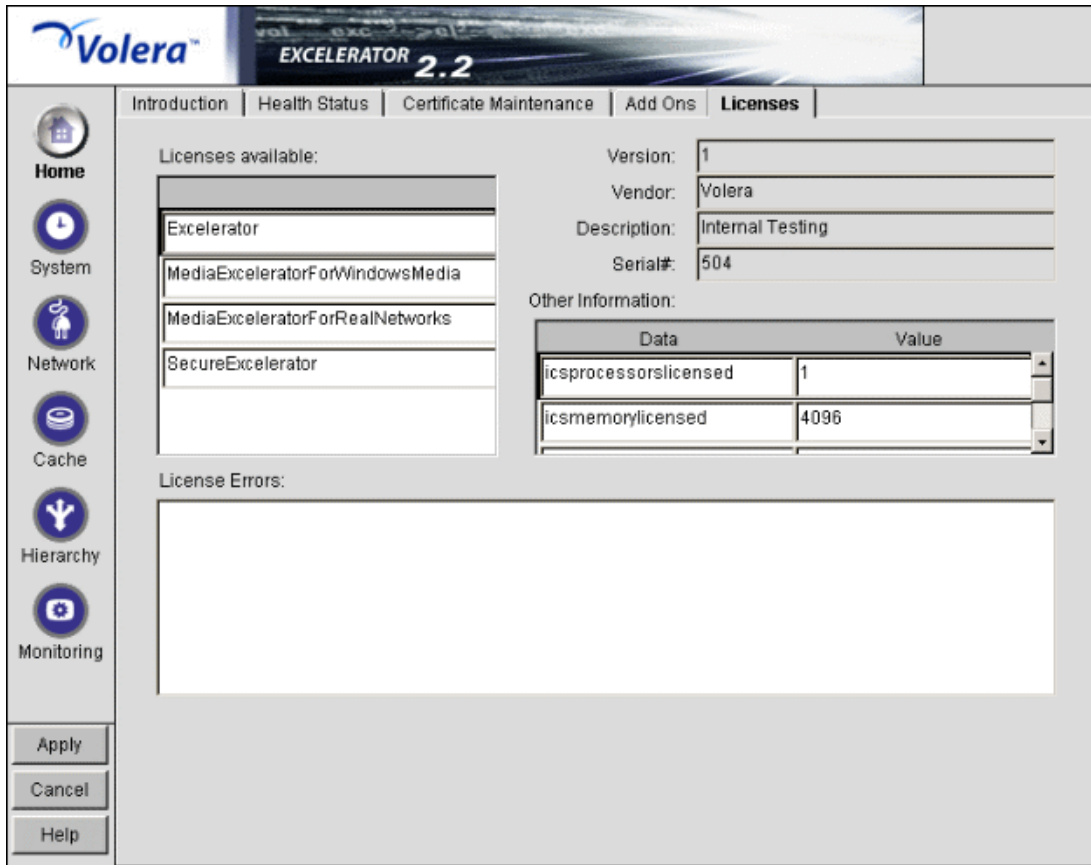


The Add Ons tab lets you launch administration utilities for third-party software components you’ve installed on your Excelsator appliance. To access the target utility, click the component name > Launch.

Licensing Tab

Path: Home > Licensing

Figure 72



The Licensing tab lets you view licensing information for licensed components installed on the appliance. Excelerator add-on products require valid licenses to run. In some cases, the browser-based tool requires a valid license before it displays add-on controls, panels, and so on.

You install add-on products and their accompanying licenses using the [Upgrade Tab](#).

IMPORTANT: If you are upgrading to from a previous version of Excelerator, you must remove all previously installed product license(s) and install new licenses before you can use the product. For further information on license management, see [Chapter 21, "Installing and Upgrading Licenses,"](#) on page 133.

Licenses Available: A list of the licensed components installed on the appliance.

Version: The version of the selected component.

Vendor: The vendor of the selected component.

Description: A brief description of the selected component.

Serial #: The serial number of the selected component.

Other Information: Details of the selected component license.

License Errors: Description of any current problems with the license. If a component is not running, this box might indicate what the problem is.

44

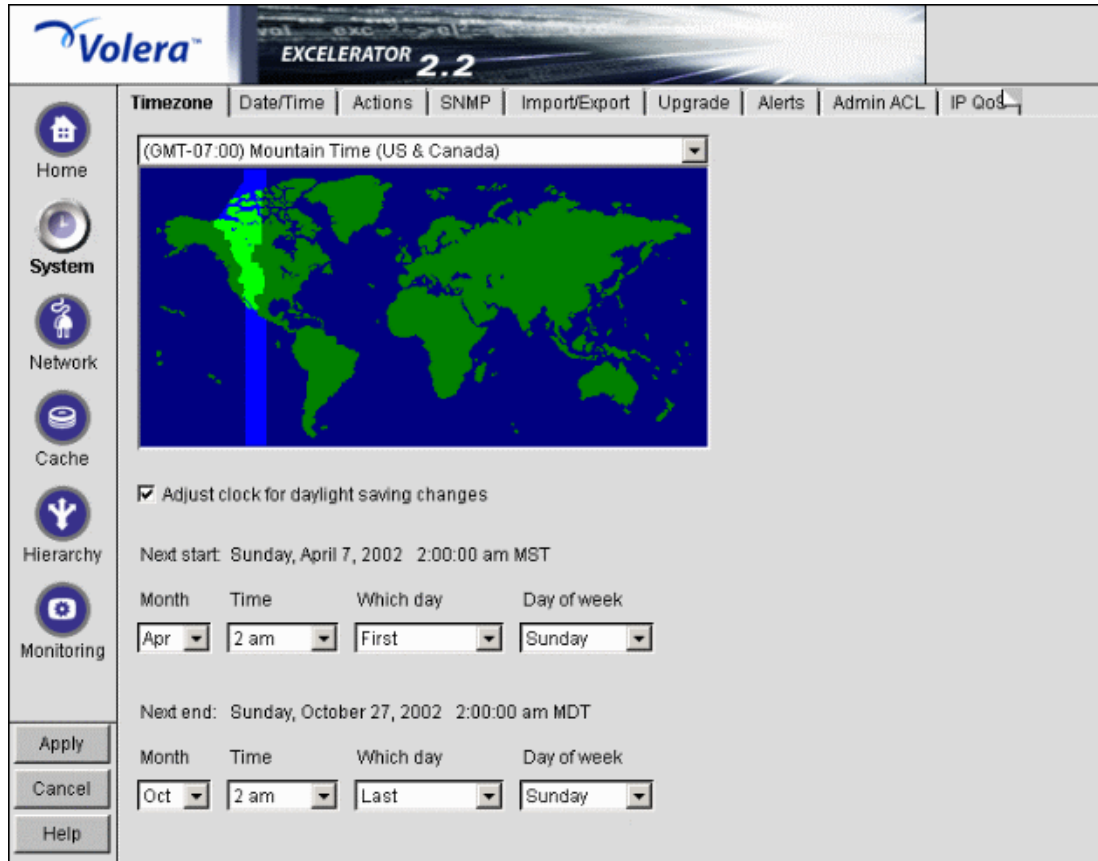
Using the System Panel

The System panel lets you perform actions that affect the appliance system in a general way. Use the tabs in this panel for changing and setting system time, changing the system password, restarting the appliance, upgrading the system, etc.

Timezone Tab

Path: System > Timezone

Figure 73



The Timezone tab lets you specify a time zone for the appliance. It also lets you specify exactly when daylight saving time begins and ends.

The Time Zone Map: Lets you select a time zone for the appliance by clicking the map. The granularity offered through this method is adequate for most appliance installations. Additional flexibility in setting time is available on this tab and from the command line.

For more information on command line options, refer to the command line help for the set command and the time zone argument. See the instructions for using command line online help in [Appendix A, “Command Line Reference,”](#) on page 431.

Adjust Clock for Daylight Saving Changes: If you check this option, the appliance clock begins daylight saving time and resumes standard time on the dates and times defined in the fields below Next Start and Next End. For example, most U.S. time zones begin daylight saving on the first Sunday of April at 2:00 a.m. and resume standard time on the last Sunday of October at 2:00 a.m.

To set nonstandard daylight saving parameters in this tab, select the start and end field values for Month, Time, Which Day, and Day of Week in their respective drop-down lists.

To set nonstandard parameters from the command line, refer to command line help for the set command and the dsstart, dsend, and dstime arguments. See the instructions for using command line online help in [Appendix A, “Command Line Reference,”](#) on page 431.

Date/Time Tab

Path: System > Date/Time

Figure 74

The screenshot shows the Volera EXCELERATOR 2.2 web interface. The top navigation bar includes tabs for Timezone, Date/Time (selected), Actions, SNMP, Import/Export, Upgrade, Alerts, Admin ACL, and IP QoS. A left sidebar contains icons for Home, System (selected), Network, Cache, Hierarchy, and Monitoring. The main content area is titled 'Date/Time' and features two radio buttons: 'Use network time protocol' (selected) and 'Set date/time manually'. Under 'Use network time protocol', there is a section for 'NTP servers' with two input fields containing '63.192.96.3' and '64.243.118.2', and buttons for 'Insert' and 'Delete'. A 'Set Date/Time' button is also present. Below this, a clock face shows the current time as approximately 10:17. To the right of the clock, the text 'Current date/time on appliance' is followed by the values: Year: 2002, Month: 2, Day: 6, and Time: 10:17:14. At the bottom left, there are buttons for 'Apply', 'Cancel', and 'Help'.

The Date/Time tab lets you set the appliance system time so that the time stamps in cache logs are accurate and valid. An ISP, for example, might bill customers based on their access to the appliance. Accurate log time stamps are essential to issuing credible billing statements.

NOTE: Excelerator stamps log entries with Greenwich Mean Time (GMT). If the appliance is using an NTP server, the GMT stamp comes from that server. If the appliance is using a manually set time, Excelerator assumes the time is accurate and calculates the GMT value based on the appliance's time zone and daylight saving settings.

Use Network Time Protocol: Checking this option turns network time protocol on or off. This enables the appliance to synchronize its system time with an NTP server. Using an NTP server makes appliance cache log time stamps as reliable as possible. This can be especially important if you use the logs for customer billing. The appliance comes with two sample NTP servers: 132.163.4.101 and 132.163.4.103. You can remove these or add other NTP servers.

IMPORTANT: When you specify an NTP server, synchronization between the NTP server clock and the appliance clock might not be immediate.

If the NTP server clock has an earlier time than the appliance clock, Excelerator will slow down the appliance clock until the two are synchronized. This provides for proper incrementation of log files and other time-sensitive information during the synchronization process.


If the NTP server clock has a later time setting than the appliance clock, synchronization between the two will generally be immediate. However, in certain situations you might observe the appliance clock incrementing by 600-minute intervals. This is normal system behavior.

The Apply button changes from Wait back to Apply. This indicates only that the NTP configuration change has been made, not that the appliance clock is fully synchronized with the NTP server.

If necessary, you can set appliance time manually to the target time and then re-enable the NTP feature.

Set Date/Time Manually: The dialog box in [Figure 75](#) appears when you select this button and click Set Time. Set the date and time using the drop-down lists. Clicking OK immediately resets the system clock.

Figure 75

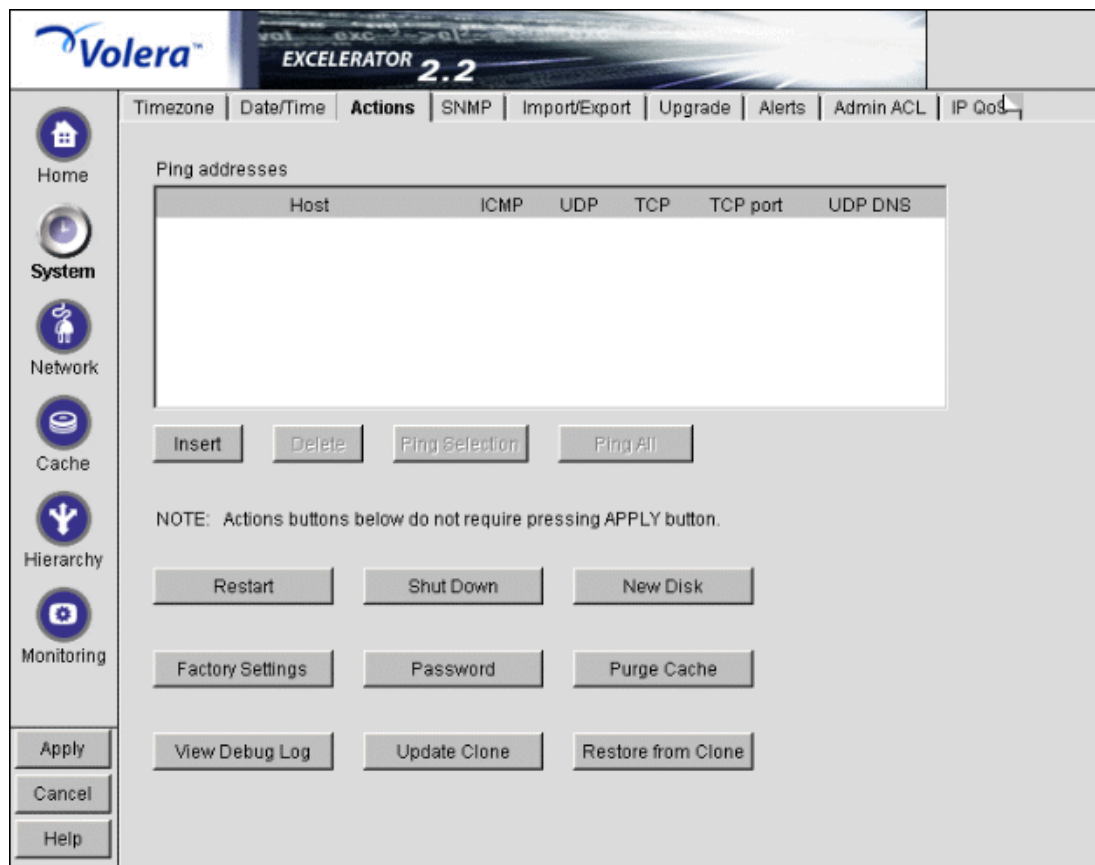


Use this option if NTP is not available to your appliance or if you need to set a specific time.

Actions Tab

Path: System > Actions

Figure 76



The Actions tab lets you perform tasks related to the appliance hardware and software.

NOTE: Most changes made in the browser-based management tool are not effective until you click Apply. Changes made in the Actions tab, however, are effective immediately.

Ping Addresses: You can check network connections using appliance ping functions by adding target hosts and port numbers to this list and then clicking Insert. Follow the address with a colon and a port number (an integer value from 0 to 65535) you want to ping. Using a port number lets you check whether a host has HTTP support (port 80), HTTP forward proxy support (port 8080), DNS support (port 53), ICP peer/parent support (port 3130), and so on.

Restart: Shuts down and then restarts the caching system. Configuration settings are retained, but cached objects are removed.

Shut Down: Shuts down the caching system. The hardware remains turned on until manually powered off.

When the appliance has been successfully shut down, a series of three beeps is repeated until the box is powered off.

New Disk: Scans for new disks that the system has not detected automatically.

Factory Settings: Resets the appliance to its original factory configuration, as explained in [“Restoring Factory Settings” on page 199](#). Passwords are retained. If you want to preserve other settings for later use on this or another appliance, see [“Import/Export Tab” on page 308](#).

Password: See [“Change Password Dialog Box” on page 305](#).

Purge Cache: See [“Purge Cache Dialog Box” on page 306](#).

View Debug Log: When an appliance experiences an abnormal shutdown due to a configuration error or other problem, Excelerator logs critical history information associated with the shutdown. Clicking this button displays the log in a separate browser window. You can then save the log file locally, print it, and e-mail it to technical support.

Update Clones: Each appliance stores a clone image that, initially, is the same as the factory image. If the appliance experiences an abnormal shutdown four times within a half-hour period, or if the appliance is restarted six times within a half-hour period, Excelerator assumes the appliance’s configuration has become unstable and automatically replaces it with the clone image.

You can overwrite the default clone image with the current configuration by selecting this option.

IMPORTANT: This process reboots the appliance, causing a temporary interruption of services.

Restore from Clones: Selecting this option restores the appliance to the configuration of the clone image (either the original factory clone image or an alternate clone image you have saved using the Update Clones option).

IMPORTANT: This process reboots the appliance, causing a temporary interruption of services. If the image being restored is the original factory clone image, you will also need to reconfigure proxy services on the appliance or use a .NAS file to restore the services. See [“Restoring the Appliance to the Clone Image” on page 200](#) and [“The AUTOLOAD.NAS File” on page 196](#).

Change Password Dialog Box

Path: System > Actions > Password

Figure 77



IMPORTANT: It is critical that you assign system passwords when initially configuring the appliance. Otherwise, access through Telnet, FTP, and the browser-based management tool is not restricted.

You can specify passwords for two users with different access privileges.

Users logging in using the View user password can view everything in the browser-based management tool

View users cannot access the command line. They can, however execute FTP get commands.

Users logging in using the Config user password have full access to the browser-based tool and the command line interface.

Change: Immediately changes the password for the selected user.

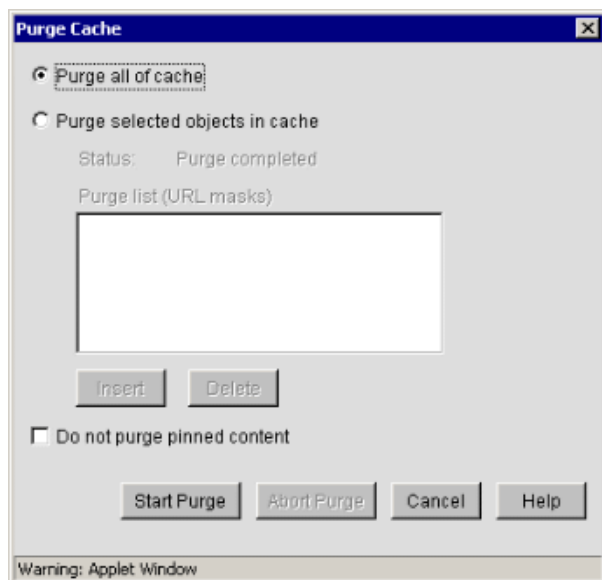
Remove: Removes (sets to null) the password for the selected user.

All ASCII-printable characters can be used in passwords, except the comma (,) and spaces at the beginning or end.

Purge Cache Dialog Box

Path: System > Actions > Purge Cache

Figure 78



You can remove all cached objects from the appliance's cache, or you can perform a limited purging of cached objects based on URL masks. Purging cannot be undone.

Purge All of Cache: This option allows you to purge everything from the appliance's cache.

Purge Selected Objects in Cache: This option allows you to specify URL patterns or masks for the pages or sites whose objects you want to purge. When defining the masks, keep in mind that the appliance interprets everything in the URL mask between the asterisk wildcard (*) and the following delimiter as a wildcard. Delimiters include the forward slash (/), the period (.), and the colon (:) characters.

This option also allows you to purge cache objects whose URL contains a specified query string or cookie. This mask is defined by placing a question mark (?) at the start of the mask, followed by text strings and wildcards as necessary. String comparisons are not case sensitive. For example, ?*=SPORTS will purge all objects with the text "=SPORTS" in any combination of upper- and lower-case letters for "=SPORTS" following the question mark in the URL.

Do Not Purge Pinned Content: You can use this option to purge everything from the appliance's cache except for the content which you have pinned.

SNMP Tab

Path: System > SNMP

Figure 79

The screenshot shows the Volera EXCELERATOR 2.2 web interface. The top navigation bar includes tabs for Timezone, Date/Time, Actions, **SNMP**, Import/Export, Upgrade, Alerts, Admin ACL, and IP QoS. A left sidebar contains icons for Home, System, Network, Cache, Hierarchy, and Monitoring. The main content area is titled 'SNMP' and contains the following sections:

- Monitor state:** Radio buttons for 'No community may read' and 'Specified community may read' (selected). A text field next to the selected option contains 'public'.
- Control state:** Radio buttons for 'No community may write' and 'Specified community may write'.
- Trap state:** Radio buttons for 'Do not send traps' (selected) and 'Trap community name'.
- IP addresses of management stations:** A text field.
- Node name for SNMP:** A text field containing 'Excelerator'.

Below these sections are three buttons: 'Hardware Description', 'Physical Location', and 'Human Contact'. At the bottom left are 'Apply', 'Cancel', and 'Help' buttons. A red note at the bottom center states: 'NOTE: The appliance will be restarted when APPLY button is pressed in order for the changes to take effect.'

The SNMP tab lets you configure the appliance with basic SNMP information so the appliance can communicate with your SNMP management workstations.

The appliance's SNMP implementation follows the ISO SNMP version 2 standard with regard to counters, 64-bit statistics, commands, configuration information, etc.

When SNMP-enabled appliance components start up, they register with the system. When the system receives a request for a specific SNMP parameter, it knows which component to contact to obtain the information.

If you purchased your appliance from a third-party vendor, you should know that each vendor has the option of customizing the SNMP capabilities on its appliance offerings, including specifying which components register with the system and which system traps are reported. If you have questions about the SNMP implementation on your appliance, contact the vendor directly.

Each appliance contains an ICS.MIB file in the SYS:\ETC\PROXY\DATA directory. To see a list of standard SNMP parameters, retrieve this file using the FTP get command and compile it for use with your SNMP management software.

If you specify a trap community name and specify an SNMP management workstation in the SNMP tab, all alerts you check in the Alerts tab (see [“Alerts Tab” on page 312](#)) are automatically sent as SNMP traps even if you have not configured syslog or e-mail alert notification on the Alerts tab.

Monitor State: Allows you to specify community Read access and the community name or password to be used. Community names must contain ASCII characters only and must not have

spaces. The default Monitor State community name of *public* allows anyone access to system configuration information. If this is a concern, change the name to *no community may read*.

Control State: Allows you to specify community Write access and the community name or password to be used. Community names must contain ASCII characters only and must not have spaces.

IMPORTANT: The default name or password for the control community is No, meaning that control access is turned off. You can reset this value. However, this is not normally recommended because the control community password is stored as clear text and could allow unauthorized write access to SNMP parameters on the appliance.

Trap State: Allows you to specify that traps are not sent or to specify a community (location, IP octets, or other identifier) from which traps are sent to the management stations you designate. Community names must contain ASCII characters only and must not have spaces.

IP Addresses of Management Stations: Allows you to specify one or more management station IP addresses, separated by semicolons.

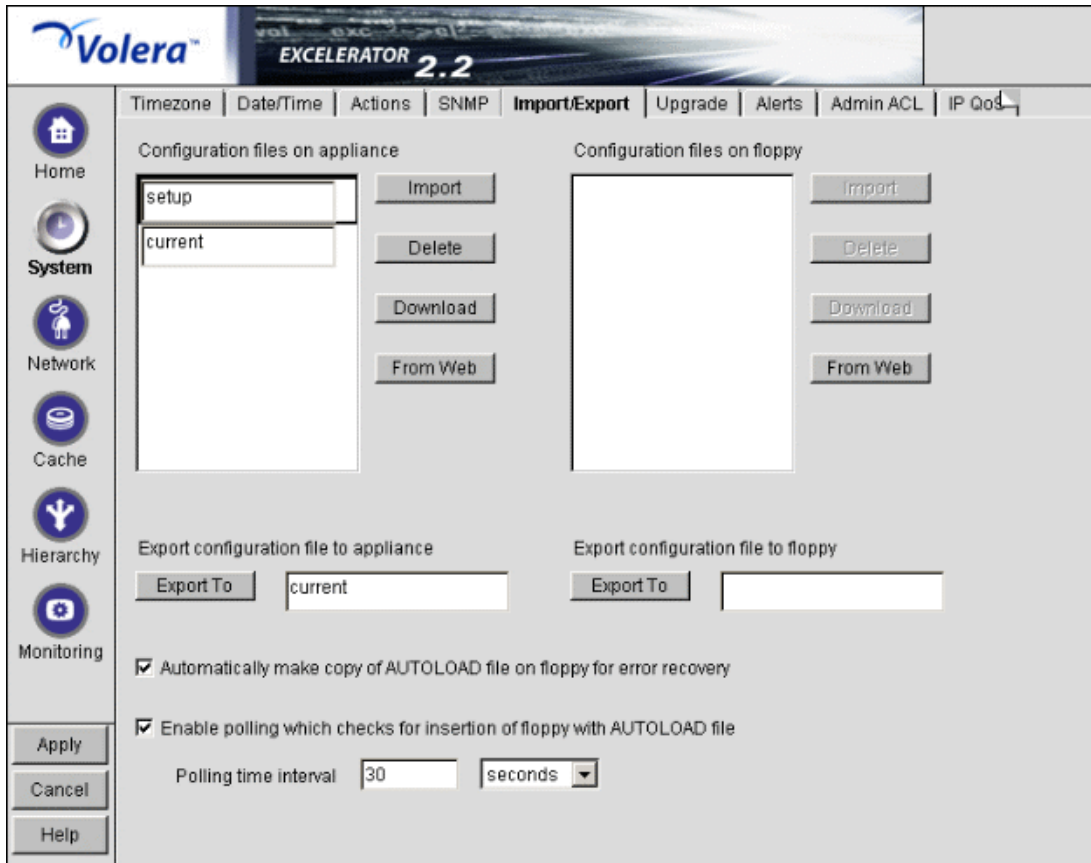
Node Name for SNMP: Allows you specify a node name for management of the appliance through SNMP.

The buttons below the Node Name for SNMP field let you enter additional information regarding the hardware, the appliance's physical location, and information regarding the person responsible for the appliance.

Import/Export Tab

Path: System > Import/Export

Figure 80



The Import/Export tab lets you manage appliance configuration files on the appliance and on floppy diskettes.

IMPORTANT: You should have a backup configuration file named something other than AUTOLOAD.NAS. For further details, see [“Backing Up the Appliance Configuration” on page 198](#).

Configuration Files on Appliance: Displays a list of all the configuration files stored on the appliance. You can use these files to configure the appliance instantly, rather than using the GUI, command line, or Telnet to make individual changes. The appliance automatically updates the configuration file, CURRENT, each time you apply a change to Excelsator. The .NAS extension of these files, which is not shown in this list, is supplied by the server.

You can download, import, and delete any file in this list. You can also copy a configuration file from any URL to the appliance. The Download option opens the file in a separate browser window. The Import option changes the appliance configuration from its current settings to those contained in the selected configuration file. The Delete option removes the selected configuration file from the appliance. The From Web option lets you specify the URL for the configuration file being copied to the appliance. If the file is in a secure area or is being downloaded using SSL (HTTPS:), you can also enter a username and password for authentication.

Configuration Files on Floppy: Displays a list of all the configuration files stored on the floppy disk located in the appliance’s floppy disk drive. You can download, import, and delete any file in this list. You can also copy a configuration file from any URL to a diskette in the appliance’s floppy disk drive. The previous section contains more detail regarding the Import, Delete, Download, and From Web options.

IMPORTANT: It is easy to confuse the diskette in the appliance's floppy disk drive with one located in your configuration workstation. Only the former is accessible through the browser-based management tool.

Export Configuration File to Appliance / Export Configuration File to Floppy: Creates a configuration file on the appliance or on the diskette in the appliance's floppy disk drive.

Files saved using the Export feature contain the complete configuration of the appliance at the time of export. The default filename is CURRENT. You can specify any DOS-style eight-character name. Names are not case-sensitive. Each file has a .NAS extension that is not displayed in the list or specified when the file is created, but is automatically appended by the system.

Automatically Make Copy of AUTOLOAD File on Floppy for Error Recovery: Creates an AUTOLOAD file on a floppy diskette in the appliance when the configuration is changed. The appliance uses the AUTOLOAD file during error recovery to restore the system configuration after a successful restoration of the clone image. The AUTOLOAD file is also used when you invoke the import floppy command from a Telnet or command line session.

You can also export or import configuration files other than AUTOLOAD to a floppy diskette from the command line or Telnet interface.

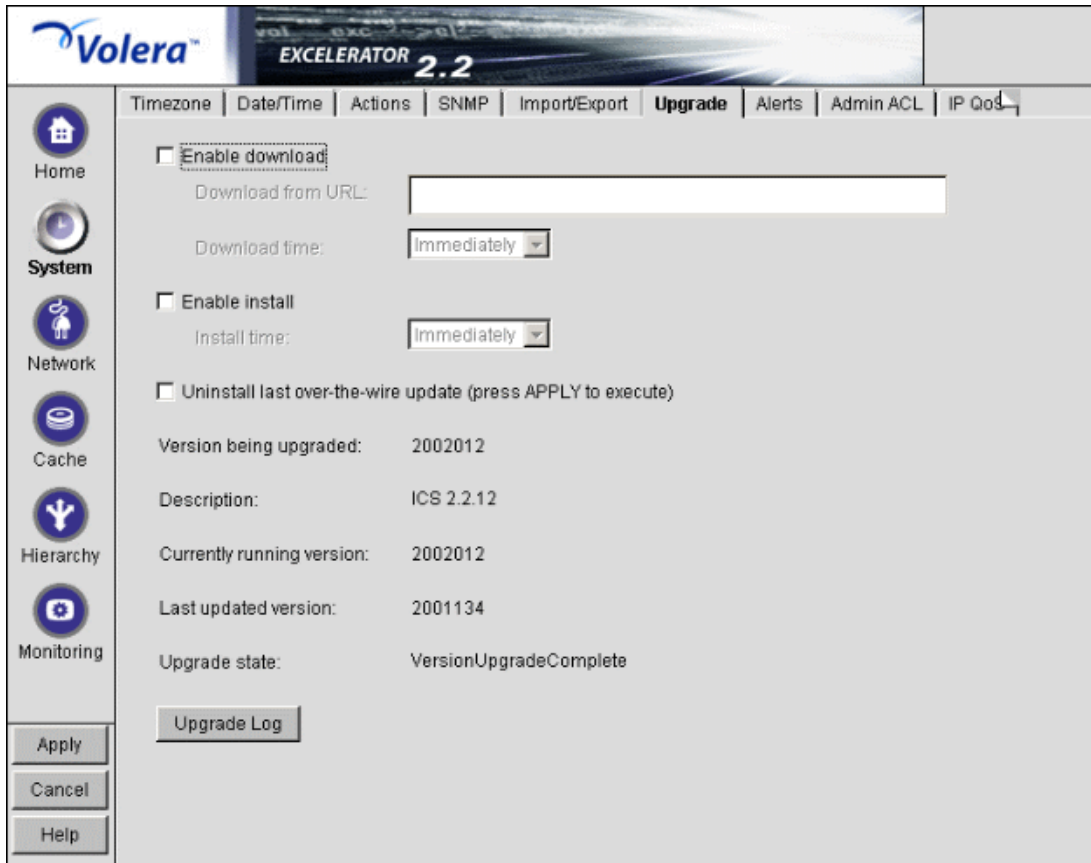
Enable Polling Which Checks for Insertion of Floppy with AUTOLOAD File: Allows the floppy diskette to be polled during normal operation of the appliance. If a floppy diskette containing an AUTOLOAD.NAS is inserted when this option is on, Excelsior automatically applies the settings in the AUTOLOAD file.

While this feature is useful, proxy activity is suspended for a short period of time while Excelsior polls the floppy drive.

Upgrade Tab

Path: System > Upgrade

Figure 81



The Upgrade tab lets you set patch and upgrade parameters so you can download and install patches to the appliance. It also lets you uninstall the most recently applied patch.

Over-the-wire upgrades are secured through signing.

IMPORTANT: Forward proxy authentication must be disabled in order to upgrade the appliance. Otherwise, the upgrade will fail.

NOTE: We recommend you update the appliance's clone image after an upgrade. See ["Restoring the Appliance to the Clone Image"](#) on page 200, ["Actions Tab"](#) on page 303, and [Appendix C, "Upgrading the Appliance,"](#) on page 441.

Enable Download: Lets you set the appliance to download updates automatically. If you check this box and enter the URL for the patch in the Install from URL field, the update is downloaded as scheduled in the Download Time field. A valid entry for Install from URL is any valid URL or DNS name for a Web site.

Enable Install: Lets you set the appliance to install patches automatically. If you check this box, patches downloaded to the appliance are automatically installed as scheduled in the Install Time field.

Uninstall Last Over-the-Wire Update (Press APPLY to Execute): If you check this box after the patch is downloaded and installed into the appliance, the patch is uninstalled or backed out when you click Apply. Only the most recently applied patch can be backed out.

Version Being Upgraded: Displays the version number of the current update. The version of the current update appears in this field the moment the update process begins. You cannot upgrade Excelsator to a lower version than the one currently installed.

Description: A text name associated with the update file.

Currently Running Version: The update version number the appliance is currently running. Before installing the first update, this number is 0.

Last Updated Version: The update version number of the last update applied. For example, if you are currently running update version 3, this number might be 2.

Upgrade State: A state value indicating upgrade status. State values include Not Started, Download Pending, and Version Download Complete. The field is updated each time you click Upgrade.

Upgrade Log: Displays the text messages that have been generated by the upgrade process.

Alerts Tab

Path: System > Alerts

Figure 82

The screenshot shows the Volera Excelsator 2.2 web interface. The top navigation bar includes tabs for Timezone, Date/Time, Actions, SNMP, Import/Export, Upgrade, Alerts (selected), Admin ACL, and IP QoS. A left sidebar contains icons for Home, System, Network, Cache, Hierarchy, and Monitoring. The main content area is titled 'Alerts' and contains the following fields and controls:

- Alert source name:** A text input field.
- Syslog:** A checkbox labeled 'Syslog'.
- Email alert:** A checkbox labeled 'Email alert'.
- Syslog servers:** A list box with 'Insert' and 'Delete' buttons below it.
- Email recipients:** A list box with 'Insert' and 'Delete' buttons below it.
- Email servers:** A list box with 'Insert' and 'Delete' buttons below it.
- Syslog port:** A text input field containing the value '514'.
- Alert types:** A list of checkboxes for various system events:
 - ☐ Disk space shortage
 - ☐ Network receive buffers shortage
 - ☐ Oversized ping packets
 - ☐ SYN packet flooding
 - ☐ Oversized UDP packets
 - ☐ ICP parent down
 - ☐ Socks server down
 - ☐ System up
 - ☐ System down
 - ☐ Login failure
 - ☐ Configuration change

At the bottom left of the main content area are three buttons: 'Apply', 'Cancel', and 'Help'.

The Alerts tab lets you configure the appliance to send notification of generated system alerts to a network server hosting a Syslog service and to a list of e-mail recipients.

Alert Source Name: Used to identify the appliance as the source of an alert. The system inserts this name in the From field of an e-mail alert and in the Syslog alert message.

Syslog: Enables syslog alerts. Alert messages are then sent to one of the syslog servers.

Email Alert: Enables e-mail alerts. Alert messages are then sent to all of the e-mail recipients.

IMPORTANT: For this feature to work, e-mail servers must be able to relay e-mail from the appliance without authentication.

Due to increasing security risks, many e-mail servers have this feature disabled.

If you plan to have the appliance use e-mail alerts, you must either ensure the e-mail server can relay unauthenticated messages, or you must configure the server to accept mail from the appliance without authentication.

Syslog Servers: Lists syslog servers to which the appliance sends alerts. The appliance pings the first server in the list. If the server doesn't respond, the appliance continues pinging other servers in this list until it receives an acknowledgment. It then sends a syslog alert using UDP to the responding server.

Email Recipients: Lists e-mail recipients to whom the appliance sends alert e-mails. The appliance sends e-mails to all addresses in the list.

Email Servers: Lists e-mail servers through which the appliance routes alert e-mails. E-mails are sent to the first e-mail server in the list. If the server doesn't respond, other servers are accessed in turn until the transmission is successful.

Syslog Port: The port the syslog server listens on for syslog alerts. The default port is 514, but it can be changed.

Alert Types: You enable or disable notification of generated alerts to the configured syslog server and e-mail recipients by checking or unchecking an alert type. The appliance generates alerts for the following conditions:

- ♦ *Disk Space Shortage:* The appliance generates this alert when disk space is low on the OS (SYS:) or log (LOG:) volumes.
- ♦ *Network Receive Buffers Shortage:* The appliance generates this alert when the network receive buffers are low.
- ♦ *Oversized Ping Packets:* The appliance generates this alert when TCP/IP receives an oversized (greater than 10K) PING packet.
- ♦ *SYN Packet Flooding:* The appliance generates this alert when TCP/IP detects a SYN packet flooding attack (half-open connections).
- ♦ *Oversized UDP Packets:* The appliance generates this alert when TCP/IP receives an oversized (greater than 16K) UDP packet.
- ♦ *ICP Parent Down:* The appliance generates this alert when an ICP parent changes from an Up to a Down status.
- ♦ *SOCKS Server Down:* The appliance generates this alert each time it cannot communicate with a SOCKS server.
- ♦ *System Up:* The appliance generates this alert each time the appliance starts.
- ♦ *System Down:* The appliance generates this alert each time the appliance is shut down properly or restarted manually.

- ♦ *Login Failure*: The appliance generates this alert each time a login failure occurs through FTP or the browser-based management tool. The alert contains the IP address of the client making the unsuccessful attempt. Unsuccessful Telnet login failures are not detected.
- ♦ *Configuration Change*: The appliance sends this alert each time the appliance's configuration is changed and each time the appliance is initialized or re-initialized.

Admin ACL Tab

Path: System > Admin ACL

Figure 83

The Admin ACL tab lets you regulate access to administrative functions in the browser-based management tool and the command line interface. You can restrict administrative client access and/or limit the appliance IP addresses through which administrative access is allowed.

Allow Administration from All Clients: This option is selected by default and allows access to administrative functions from any IP address.

Allow Administration from Specified Clients: When you select this option, you must also insert at least one IP address from which IP administrative access is allowed. Otherwise, the system will deselect the option to prevent a global lockout. The limit is eight IP addresses. The field accommodates nine addresses, but the ninth address won't work.

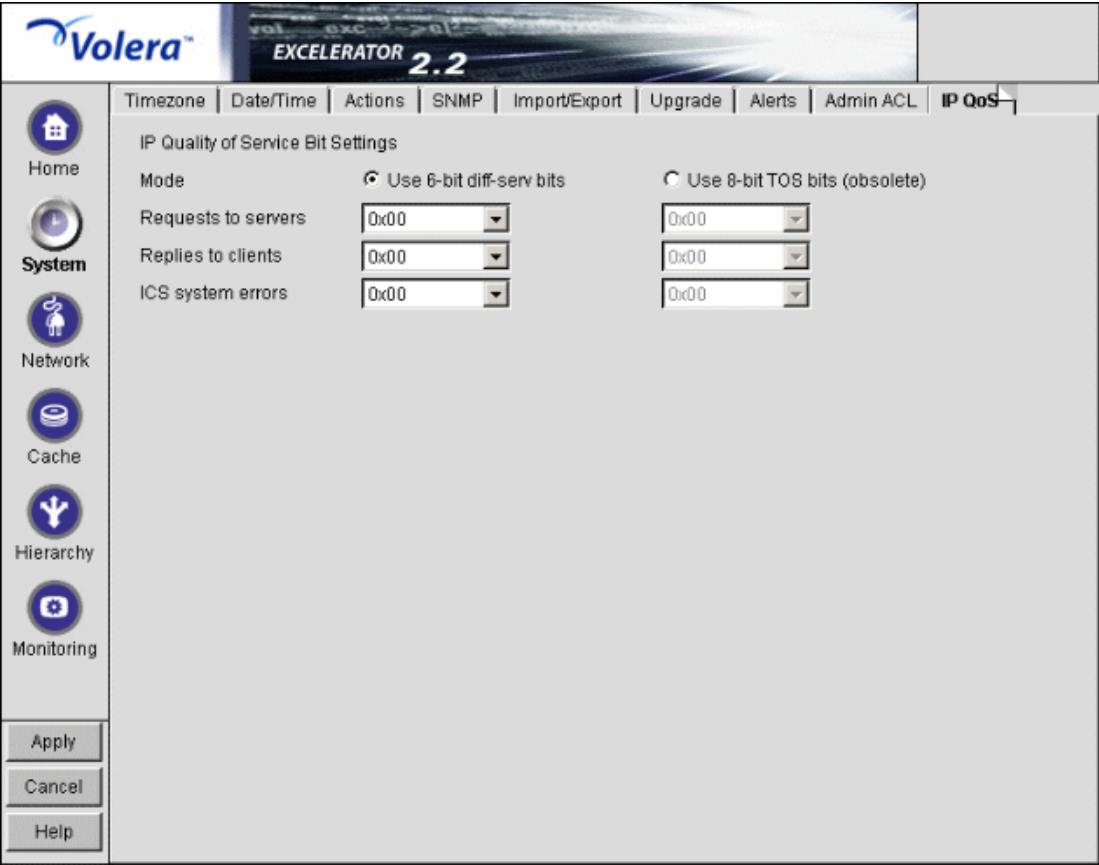
NOTE: If you do not include the IP address from which you are specifying client access, and you click Apply, the address will not be available for future administration sessions unless it is added later.

Allow Administration on Specified Server Addresses: This list contains all appliance IP addresses and indicates which are enabled for administrative access. The first addresses assigned to each network adapter are enabled for administration access by default. You can change administrative access by checking and unchecking addresses in the list. You cannot uncheck all addresses. If you attempt to uncheck all the addresses in the list, the system reverts to the default setting by re-checking all first-assigned addresses.

IP QoS Tab

Path: System > IP QoS

Figure 84



The IP QoS tab lets you specify how Excelsator sets Quality of Service (QoS) parameters in requests to servers, replies to clients, and appliance error pages.

NOTE: Excelsator only sets QoS bits. It doesn't provide QoS services.

The available options are summarized in the following table:

Option	Effect	Comments
pass through	Excelsator doesn't change the QoS bit setting in the packet.	This option is not available for appliance error pages.

Option	Effect	Comments
0X00	Exceleator changes the QoS bit setting in the packet to all zeros—no QoS priority.	The default behavior.
0X01 through 0X3F	Exceleator changes the QoS bit setting in the packet to the QoS priority selected.	
0X40 through 0XFF	Exceleator changes the QoS bit setting in the packet to the QoS priority selected.	These values are available in 8-bit mode only.

Mode: Select either the 6-bit diff-serv (differentiated services) bits or the obsolete 8-bit TOS bits method, depending on the TCP/IP environment in which your appliance is running.

Requests to Servers: Specify how Exceleator handles QoS parameters in requests to servers according to the options summarized in the table above.

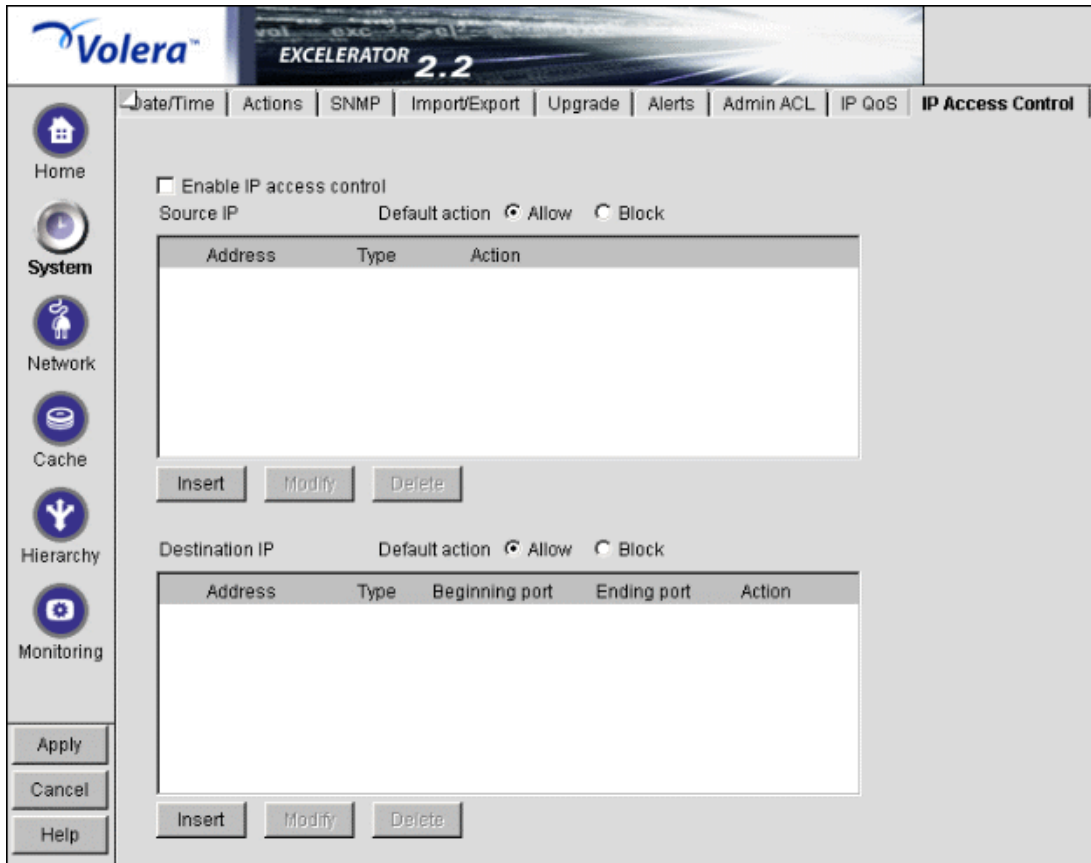
Replies to Clients: Specify how Exceleator handles QoS parameters in replies to clients according to the options summarized in the table above.

ICS System Errors: Specify how Exceleator sets QoS parameters in packets containing appliance alerts and error messages. Because the QoS tag is generated by Exceleator, the pass-through option is not available.

IP Access Control Tab

Path: System > IP Access Control

Figure 85



The IP Access Control tab lets you allow or block browser access to the appliance using the Source IP list. It also lets you allow or block IP addresses from which the appliance will fill its cache using the Destination IP list.

Both the Source IP and Destination IP lists have default actions associated with them. These let you either allow or block all IP addresses. After setting the default action, you create exceptions to the action by adding specific IP addresses, specific IP address subnets, or arbitrary address ranges to the lists.

In most cases you should populate the list with exceptions to the default action. However, you might find it convenient to use a combination of actions within the list itself.

For example, you might need to block all browsers except those on a specific subnet and also block some browsers within the allowed subnet. You could accomplish this by doing the following:

1. Checking the Block default action for the Source IP list.
2. Inserting the allowed subnet in the Source IP list with the Allow action selected.
3. Inserting each of the specific browser IP addresses in the Source IP list with the Block action selected for each entry.

When address ranges overlap and actions specified for affected addresses conflict, the most granular specification takes precedence. For example, if a subnet of addresses is blocked and a specific address within the subnet is allowed, the specific address is allowed. By the same token, if a subnet of addresses is allowed and a specific address within the subnet is blocked, the specific address is blocked.

Enable IP Access Control: Checking this box activates the Source IP and Destination IP restrictions you have specified for the appliance.

Default Action: Allows you to set an action for Source IP addresses or Destination IP addresses. You can then build exceptions to this action by populating the list below each default action setting.

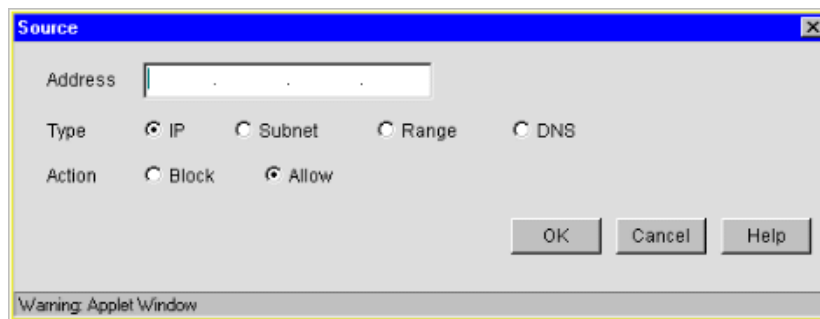
Source IP Settings: Summarizes the exceptions to the default action specified for IP addresses attempting to access the appliance. After IP access control has been enabled, only IP addresses that are explicitly allowed can access the appliance. For more information on how the list works, see the explanation at the start of this section.

Destination IP Settings: Summarizes the exceptions to the default action specified for IP addresses from which the appliance fills its cache. After IP access control has been enabled, the appliance will fill cache only from those addresses explicitly allowed. For more information on how the list works, see the explanation at the start of this section.

Source IP or Subnet Dialog Box

Path: System > IP Access Control > Insert under the Source IP Settings list

Figure 86



The Source IP or Subnet dialog box lets you specify an IP address, a subnet, or an arbitrary range of IP addresses that are either allowed access to the appliance or blocked from such access.

Address: You can enter a valid IP address, a valid subnet address (containing one or more zeros), a DNS name, or an IP address range.

For example, 20.30.40.50 is a valid IP address, and 20.30.0.0 is a valid subnet address that includes all IP addresses from 20.30.1.1 through 20.30.255.255.

To enter a range, you must first select the Range radio button below the Address field to display the To field and then type the first and last IP addresses in the range in the Address and To fields, respectively. You should avoid entering ranges that overlap subnet boundaries.

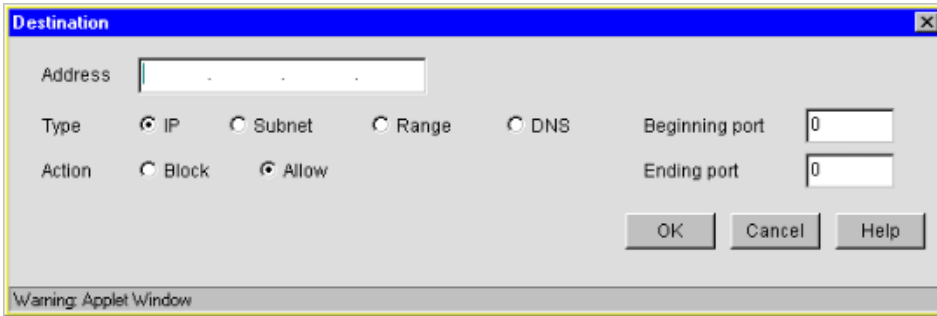
Type: The system selects the correct type for IP and Subnet, depending on the value you type in the Address field. To enter an address range or DNS name, you must first select the appropriate radio button.

Action: The action value you select determines whether Excelerator allows access to or blocks access from the IP address or addresses specified. This action works in combination with the default action, as explained in [“IP Access Control Tab” on page 316](#).

Destination IP or Subnet Dialog Box

Path: System > IP Access Control > Insert under the Destination IP Settings list

Figure 87

The image shows a Java applet window titled "Destination". It contains a text field for "Address" with a placeholder " ". Below the address field are four radio buttons for "Type": "IP" (selected), "Subnet", "Range", and "DNS". To the right of these are two text fields for "Beginning port" and "Ending port", both containing the value "0". Below the type radio buttons are two radio buttons for "Action": "Block" and "Allow" (selected). At the bottom right are three buttons: "OK", "Cancel", and "Help". A status bar at the bottom left of the window says "Warning: Applet Window".

The Destination IP or Subnet dialog box lets you specify an IP address, a subnet, or an arbitrary range of IP addresses that the appliance will either specifically access or specifically not access when filling its cache.

Address: You can enter a valid IP address, a valid subnet address (containing one or more zeros), a DNS name, or an IP address range.

For example, 20.30.40.50 is a valid IP address, and 20.30.0.0 is a valid subnet address that includes all IP addresses from 20.30.1.1 through 20.30.255.255.

To enter a range, you must first select the Range radio button below the Address field to display the To field and then type the first and last IP addresses in the range in the Address and To fields, respectively. You should avoid entering ranges that overlap subnet boundaries.

Type: The system selects the correct type for IP and Subnet, depending on the value you type in the Address field. To enter an address range or DNS name, you must first select the appropriate radio button.

Action: The action value you select determines whether Excelsior allows access to or blocks access from the IP address or addresses specified. This action works in combination with the default action as explained in [“IP Access Control Tab” on page 316](#).

Beginning Port: The first port number of a range of port numbers to be included in the block or allow action.

Ending Port: The last port number of a range of port numbers to be included in the block or allow action.

45

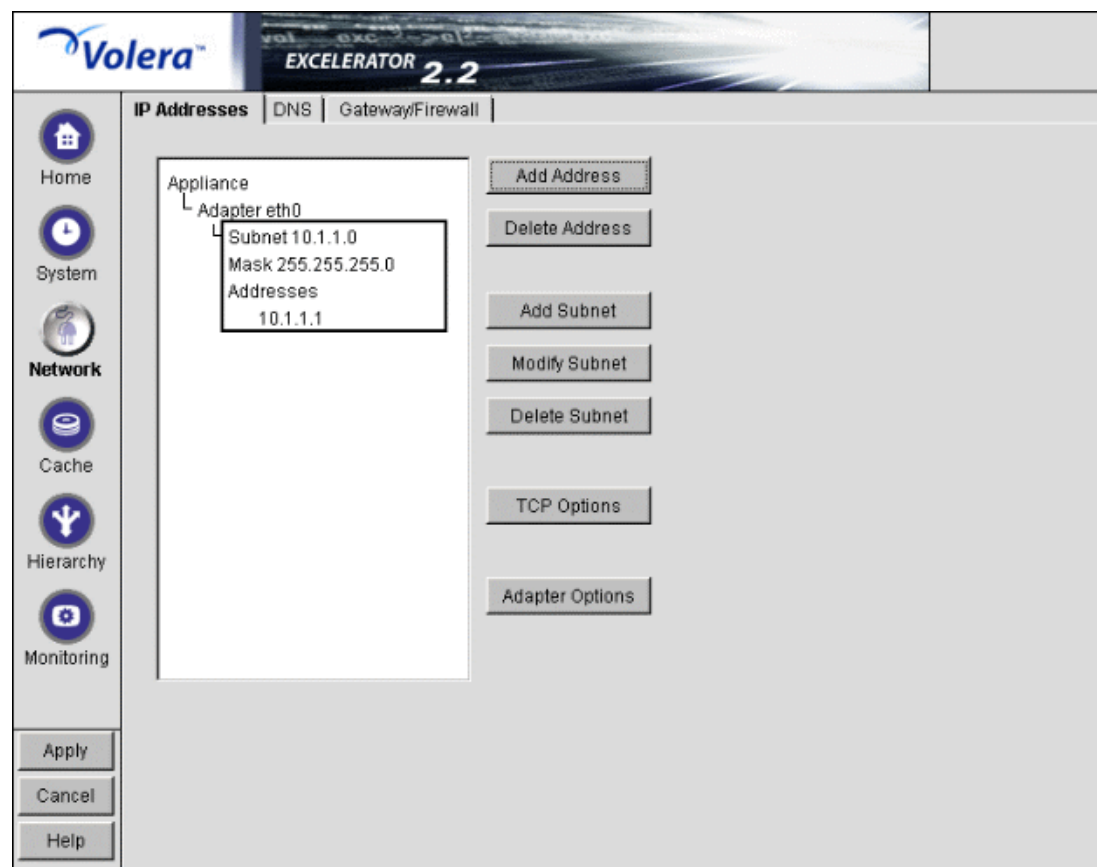
Using the Network Panel

The Network panel lets you configure the appliance to function on the network on which it is installed.

IP Addresses Tab

Path: Network > IP addresses

Figure 88



The IP Addresses tab displays the network adapters, which are the physical connectors into the appliance, and the IP addresses associated with each adapter. The list reflects the current appliance hardware configuration.

Using the buttons to the right of the list, you can associate IP addresses with adapters and change IP address information. Each adapter can have multiple subnets associated with it, and each subnet

can have one or more IP addresses associated with it. You can either define individual IP addresses and masks, or you can add a subnet address and mask and then add multiple IP addresses from that subnet range.

IMPORTANT: If you plan to use an appliance cluster, do not assign the IP addresses for clustered services to network adapters. IP addresses used for clustered services are assigned during the clustered service creation process.

The IP address and the mask define a subnet. You cannot use the first or last address in any given subnet. You cannot create a subnet that collides with another subnet. You cannot create a subnet which spans multiple adapters.

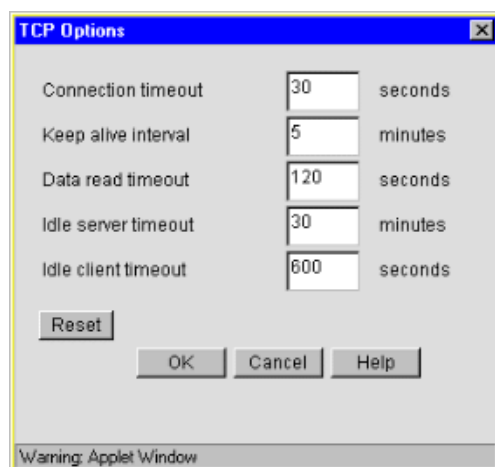
The following are valid appliance subnet masks (representing /1 through /31 in common router notation):

128.0.0.0	192.0.0.0	224.0.0.0	240.0.0.0	248.0.0.0
252.0.0.0	254.0.0.0	255.0.0.0	255.128.0.0	255.192.0.0
255.224.0.0	255.240.0.0	255.248.0.0	255.252.0.0	255.254.0.0
255.255.0.0	255.255.128.0	255.255.192.0	255.255.224.0	255.255.240.0
255.255.248.0	255.255.252.0	255.255.254.0	255.255.255.0	255.255.255.128
255.255.255.192	255.255.255.224	255.255.255.240	255.255.255.248	255.255.255.252
255.255.255.254				

TCP Options Dialog Box

Path: Network > IP Addresses > TCP Options

Figure 89



The parameters displayed in the TCP Options dialog box are standard TCP configuration settings. For more information on adjusting these parameters, see one of the TCP/IP references available at any bookstore carrying computer reference manuals.

Connection Timeout: The number of seconds the proxy server attempts to establish a connection before timing out because the other side has not responded. You might want to increase this value if you notice that the remote server is reachable (the ping succeeds), but the load is heavy.

Keep Alive Interval: The number of minutes a connection is idle before the proxy server queries the other server.

Data Read Timeout: The number of seconds the proxy server waits for expected data to begin arriving before it times out. You might want to increase this value if you notice that the browser receives incomplete data or the connection is disconnected in the middle of data transfer.

Idle Server Timeout: The number of minutes the proxy server keeps the TCP connection between the browser and the proxy server active, even if there is no data flow.

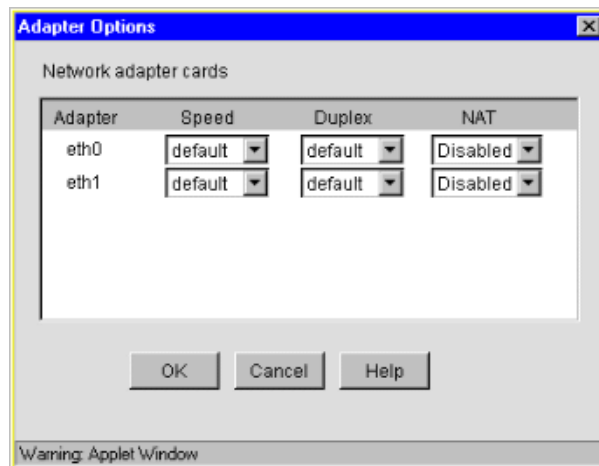
Idle Client Timeout: The number of seconds the proxy server keeps the connection to the origin Web server or another proxy server active, even if there is no data flow.

Reset: Resets the TCP configuration settings to the default values.

Adapter Options Dialog Box

Path: Network > IP Addresses > Adapter Options

Figure 90



The Adapter Options dialog box lets you change settings for the network adapters on the appliance to ensure compatibility with an existing LAN. Modify the default settings only if your LAN requires specialized adapter card changes.

Speed: Options include Default, 10 M, and 100 M.

Duplex: Options include Default, Half, and Full.

IMPORTANT: Some network adapter drivers do not detect duplex settings correctly. This is a general industry problem with Fast Ethernet technology.

If your appliance isn't performing as expected, make sure that the duplex settings for its network adapters match your network configuration. It might be necessary to manually configure the duplex settings on both your appliance and your Ethernet switch or hub.

NAT: Options include Dynamic and Disabled.

If the appliance is serving as a router and your network employs non-unique private IP addresses, you can configure the appliance to provide Network Address Translation (NAT) services.

For example, if you have a 10.0.0.0 private network on eth0 and a registered public network such as 130.0.0.0 on eth1, the clients on the private network can access the Internet through the appliance, if the Dynamic option has been selected in the NAT drop-down list for the eth1 adapter.

The appliance then functions as a network address translator and dynamically maps the private, non-routable 10-net addresses to the registered public address assigned to eth1.

IMPORTANT: You cannot configure a transparent proxy service on an IP address assigned to a card that has the Dynamic option set for NAT. NAT and transparent proxy cannot coexist on the same card.

DNS Tab

Path: Network > DNS

Figure 91

The screenshot shows the Volera EXCELERATOR 2.2 web interface. The top header includes the Volera logo and the product name 'EXCELERATOR 2.2'. Below the header is a navigation bar with tabs: 'IP Addresses', 'DNS' (which is selected), and 'Gateway/Firewall'. On the left side, there is a vertical menu with icons and labels: 'Home', 'System', 'Network' (highlighted), 'Cache', 'Hierarchy', and 'Monitoring'. The main content area of the 'DNS' tab contains the following fields and controls:

- Domain:** A text input field containing 'orem.volera.com'.
- DNS server IP addresses:** A table with three rows. The first row contains '10', '.1', '.1', and '.250'. The second and third rows contain dashes.
- Appliance domain name or alias:** An empty text input field.
- Enable DNS proxy:** A checkbox that is currently unchecked.
- Advanced DNS Options:** A button.
- DHCP server IP addresses:** A table with four rows, each containing dashes.

At the bottom left of the main content area, there are three buttons: 'Apply', 'Cancel', and 'Help'.

The DNS tab lets you configure the domain name service that the appliance will use, including setting a domain name for domain-relative address resolution.

DNS servers are searched in the order listed.

You must specify a domain name for the appliance to use relative domain names.

Domain: Allows you to specify the domain of your appliance. Valid ranges include all valid domain names.

DNS Server IP Addresses: Allows you to specify the IP addresses of the DNS servers you are using. You can enter up to three.

Appliance Domain Name or Alias: (Optional) Allows you to specify a unique domain name or alias for the appliance. This name is used in the Via headers that track packet routes across the network.

Enable DNS Proxy: Allows you to enable a DNS proxy. Versions earlier than 1.3 included an active DNS proxy. Due to a potential security risk through the DNS port, version 1.3 and later have the DNS proxy disabled by default. You can enable the DNS proxy by checking this box.

Advanced DNS Options: See “[Advanced DNS Options Dialog Box](#)” on page 325.

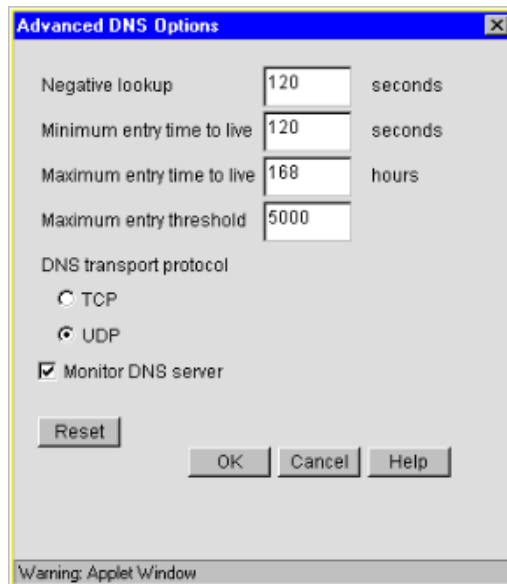
DHCP Server IP Addresses: Allows you to specify a list of DHCP servers to which the appliance will forward client DHCP requests.

This is critical if DHCP clients cannot directly access their designated DHCP servers. The appliance forwards the DHCP requests from the clients to the servers and forwards the replies back to clients. The appliance does not have to be enabled as a router to forward DHCP requests. However, the DHCP Server IP list must be filled in.

Advanced DNS Options Dialog Box

Path: Network > DNS > Advanced Options

Figure 92



The parameters displayed in the DNS Advanced Options dialog box are standard DNS configuration settings. For more information on adjusting these parameters, see one of the TCP/IP references available at any bookstore carrying computer reference manuals.

Negative Lookup: The number of seconds a failed DNS lookup domain name remains in the proxy server cache. If the proxy server cannot resolve a domain name, it stores that information in

its cache for the specified amount of time. If the proxy server receives requests for that domain name within this period, it sends a Bad Gateway error message to the browser and does not resolve the domain name again. Valid field values include 0 – 3600 seconds.

Minimum Entry Time to Live: The minimum amount of time that DNS entries remain in cache before they expire. This is the minimum value the appliance uses regardless of the value returned by the DNS name server. Valid field values include 0 – 3600 seconds.

Maximum Entry Time to Live: The maximum amount of time that DNS entries remain in cache before they expire. This is the maximum value the appliance uses regardless of the value returned by the DNS name server. Valid field values include 0 – 744 hours.

Maximum Entry Threshold: The maximum number of DNS cache entries. When this number is reached, the proxy server deletes old entries to make room for newer ones. The default is 5000. Valid field values include 2000 – 100000.

DNS Transport Protocol: The transport protocol DNS uses on the network where the appliance is installed.

Monitor DNS Server: The appliance normally monitors DNS server availability by pinging the configured servers every minute. This ensures timely handling of DNS requests. You should uncheck this item if the appliance accesses DNS through a connection, such as a dial-up phone line or ISDN connection, that should not be kept continually open. Keep in mind, however, that unchecking the option will cause the DNS configuration on the **Health Status Tab** to fail.

Gateway/Firewall Tab

Path: Network > Gateway/Firewall

Figure 93

The screenshot shows the Volera EXCELERATOR 2.2 web interface. The 'Gateway/Firewall' tab is selected. The left sidebar contains icons for Home, System, Network, Cache, Hierarchy, and Monitoring, with 'Network' currently active. The main content area has the following fields and controls:

- Default gateway IP address:** A text box containing '10 . 1 . 1 . 254' and an 'Additional Gateways' button.
- Enable RIP:** A checkbox that is unchecked.
- Act as router:** A checkbox that is unchecked.
- Enable gateway monitoring:** A checkbox that is checked.
- Enable SOCKS client:** A checkbox that is unchecked.
- Server IP address:** A text box with a placeholder ' . . . '.
- Server port:** A text box containing '1080'.
- SOCKS V4:** A radio button that is selected.
- Username:** A text box.
- SOCKS V5:** A radio button that is unselected.
- No authentication:** A radio button that is unselected.
- User name/password authentication:** A radio button that is unselected.
- Username:** A text box.
- Password:** A text box.
- SOCKS bypass Web server list:** A large empty text area.
- Buttons:** 'Show Routes', 'Reset Learned Routes', 'Insert', and 'Delete'.

At the bottom left of the main area are 'Apply', 'Cancel', and 'Help' buttons.

The Gateway/Firewall tab lets you set up both default gateways as well as additional gateways for specific routing to hosts or networks. It also lets you specify RIP and SOCKS information for firewalls.

To let the appliance function, you must specify a default gateway (router) whether the appliance is originating packets that need to be routed (from proxy requests or scheduled downloads) or is serving as a router for packets that need to be routed externally.

Default Gateway IP Address: You must have at least one gateway defined for the appliance to function. This is the IP address of the gateway or router being used by the appliance.

Additional Gateways: The appliance uses these only if the Act As Router option is checked. See [“Additional Gateways Dialog Box” on page 328](#).

Enable RIP: Allows you to turn on Routing Information Protocol 1. Through this protocol, the appliance is able to learn routes.

The appliance can also work in a network that uses RIP 2, but you must manually add static routes using the [Routes Dialog Box](#).

Show Routes: See [“Routes Dialog Box” on page 330](#).

Reset Learned Routes: Throws away all information acquired through RIP. RIP must be turned on for this to have any effect.

Act as Router: Check this box if the appliance will function as the default gateway for clients on the network. See [“Transparent Proxy as a Default Gateway \(Router\)” on page 41](#) and [“Transparent](#)

Proxy as an Inline Router (Network Gateway)” on page 42. If you check this option, you can specify additional gateways.

Enable Gateway Monitoring: The appliance normally monitors gateway availability by pinging the configured gateways every minute. You should uncheck this item if the appliance accesses its gateways through a connection, such as a dial-up phone line or ISDN connection, that should not be kept continually open. Keep in mind, however, that unchecking the option will cause the gateway configuration on the **Health Status Tab** to fail.

Enable SOCKS Client: SOCKS is a firewall communication protocol. If a firewall prevents the appliance from communicating directly, you can specify information for SOCKS4 or SOCKS5 servers.

Server IP Address: The address of the SOCKS server you want to use.

Server Port: The port number for SOCKS traffic on the network.

SOCKS V4: Enables the SOCKS4 protocol.

Username: Specify a username if the SOCKS4 server requires one for communication.

SOCKS V5: Enables the SOCKS5 protocol. The appliance currently supports only NULL and Username/Password authentications.

No Authentication: If you use SOCKS5 without verification, this box must be checked (where there is no username or password required).

Username/Password Authentication: Enables the entry of a SOCKS5 username and password if your SOCKS server requires authentication.

Username: Enter your SOCKS username.

Password: Enter your SOCKS password.

SOCKS Bypass Web Server List: If the SOCKS client is enabled, all HTTP and FTP server traffic is redirected to the SOCKS firewall. However, requests to origin servers on an intranet within the firewall should not be routed through the SOCKS server. Requests to servers whose IP addresses are inserted into this list will not be sent to the SOCKS server.

Additional Gateways Dialog Box

Path: Network > Gateway/Firewall > Additional Gateways

Figure 94

Additional Gateways

Default gateway

Next hop address	Metric	Type
10.1.1.1	1	Passive

Host gateways

Next hop address	Host address	Metric	Type
------------------	--------------	--------	------

Insert Delete

Network gateways

Next hop address	Subnet base address	Mask	Metric	Type
------------------	---------------------	------	--------	------

Insert Delete

OK Cancel Help

Warning: Applet Window

This dialog box lets you specify additional gateways. The appliance routes requests to specific destinations through these gateways. If a request could be routed through multiple gateways, the appliance chooses the gateway associated with the most restrictive mask (the smallest range of destination addresses). The default gateway is used only when no other routes apply.

IMPORTANT: The appliance uses additional gateways only when the Act As Router option is checked on the Gateway/Firewall tab.

Gateways fall within three basic groups:

- ♦ Host gateways for specific destination addresses
- ♦ Network gateways for destination addresses that fall within specific subnets
- ♦ The default gateway for destination addresses that aren't covered by host or network gateways

The syntax for this gateway is often expressed in router configuration tables as follows:

`0.0.0.0 / 0.0.0.0 / iii.iii.iii.iii`

The variable *i* represents the IP address of the default gateway.

IMPORTANT: If the appliance is acting as a router and you don't specify a default gateway, the appliance routes only those requests whose destination addresses are covered by a host or network gateway. Other requests are not routed.

The appliance uses Metric field values to alter normal (most restrictive) routing. The default field value is 1. A higher number indicates a higher cost associated with the gateway being referenced. This lets you configure the appliance in such a way that more expensive gateways are not used unless the default or less specific gateway is unavailable.

The appliance conveniently determines masking information when you enter the host or network information.

Default Gateway: The default gateway entered on the gateway panel. You can add a metric and specify whether the gateway is active or passive.

- ◆ *Next Hop Address:* The IP address of the gateway.
- ◆ *Metric:* A relative number indicating the bias one wants to add to the normal flow of gateway logic. Entering a number higher than 1 makes this resource more expensive and alters the gateway logic used. Valid numbers include 1 through 16.
- ◆ *Type:* Active or passive. Gateways can be active where they publish their presence or passive where they do not.

Host Gateways: You can define one or more gateways to be used for packets being sent to specific hosts:

- ◆ *Next Hop Address:* The address of the host gateway that is to be used.
- ◆ *Host Address:* The IP address of the destination host. The address cannot be the first or last address of a class and must be unique.
- ◆ *Metric:* A value that alters the normal gateway-use logic, depending on a relative cost factor for using the gateways.
- ◆ *Type:* Active or passive. Gateways can be active where they publish their presence or passive where they do not.

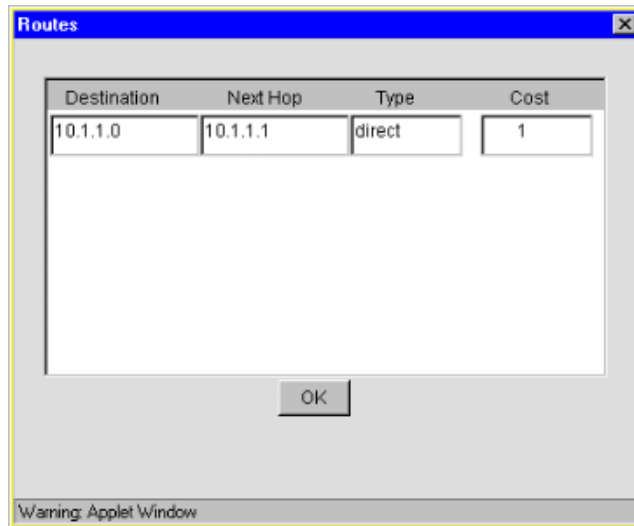
Network Gateways: You can define one or more gateways to be used for packets being sent to specific subnets.

- ◆ *Next Hop Address:* The address of the gateway that is to be used.
- ◆ *Subnet Base Address:* The subnet address for the destination IP address range. You can also enter a specific IP address on a given subnet and the appliance will calculate the subnet address using the mask.
- ◆ *Mask:* The subnet mask for the subnet or IP address above. A valid entry must be at least as large as a class mask where Class A Mask is 255.0.0.0, Class B Mask is 255.255.0.0, and Class C, D, and E Masks are 255.255.255.0.
- ◆ *Metric:* A value that alters the normal gateway-use logic, depending on a relative cost factor for using the gateways.
- ◆ *Type:* Active or passive. Gateways can be active where they publish their presence or passive where they do not.

Routes Dialog Box

Path: Network > Gateway/Firewall > Show Routes

Figure 95



This dialog box is useful for viewing and troubleshooting the routes the appliance is using. The list contains an entry for each defined gateway, each IP address assigned to an appliance network adapter, and routes discovered through RIP if the Enable RIP box is checked. Click Reset Learned Routes to clear RIP entries from the list.

Destination: The default route is named and listed first. The subnet address is shown for other routes.

Next Hop: This is the IP address of appliance network adapters, or the gateway address for all routes that are external to the appliance.

Type: Appliance network adapter routes are direct; all others are remote.

Cost: This is either the metric value you assigned to manually configured additional gateways (including the default gateway) or a relative cost factor assigned by the RIP function if the Enable RIP box is checked.

46

Using the Cache Panel

The cache panel lets you set up client acceleration (forward and transparent proxy) and Web acceleration (reverse proxy) for HTTP and FTP requests. It also lets you set up clusters of cooperating appliances, configure batch downloads into cache, set up filtering to block certain content, perform cache management, fine tune caching services, and specify how error pages are vended to browsers.

Client Accelerator Tab

Path: Cache > Client Accelerator

Figure 96

The screenshot shows the Volera EXCELERATOR 2.2 web interface. The 'Client Accelerator' tab is selected in the top navigation bar. On the left, a sidebar contains icons for Home, System, Network, Cache (highlighted), Hierarchy, and Monitoring. The main configuration area includes the following options:

- ☒ Enable client acceleration (forward proxy):
 - Proxy IP Addresses: A text box containing '10.1.1.1'.
 - Proxy port: A text box containing '8080'.
- ☐ Enable automatic proxy configuration (WPAD)
- ☐ Enable logging for client acceleration (with a 'Log Options' button)
- ☐ Enable authentication (with an 'Authentication Options' button)
- SSL listening port: A text box containing '443'.
- Certificate: A dropdown menu set to 'Auto'.
- ☐ Enable X-Forwarded-For
- ☐ Enable custom cache control header (with a 'Header Options' button)
- ☒ Allow HTTP CONNECT method
- ☒ Allow only SSL CONNECT traffic
- ☐ Enable access control (with an 'Access Control Options' button)

At the bottom right, there is an 'Advanced Options' button. At the bottom left, there are 'Apply', 'Cancel', and 'Help' buttons.

The Client Accelerator tab lets you configure the appliance as a forward proxy server. It lets you specify which IP addresses receive forward proxy requests from browsers and the ports that the appliance listens on for forward requests.

Enable Client Acceleration (Forward Proxy): Enables the appliance to handle forward proxy services. Browsers using this service must be configured with the appliance as a proxy server or must be enabled to obtain the proxy address automatically using WPAD. To activate the service, you must check one or more IP addresses for this service in the Proxy IP Addresses list.

Proxy IP Addresses: Displays all appliance IP addresses. You must check all addresses that you want to have forward proxy services on.

Proxy Port: Identifies the port from which the appliance will receive forward proxy requests and send requested data back to the requesting browsers.

Enable Automatic Proxy Configuration (WPAD): Enables WPAD listening on all IP addresses configured for client acceleration.

IMPORTANT: These IP addresses must not be used by any other appliance service that uses port 80. For example, a single IP address cannot provide both WPAD and Web server acceleration services that are configured to use port 80. The latter will always override WPAD. The same is true for any other service configured to use port 80.

The WPAD standard provides automatic configuration of browsers to use proxy services on a network. It requires browsers on the network to be configured to request and use WPAD information. Configuration procedures vary for each browser.

For more information on configuring your network for WPAD, see [“Setting Up Forward Proxy with WPAD” on page 287](#).

For information on customizing the WPAD implementation on your appliance, see [“Customizing Web Proxy Auto-Discovery” on page 285](#).

Enable Logging for Client Acceleration: Enables logging of forward activity.

Log Options: Lets you specify how often new log files are started and how long log files are retained. See [“Using Appliance Logging Services” on page 237](#) and [“Log Options Dialog Box” on page 335](#).

Enable Authentication: Causes the appliance to require authentication of users wanting to use its client accelerator services. Clicking Authentication displays the Add Authentication Profiles dialog box. For more information, see [“Add Authentication Profiles Dialog Box” on page 338](#).

IMPORTANT: Excelerator requires that each service (including authentication) uses a unique IP address and port combination. The default authentication port is 443. Attempts to enable authentication for more than one service on the same IP address and port will result in a TCP bind error.

DNS Name: If the authentication profiles selected for this service are not Basic or NTLM profiles, the initial information exchanges between the requesting browsers and the cache device are SSL-encrypted.

If the Subject Name in the SSL certificate on the cache device is not the IP address of the cache device, the value of this field must match both the DNS name returned from the browser’s DNS lookup and the Subject Name in the cache device’s certificate. Otherwise, browsers will issue security alerts to users.

SSL Listening Port: The port on which the appliance listens for authentication requests.

Certificate: This drop-down list displays any certificates you have stored on your appliance. System-generated certificates do not appear in the list.

Use this field to select the certificate you created specifically for the appliance’s client accelerator services. This will prevent browsers from receiving certificate confirmation messages each time

they access the appliance. For more information, see [“Managing Appliance Certificates” on page 165](#).

Enable X-Forwarded-For: Headers used to pass browser ID information along with browser request packets. If the headers are included, Web servers can determine the origin of browser requests they receive. If the headers are not included, browser requests have anonymity.

Checking the X-Forward-For option causes the appliance either to add information to an existing X-Forwarded-For or Forwarded-For header, or to create a header if one doesn’t already exist.

Leaving the option unchecked causes the appliance to remove X-Forwarded-For headers from any forward proxy requests passing through the appliance.

You must weigh the desires of browser users to remain anonymous against the desire of Web server owners (e-commerce sites, for example) to collect data about who is accessing their sites.

Enable Custom Cache Control Header: Lets you enable the caching of objects on the appliance while preventing caching by requesting browsers.

Custom cache control headers are designed primarily to be used in Web server accelerators. However, very large hosting sites sometimes prefer to place forward or transparent accelerators in front of their server farms rather than creating hundreds or even thousands of Web server accelerator services.

For details on how the headers work, see [“Custom Cache Control Header Dialog Box” on page 340](#).

Allow HTTP CONNECT Method: Lets you enable the forward proxy service to use the HTTP CONNECT method. For details, see [“Managing HTTP CONNECT Method Support” on page 191](#).

Allow Only SSL CONNECT Traffic: Lets you have Excelerator check to ensure that HTTP CONNECT requests to the forward service contain SSL-related traffic. For details, see [“Managing HTTP CONNECT Method Support” on page 191](#).

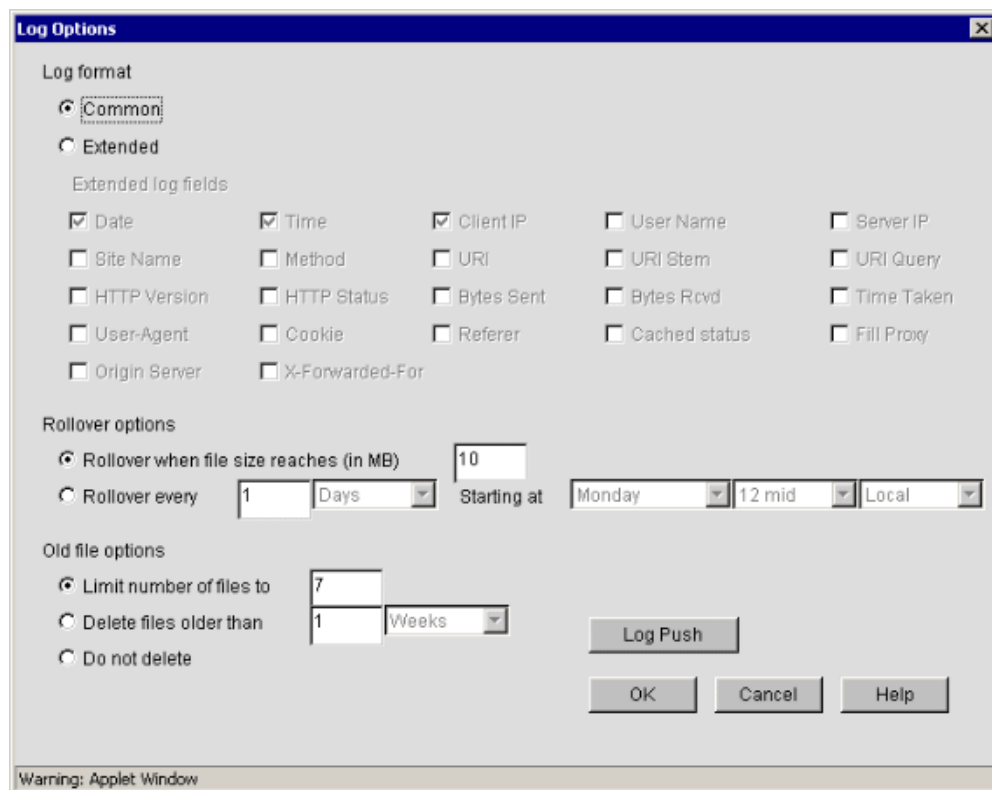
Enable Access Control: . For details, see [“Access Control Options Dialog Box” on page 341](#).

Advanced Options: Lets you control the TCP receive window size for cache device fills from origin Web servers, the caching of objects on the cache device that would not normally be cached, and the filling and vending of browser no-cache requests by the cache device. For details, see [“Advanced Options \(Tuning\) Dialog Box” on page 341](#).

Log Options Dialog Box

Path: Cache > Client Accelerator > Enable logging for Client Acceleration > Log Options

Figure 97

The image shows a Java applet window titled "Log Options". It contains three main sections: "Log format", "Rollover options", and "Old file options". In the "Log format" section, the "Common" radio button is selected. Below it, a grid of checkboxes lists various log fields: Date, Time, Client IP, User Name, Server IP, Site Name, Method, URI, URI Stem, URI Query, HTTP Version, HTTP Status, Bytes Sent, Bytes Rcvd, Time Taken, User-Agent, Cookie, Referer, Cached status, Fill Proxy, Origin Server, and X-Forwarded-For. In the "Rollover options" section, the "Rollover when file size reaches (in MB)" option is selected with a value of 10. The "Rollover every" option is also visible with a value of 1, and the "Starting at" section shows "Monday", "12 mid", and "Local". In the "Old file options" section, the "Limit number of files to" option is selected with a value of 7. At the bottom right are buttons for "Log Push", "OK", "Cancel", and "Help". A warning bar at the bottom left says "Warning: Applet Window".

The Log Options dialog box lets you set logging format and other options for a proxy service.

Common: For information, see the [Common Log File Format Web site \(http://www.w3.org/daemon/user/config/logging.html\)](http://www.w3.org/daemon/user/config/logging.html).

Extended: For information, see the [Extended Log File Format Web site \(http://www.w3.org/TR/WD-logfile\)](http://www.w3.org/TR/WD-logfile). For information on the appliance-specific extended log format, see “[About Extended Log Field Headers](#)” on page 249.

The list of possible fields to log includes User Name, Site Name, Fill Proxy, and Origin Server. These fields are not part of the extended log file format definition.

Rollover Options: Lets you specify how often new log files are started or rolled over. You can use either periods of time or log file size to trigger the start of a new file.

NOTE: If you specify file size as the trigger to start a new file and the appliance is shut down before and restarted after midnight, Excelerator will start a new log file automatically.

Old File Options: Lets you specify the disposition of old log files. The Do Not Delete option does not prevent the manual deletion of files nor the deletion of files specified in the [FTP Log Push Configuration Dialog Box](#) dialog box.

FTP Log Push Configuration Dialog Box

Path: Any Log Options dialog box > Log Push

Figure 98

Warning: Applet Window

The FTP Log Push Configuration dialog box lets you schedule regular downloading and deleting of appliance log files to an FTP server, thus preventing a disk full condition and the loss of logging data.

IMPORTANT: Although this dialog box is accessed through the service-specific Log Options dialog boxes, FTP log push is a global feature that affects all log files on the appliance.

For more information on log file management, see [“Using Appliance Logging Services” on page 237](#).

FTP Log Push Enable: Checking this box enables FTP Log Push for the appliance. All log files for the services checked under Log Types are affected.

Host Server: The DNS name or IP address of the FTP server to which log files are to be downloaded.

Login Name: A valid FTP login name.

Password: The password of the FTP login name.

Default Directory: A subdirectory located in the default directory of the FTP login name. If this directory doesn’t exist, FTP Log Push will create it. All log files are downloaded in a subdirectory structure that mirrors the ETC/PROXY/DATA/LOGS directory on the appliance. For log file locations, see [“Getting Log Filenames” on page 246](#).

IP Address: An appliance IP address through which FTP Log Push will send the log files.

FTP Log Push does not require that the appliance’s mini FTP server be activated.

Delete Log Files from Exceleator Server after Successful Push: When selected, this option causes log files to be deleted immediately after they have been downloaded to the FTP server. This option overrides rollover and old file settings in the Log Options dialog box.

Log Push Result: A non-editable status box indicating if push operations were successful.

Push Logs When the Logs Roll Over: When this option is selected, log files are immediately downloaded and deleted when they are closed by a rollover operation.

Days to Push the Logs: Allows you to specify how often to push logs. You can push logs monthly, weekly, only on specific days, or every day. You can specify frequent pushing through the Push Logs When the Logs Roll Over option.

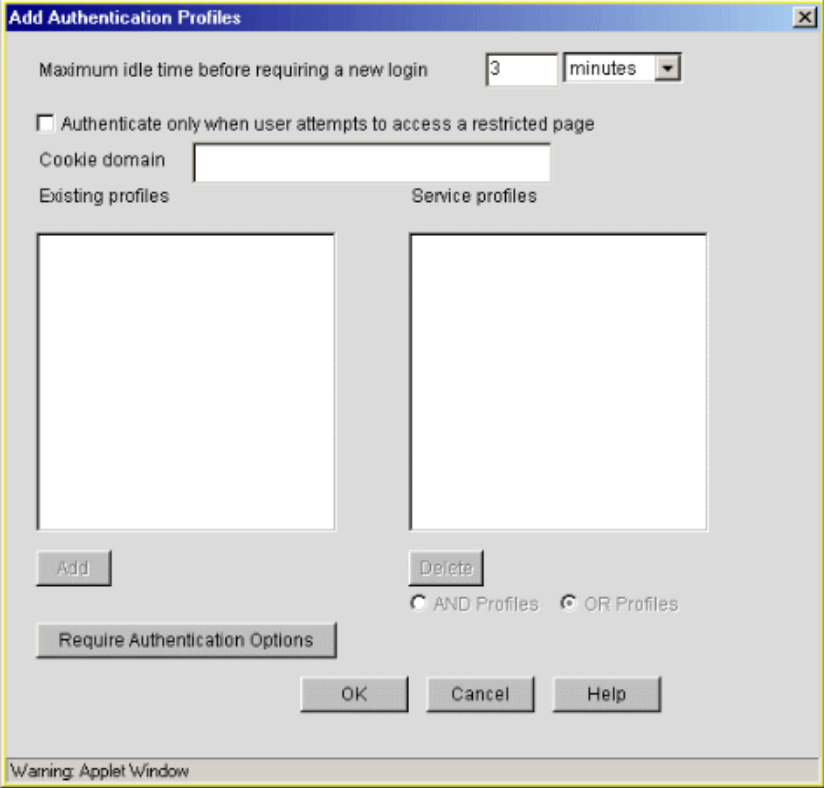
Time to Push the Logs: Specifies the hour at which the push will occur.

Log Types: Identifies the services whose logs are to be pushed. The Cluster option pushes all logs for any clustered services running on the appliance.

Add Authentication Profiles Dialog Box

Path: Cache > Client Accelerator or Transparent Handling or Web Server Accelerator > Enable Authentication > Authentication Options

Figure 99



The dialog box titled "Add Authentication Profiles" contains the following elements:

- A text input field for "Maximum idle time before requiring a new login" with the value "3" and a "minutes" dropdown menu.
- An unchecked checkbox labeled "Authenticate only when user attempts to access a restricted page".
- A text input field for "Cookie domain".
- Two empty rectangular boxes labeled "Existing profiles" and "Service profiles".
- "Add" and "Delete" buttons below the profile boxes.
- Radio buttons for "AND Profiles" and "OR Profiles", with "OR Profiles" selected.
- A "Require Authentication Options" button.
- "OK", "Cancel", and "Help" buttons at the bottom.
- A status bar at the bottom left that reads "Warning: Applet Window".

The Add Authentication Profiles dialog box lets you select one authentication profile, which you can use to authenticate the users of the proxy service from which you accessed this dialog box.

Maximum Idle Time Before Requiring a New Login: The period of browser inactivity allowed before the appliance requests a new login.

Authenticate Only When User Attempts to Access a Restricted Page: This option is used in conjunction with access control policies created for the proxy service you are configuring. For information on access control, see [“Access Control” on page 161](#).

If this option is checked, users are prompted to authenticate only when accessing pages that are covered by an access control policy. If this is not checked, users must authenticate whenever they attempt to use the proxy service being configured.

Cookie Domain: This optional field lets you explicitly specify the domain for which the authentication cookie will be set.

Excelerator's default authentication cookie generation works for domains of the form `www.companyname.com`.

For other domains such as international domain names wherein the right two fields are often top-level domains, the default generation can result in invalid cookie domain names, such as `co.uk`.

If you need to specify an authentication cookie domain, the following rules apply:

- ♦ You can't specify top level domains, for example `.com` or `.co.uk`
- ♦ You can't include the left-most dot-delimited field in the domain name. For example, `setup.foo.net` is not allowed for the URL `http://setup.foo.net`, and `www.foo.net` is also not allowed for the URL `http://www.foo.net`.
- ♦ The cookie domain name must be a valid, dot-delimited substring of the DNS name specified in the service definition you are configuring (forward, transparent, or reverse).

For example, `foo.net` is a valid, dot-delimited substring of `setup.foo.net`.

Existing Profiles: A list of the authentication profiles you have created in Cache > Authentication. For more information, see [“Authentication Tab” on page 377](#).

Service Profiles: The authentication profiles that are active for the proxy service you are configuring. You add a profile to the list by clicking a profile in the Existing Profiles list and then clicking Add. You can remove the profile from this list using the Delete button.

The list can contain one or two profiles as follows:

- ♦ A single mutual authentication, username/password, or basic authentication profile
 - ♦ One mutual authentication and one username/password authentication profile
- If you choose this combination, the AND Profiles and OR Profiles options are activated.
- ♦ One username/password and one basic authentication profile

AND Profiles: If this option is selected, users must pass both the mutual authentication and the username/password authentication criteria to access the service.

OR Profiles: If this option is selected, users can use the service after passing either the mutual or the username/password profile's criteria.

Require Authentication Options: Clicking this button displays the Require Authentication on Request Header dialog box. This dialog box contains the following options.

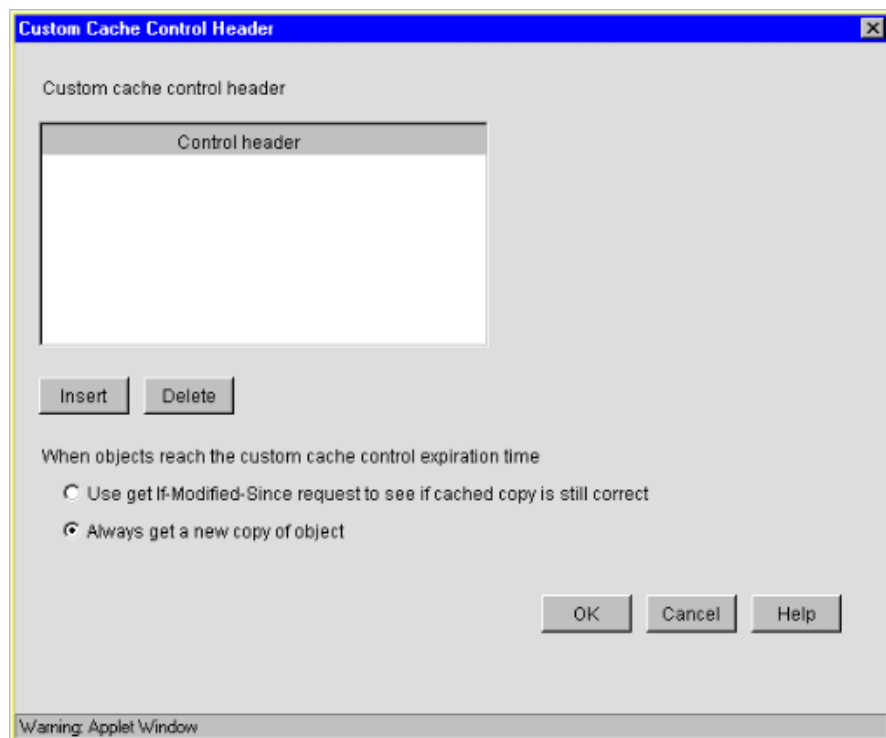
- ♦ *GET and HEAD:* If this option is checked, all GET and HEAD requests will require authentication before vending content. This option is enabled by default.
- ♦ *Non-HTTP Schemes:* If this option is checked, GET and HEAD requests for non-HTTP schemes (FTP, etc.) will require authentication before vending content. This option is enabled by default.
- ♦ *POST:* If this option is checked, HTTP POST requests will require authentication.
- ♦ *PUT:* If this option is checked, HTTP PUT requests will require authentication.

- ♦ *CONNECT*: If this option is checked, HTTP CONNECT requests will require authentication.
- ♦ *TRACE*: If this option is checked, HTTP TRACE requests will require authentication.
- ♦ *OPTIONS*: If this option is checked, HTTP OPTIONS requests will require authentication.
- ♦ *DELETE*: If this option is checked, HTTP DELETE requests will require authentication.
- ♦ *Other*: If this option is checked, all HTTP requests of a type not listed above will require authentication.

Custom Cache Control Header Dialog Box

Path: Cache > Web Server Accelerator > Enable Custom Cache Control Header > Header Options

Figure 100



This dialog box lets you specify object headers that the appliance recognizes as overriding standard HTTP cache directives. The dialog box has two options that let you specify how the appliance refills objects with custom cache control headers when the objects expire in cache.

For more information on how objects expire and what happens when they do, see [“Overview” on page 185](#) and [“Managing Cache Freshness” on page 185](#).

For more information on custom cache control headers, see [“Using Custom Cache Control Headers” on page 187](#).

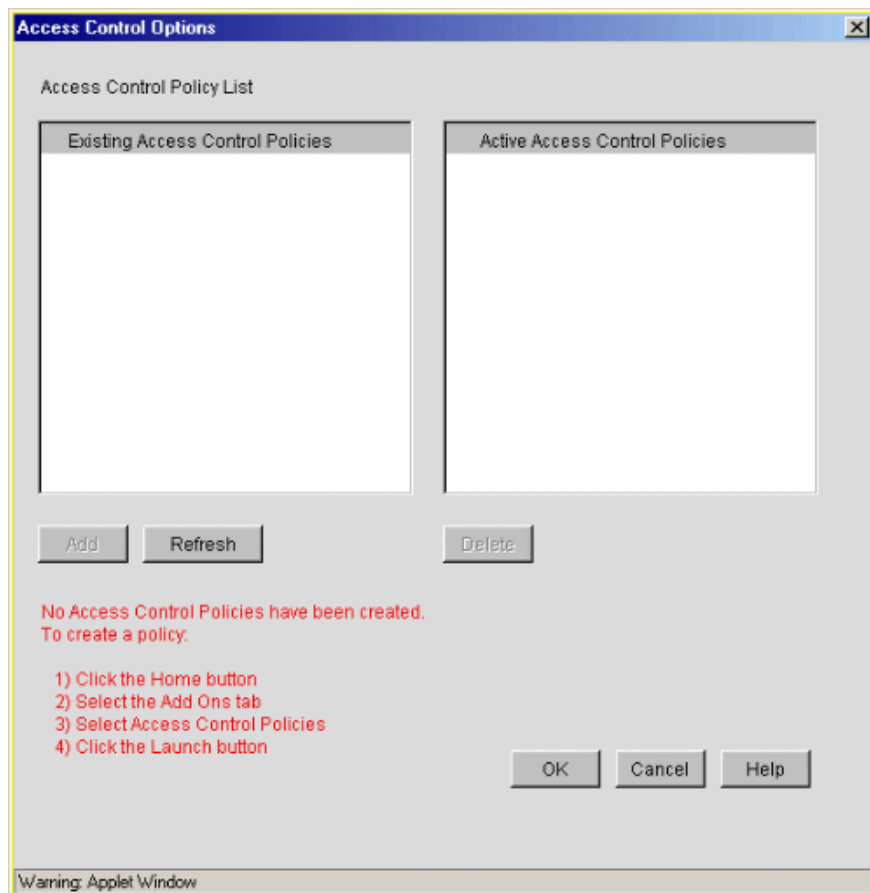
Use Get If-Modified-Since Request to See If Cached Copy Is Still Correct: When objects expire in cache and are subsequently requested by a browser, they are refilled in cache only if they have changed on the origin Web server. Otherwise, they are retained in cache and their expiration timer is set to the custom cache control header value.

Always Get a New Copy of Object: When objects expire in cache and are subsequently requested by a browser, they are refilled in cache and their expiration timer is set to the custom cache control header value. Excelerator doesn't check to see whether objects have changed on the origin Web server.

Access Control Options Dialog Box

Path: Cache > Client Accelerator or Transparent Handling or Web Server Accelerator > Enable Access Control > Access Control Options

Figure 101



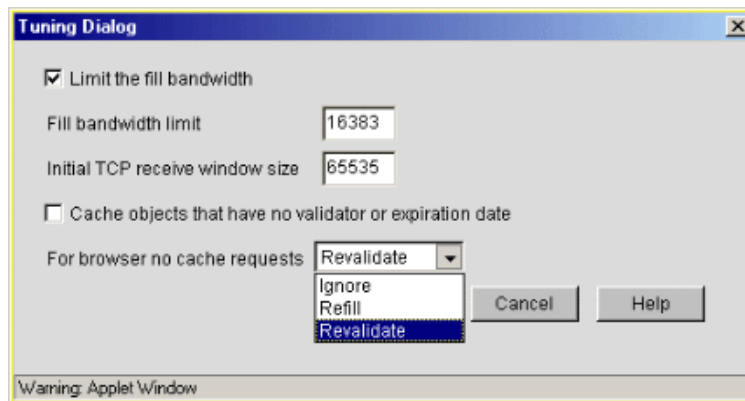
The Access Control Options dialog box lets you assign access control policies to the proxy service through which you accessed the box. You assign policies to the service by selecting a profile in the Existing Access Control Policies list and clicking Add. The policy then moves to the Active Access Control Policies list. Use the Delete button to remove policies from the active list.

For more information on how access control policies function and how to create them, see [Chapter 23, “Access Control,”](#) on page 161.

Advanced Options (Tuning) Dialog Box

Path: Cache > Client Accelerator or Transparent Handling or Web Server Accelerator > Advanced Options

Figure 102



The Tuning dialog box lets you control the initial and subsequent TCP receive window size for cache device fills from origin Web servers, the caching of objects on the cache device that would not normally be cached, and the filling and vending of browser no-cache requests by the cache device.

Limit the Fill Bandwidth: This enables limits on the TCP receive window size for browser requests using the service.

Fill Bandwidth Limit: This sets the TCP receive window size for cache fills that occur after the initial fill request.

Initial TCP Receive Window Size: This sets the TCP receive window size for initial cache fill requests.

Cache Objects That Have No Validator or Expiration Date: This enables caching of objects that would not normally be cached because they have no validator or expiration date set.

For Browser No-Cache Requests: This controls the handling of browser requests that specify the request should not be filled from cache. Requests can be ignored (vended from cache without checking content freshness), refilled (automatically refilled from the origin Web server before vending), or revalidated (checked against the object on the origin Web server and refilled prior to vending if that object is fresher).

Transparent Handling Tab

Path: Cache > Transparent Handling

Figure 103

The screenshot displays the Volera EXCELERATOR 2.2 web interface. The top navigation bar includes tabs for Client Accelerator, Transparent Handling (selected), Web Server Accelerator, FTP, Streaming, and Cluster. The left sidebar contains icons and labels for Home, System, Network, Cache, Hierarchy, and Monitoring. The main content area is titled 'Transparent Handling' and features a checkbox for 'Enable transparent client acceleration (transparent proxy - L4 switch support)'. Below this are three lists: 'Proxy ports' (containing 80), 'Exception IP addresses' (empty), and 'Proxy IP Addresses' (containing 10.1.1.1). Each list has 'Insert' and 'Delete' buttons. Further down, there are checkboxes for 'Router Options', 'Enable WCCP', 'Enable authentication', and 'SSL listening port' (set to 443). There are also buttons for 'Router Options', 'WCCP Options', 'Authentication Options', and 'Access Control Options'. The 'Error Handling Method' is set to 'SendCacheError'. At the bottom, there are checkboxes for 'Enable X-Forwarded-For', 'Enable logging for transparent handling', 'Enable custom cache control header', 'Allow HTTP CONNECT method', and 'Allow only SSL CONNECT traffic'. There are also buttons for 'Log Options', 'Header Options', and 'Advanced Options'. At the very bottom are 'Apply', 'Cancel', and 'Help' buttons.

The Transparent Handling tab lets you configure the appliance as a transparent proxy server. It lets you specify which ports and IP addresses the appliance listens on for transparent requests.

Enable Transparent Client Acceleration (Transparent Proxy—L4 Switch Support): Enables the appliance to handle transparent proxy services. You must also check the IP addresses for this service in Cache > Client Accelerator > the Proxy IP Addresses list.

Proxy Ports: The ports from which the appliance receives transparent proxy requests and sends requested data back to the requesting browsers.

Exception IP Addresses: A list of origin server IP addresses. Browser requests to these addresses will bypass Excelsator's transparent handling service and be sent directly to the origin server.

Proxy IP Addresses: The addresses that are enabled for transparent proxy services and to which HTTP requests are forwarded by the L4 switch, WCCP-capable router, or appliance internal router.

IMPORTANT: Two restrictions apply:

You cannot configure a transparent proxy service on an IP address assigned to a card that has the Dynamic option set for NAT. NAT and transparent proxy cannot coexist on the same card.

You must not enable more than one address on each network card for transparent proxy services.

Router Options: Enables the appliance to act as a router. The appliance provides basic routing without additional configuration. For information about setting up alternate routes and so on, see [“Router Options Dialog Box” on page 347](#).

Enable Access Control: . For details, see [“Access Control Options Dialog Box” on page 341](#).

Enable WCCP: Enables WCCP-capable routers to route HTTP requests to the appliance. You must also select the version of WCCP used by the routers and configure WCCP options so that routers can recognize the appliance and know how to work with it. For more information, see [“WCCP Version 1.0 Options Dialog Box” on page 345](#) and [“WCCP Version 2.0 Options Dialog Box” on page 345](#).

Enable Authentication: Checking this box causes the appliance to require authentication of users wanting to use its transparent proxy services. Click Authentication to display the Add Authentication Profiles dialog box. For more information, see [“Add Authentication Profiles Dialog Box” on page 338](#).

SSL Listening Port: The port on which the appliance listens for authentication requests.

IMPORTANT: Excelerator requires each service (including authentication) to use a unique IP address and port combination. The default authentication port is 443. Attempts to enable authentication for more than one service on the same IP address and port will result in a TCP bind error.

Certificate: Drop-down list that displays any certificate you have stored on your appliance. System-generated certificates do not appear in the list.

Use this field to select the certificate you created specifically for the appliance’s transparent services. This will prevent browsers from receiving certificate confirmation messages each time they access the appliance. For more information, see [“Managing Appliance Certificates” on page 165](#).

Error Handling Method: Lets you select a method for handling origin server error pages. You can have the appliance send a cache error to the browser, reset the connection with the browser, or transparently pass through the message from the origin server.

Enable X-Forwarded-For: Headers used to pass browser ID information along with browser request packets. If the headers are included, Web servers can determine the origin of browser requests they receive. If the headers are not included, browser requests have anonymity.

Checking the X-Forwarded-For option causes the appliance to either add information to an existing X-Forwarded-For or Forwarded-For header, or to create a header if one doesn’t already exist.

Leaving the option unchecked causes the appliance to remove X-Forwarded-For headers from any transparent proxy requests passing through the appliance.

You must weigh the desires of browser users to remain anonymous against the desires of Web server owners (e-commerce sites, for example) to collect data about who is accessing their site.

Enable Logging for Transparent Handling: Enables logging of transparent activity.

Log Options: Lets you specify how often new log files are started and how long log files are retained. See [“Using Appliance Logging Services” on page 237](#) and [“Log Options Dialog Box” on page 335](#).

Enable Custom Cache Control Header: Lets you enable the caching of objects on the appliance while preventing caching by requesting browsers.

Custom cache control headers are designed primarily to be used in Web server accelerators. However, very large hosting sites sometimes prefer to place forward or transparent accelerators in front of their server farms rather than creating hundreds or even thousands of Web server accelerator services.

For details on how the headers work, see [“Custom Cache Control Header Dialog Box” on page 340](#).

Allow HTTP CONNECT Method: Lets you enable the transparent proxy service to use the HTTP CONNECT method. For details, see [“Managing HTTP CONNECT Method Support” on page 191](#).

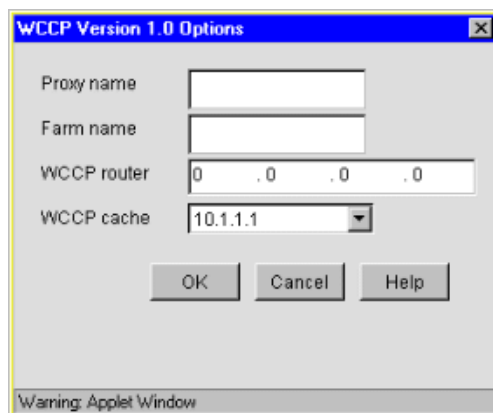
Allow Only SSL CONNECT Traffic: Lets you have Excelerator check to ensure that HTTP CONNECT requests to the transparent service contain SSL-related traffic. For details, see [“Managing HTTP CONNECT Method Support” on page 191](#).

Advanced Options: Lets you control the TCP receive window size for cache device fills from origin Web servers, the caching of objects on the cache device that would not normally be cached, and the filling and vending of browser no-cache requests by the cache device. For details, see [“Advanced Options \(Tuning\) Dialog Box” on page 341](#).

WCCP Version 1.0 Options Dialog Box

Path: Cache > Transparent Handling > Enable WCCP > WCCP V1 Options > WCCP Options

Figure 104

The image shows a Java applet window titled "WCCP Version 1.0 Options". It contains four input fields: "Proxy name" (a text box), "Farm name" (a text box), "WCCP router" (a text box with the value "0 . 0 . 0 . 0"), and "WCCP cache" (a dropdown menu with the value "10.1.1.1"). At the bottom are three buttons: "OK", "Cancel", and "Help". A warning bar at the bottom of the window reads "Warning: Applet Window".

The WCCP Version 1.0 Options dialog box lets you configure the appliance to provide configuration information to the router that uses WCCP 1.0. The router can work with multiple appliances, but an appliance can work with only one WCCP 1.0 router.

Proxy Name: Not required for router configuration. This is provided for your reference only. You can use text of any length and content for the name.

Farm Name: Not required for router configuration. This is provided for your reference only. You can use text of any length and content for the name.

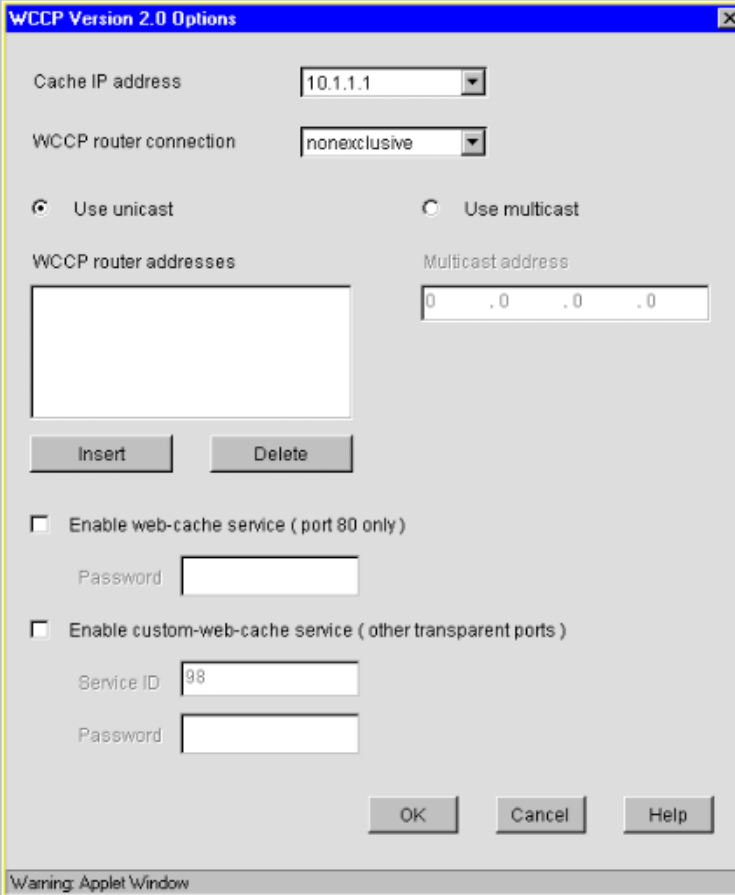
WCCP Router: The address of the WCCP-capable router. The appliance uses this address to request that the router route HTTP traffic to the appliance.

WCCP Cache: The address of the transparent proxy service on the appliance. The router routes HTTP traffic to this address.

WCCP Version 2.0 Options Dialog Box

Path: Cache > Transparent Handling > Enable WCCP > WCCP V2 Options > WCCP Options

Figure 105



The image shows a Java applet window titled "WCCP Version 2.0 Options". It contains the following fields and controls:

- Cache IP address:** A text box containing "10.1.1.1".
- WCCP router connection:** A dropdown menu set to "nonexclusive".
- Use unicast:** A radio button that is selected.
- Use multicast:** An unselected radio button.
- WCCP router addresses:** An empty list box with "Insert" and "Delete" buttons below it.
- Multicast address:** A text box containing "0 . 0 . 0 . 0".
- Enable web-cache service (port 80 only):** An unchecked checkbox. Below it is a "Password" text box.
- Enable custom-web-cache service (other transparent ports):** An unchecked checkbox. Below it are "Service ID" (containing "98") and "Password" text boxes.
- Buttons:** "OK", "Cancel", and "Help" buttons at the bottom right.
- Warning:** A small "Warning: Applet Window" message at the bottom left.

The WCCP Version 2.0 Options dialog box lets you configure the appliance to provide configuration information to the routers that use WCCP 2.0.

Cisco* recommends that you use IOS 12.05t or later for WCCP 2.0 to work as planned.

Cache IP Address: The address of a transparent proxy service on the appliance. The router routes HTTP traffic to this address. This address should also be the appliance's main connection to the Internet.

WCCP Router Connection: Option used to specify how the router and appliance are connected:

- ♦ *Non-exclusive:* The interface connecting the appliance to the router is redirecting traffic. The appliance encapsulates the return packet so the router won't redirect it, and the router unencapsulates the packet for processing.
- ♦ *Exclusive:* The interface connecting the appliance to the router is not redirecting traffic. The appliance doesn't encapsulate the return packet. The router processes the packet as is.

Use Unicast: Option which causes the appliance to communicate with the configured WCCP 2.0 routers using UDP packets and enables the WCCP Router Addresses list.

WCCP Router Addresses: The addresses of one or more WCCP-capable routers. The appliance uses these addresses to request that the routers route HTTP traffic to the appliance.

Use Multicast: Option which causes the appliance to use multicast packets for requesting that WCCP 2.0 routers route traffic to the appliance.

Multicast Address: Multicast address for your network. The appliance verifies that the address is valid when you click OK.

Enable Web-Cache Service (Port 80 Only): Checking this option causes the appliance to register with the WCCP-capable routers for receiving port 80 transparent proxy traffic.

Password: If you enter a password, the appliance signs WCCP version 2 communication packets with an MD5 hash or “signature” of the password you enter. The valid string range is one to seven characters.

Enable Custom-Web-Cache Service (Other Transparent Ports): Check this option if your WCCP-capable routers handle transparent proxy traffic on multiple ports.

Service ID: The service ID that Cisco routers use for the multiple port service. Currently, the only valid value is 98.

Password: If you enter a password, the appliance signs WCCP 2.0 communication packets with an MD5 hash or signature of the password you enter. The valid string range is one to seven characters.

Router Options Dialog Box

Path: Cache > Transparent Handling > Router Options

Figure 106

Router Options

☐ Enable RIP

Default gateway

Next hop address	Metric	Type
10 . 1 . 1 . 1	1	Passive

Host gateways

Next hop address	Host address	Metric	Type
------------------	--------------	--------	------

Network gateways

Next hop address	Subnet base address	Mask	Metric	Type
------------------	---------------------	------	--------	------

Warning: Applet Window

The Router Options dialog box lets you set up static routes and cause the appliance to maintain and use RIP tables.

For an overview of appliance routing capabilities, see [“Router Capabilities” on page 277](#).

Enable RIP: Causes the appliance to use the Routing Information Protocol (RIP) to build and maintain a table of the shortest routes to destinations.

Show Routes: See [“Routes Dialog Box” on page 330](#).

Reset Learned Routes: Causes the RIP table to be cleared. The appliance then begins building the table from scratch.

Default Gateway: The gateway for requests to destinations not covered by host or network gateways.

Host Gateways: Lets you set routing for requests to specific hosts. Packets are checked against these routes first. Packets not routed are then checked against the network gateways table.

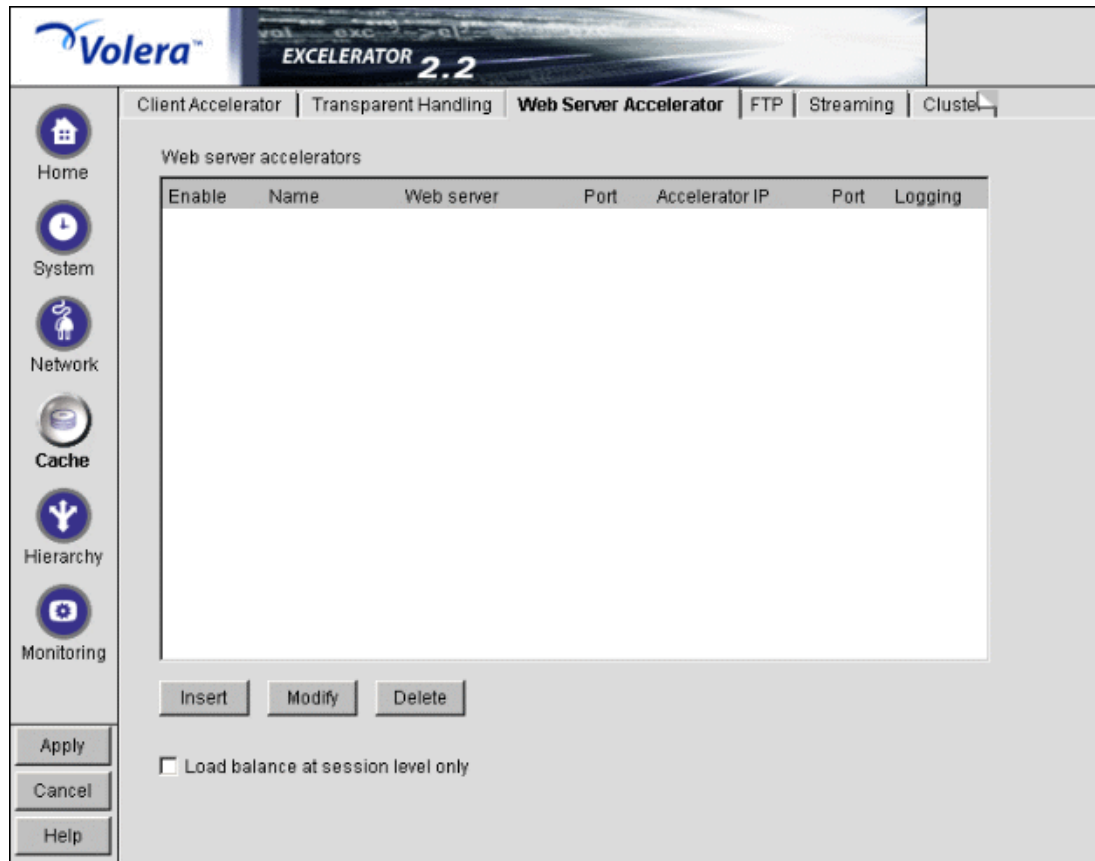
Network Gateways: Lets you set routing for requests to specific networks. Packets not fitting the host gateways criteria are checked against the routes in these tables, beginning with the most restrictive routes. Packets not routed are then sent to the default gateway.

For more detail on the fields in this dialog box, see [“Additional Gateways Dialog Box” on page 328](#).

Web Server Accelerator Tab

Path: Cache > Web Server Accelerator

Figure 107



The Web Server Accelerator tab lets you add one or more Web server (reverse proxy) accelerators. The appliance acts as the front end to Web servers on your Internet or intranet and off-loads frequent requests, thereby freeing up bandwidth. Using a Web server accelerator also increases security because the IP addresses of your Web servers are hidden from the Internet. For more information, see [“Overview of Web Server Acceleration” on page 51](#).

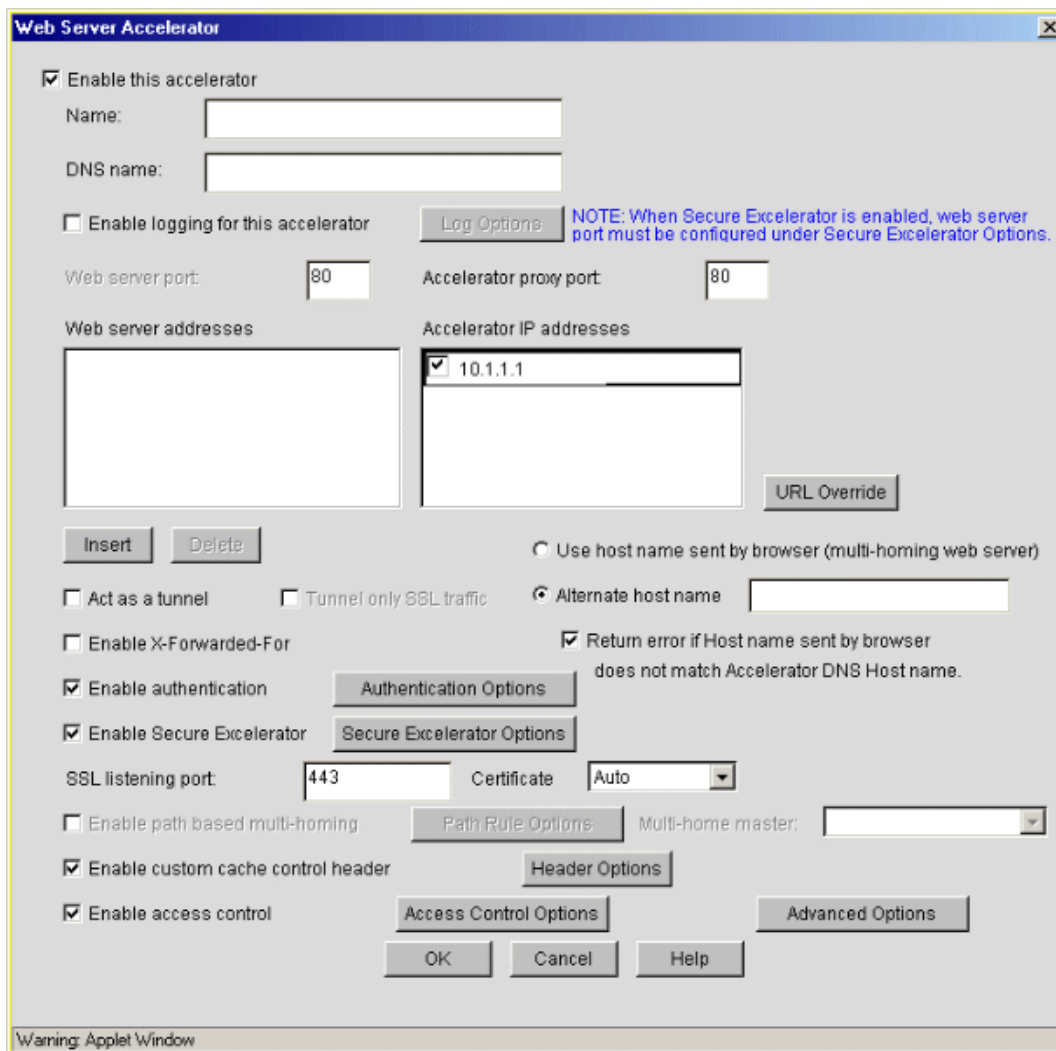
Field names shown in the Web Server Accelerators list are explained in [“Web Server Accelerator Dialog Box” on page 349](#).

Load Balance at Session Level Only: Checking this box causes the appliance to use the same Web server for all fills during a session. This prevents e-business users from having to log in multiple times. This setting affects all Web server accelerators configured on the appliance.

Web Server Accelerator Dialog Box

Path: Cache > Web Server Accelerator > Insert

Figure 108



The screenshot shows the 'Web Server Accelerator' dialog box. It has a title bar with a close button. The main area contains various configuration options:

- ☒ Enable this accelerator. Below it are text boxes for 'Name:' and 'DNS name:'.
- ☐ Enable logging for this accelerator. Next to it is a 'Log Options' button. A blue note states: 'NOTE: When Secure Accelerator is enabled, web server port must be configured under Secure Accelerator Options.'
- Text boxes for 'Web server port:' (containing 80) and 'Accelerator proxy port:' (containing 80).
- Text boxes for 'Web server addresses' and 'Accelerator IP addresses'. The 'Accelerator IP addresses' box contains '10.1.1.1' with a checkmark. A 'URL Override' button is to the right.
- 'Insert' and 'Delete' buttons.
- Radio buttons for 'Use host name sent by browser (multi-homing web server)' and 'Alternate host name' (selected). The 'Alternate host name' has a text box.
- ☐ Act as a tunnel, ☐ Tunnel only SSL traffic.
- ☐ Enable X-Forwarded-For.
- ☒ Enable authentication. Next to it is an 'Authentication Options' button.
- ☒ Enable Secure Accelerator. Next to it is a 'Secure Accelerator Options' button.
- Text boxes for 'SSL listening port:' (containing 443) and 'Certificate:' (containing Auto).
- ☐ Enable path based multi-homing. Next to it is a 'Path Rule Options' button.
- ☒ Enable custom cache control header. Next to it is a 'Header Options' button.
- ☒ Enable access control. Next to it is an 'Access Control Options' button.
- An 'Advanced Options' button.
- 'OK', 'Cancel', and 'Help' buttons at the bottom.

A warning bar at the bottom left says 'Warning: Applet Window'.

The Web Server Accelerator dialog box lets you create Web server accelerator services for handling requests to Web servers.

IMPORTANT: Accelerators and tunnels are incompatible. You can either configure an accelerator (includes multiple multihoming accelerators) on an IP address and port or configure one tunnel on an IP address and port.

Enable This Accelerator: Specifies whether to enable the defined Web server accelerator after you have configured it. The default is Enabled.

Name: Name of the Web server accelerator service. Each Web server accelerator service requires a name you create. For example, you can select a name that indicates which Web server is being serviced by the appliance, or alternately, a set of browsers configured to access the Web server accelerator as a proxy server. A valid name consists of a DOS-style, eight-character name with no special characters or spaces.

If logging is enabled, the appliance uses the Web server accelerator name as the name of the directory in which the log files are kept.

DNS Name: The contents of this field depend on the type of accelerator you are using.

If you are accelerating multiple Web servers for multiple Web sites on the same IP address, you must create a Web server accelerator definition for each DNS name that is used in browser requests. This name must exactly match one of the names in the requests. See [“Standard Multihoming for Multiple Web Sites” on page 122](#).

If you are accelerating multiple Web servers for a single Web site and plan to use path-based multihoming, you must use the same DNS name in every accelerator definition.

If you specify one or more DNS names in the Web Server Addresses list and those names are different than this DNS name, the system creates an automatic DNS name override that replaces those names with the service’s DNS name in all object references that point to Web server’s accelerated by the service. For more information, see [“Understanding When Automatic URL Overrides Are Created” on page 174](#).

This name is also used if the appliance is part of an ICP hierarchy that needs to resolve relative URLs.

Enable Logging for This Accelerator: Causes log files to be kept for this Web server accelerator.

Web Server Port: The port number that the origin Web server is listening on for incoming connections. The default for HTTP is 80. The valid port range is 1 through 65535.

Accelerator Proxy Port: The port number that the proxy server is listening on for incoming connections. The default for HTTP is 80. The valid port range is 1 through 65535.

IMPORTANT: Configuring both a tunnel and an accelerator on the same IP address and port is not supported.

Web Server Addresses: The IP address or local DNS name of each Web server from which the appliance fills the cache for this Web server accelerator. The cache must be able to fill all requests through any of these names or addresses unless path-based multihoming is being used, in which case the device will refer first to the parent accelerator and then to the appropriate child accelerator based on the Sub-Path Match String entered in the [Path Rule Options Dialog Box](#).

Accelerator IP Addresses: For normal accelerator situations, non-path-based multihoming configurations, and accelerators configured as multihoming masters, this is the appliance’s IP addresses to which DNS resolves the Web server’s (or Web site’s) DNS name and on which the Web server accelerator listens for incoming connections from the Internet.

For child accelerators in path-based multihoming configurations, this is the IP address or addresses to which the multihoming master forwards browser requests that match the specified path rule.

URL Override: This dialog box lets you create additional URL overrides for the accelerator service you are configuring.. For details, see [“URL Override Dialog Box” on page 353](#).

Use Hostname Sent by Browser (Multi-homing Web Server): When selected, this option preserves the hostname in the HTTP header exactly as it came in the browser request. This disables any DNS name URL overrides that would normally apply.

Act as a Tunnel: Lets you create one or more accelerator services for the specific purpose of tunneling non-HTTP traffic through the appliance to the origin Web server. Normally an accelerator service processes HTTP requests in order to fill them. However, it is not unusual that some of the traffic coming through the appliance is not HTTP-based.

Web servers often handle SSL connections, and less frequently they might need to let Telnet, FTP, chat, or other kinds of traffic through without attempting to process it.

When the Act as a Tunnel option is checked, the accelerator sets up a tunnel for all incoming traffic. Because tunnels are already limited (they can only connect to the fill Web server on the specified fill port), they are much more secure than connections to forward or transparent services that use the CONNECT method. For more information on CONNECT security concerns, see [“Managing HTTP CONNECT Method Support” on page 191](#).

When you check the Act as a Tunnel option, you have the additional option of having the accelerator service tunnel only SSL traffic.

IMPORTANT: Configuring both a tunnel and an accelerator on the same IP address and port is not supported.

Tunnel Only SSL Traffic: If you decide to have the accelerator act as a tunnel, you can elect to have it tunnel only SSL traffic. The service will then verify that the address and port being accessed are actually an SSL Web site. If verification fails, the service will tear down the connection.

Alternate Hostname: Checking this option causes the string specified to be substituted for the hostname in the HTTP header before the request is forwarded to the Web server.

This field is required in two circumstances:

- ♦ The Web server being accelerated will only answer requests when a specific DNS hostname is included in the request header.
- ♦ The Web server being accelerated is multihomed and needs to know which DNS name home to route the request to.

Enable X-Forwarded-For: Headers used to pass browser ID information along with browser request packets. If the headers are included, Web servers can determine the origin of browser requests they receive. If the headers are not included, browser requests have anonymity.

Checking the X-Forwarded-For option causes the appliance to either add information to an existing X-Forwarded-For or Forwarded-For header, or to create a header if one doesn't already exist.

Leaving the option unchecked causes the appliance to remove X-Forwarded-For headers from any Web accelerator requests passing through the appliance.

You must weigh the desires of browser users to remain anonymous against the desires of Web server owners (e-commerce sites, for example) to collect data about who is accessing their site.

Return Error If Hostname Sent by Browser Does Not Match Accelerator DNS Hostname: Checking this option causes Accelerator to match the hostname in the DNS header that came from

the browser against the DNS name specified in this accelerator definition. If the names don't match, the request is not forwarded to the Web server. Instead, Excelsator returns an error to the requesting browser.

Enable Authentication: Checking this box causes the appliance to require authentication of users wanting to use its Web server accelerator services. Clicking Authentication Options displays the Add Authentication Profiles dialog box. For more information, see [“Add Authentication Profiles Dialog Box” on page 338](#).

Enable Secure Excelsator: If you have installed Secure Excelsator on this cache device, you can check this option to enable the service to work with Secure Excelsator. Click Secure Excelsator Options to configure the Secure Excelsator options. For more information, see the [Volera Secure Excelsator 1.1 Administration Guide](#).

SSL Listening Port: The port on which the appliance listens for authentication requests.

IMPORTANT: Each service (including authentication) must use a unique IP address and port combination. The default authentication port is 443. Attempts to enable authentication for more than one service on the same IP address and port will result in a TCP bind error.

Certificate: This drop-down list displays any certificate you have stored on your appliance. System-generated certificates do not appear in the list.

Use this field to select the certificate you created specifically for the Web server accelerator you are creating. This will prevent browsers from receiving certificate confirmation messages each time they access the appliance. For more information, see [“Managing Appliance Certificates” on page 165](#).

Enable Path-Based Multihoming: Allows you to create child accelerators for path-based multihoming configurations. This option is enabled only when you have created another accelerator definition and have not created a standard multihoming relationship between previously defined accelerators on the appliance. In other words, it is enabled if you don't have multiple accelerators sharing the same accelerator IP address and port.

When you enable path-based multihoming for the accelerator you are defining, you must also click the Path-Based Options button and specify a path rule that the multihoming master can use to route traffic to the accelerator you are defining.

You will also notice that if you have created multiple accelerators that can function as multihoming masters, when you select a name in the Multihoming Master drop-down list, the DNS Name, Accelerator Proxy Port, and Accelerator IP Addresses selections dynamically change to match the accelerator whose name you have selected. For more information regarding path-based multihoming, see [“Multihoming and Path-Based Support” on page 123](#).

Multihome Master: This drop-down list contains the names of accelerators you have defined that can function as multihoming masters, meaning they are not configured as child accelerators to other multihoming masters.

Enable Custom Cache Control Header: Checking this option and clicking Header Options provides access to the Custom Cache Control Header dialog box.

For more information, see [“Custom Cache Control Header Dialog Box” on page 340](#).

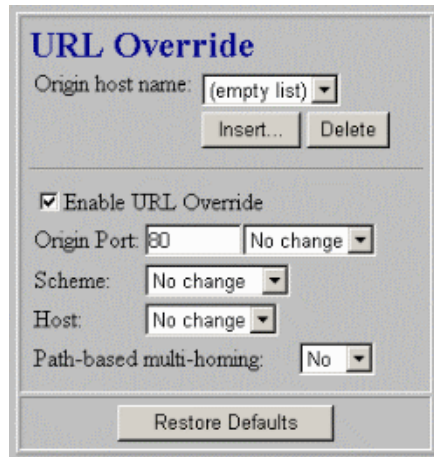
Enable Access Control: . For details, see [“Access Control Options Dialog Box” on page 341](#).

Advanced Options: Lets you control the TCP receive window size for cache device fills from origin Web servers, the caching of objects on the cache device that would not normally be cached, and the filling and vending of browser no-cache requests by the cache device. For details, see [“Advanced Options \(Tuning\) Dialog Box” on page 341](#).

URL Override Dialog Box

Path: Cache > Web Server Accelerator > Insert > URL Override

Figure 109



This dialog box lets you create additional URL overrides for the accelerator service you are configuring. For more information regarding URL overrides, see [Chapter 25, “Transforming Content for Internet Delivery,” on page 173](#).

Origin Host Name: The internal DNS name of the origin Web server for which you are creating the override.

IMPORTANT: The Origin Host Name cannot be a DNS name listed in the Web Server Addresses list of the service. Automatic DNS name overrides are already created for these names. However, if the entry in the Web Server Addresses list is an IP address, this name could be the internal DNS name associated with that address since an automatic DNS name override would not be created for the address entry.

Enable URL Override: Selecting this option enables the override for the service.

Origin Port: The internal port number you want to change.

Scheme: You use this option only in connection with Secure Excelerator. For more information see, the [Volera Secure Excelerator 1.1 Administration Guide](#) and [“Understanding URL Overrides” on page 173](#).

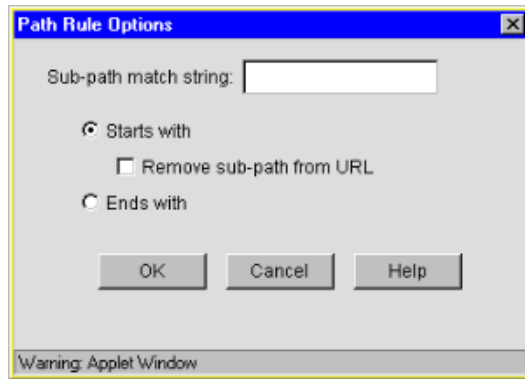
Host: You select this option to create additional DNS name overrides. For more information, see [“Understanding URL Overrides” on page 173](#).

Path-Based Multihoming: You use this option to create additional child accelerator service overrides. For more information, see [“Understanding URL Overrides” on page 173](#).

Path Rule Options Dialog Box

Path: Cache > Web Server Accelerator > Insert > Enable Path-Based Multihoming > Path Rule Options

Figure 110



This dialog box lets you specify a string that, if present in the browser request, will cause the multihoming master accelerator to route the request to the child accelerator being defined.

The string match can occur immediately following the DNS name (the Starts With option) or at the end of the URL (the Ends With option).

Sub-Path Match String: The string the multihoming master will compare against the browser request. If the string is not found at the beginning of the URL path, the multihoming master accelerator attempts to fill the request through the Web server addresses in its accelerator definition. If the string is found, the multihoming master accelerator routes the request to the child accelerator defined for the matching string. For more information on path-based multihoming, see [“Multihoming and Path-Based Support” on page 123](#). For more information on creating additional match strings for other child accelerators whose objects references also need to be transformed, see [“Understanding URL Overrides” on page 173](#).

Starts With: When checked, this option indicates to the multihoming master that the Sub-Path Match String field contains a path that might immediately follow the DNS name in the browser request. If the string matches the path in the request, the multihoming master forwards the request to the child accelerator being defined.

IMPORTANT: The initial forward slash (/) must not be included in the string because Excelerator automatically inserts it into the rewritten path.

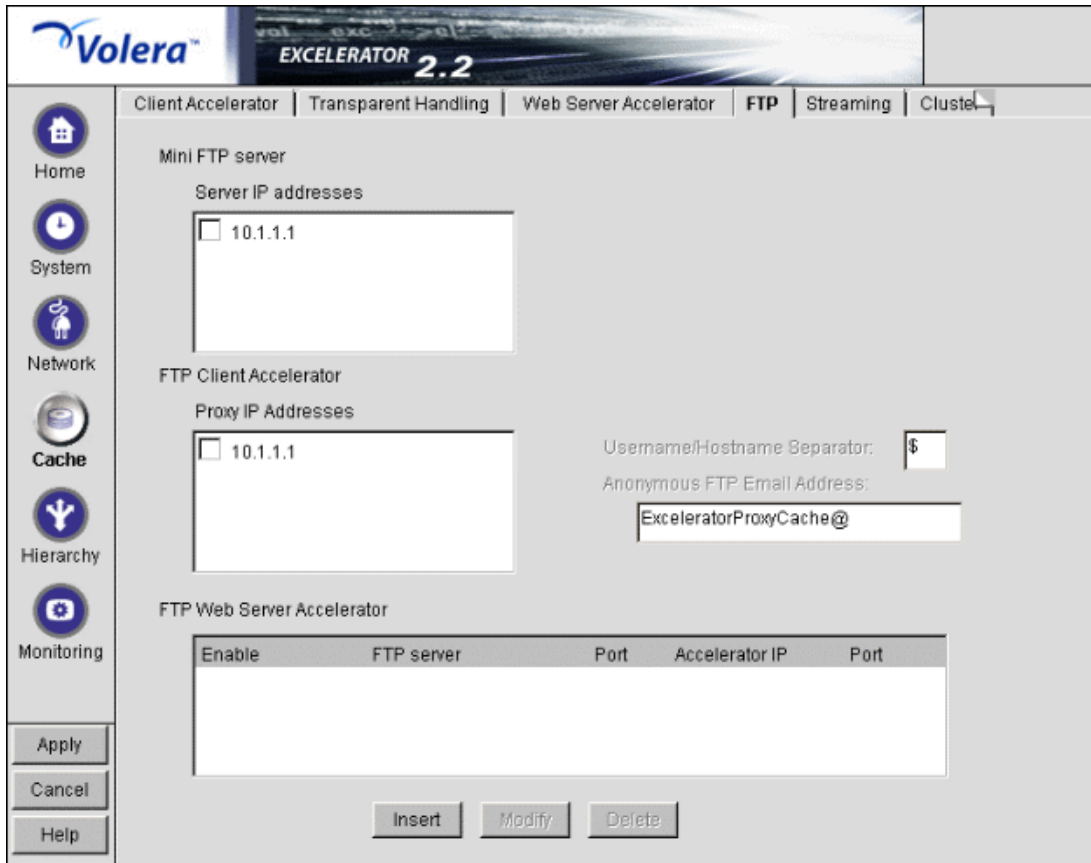
Remove Sub-Path from URL: Check this option if the path string doesn’t actually appear at the root of the Web server. If this option is checked, the string is stripped from the request before the request is sent to the Web server. This probably indicates that the object is at the root of the Web server. If this option is not checked, the matched string is retained in the request sent to the Web server.

Ends With: When checked, this option indicates to the multihoming master accelerator that the Sub-Path Match String field contains a file extension, such as gif, jpg, mpg, or cgi. If a match is found at the end of the browser request, the multihoming master will route the request to the child accelerator being defined.

FTP Tab

Path: Cache > FTP

Figure 111



The FTP tab lets you configure the appliance to provide an FTP listening address (Mini FTP Server) for appliance management, FTP forward proxy (client acceleration), and FTP reverse proxy (Web server acceleration) services.

IMPORTANT: All appliance FTP services listen on port 21. Therefore, each IP address on the appliance can be configured for only one FTP service (Mini FTP, Forward Proxy, or FTP Accelerator).

Mini FTP Server: Check an address in the Server IP Addresses list to enable it for FTP listening. If an address is not checked, you cannot upload or download files using FTP.

FTP Client Accelerator: Check an address in the Proxy IP Addresses list to enable FTP forward proxy services on that address.

- ♦ *Username/Hostname Separator:* Use this character on the FTP command line to indicate where the username ends and the hostname begins.
 - ♦ Example: From a DOS window using native FTP mode (an FTP/FTP, client to an appliance using the FTP protocol, or an appliance to the origin FTP host using the FTP protocol), enter

```
syntax: ftp appliance_proxy_address
username: anonymous$ftp_host
password:
```
 - ♦ Example: From a browser using native FTP mode (FTP/FTP), enter

```
ftp://anonymous$ftp_host@appliance_proxy_address
```


- ♦ Example: From a browser configured to route FTP traffic to an appliance via HTTP (HTTP/FTP) (the appliance retrieves data from the origin FTP host using the FTP protocol), enter

```
ftp://anonymous@ftp_host
```

- ♦ Example: Non-anonymous FTP/FTP can also be used. From a DOS window, enter:

```
syntax: ftp appliance_proxy_address
```

```
username: username$ftp_host
```

```
password:
```

IMPORTANT: The default separator is the dollar sign (\$). This character should be used whenever possible. The pound sign (#) and the at sign (@) are also known to work as separators.

Although the browser-based tool will let you specify other characters such as the comma (,), the percent sign (%), and the colon (:), these characters will not work as FTP separators, and they will not appear in command line interface queries.

If you are using an FTP client, you might need to configure the client to recognize the dollar sign (\$).

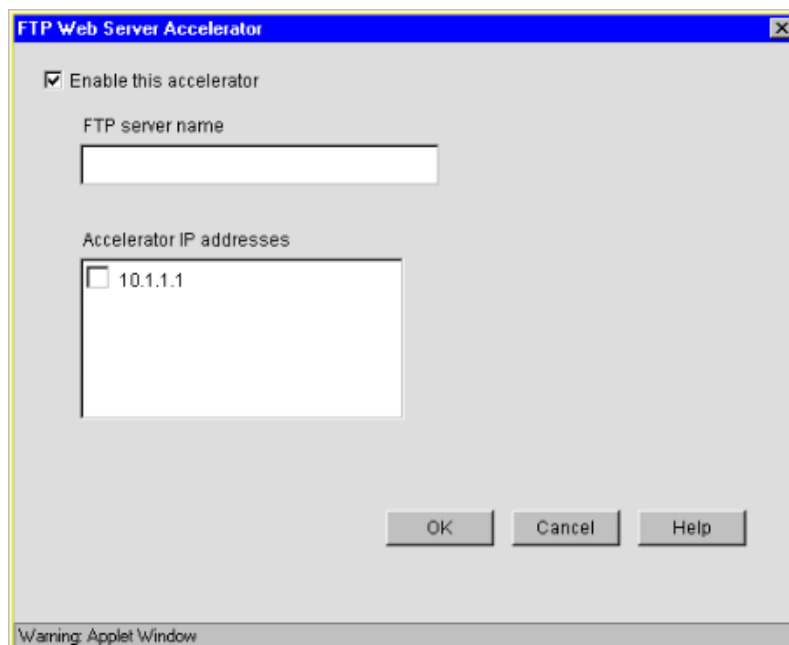
- ♦ *Anonymous FTP Email Address:* When the appliance's FTP server prompts you for a password, if you enter a blank, this e-mail address will be used as the password.

FTP Web Server Accelerator: Summary of the FTP accelerators created in the FTP Accelerator dialog box. You can only enable or disable the service represented by an item in the list. To modify a list item, click the item > click Modify.

FTP Accelerator Dialog Box

Path: Cache > FTP > Insert under FTP Accelerator list

Figure 112



The FTP Accelerator dialog box lets you define an FTP accelerator service on the appliance.

Enable This Accelerator: Activates the defined FTP accelerator service.

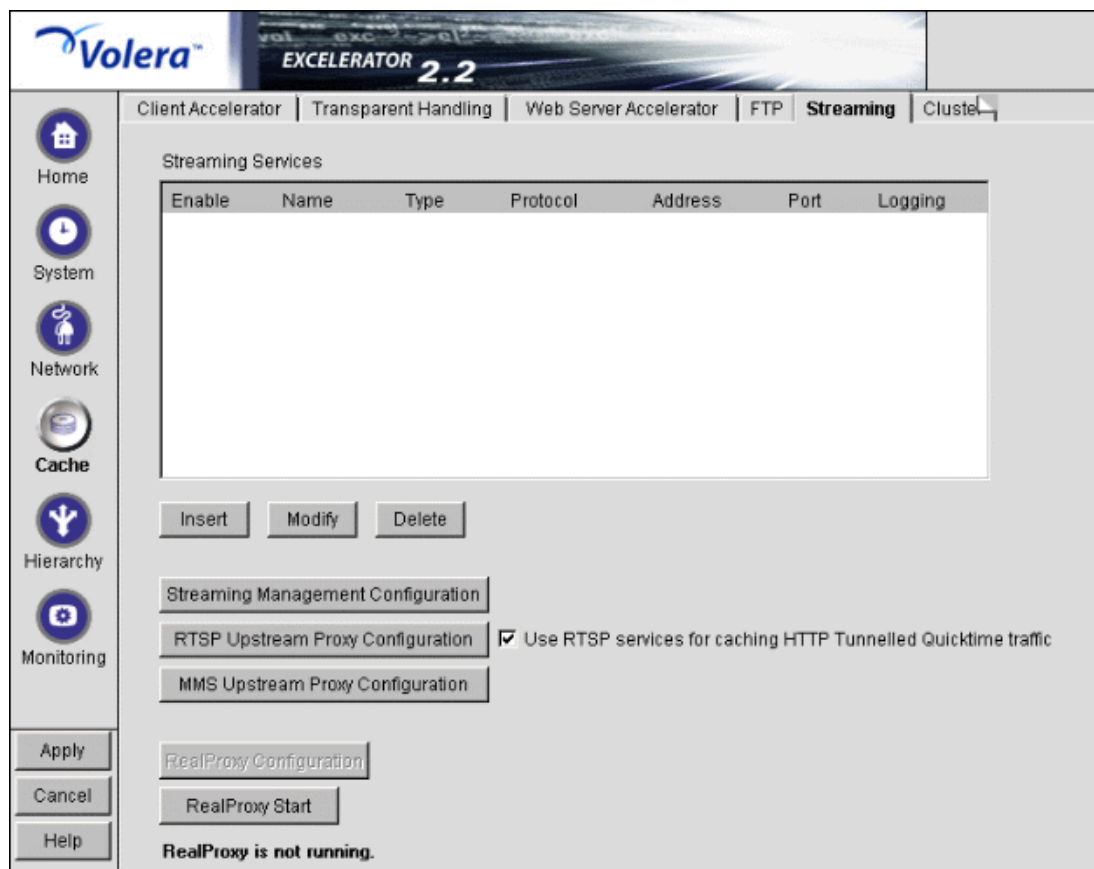
FTP Server Name: Identifies either the IP address or DNS name of the FTP server being accelerated.

Accelerator IP Addresses: Lists all appliance IP addresses. Checking the box to the left of an address enables FTP acceleration through the address for the server listed in the FTP Server Name field.

Streaming Tab

Path: Cache > Streaming

Figure 113



The Streaming tab lets you configure the appliance with both accelerator media services for Web servers and forward media services for browsers.

Streaming Services: Lists the streaming media services you have defined for (inserted into) the appliance.

Streaming Management Configuration: Lets you access the Streaming Management Configuration dialog box, in which you can set various parameters for tuning bandwidth management, storage management, streaming cache freshness, stream filling parameters, and

RTSP protocol settings. For more information, see [“Streaming Management Configuration Dialog Box” on page 358](#).

RTSP Upstream Proxy Configuration: Checking this box and clicking Upstream Proxy lets you specify another proxy server to which the appliance should look for filling RTSP requests. The relationship created between the appliance and the other server is similar to CERN hierarchical relationships that service HTTP requests. However, only one parent is supported.

MMS Upstream Proxy Configuration: Checking this box and clicking Upstream Proxy lets you specify another proxy server to which the appliance should look for filling MMS and streaming HTTP requests. The relationship created between the appliance and the other server is similar to CERN hierarchical relationships that service HTTP requests. However, only one parent is supported.

RealProxy Configuration: This button provides access to the RealProxy configuration dialogs. It is only active when the RealProxy service is running.

RealProxy Start: Clicking this option starts the RealProxy service. License requirements and CLI setup instructions are contained in the [Volera Media Excelerator 1.2 for RealSystem Proxy 8 Startup Guide](#).

Streaming Management Configuration Dialog Box

Path: Cache > Streaming > Streaming Management Configuration

Figure 114

Bandwidth Management		MMS	RTSP	
Max bandwidth per stream			0	kbps
Max upstream bandwidth			0	kbps
Max downstream bandwidth	0		0	kbps

Storage Management		RTSP Cache Freshness	
Max object size	0 Mbytes	Max TTL	48 hours
Max object duration	0 hours	Min TTL	3600 seconds
Max live object duration	5 seconds	Default TTL	24 hours
Max disk usage	0 Gbytes		

RTSP Stream Filling		RTSP Protocol Settings	
Continue fill time	0 minutes	RTSP idle timeout	60 seconds
Max sessions	10000	RTSP keep alive timeout	45 seconds

Reset OK Cancel Help

Warning: Applet Window

The Streaming Management Configuration dialog box lets you specify how the caching system handles streaming media objects.

Bandwidth Management

Max Bandwidth Per Stream: The maximum bandwidth that SMC allows a single stream to consume. Streams that require more bandwidth are rejected. Valid field values are 0 through 10,000 Kbps. The default value is 0, which disables single stream bandwidth restrictions.

Max Upstream Bandwidth: The total bandwidth allowed for all streams coming from origin Web servers. All requests that would cause this value to be exceeded are denied. Valid field values are 0 through 1,000,000 Kbps. The default value is 0, which disables total upstream bandwidth restrictions.

Max Downstream Bandwidth: The total bandwidth allowed for all streams flowing to client players. All requests that would cause this value to be exceeded are denied. Valid field values are 0 through 1,000,000 Kbps. The default value is 0, which disables total downstream bandwidth restrictions.

IMPORTANT: You must set this parameter to match the capacity of your network hardware.

The appliance's high-performance capabilities can easily cause streaming media output to exceed the capacity of the network and the appliance's network cards. The resulting backlog can result in lost packets and other more severe problems, such as an appliance system crash.

Storage Management

Max Object Size: The largest non-live streaming object that the appliance will cache. The appliance will pass through but will not cache non-live objects that exceed this value.

This value does not apply to live streaming objects.

Setting the value to 0 disables object size restrictions. However, objects will not exceed the system-calculated largest cachable object size (see the Note below) nor will this value cause the Max Disk Usage value to be exceeded. The default field value is 0.

NOTE: This field value does not apply to tunneled content, which Excelsior treats as HTTP objects. The maximum object size for HTTP objects is equal to approximately one fourth of the smallest disk drive in the appliance. This means that if the smallest drive is 18 GB, the largest cached object size allowed by the system would be 4.5 GB.

Max Object Duration: The playing-time threshold beyond which SMC will not cache a non-live streaming object.

When this threshold is reached, SMC continues to pass through the stream, but any streaming object whose playing time exceeds this limit is removed from cache. In other words, SMC does not cache objects whose total playing time exceeds this value.

When determining whether to continue caching an object, SMC considers the Max Object Size and Max Object Duration values and uses the more restrictive of the two.

This value does not apply to live streaming objects.

Setting the value to 0 disables playing time restrictions on object cacheability. However, the same restrictions explained under Max Object Size still apply. Valid field values are 0 through 24 hours. The default value is 0.

Max Live Object Duration: For future use. Not currently implemented.

Max Disk Usage: The total amount of appliance disk space that can be used for caching streaming objects. Valid field values are 0 through 10,000 MB. The default value is 0, which means there are no disk usage restrictions that apply to streaming objects.

RTSP Stream Filling

Continue Fill Time: The length of time that Excelsator will continue to fill a streaming request after there are no consumers for the stream. Valid values range from 0 through 600 seconds. The default value is 0, which stops filling immediately.

This value applies only when No Consumer Abort Fill is set to Yes.

Max Sessions: The total number of user sessions that SMC will service. SMC denies requests beyond this value.

Valid field values are 1 through 10,000. The default is 10,000.

RTSP Cache Freshness

Max TTL: The maximum number of hours that SMC will serve a streaming media object from cache before refilling from the origin Web server. No streaming media object is served from cache after this value expires without the object being refilled first.

This value overrides the freshness or Time to Expire header value specified by the Webmaster if he or she specified a longer time.

Use this field to reduce the maximum time that SMC waits before refilling requested objects. Valid values range from 1 through 168 hours. The default value is 48.

NOTE: This does not apply to tunneled content. The configured limits on HTTP freshness apply to all HTTP objects, including tunneled streaming content. For more information, see the [“Cache Freshness Dialog Box” on page 401](#).

Min TTL: The minimum number of seconds that SMC will serve a streaming media object from cache before refilling it from the origin Web server. No requested object will be refilled sooner than specified by this value.

This overrides the freshness or Time to Expire header value specified by the Webmaster if he or she specified a shorter time.

Use this field to increase the minimum time Excelsator waits before refilling requested objects. This parameter does not override No Cache or Must Revalidate directives from the origin Web server. Valid values range from 0 through 86,400 seconds (24 hours). A value of 0 means that Excelsator will use the value assigned the object by the Webmaster. The default value is 3600, which means Excelsator will not request that an object be refilled sooner than one hour.

NOTE: This does not apply to tunneled content. The configured limits on HTTP freshness apply to all HTTP objects, including tunneled streaming content. For more information, see the [“Cache Freshness Dialog Box” on page 401](#).

Default TTL: The number of hours SMC waits before refilling requested streaming media objects for which Webmasters have not specified a freshness or Time to Expire header value.

NOTE: This does not apply to tunneled content. The configured limits on HTTP freshness apply to all HTTP objects, including tunneled streaming content. For more information, see the [“Cache Freshness Dialog Box” on page 401](#).

RTSP Protocol Settings

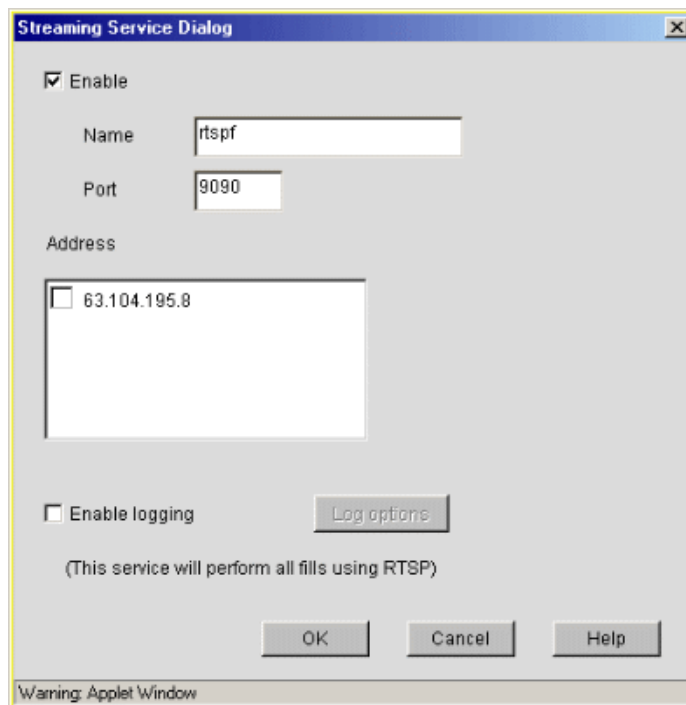
RTSP Idle Timeout: Do not change this value. If the media player doesn't send any messages for this period of time, the session is terminated. Valid field values are 10 through 600 seconds. The default value is 60 seconds.

RTSP Keep Alive Timeout: Do not change this value. This is the interval at which the system sends messages to the origin Web servers to keep streaming media sessions alive. Valid field values are 10 through 600 seconds. The default value is 45 seconds.

Forward Streaming Services (RTSP)

Path: Cache > click Streaming > click Insert > set Type = Forward and Protocol = RTSP > click OK

Figure 115



The Forward Streaming Service dialog box lets you set up a forward proxy service for streaming media on the appliance. After the service is defined, RTSP/RTP-compatible media players can designate the appliance as their streaming proxy server. Any appliance IP addresses can be used. Enabling specific addresses is not required for forward streaming services.

Enable: Enables the defined service to provide forward proxy services for streaming media content.

Name: Name you assign to the streaming media accelerator. For example, you can select a name that indicates the group of client players that will be accessing the service or another name that helps you differentiate between different forward services. A valid name is a DOS-style, eight-character, alphanumeric string with no special characters or spaces.

Address: Lists IP addresses of available servers.

Port: The port from which the appliance will receive forward proxy requests for streaming media content and on which it will send the content back to the client players. The default port value is 9090. Valid port numbers range from 1 through 65535.

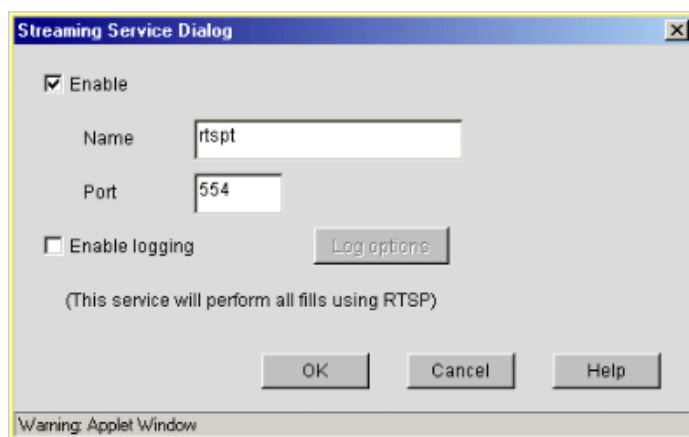
Enable Logging: Enables logging of events associated with the service.

Log Options: Lets you specify the information that is logged, how often new log files are started, and how long log files are retained. See the “[Streaming Media Log Options Dialog Box \(RTSP\)](#)” on page 364.

Transparent Streaming Service Dialog Box (RTSP)

Path: Cache > click Streaming > click Insert > set Type = Transparent and Protocol = RTSP > click OK

Figure 116



The Transparent Streaming Service dialog box lets you set up a transparent proxy service for streaming media on the appliance. After the service is defined, RTSP/RTP requests can be automatically routed to the appliance by a router or switch on your network. Any appliance IP addresses can be used. Enabling specific addresses is not required for transparent streaming services.

Enable: Enables the defined service to provide transparent proxy services for streaming media content.

Name: Name you assign to each streaming media accelerator. For example, you can select a name that indicates the location of the client players that will be accessing the service or another name that helps you differentiate between different transparent services. A valid name is a DOS-style, eight-character, alphanumeric string with no special characters or spaces.

Port: The port from which the appliance will receive transparent proxy requests for streaming media content and on which it will send the content back to the client players. The default port value is 554. Valid port numbers range from 1 through 65535.

Enable Logging: Enables logging of events associated with the service.

Log Options: Lets you specify the information that is logged, how often new log files are started, and how long log files are retained. See the “[Streaming Media Log Options Dialog Box \(RTSP\)](#)” on page 364.

Reverse Streaming Service Dialog Box (RTSP)

Path: Cache > click Streaming > click Insert > set Type = Reverse and Protocol = RTSP > click OK

Figure 117

The image shows a 'Streaming Service Dialog' window. At the top, there is a title bar with the text 'Streaming Service Dialog' and a close button. Below the title bar, there is a section with a checked checkbox labeled 'Enable'. Underneath, there are two text input fields: 'Name' with the value 'rtspa' and 'DNS name' which is empty. Below these are two more text input fields: 'Streaming server port' with the value '554' and 'Accelerator port' with the value '554'. Below these are two text area fields: 'Streaming server addresses' which is empty, and 'Accelerator addresses' which contains the IP address '63.104.195.8'. Below the text area fields are two buttons: 'Insert' and 'Delete'. Below these are two checkboxes: 'Enable logging' (unchecked) and 'Log options' (disabled). Below the checkboxes is a note: '(This service will perform all fills using RTSP)'. At the bottom of the dialog are three buttons: 'OK', 'Cancel', and 'Help'. At the very bottom of the window, there is a warning message: 'Warning: Applet Window'.

The Reverse Streaming Service dialog box lets you create accelerator services for handling streaming media requests to origin streaming servers.

Enable: Specifies whether the service is enabled.

Name: Name you assign to each streaming media accelerator. For example, you can select a name that indicates the streaming media source being serviced by the appliance. Alternatively, you might choose a name that matches the clients (players) that will be accessing the service. A valid name is a DOS-style, eight-character, alphanumeric string with no special characters or spaces.

DNS Name: The DNS name of the streaming server you are accelerating.

Streaming Server Port: The port number that the SMC accelerator service is listening on for incoming connections. The default is 554. The valid port range is 1 through 65535.

Accelerator Port: The port number that the origin streaming server is listening on for incoming streaming media connections. The default is 554. The valid port range is 1 through 65535.

Streaming Server Address: The appliance's IP addresses to which DNS resolves the Web site's DNS name and on which the SMC accelerator service listens for incoming connections from the Internet.

Accelerator Address: The IP address of the streaming server from which the appliance fills the cache for this SMC accelerator service.

Enabling Logging: Enables logging of SMC-related events associated with the service.

Log Options: Lets you specify the information that is logged, how often new log files are started, and how long log files are retained. See the “[Streaming Media Log Options Dialog Box \(RTSP\)](#)” on page 364.

Streaming Media Log Options Dialog Box (RTSP)

Path: Cache > click Streaming > click Insert > set Type > set Protocol = RTSP > click OK > check Enable Logging > click Log Options

Figure 118

The screenshot shows the 'Log Options' dialog box. Under 'Extended log fields', 'Date', 'Time', and 'Client IP' are checked. Under 'Rollover options', the first option is selected with a value of 10 MB. Under 'Old file options', the second option is selected with a value of 168 hours. The 'Log Push' button is visible, along with 'OK' and 'Cancel' buttons at the bottom right. A warning bar at the bottom left indicates 'Warning: Applet Window'.

The Streaming Media Log Options dialog box lets you specify the SMC-specific information that is logged, how often new log files are started, and how long log files are retained.

Extended Log Fields: These options let you specify which SMC-specific information is logged. Fields that are inactive are not selectable. The following are brief explanations of each log field:

- ♦ *Date:* The appliance date the streaming media data was requested.
- ♦ *Time:* The appliance time when the streaming media data was requested.
- ♦ *Client IP:* The IP address of the requesting player.
- ♦ *Server IP:* The IP address of the origin Web server containing the streaming object.
- ♦ *Method:* The RTSP method that the client player sent to SMC.
- ♦ *URL:* The RTSP URL that the client player sent to SMC.
- ♦ *RTSP Version:* The RTSP version specified in the URL that the client player sent to SMC.
- ♦ *RTSP Status:* The RTSP status code that SMC sent to the client player.
- ♦ *Cached Status:* Whether the object requested in the URL was retrieved from appliance cache.

Rollover Options: These options let you specify the method the appliance uses to determine when to start new log files. These fields should be set as you plan your appliance’s logging strategy (see [“Using Appliance Logging Services” on page 237](#)).

- ♦ *Rollover When File Size Reaches (in MB):* Starts new log files each time the size limit you specify is reached. The default is 10 MB.
- ♦ *Rollover Every:* Starts new log files for each time increment (days or hours) you specify. Logging begins on the day and time you specify using either appliance or GMT time.

Old File Options: These options let you specify the automatic disposition of older files. Because the disk space available to log files is limited and can become full fairly quickly, setting these options is an important part of defining your appliance logging strategy. See [“Using Appliance Logging Services” on page 237](#).


- ♦ *Limit Number of Files To:* After this limit is reached, the oldest file is deleted each time a new file is created.
- ♦ *Delete Files Older Than:* Excelsior automatically deletes files when they are older than you specify.
- ♦ *Do Not Delete:* This option is not normally recommended because it can lead to a disk full condition. For more information, see [“Using Appliance Logging Services” on page 237](#).

Log Push: For more information, see [“FTP Log Push Configuration Dialog Box” on page 336](#)

Upstream Proxy Dialog Box (RTSP)

Path: Cache > Streaming > RTSP Upstream Proxy Configuration

Figure 119



The Upstream Proxy dialog box lets you designate another proxy server as a hierarchical parent to which the appliance should look for filling RTSP requests. The relationship created between the appliance and the other server is similar to the CERN hierarchical relationship for HTTP requests.

IMPORTANT: Although the dialog box allows you to insert multiple IP addresses, only one parent is supported.

Upstream Proxy Enable: Causes the service to request fills through the Upstream Proxy Address specified.

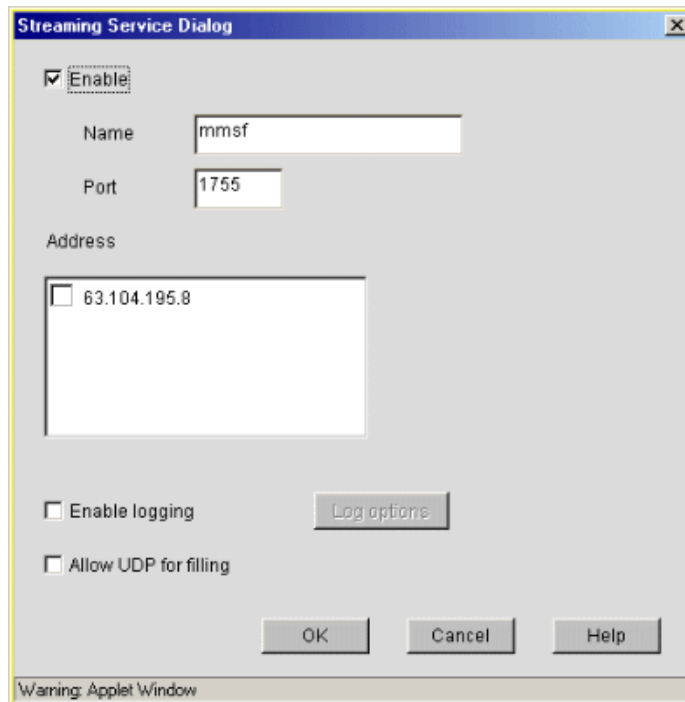
RTSP Port: The port on which the service will request the RTSP streaming objects.

Upstream Proxy Address: The streaming proxy through which the service will request the RTSP streaming object. (You can specify only one upstream proxy per service.)

Forward Streaming Services (MMS)

Path: Cache > click Streaming > click Insert > set Type = Forward and Protocol = MMS > click OK

Figure 120



The Forward Streaming Service dialog box lets you set up a forward proxy service for streaming media on the appliance. After the service is defined, Windows Media players can designate the appliance as their streaming proxy server. Any appliance IP addresses can be used. Enabling specific addresses is not required for forward streaming services.

Enable: Enables the defined service to provide forward proxy services for streaming media content.

Name: Name you assign to the streaming media accelerator. For example, you can select a name that indicates the group of client players that will be accessing the service or another name that helps you differentiate between different forward services. A valid name is a DOS-style, eight-character, alphanumeric string with no special characters or spaces.

Address: Lists IP addresses of available servers.

Port: The port from which the appliance will receive forward proxy requests for streaming media content and on which it will send the content back to the client players. The default port value is 1755. Valid port numbers range from 1 through 65535.

Enable Logging: Enables logging of events associated with the service.

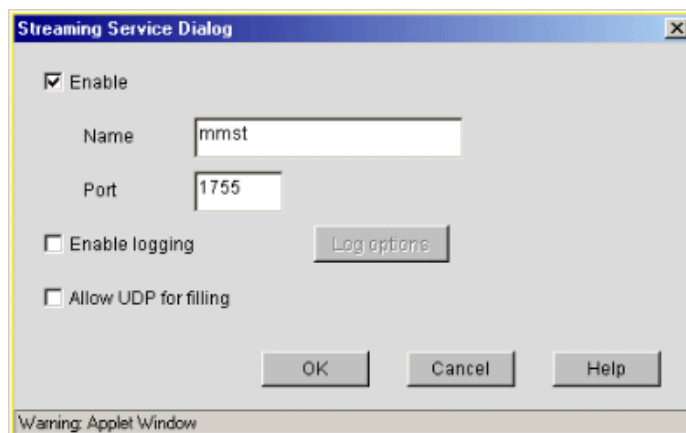
Log Options: Lets you specify the information that is logged, how often new log files are started, and how long log files are retained. See the “[Streaming Media Log Options Dialog Box \(MMS\)](#)” on page 369.

Allow UDP for Filling: To ensure quality of cached streams, Excelerator fills all MMS stream requests using TCP connections by default, even when the initial request is for UDP. Checking this option causes Excelerator to fill MMS in UDP stream requests associated with this service using UDP.

Transparent Streaming Service Dialog Box (MMS)

Path: Cache > click Streaming > click Insert > set Type = Transparent and Protocol = MMS > click OK

Figure 121



The Transparent Streaming Service dialog box lets you set up a transparent proxy service for streaming media on the appliance. After the service is defined, MMS requests can be automatically routed to the appliance by a router or switch on your network. Any appliance IP addresses can be used. Enabling specific addresses is not required for transparent streaming services.

Enable: Enables the defined service to provide transparent proxy services for streaming media content.

Name: Name you assign to each streaming media accelerator. For example, you can select a name that indicates the location of the client players that will be accessing the service or another name that helps you differentiate between different transparent services. A valid name is a DOS-style, eight-character, alphanumeric string with no special characters or spaces.

Port: The port from which the appliance will receive transparent proxy requests for streaming media content and on which it will send the content back to the client players. The default port value is 1755. Valid port numbers range from 1 through 65535.

Enable Logging: Enables logging of events associated with the service.

Log Options: Lets you specify the information that is logged, how often new log files are started, and how long log files are retained. See the “[Streaming Media Log Options Dialog Box \(MMS\)](#)” on page 369.

Allow UDP for Filling: To ensure quality of cached streams, Excelerator fills all MMS stream requests using TCP connections by default, even when the initial request is for UDP. Checking this option causes Excelerator to fill MMS in UDP stream requests associated with this service using UDP.

Reverse Streaming Service Dialog Box (MMS)

Path: Cache > click Streaming > click Insert > set Type = Reverse and Protocol = MMS > click OK

Figure 122

The screenshot shows the 'Streaming Service Dialog' window. It has a title bar with 'Streaming Service Dialog' and a close button. The main area contains the following elements:

- ☒ Enable
- Name:
- DNS name:
- Streaming server port:
- Accelerator port:
- Streaming server addresses:
- Accelerator addresses:

☐ 63.104.195.8
- Buttons: Insert, Delete
- ☐ Enable logging
- Log options (button)
- ☐ Allow UDP for filling
- Buttons: OK, Cancel, Help

Warning: Applet Window

The Reverse Streaming Service dialog box lets you create accelerator services for handling streaming media requests to Windows Media servers.

Enable: Specifies whether the service is enabled.

Name: Name you assign to each streaming media accelerator. For example, you can select a name that indicates the streaming media source being serviced by the appliance. Alternatively, you might choose a name that matches the clients (players) that will be accessing the service. A valid name is a DOS-style, eight-character, alphanumeric string with no special characters or spaces.

DNS Name: The DNS name of the streaming server you are accelerating.

Streaming Server Port: The port number that the SMC accelerator service is listening on for incoming connections. The default is 1755. The valid port range is 1 through 65535.

Accelerator Port: The port number that the origin streaming server is listening on for incoming streaming media connections. The default is 1755. The valid port range is 1 through 65535.

Streaming Server Address: The appliance's IP addresses to which DNS resolves the Web site's DNS name and on which the MMS accelerator service listens for incoming connections from the Internet.

Accelerator Address: The IP address of the streaming server from which the appliance fills the cache for this MMS accelerator service.

Enabling Logging: Enables logging of events associated with the service.

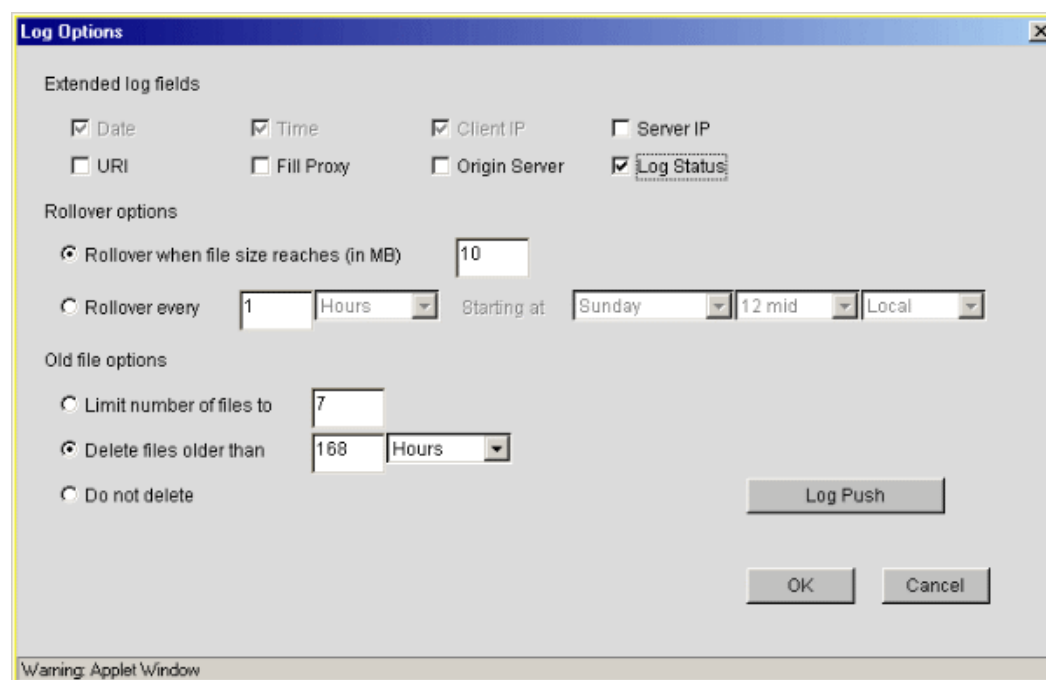
Log Options: Lets you specify the information that is logged, how often new log files are started, and how long log files are retained. See the [“Streaming Media Log Options Dialog Box \(MMS\)” on page 369](#).

Allow UDP for Filling: To ensure quality of cached streams, Exceleator fills all MMS stream requests using TCP connections by default, even when the initial request is for UDP. Checking this option causes Exceleator to fill MMS in UDP stream requests associated with this service using UDP.

Streaming Media Log Options Dialog Box (MMS)

Path: Cache > click Streaming > click Insert > set Type > set Protocol = MMS > click OK > check Enable Logging > click Log Options

Figure 123



The Streaming Media Log Options dialog box lets you specify the SMC-specific information that is logged, how often new log files are started, and how long log files are retained.

Extended Log Fields: These options let you specify which MMS information is logged. Fields that are inactive are not selectable. The following are brief explanations of each log field:

- ◆ *Date*: The appliance date the streaming media data was requested.
- ◆ *Time*: The appliance time when the streaming media data was requested.
- ◆ *Client IP*: The IP address of the requesting player.
- ◆ *Server IP*: The IP address of the origin server specified in the original player request. For forward and transparent requests, this is the IP address of the origin server. For requests to a reverse accelerator, the IP address actually belongs to the cache device, not the origin streaming server.
- ◆ *URL*: The URL that the client player sent to the service.
- ◆ *Fill Proxy*: The address of the upstream proxy (if applicable).
- ◆ *Origin Server*: The IP address of the streaming server containing the original streaming object.
- ◆ *Logged Status*: The HTTP status code associated with the reply (if applicable).

Rollover Options: These options let you specify the method the appliance uses to determine when to start new log files. These fields should be set as you plan your appliance's logging strategy (see [“Using Appliance Logging Services” on page 237](#)).

- ◆ *Rollover When File Size Reaches (in MB)*: Starts new log files each time the size limit you specify is reached. The default is 10 MB.
- ◆ *Rollover Every*: Starts new log files for each time increment (days or hours) you specify. Logging begins on the day and time you specify using either appliance or GMT time.

Old File Options: These options let you specify the automatic disposition of older files. Because the disk space available to log files is limited and can become full fairly quickly, setting these options is an important part of defining your appliance logging strategy. See [“Using Appliance Logging Services” on page 237](#).


- ◆ *Limit Number of Files To*: After this limit is reached, the oldest file is deleted each time a new file is created.
- ◆ *Delete Files Older Than*: Excelsior automatically deletes files when they are older than you specify.
- ◆ *Do Not Delete*: This option is not normally recommended because it can lead to a disk full condition. For more information, see [“Using Appliance Logging Services” on page 237](#).

Log Push: For more information, see [“FTP Log Push Configuration Dialog Box” on page 336](#)

Upstream Proxy Dialog Box (MMS)

Path: Cache > Streaming > MMS Upstream Proxy Configuration

Figure 124



The Upstream Proxy dialog box lets you designate another proxy server as a hierarchical parent to which the appliance should look for filling MMS requests. The relationship created between the appliance and the other server is similar to the CERN hierarchical relationship for HTTP requests.

IMPORTANT: Although the dialog box allows you to insert multiple IP addresses, only one parent is supported.

Upstream Proxy Enable: Causes the service to request fills through the Upstream Proxy Address specified.

MMS Port: The port on which the service will request the MMS streaming objects.

HTTP Port: The port on which the service will request the MMS in HTTP streaming objects.

Upstream Proxy Address: The streaming proxy through which the service will request the streaming objects. (You can specify only one upstream proxy per service.)

Cluster Tab

Path: Cache > Cluster

Figure 125

The screenshot displays the Volera EXCELERATOR 2.2 software interface. The top navigation bar includes tabs for Client Accelerator, Transparent Handling, Web Server Accelerator, FTP, Streaming, and Cluster. The Cluster tab is selected. On the left, a sidebar contains icons for Home, System, Network, Cache, Hierarchy, and Monitoring. The main content area for the Cluster tab includes an 'Enable cluster' checkbox, a 'Name' text field containing 'Cluster', and a 'Subnet' dropdown menu showing '63.104.192.0'. Below these are two empty tables: 'Servers' with headers 'Name', 'IP address', 'Role', and 'Capacity'; and 'Services' with headers 'Name' and 'Type'. At the bottom of the main area are buttons for 'Insert', 'Delete', 'Forward', 'Transparent', 'Accelerator', 'Modify', and 'Delete'. A vertical sidebar on the far left contains buttons for 'Apply', 'Cancel', and 'Help'.

The Cluster tab lets you set up a group of appliances to handle one or more forward, transparent, or reverse proxy services. The cluster of appliances provides fail-over support for all defined services. Distribution of services among the appliances is automatically negotiated within the cluster. If one of the appliances shuts down for any reason, the services that were previously assigned to it are redistributed among the remaining appliances.

Enable Cluster: Enables clustering.

Name: Name of cluster. Supply a descriptive name for the cluster. This can be any text string.

Subnet: Drop-down list containing all the subnets currently bound to the appliance's network interface cards. Select the appropriate subnet. All members of the cluster must be on the same subnet.

Servers: Click Insert to add appliances to the Servers list. All cluster members must be listed in the Servers list on each appliance. You then need to fill in the following fields:

- ◆ *Name:* A descriptive text label for the cluster member.
- ◆ *IP Address:* One of the IP addresses bound to the cluster member's network adapter (the adapter associated with the cluster). All IP addresses associated with the cluster, both those assigned to the appliances' network adapters and those for clustered services, must be on the same subnet.

IMPORTANT: Do not confuse this address with clustered service IP addresses. You assign the latter within the service definitions. They must not be bound to appliance network adapters.

- ♦ **Role:** Cluster member's role. Specify whether this cluster member's role is Active, Off-line, or Standby. Active means it is an available member of the cluster for proxy services. Standby means that the member becomes available for proxy services if another active member of the cluster goes down or is removed. Off-line means that this cluster member does not participate in providing services defined for the cluster.
- ♦ **Capacity:** An arbitrary number that you assign to rank each cluster member by its relative ability to handle services. Use the lowest numbers possible. Higher numbers suggest relatively superior memory, disk space, speed, and so on. The appliance uses these numbers to assign services and balance the load among cluster members. For more information, see [“About Capacity and Weight” on page 118](#). Valid field values include 1 – 100.

Services: This list shows the name you enter for each service you define in the cluster and the type of service it is. Click a service type (Forward, Transparent, Accelerator) below the list to insert a new entry into the list.

Insert Client Accelerator Service (Forward Proxy) Dialog Box

Path: Cache > Cluster > Forward

Figure 126

The Insert Client Accelerator Service (Forward Proxy) dialog box lets you define a forward proxy service to be handled by a cluster of appliances. You can define as many forward services for the cluster as makes sense for your network.

Name: A unique DOS 8.3 name assigned to the forward service.

Weight: A number between 1 and 100 representing the traffic this service needs to accommodate. See [“About Capacity and Weight” on page 118](#).

Forward Port: The port number that browsers using the forward service send their HTTP requests to.

IP Addresses: The IP addresses that browsers using the forward service send their HTTP requests to. All IP addresses associated with a cluster must be on the same subnet.

Enable X-Forwarded-For: Headers used to pass browser ID information along with browser request packets. If the headers are included, Web servers can determine the origin of browser requests they receive. If the headers are not included, browser requests have anonymity.

Checking the X-Forwarded-For option causes the appliance to either add information to an existing X-Forwarded-For or Forwarded-For header, or to create a header if one doesn't already exist.

Leaving the option unchecked causes the appliance to remove X-Forwarded-For headers from any forward proxy requests passing through the appliance.

You must weigh the desires of browser users to remain anonymous against the desires of Web server owners (e-commerce sites, for example) to collect data about who is accessing their site.

Enable Logging: Enables logging for the forward service.

Cache Objects That Have No Validator or Expiration Date: This enables caching of objects that would not normally be cached because they have no validator or expiration date set.

Insert Transparent Client Accelerator Service (Transparent Proxy) Dialog Box

Path: Cache > Cluster > Transparent

Figure 127

Insert Transparent Client Accelerator Service (transparent proxy)

Name: cluster

Weight: 1

Transparent ports: 80

Transparent IP addresses

Exception IP addresses

Insert Delete Insert Delete Insert Delete

☐ Enable X-Forwarded-For ☐ Cache objects that have no validator or expiration date

☐ Enable logging Log Options

OK Cancel Help

Warning: Applet Window

The Insert Transparent Client Accelerator Service (Transparent Proxy) dialog box lets you define a transparent proxy service to be handled by a cluster of appliances. You can define only one (1) transparent service for the network; however, you can define as many IP addresses as makes sense for the service and for your network.

IMPORTANT: For transparent client acceleration services to work properly in a cluster, each member of the cluster must have the Act As Router option checked in Network > Gateway/Firewall.

WCCP routing is not supported as a transparent service in clusters.

If you plan to use the clustered appliances in an inline routing configuration, see the Notes column in the table under “[Cluster Setup](#)” on page 115.

Name: A unique name for the transparent service. It must conform with DOS 8.3 naming conventions. The directory where logs for this service are stored uses this name.

Weight: A number between 1 and 100 representing the traffic this service needs to accommodate. See “[About Capacity and Weight](#)” on page 118.

Transparent Ports: The ports on which the transparent service handles HTTP requests.

Transparent IP Addresses: The service’s IP addresses to which the L4 switch or WCCP-capable router forward HTTP requests. All IP addresses associated with a cluster must be on the same subnet.

Exception IP Addresses: A list of origin server IP addresses. Browser requests to these addresses will bypass Excelsior’s transparent handling service and be sent directly to the origin server.

Enable X-Forwarded-For: Headers used to pass browser ID information along with browser request packets. If the headers are included, Web servers can determine the origin of browser requests they receive. If the headers are not included, browser requests have anonymity.

Checking the X-Forwarded-For option causes the appliance to either add information to an existing X-Forwarded-For or Forwarded-For header, or to create a header if one does not exist.

Leaving the option unchecked causes the appliance to remove X-Forwarded-For headers from any transparent proxy requests passing through the appliance.

You must weigh the desires of browser users to remain anonymous against the desire of Web server owners (e-commerce sites, for example) to collect data about who is accessing their site.

Cache Objects That Have No Validator or Expiration Date: This enables caching of objects that would not normally be cached because they have no validator or expiration date set.

Enable Logging: Enables logging for the service.

IMPORTANT: If transparent client acceleration services are defined for both an appliance and a cluster of which the appliance is a member, the appliance logs cluster-based service entries in the appliance-based forward log file.

If you define transparent client acceleration services for a cluster, and you want to ensure that the cluster’s transparent service log entries are written only to the cluster-transparent log files, you must ensure that none of the cluster members have appliance-based transparent services enabled in their Transparent Handling tabs.

Insert Web Server Accelerator Service Dialog Box

Path: Cache > Cluster > Accelerator

Figure 128

The screenshot shows a Windows-style dialog box titled "Insert Web Server Accelerator Service". It contains the following elements:

- Name:** A text input field.
- DNS name:** A text input field.
- Weight:** A text input field containing the number "1".
- Accelerator port:** A text input field containing "80".
- Web server port:** A text input field containing "80".
- Accelerator IP addresses:** A large empty text area with "Insert" and "Delete" buttons below it.
- Web server name/IP addresses:** A large empty text area with "Insert" and "Delete" buttons below it.
- Enable X-Forwarded-For:** An unchecked checkbox.
- Enable logging:** An unchecked checkbox, with a "Log Options" button to its right.
- Cache objects that have no validator or expiration date:** An unchecked checkbox.
- Buttons:** "OK", "Cancel", and "Help" buttons at the bottom.
- Status Bar:** A warning message "Warning: Applet Window" at the very bottom.

The Insert Web Server Accelerator Service dialog box lets you define a reverse acceleration service to be handled by a cluster of appliances. You can define as many reverse services for the cluster as makes sense for your network.

Name: A unique name that conforms with DOS 8.3 naming conventions. The directory where logs for this service are stored uses this name.

DNS Name: If you are accelerating multiple Web servers on the same IP address, you must create a Web server accelerator definition for each DNS name that is used in browser requests. This name must exactly match one of the names in the requests. See [“Standard Multihoming for Multiple Web Sites” on page 122](#).

This name is also used if the appliance is part of an ICP hierarchy that needs to resolve relative URLs.

Weight: A number between 1 and 100 representing the traffic this service needs to accommodate. See [“About Capacity and Weight” on page 118](#).

Accelerator Port: The port on which the service communicates with browsers requesting data from accelerated Web servers.

Accelerator IP Addresses: The service’s IP address to which DNS resolves the Web server’s DNS name to and on which the service communicates with browsers requesting Web server data. All IP addresses associated with a cluster must be on the same subnet.

Web Server Port: The port that the Web server uses for HTTP traffic.

Web Server Name/IP Addresses: The IP address or local DNS name of each Web server from which the appliance fills the cache for this Web server accelerator service.

Enable X-Forwarded-For: Headers used to pass browser ID information along with browser request packets. If the headers are included, Web servers can determine the origin of browser requests they receive. If the headers are not included, browser requests have anonymity.

Checking the X-Forwarded-For option causes the appliance to either add information to an existing X-Forwarded-For or Forwarded-For header, or to create a header if one doesn't already exist.

Leaving the option unchecked causes the appliance to remove X-Forwarded-For headers from any requests passing through the appliance.

You must weigh the desires of browser users to remain anonymous against the desires of Web server owners (e-commerce sites, for example) to collect data about who is accessing their site.

Enable Logging: Enables logging for the Web server accelerator.

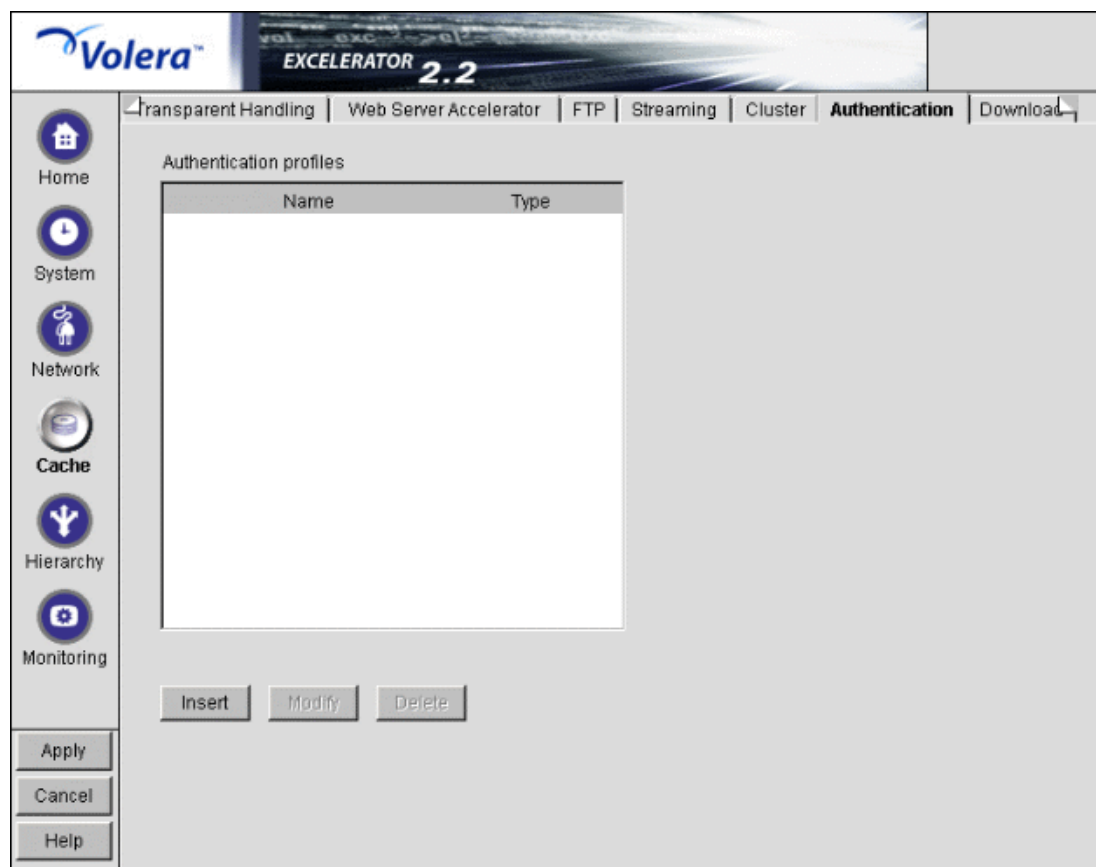
For more information, see [“Overview of Web Server Acceleration” on page 51](#).

Cache Objects That Have No Validator or Expiration Date: This enables caching of objects that would not normally be cached because they have no validator or expiration date set.

Authentication Tab

Path: Cache > Authentication

Figure 129



The Authentication tab lets you control access to proxy services by creating authentication profiles and assigning them to the services.

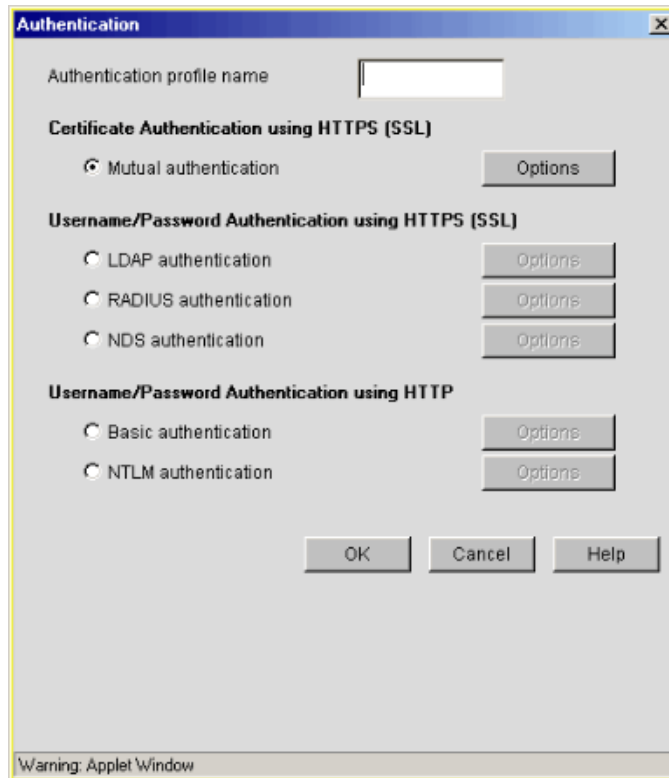
For more information, see [Chapter 22, “Authentication Services,” on page 135](#).

Authentication Profiles: List of the authentication profiles you have configured on the appliance using the Authentication dialog box.

Authentication Dialog Box

Path: Cache > Authentication > Insert under the Authentication Profiles list

Figure 130



The Authentication dialog box lets you assign an authentication profile name and selected the desired authentication method.

For more information, see [Chapter 22, “Authentication Services,” on page 135](#).

IMPORTANT: Excelerator doesn’t recognize case differences in profile names. MyProfile and myprofile are, effectively, the same profile name.

Also, Excelerator partially overwrites and concatenates previously created profiles without warning if a duplicate name is used. Therefore, if you create a profile named MyProfile and later create another profile named myprofile, Excelerator will remove the first name, concatenate parts the first profile with the second, and use the second name.

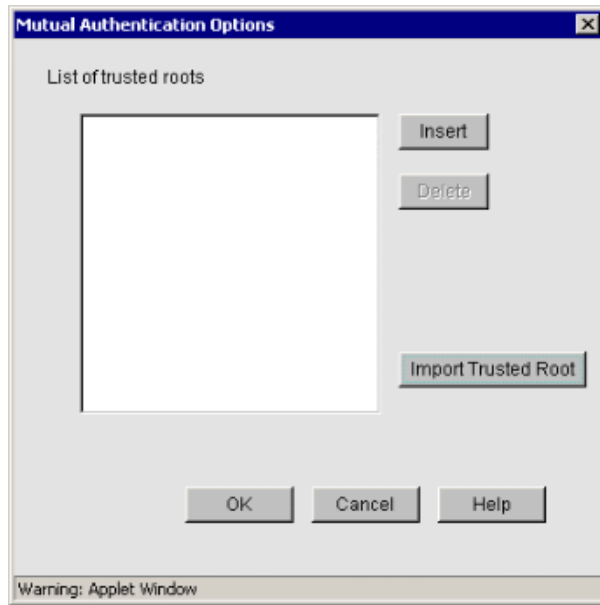
To avoid these problems, ensure that each profile has a unique name.

After selecting the authentication source, you must configure the source by clicking its respective Options button.

Mutual Authentication Options Dialog Box

Path: Cache > Authentication > Insert > Mutual Authentication > Mutual Authentication Options

Figure 131



Use the Mutual Authentication Options Dialog Box to create a mutual authentication profile.

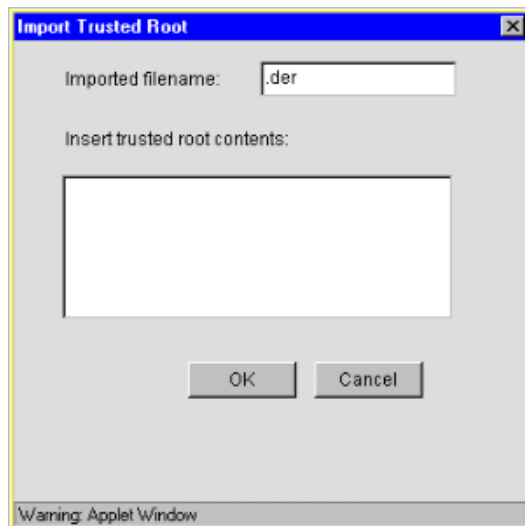
List of Trusted Roots: Displays trusted root certificates already installed. Click Insert to add trusted root certificates; click delete to remove existing certificates.

For more information regarding mutual authentication profiles, see [“Using Mutual \(Certificate-Based\) Authentication” on page 141](#).

Import Trusted Root Dialog Box (Mutual Authentication)

Path: Cache > Authentication > Insert > Mutual Authentication > Mutual Authentication Options > Import Trusted Root

Figure 132



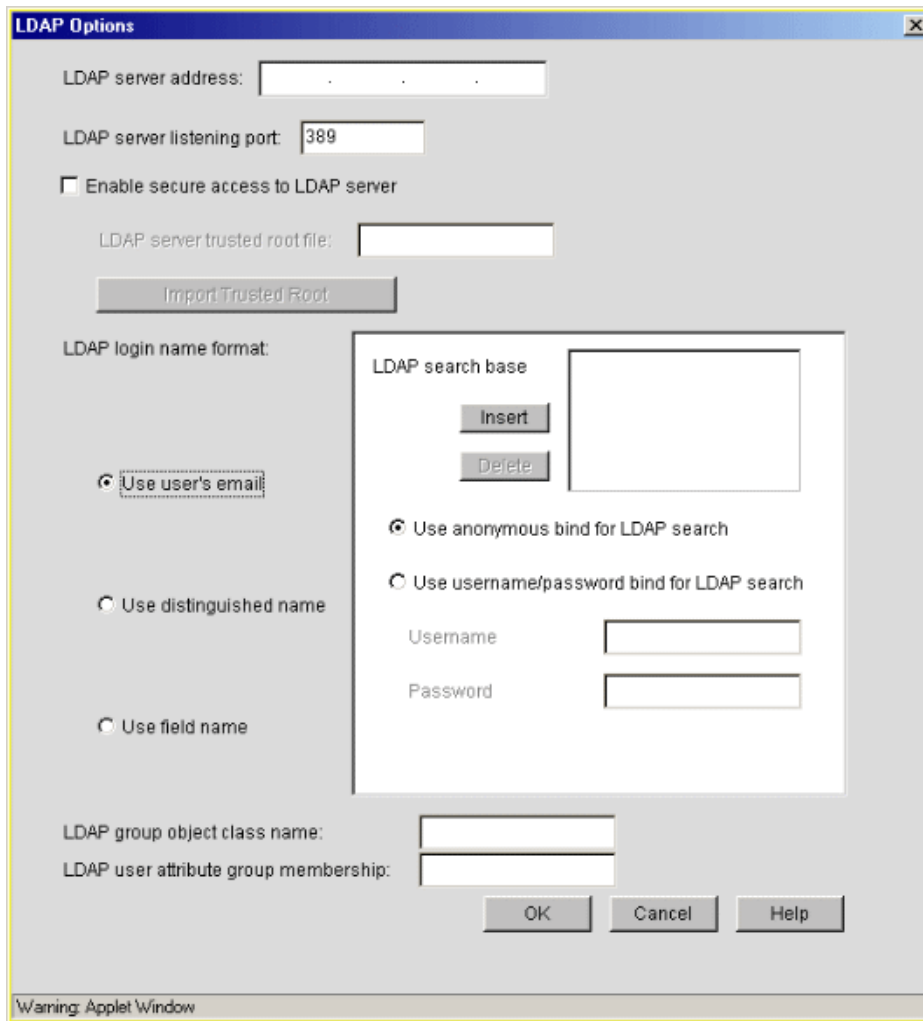
The Import Trusted Root dialog box lets you create a trusted root file that contains information identifying the Certificate Authority used by the server for the profile you are creating.

To create a trusted root file, see the instructions in [“Importing a Trusted Root to a Cache Device” on page 169](#).

LDAP Options Dialog Box

Path: Cache > Authentication > Insert > LDAP Authentication > Options

Figure 133



The LDAP Options dialog box is a Java applet window with a title bar that says "LDAP Options". It contains several input fields and buttons. At the top, there is a text field for "LDAP server address:" followed by a dotted placeholder. Below it is a text field for "LDAP server listening port:" with the value "389". A checkbox labeled "Enable secure access to LDAP server" is currently unchecked. Below that is a text field for "LDAP server trusted root file:" and a button labeled "Import Trusted Root". Under the heading "LDAP login name format:", there are three radio buttons: "Use user's email:" (selected), "Use distinguished name", and "Use field name". To the right of these is a sub-dialog box titled "LDAP search base" containing an empty text field, "Insert" and "Delete" buttons, and two radio buttons: "Use anonymous bind for LDAP search" (selected) and "Use username/password bind for LDAP search". Below these are text fields for "Username" and "Password". At the bottom of the main dialog are text fields for "LDAP group object class name:" and "LDAP user attribute group membership:". At the very bottom are "OK", "Cancel", and "Help" buttons. A warning bar at the bottom left says "Warning: Applet Window".

Use the LDAP Options dialog box to configure the appliance for users who authenticate through an LDAP database.

LDAP Server Address: The IP address of the LDAP server.

LDAP Server Listening Port: The port number on which the LDAP server is listening for requests from LDAP clients. The default is 389 for normal access. Use 636 for secure access.

Enable Secure Access to LDAP Server: Causes the data sent between the LDAP client and the LDAP server to be sent using SSL.

LDAP Server Trusted Root File: The path to a trusted root file that contains the Certificate Authority (CA) used by the LDAP server in the profile you are creating.

Excelerator fills this field with information for the trusted root file you create using the Import Trusted Root button. See the instructions found in “**Import Trusted Root Dialog Box (LDAP Authentication)**” on page 384.

If the LDAP server uses a CA for which you have previously created a trusted root file, you can manually type the path and filename in this field. For example, you might be using the same LDAP server for multiple authentication profiles.

Import Trusted Root: Opens the Import Trusted Root dialog box. See **Import Trusted Root Dialog Box (LDAP Authentication)**.

LDAP Login Name Format

The contents of this box change depending on the option selected.

Use User's E-Mail

(See **Figure 133**.) Select this option to have users log in using their e-mail name field in the LDAP database. You must provide one or more contexts in which the LDAP server will search for the e-mail name.

This option is somewhat redundant with Use Field Name because the e-mail name is simply an LDAP field name. E-mail is offered separately because it is used so often.

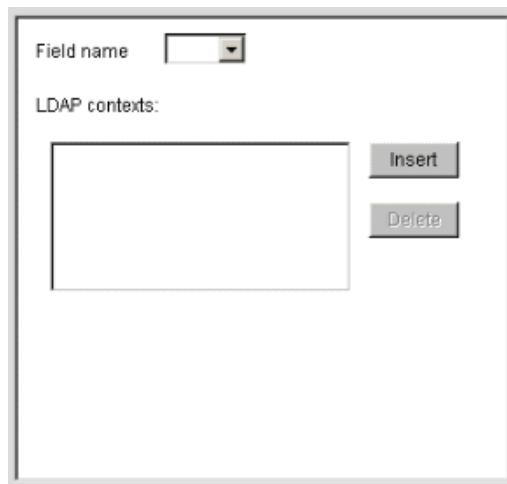
LDAP Search Base: Click Insert to enter the context of one or more LDAP containers from which the search for the e-mail name should begin.

You must also provide authentication information for the appliance to access the LDAP server using one of the following options:

- ◆ *Use Anonymous Bind for LDAP Search:* Select this option if the appliance can authenticate to the LDAP server using anonymous bind.
- ◆ *Use User Name/Password Bind for LDAP Search:* Select this option if anonymous bind is not enabled on the LDAP server > enter the username and password pair through which the appliance authenticates to use the LDAP server's authentication services.

Use Distinguished Name

Figure 134



The interface for configuring LDAP contexts. It features a 'Field name' dropdown menu at the top. Below it, the text 'LDAP contexts:' is followed by a large empty rectangular box for listing contexts. To the right of this box are two buttons: 'Insert' and 'Delete'.

Select this option to allow users to authenticate using their LDAP usernames. Users can use either their fully distinguished LDAP (full LDAP contexts) usernames, or you can provide a list of LDAP contexts so users only need to type their usernames.

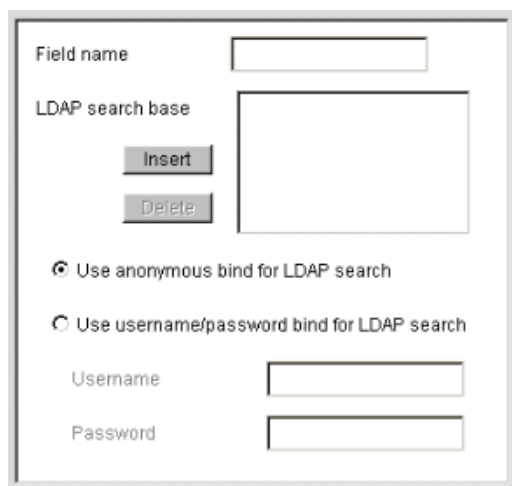
IMPORTANT: Using this option with Netscape's LDAP server requires a special setup procedure. For more information, see [“Use Distinguished Name” on page 146](#).

LDAP Contexts: Specific contexts in which the LDAP server will look for usernames. This provides a shortcut to authentication of users by allowing them to type only their LDAP usernames.

The appliance searches each context until it either locates the name or exhausts the search. If duplicate names exist in different contexts, the appliance searches until the correct name/password match is found.

Use Field Name

Figure 135



The interface for configuring LDAP search settings. It includes a 'Field name' text input field. Below it is the 'LDAP search base' section, which contains a large text area and two buttons: 'Insert' and 'Delete'. Further down are two radio button options: 'Use anonymous bind for LDAP search' (which is selected) and 'Use username/password bind for LDAP search'. At the bottom, there are two text input fields labeled 'Username' and 'Password'.

Select this option to require that users enter a specific LDAP field name.

Field Name: The LDAP field name (such as CN or UID) through which users can authenticate. If the field is left blank, the system automatically uses CN as the field name.

LDAP Search Base: Click Insert to enter the context of one or more LDAP containers. The appliance will perform a subtree search in all containers in the list. The subcontainers of the listed containers will also be searched.

Use Anonymous Bind for LDAP Search: Select this option if the appliance can authenticate to the LDAP server using anonymous bind.

Use User Name/Password Bind for LDAP Search: Select this option if anonymous bind is not enabled on the LDAP server > enter the username and password pair through which the appliance authenticates to use the LDAP server's authentication services.

LDAP Group Fields

LDAP Group Object Class Name: The mechanisms the target directory's schema uses to designate an LDAP group.

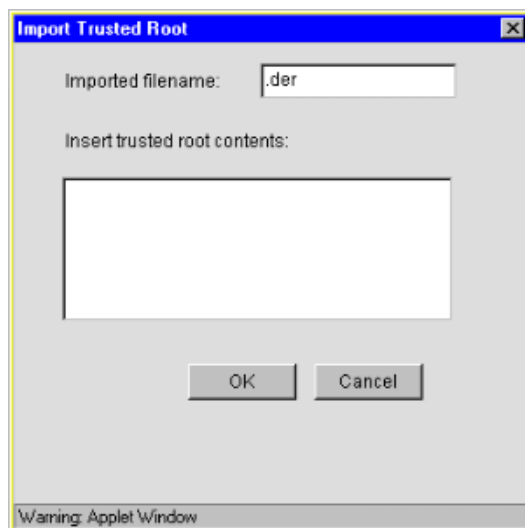
LDAP User Attribute Group Membership: The user object attribute used by the target directory to designate group membership.

For more information, see [“Enabling and Using LDAP Groups” on page 147](#) and [“Designating the Group Class and/or Attribute Name” on page 147](#).

Import Trusted Root Dialog Box (LDAP Authentication)

Path: Cache > Authentication > Insert > LDAP Authentication > LDAP Options > Import Trusted Root

Figure 136



The Import Trusted Root dialog box lets you create a trusted root file that contains information identifying the Certificate Authority used by the LDAP server for the profile you are creating.

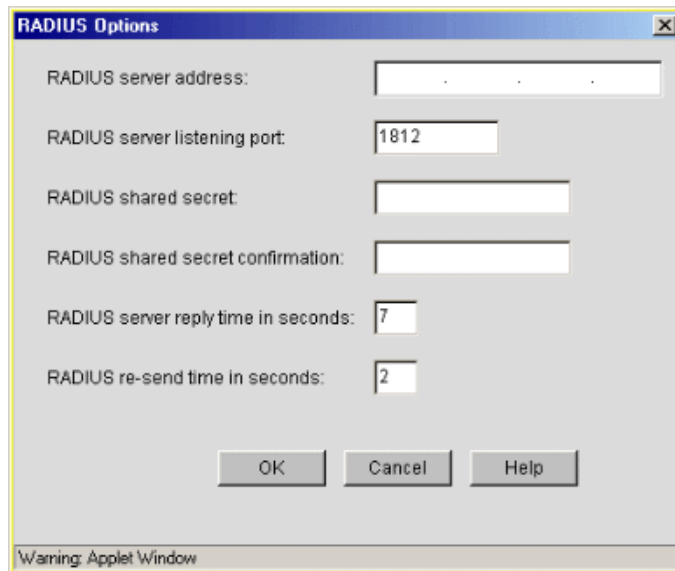
NOTE: Importing a trusted root file using this dialog box does not affect the list of trusted roots available for mutual authentication profiles, which are imported using the dialog box explained in [“Import Trusted Root Dialog Box \(Mutual Authentication\)”](#) on page 380.

For more information, see [“Importing a Trusted Root to a Cache Device”](#) on page 169.

RADIUS Options Dialog Box

Path: Cache > Authentication > Insert > RADIUS Authentication > Options

Figure 137

The image shows a dialog box titled "RADIUS Options" with a standard Windows-style title bar (blue with a close button). The dialog box has a light gray background and contains several input fields and buttons. The fields are: "RADIUS server address:" with a text box containing three dots; "RADIUS server listening port:" with a text box containing "1812"; "RADIUS shared secret:" with a text box; "RADIUS shared secret confirmation:" with a text box; "RADIUS server reply time in seconds:" with a text box containing "7"; and "RADIUS re-send time in seconds:" with a text box containing "2". At the bottom of the dialog box are three buttons: "OK", "Cancel", and "Help". A small warning message "Warning: Applet Window" is visible at the very bottom of the dialog box.

Use this dialog box to specify a RADIUS server the appliance can use for authentication. For more information regarding RADIUS authentication, see [“Using RADIUS Authentication”](#) on page 148.

RADIUS Server Address: The IP address of the RADIUS server.

RADIUS Server Listening Port: The port number on which the RADIUS server listens for incoming authentication requests.

RADIUS Shared Secret: The string the RADIUS server uses to verify that the appliance can request authentication of users.

RADIUS Shared Secret Confirmation: Confirmation string the system will compare with the RADIUS Shared Secret. The system compares the strings to ensure they match before accepting the configuration.

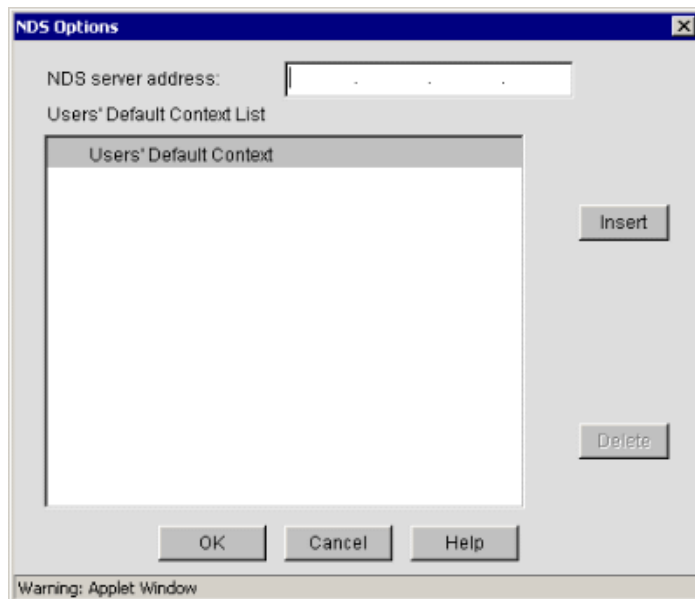
RADIUS Server Reply Time in Seconds: The total time the appliance will wait for a response from the RADIUS server before authentication fails. The default is 7 seconds.

RADIUS Re-send Time in Seconds: The interval in seconds between appliance requests to the RADIUS server. The default is 2 seconds. This means that the appliance could send three requests before the 7-second default limit expires and the authentication request fails.

NDS Options Dialog Box

Path: Cache > Authentication > Insert > NDS Authentication > Options

Figure 138



Use the NDS Options dialog box to configure the appliance for having users authenticate through an NDS database. For more information regarding NDS authentication, see [“Using NDS \(eDirectory\) Authentication” on page 150](#).

NDS Server Address: The IP address of the NDS server.

Users’ Default Context List: Displays the defined NDS context(s).

To add an NDS context, click Insert. The following dialog displays:

Figure 139



Enter the appropriate NDS context and tree name and click OK.

Basic Authentication Options Dialog

Path: Cache > Authentication > Insert > Basic Authentication > Basic Authentication Options

Figure 140



Use this dialog to set up basic authentication. With basic authentication, usernames and passwords are lightly encrypted (low security).

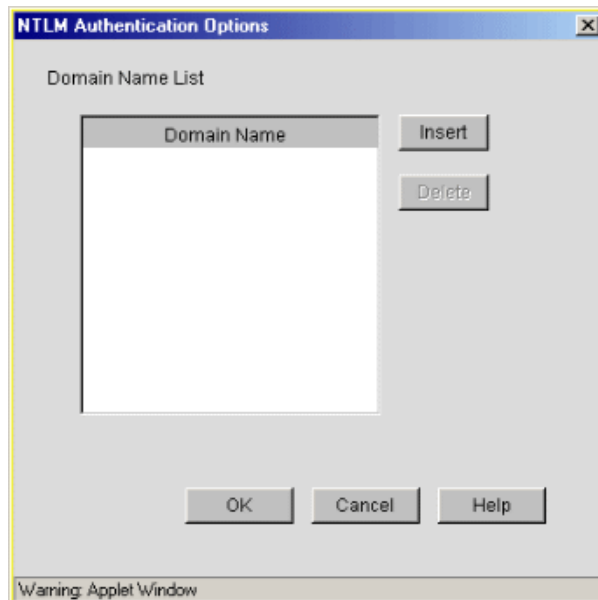
To use basic authentication, you must already have established an authentication method with at least one of the existing authentication options. Select the desired profile from the drop-down menu and click OK.

For more information regarding Basic authentication, see [“Using Basic Authentication” on page 155](#).

NTLM Authentication Options Dialog Box

Path: Cache > Authentication > Insert > NTLM Authentication > Options

Figure 141



Use this dialog box to create NTLM-based authentication profiles for forward proxy services.

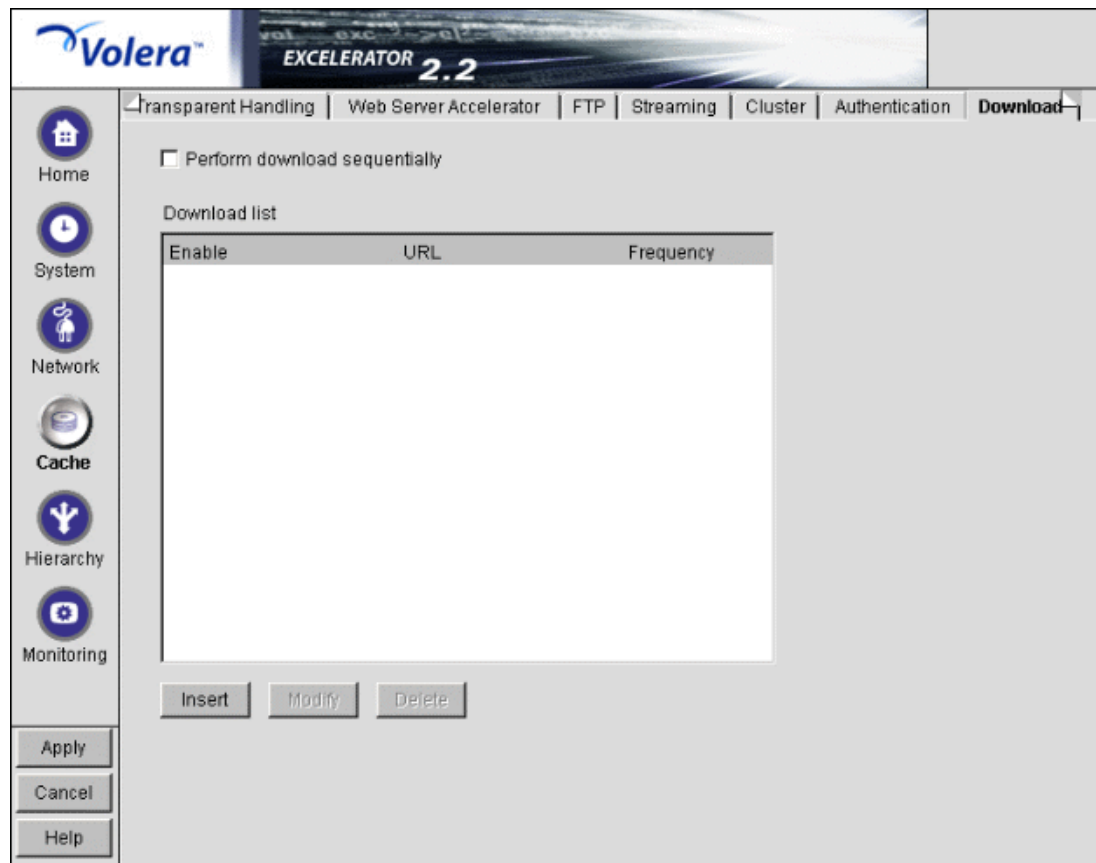
Address List: This list contains the IP address of the Domain Controller used by the profile.

For more information regarding NTLM authentication profiles, see [“Using NTLM Authentication” on page 157](#).

Download Tab

Path: Cache > Download

Figure 142



The Download tab lets you keep the cache of objects up to date for your users. You can schedule downloads of files from a Web site to the local cache. You can download a single URL, multiple URLs up to a specified number of links, or an entire Web site. You can download for both forward and reverse proxies. However, reverse proxy does not download links that are external to a site.

To conserve network resources, consider scheduling your downloads during times of low network usage.

Perform Download Sequentially: To conserve network resources, specify the downloads to occur sequentially rather than concurrently.

Download List: Create multiple lists of URLs for download by inserting them into this list.

Scheduled Download Dialog Box

Path: Cache > Download > Insert

Figure 143

Scheduled Download

☒ Enable this download

HTTP URL:

Levels to download from web site:

☐ Follow links to other hosts

Maximum number of concurrent requests:

Maximum number of objects to download:

Maximum amount of data to download: MB

☐ This site requires authentication

Username:

Password:

Frequency

☐ Immediately

☒ One time only:

☐ Once a day at:

☐ Daily from: to every

OK Cancel Help

Warning: Applet Window

The Scheduled Download dialog box lets you set up download options for a specific URL.

Enable This Download: Lets you edit and enable a specific download. You can uncheck this if you want to turn various download categories off without deleting them.

HTTP URL: The URL or IP address of the site you want to use as the top level of the download.

Levels to Download from Web Site: Each link takes you one level away from the URL you specified. By default, you get one level—the top page (graphics and text objects). (Valid numbers are 1 – 999.) Increase this number for each level you want to download. For example, selecting a 2 will download the initial page (graphics and text objects) and all its links.

Follow Links to Other Hosts: By default, only links that are on the host of the HTTP URL specified above are followed. Check this option to allow the download to follow external links as well.

Maximum Number of Concurrent Requests: By default, Excelerator limits the concurrent requests for a specific download to 6. You can specify a larger or smaller number to increase or decrease the impact on network traffic. Valid numbers are 1 – 50.

IMPORTANT: If this number is set to a larger value than the default, it is possible to exceed the Maximum Amount of Data to Download setting.

Maximum Number of Objects to Download: You can specify a maximum number of objects (graphics and text) that are downloaded. By default, this number is 5000 (1 – 100,000).

Maximum Amount of Data to Download: You can specify a maximum MB size in the download. By default, this is 200 MB (1 – 4095).

IMPORTANT: This is a soft limit. The accuracy of this setting is dependent upon the Maximum Number of Concurrent Requests setting.

The reason for this is that the size of the downloaded object is not compared to the limit until after it has successfully downloaded. If the maximum limit has not been reached, then a new request to download an object is initiated. If that limit has been reached, then no new requests are initiated but all current requests are allowed to complete. If each object currently being downloaded is very large, then the maximum amount of data will be exceeded by a significant amount.

This Site Requires Authentication: If the origin server requires authentication, check this option and type the authentication username and password in their respective fields.

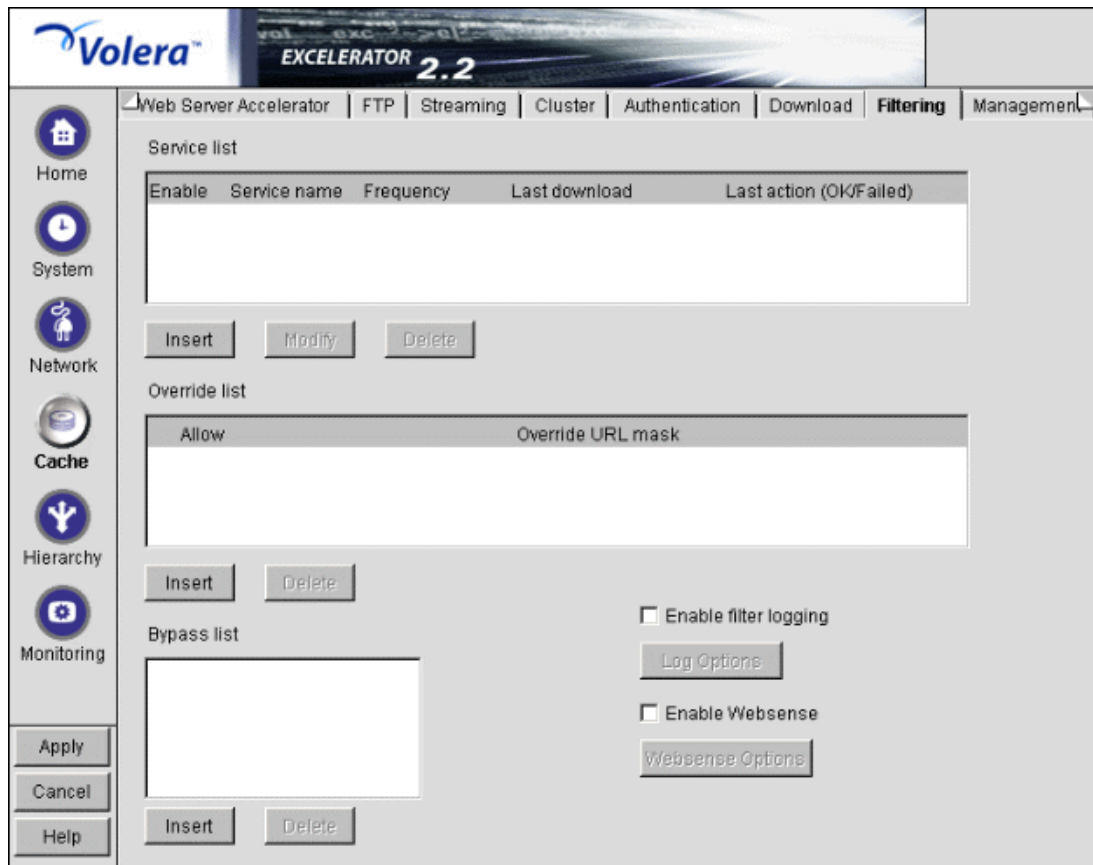
Frequency: Specify when this download should occur. Select one of the following options to establish download frequency:

- ♦ *Immediately:* Selecting this option causes the download to begin when you click Apply.
- ♦ *One Time Only:* You can specify a specific instance for a download (day, month, year, hour, time zone).
- ♦ *Once a Day At:* Specify a time and time zone for a daily download.
- ♦ *Daily From:* Specify a daily time frame for a scheduled download. You can also make this repeat every 1 to 24 hours.

Filtering Tab

Path: Cache > Filtering

Figure 144



The Filtering tab lets you do the following:

- ◆ Subscribe to a cooperating filtering service.
- ◆ Create your own always- or never-vend instructions for the URLs or URL patterns you create.
- ◆ Create a bypass list for appliance IP addresses and ports whose requests need to bypass appliance content filtering.
- ◆ Enable logging of filter activity.

When you have an account with the filtering service, you can perform scheduled downloads of the filter service files. You do not need a filtering service subscription to use your own override or bypass lists, however.

Service List: Contains all the filtering services you are currently subscribed to. It has the following display fields:

- ◆ *Enable:* Turn the service (including the scheduled download of the list) on and off.
- ◆ *Service Name:* The name you created for the service.
- ◆ *Frequency:* How often a scheduled download of the service files occurs. Options are Daily, Weekly, Monthly, and Bimonthly.
- ◆ *Last Download:* When the filtering service files were last updated.

- ♦ *Last Action*: The last action performed relative to downloading the service files. The state is OK for a successful download or Failed for an unsuccessful one. Text information is also displayed about the last action performed or pending.

You can insert a new service, or modify or delete an existing one. See “[Insert Filter Service Dialog Box](#)” on page 393 and “[Modify Filter Service Dialog Box](#)” on page 393.

Override List: A list of URLs you can create that overrides any filtering described in the service list. You can specify to never or always allow a specific URL.

You can use wildcards in describing the URL. See “[Critical Information about Wildcards in the Override List](#)” on page 206.

You can also select Insert to add members or Delete to remove an existing URL from the list.

- ♦ *Allow*: Select the Always option to always allow the specified URL or select Never to never allow the specified URL to vend.
- ♦ *Override URL Mask*: Enter the URL address (DNS name or IP address) of the site you want to control access to. You can use the asterisk (*) wildcard character or any valid URL scheme. See “[Critical Information about Wildcards in the Override List](#)” on page 206.

Be aware that the appliance applies the action specified for the most specific or granular URL.

For example, if you have specified to always allow AnEducationalSite.org but to never allow AnEducationalSite.org/snakes, the appliance always vends everything on the site, except for objects with /snakes in the path because /snakes is a more granular URL.

IMPORTANT: When you enter a URL (hostname) in the Override list, Excelerator checks to see whether the hostname resolves to a single IP address. If so, a second entry is added to the list for the IP address that corresponds to the hostname entry. The same Allow status you specify for the hostname entry is applied to the IP address entry.

These entries are not linked together. If you delete one of the entries (hostname or IP address), the corresponding entry is not automatically deleted. Therefore, the override might continue to be in effect after you think it has been removed.

Also, if you change the Allow status of one of the entries, the other entry is not automatically changed. Excelerator processes IP address entries before hostname entries. Therefore, if you change the Allow status of the hostname entry and don't change the IP address entry, the previous override behavior will continue unchanged.

Bypass List: A list of appliance IP addresses and ports whose requests need to bypass appliance content filtering. Clicking Insert opens the Enter Filter Bypass dialog box wherein you can select one of the appliance's IP addresses and enter the port number.

Once an IP address and port have been inserted in the list, all requests pass through the specified appliance IP address and port unaffected by filtering.

Enable Filter Logging: Turns on logging of filter activity. Excelerator supports only appliance-specific filtering common log format for filter logging, which has different fields than the common log format used to log caching transactions.

Enable Websense: This option is not activated until you have installed a valid Websense license on the cache device. You can download a valid license from [Volera's support pages on Novell.com](http://www.volera.com/support/patches/WEBSENSE.LIC) (<http://www.volera.com/support/patches/WEBSENSE.LIC>) or contact your Novell Authorized Reseller. For more information on installing licenses, see [Chapter 21, “Installing and Upgrading Licenses,”](#) on page 133.

After installing a valid Websense license, click Websense Options to complete the Websense integration. For more information, see “[Websense Options Dialog Box](#)” on page 395.

Insert Filter Service Dialog Box

Path: Cache > Filtering > Insert under Service List

Figure 145

Insert Filter Service

Service name

Configuration file URL

Account name

Account password

Disclaimer

The caching server vendor does not control the content, software, services or experience associated with any of the offered filtering services. The caching server vendor does not affirm or warrant the integrity or quality of the providers or their goods and services. No filtering service is perfect. Even when using one of the offered filtering services, you may be exposed to content that is offensive to you or content that you may consider explicit, indecent or otherwise objectionable. IN NO EVENT SHALL THE CACHING SERVER VENDOR BE LIABLE IN ANY WAY FOR ANY DAMAGE THAT OCCURS RELATED TO THE OFFERED FILTERING SERVICES, INCLUDING BUT NOT LIMITED TO, CONTENT RELATED DAMAGE, ERRORS OR OMISSIONS, LOSS OF DATA, DAMAGE TO EQUIPMENT OR FACILITIES, DAMAGES RELATED TO THE RELEVANCE OR USEFULNESS OF THE OFFERED SERVICES AND THEIR CONTENT, OR DAMAGES CAUSED BY EMOTIONAL DISTRESS.

OK Cancel Help

Warning: Applet Window

The Insert Filter Service dialog box lets you configure the appliance to download a filter service file.

Service Name: Enter the name you have selected to identify this filter service. Enter any text you want, with no spaces or special characters.

Configuration File URL: Enter the complete URL supplied by your filtering service for the filter service file.

Account Name: Enter the account name provided by your filtering service.

Account Password: Enter the password provided for your filtering service account.

To add a new service to the list, click OK > Apply. To ignore changes, click Cancel.

Modify Filter Service Dialog Box

Path: Cache > Filtering > Service in Service List > Modify

Figure 146

Help	Category name	Filter rule	Block unrated
?	Alcohol	Block	<input type="checkbox"/>
?	Alternative Journal	Block	<input type="checkbox"/>
?	Anarchy	Block	<input type="checkbox"/>
?	Automobile	Block	<input type="checkbox"/>

The Modify Filter Service dialog box allows you to modify the configuration of a filter service, including the service name, URL, account name, and account password, after the PICS description file has been downloaded.

Download Frequency

Most filter services allow you to refresh their rating list as often as you like. When you initially configure a filter service, check the Next Download information to ensure timely activation of filtering. (Filtering doesn't start until a rating list has been successfully downloaded.)

The Categories List

The Categories list lets you configure filtering behavior for each category defined by your filter service.

Help: Click the question mark button to the left of the category name to display a category help file. If a category help file is not available, a generic help page appears.

Category Name: The category name is created by the service provider and identifies some aspect of Web data. Names might include Violence, Language, Adult Oriented, and Education. The appliance downloads the category names and associated help files when you specify the service.

Service providers rate Web pages for each category as either true or false. A page rated true for Violence has violent content as defined by the service provider. A page rated true for Education meets the service provider's criteria for the Education category.

Filter Rule: The appliance uses filter rules you create in combination with category ratings from your service provider to determine whether to vend, monitor, or block a page. You specify one of the following filter rules for each category.

- ♦ *Ignore*: The appliance ignores the category when determining what to do with a page. For example, if you specify that the Violence category should be ignored, the appliance ignores Violence ratings when determining whether to vend or block a page.
- ♦ *Always Allow*: The appliance immediately vends a page rated true for the category and ignores other filter rules that might require the page to be blocked. For example, if a page is rated true for both the Education and Violence categories, and you specify that pages rated true for Education are always allowed, the appliance will immediately vend the page, ignoring the rating for Violence.
- ♦ *Monitor*: The appliance vends a page rated true for the category and logs the event, unless the page is blocked. If another category blocks vending the page, the event is logged, but the page is blocked. The appliance provides this feature for administrators who want to monitor activity for specific categories such as Educational Content.
- ♦ *Block*: The appliance blocks a page rated true for the category unless another category with an Always Allow filter rule set has been rated as true. See the example under Always Allow above.

Block Unrated: The Block Unrated check box appears only when the Block filter rule is selected. Checking the box causes the appliance to block all pages that the service provider has not rated for the category.

IMPORTANT: Checking this option sometimes causes pages to be blocked that appliance administrators hadn't anticipated. As you use this option, remember that many sites are not rated for all categories. The appliance will block all pages that are not rated for a category that is checked.

For example, a service provider has rated a page on trees as true for Education but has not assigned a rating for the Violence category since the page has no violent content. If an appliance administrator checks Block Unrated for the Violence category, the appliance will block the tree page even though it has no violent content.

Websense Options Dialog Box

Path: Cache > Filtering > Enable Websense > Websense Options

Figure 147



Use this dialog box to configure Websense Enterprise software.

IMPORTANT: This dialog box cannot be displayed until a Websense license is installed on the cache device. For more information on installing licenses, see [Chapter 21, “Installing and Upgrading Licenses,”](#) on page 133.

Management Tab

Path: Cache > Management

Figure 148

The screenshot shows the 'Management' tab of the Volera EXCELERATOR 2.2 interface. The left sidebar contains navigation icons for Home, System, Network, Cache, Hierarchy, and Monitoring. The main content area is titled 'Management' and includes the following sections:

- Enable pin list:** A checkbox to enable pinning. Below it, 'Default refresh frequency' is set to 'Once Immediately' and 'Default refresh time' is set to '4 am :00'.
- Pin list:** A table with columns: URL Mask, Pin Type, Pin Links, Pin Images, and Refresh Frequency/Time. Below the table are 'Insert', 'Modify', and 'Delete' buttons.
- Do not purge pinned content:** A checkbox.
- Enable dynamic bypass (Requires router to be enabled):** A checkbox. Below it, 'Dynamic bypass duration' is set to '0 :05'. 'Dynamic bypass error codes to enable' has 'All' and 'None' buttons.
- Error codes:** A grid of checkboxes for error codes 400 through 505.
- Enable dynamic bypass logging:** A checkbox with a 'Log Options' button.
- Apply/Cancel/Help buttons:** Located at the bottom left.
- Reset button:** Located at the bottom center.
- Cache settings:** Two dropdown menus for 'Allow caching of objects with a ? in the URL' and 'Allow caching of objects with /cgi in the path'.

The Management tab lets you identify objects that will be either pinned (explicitly downloaded and retained in cache as long as possible) or bypassed (explicitly not cached when requested by users). It also lets you specify how pinned objects are stored on the appliance (pin type).

Enable Pin List: Activate pinning on the appliance. For information regarding what pin lists are and how they function, see [“The Pin List” on page 269](#) and [“Pin List Examples” on page 274](#).

NOTE: This option affects only the pinning of objects on the appliance. It has no effect on whether objects are cached unless the pin type is set to Bypass.

Default Refresh Frequency: Specify when the appliance checks to see if items should be added to or removed from the list of objects being pinned. At refresh time the appliance re-evaluates the objects in cache for the URL list and downloads those objects that have changed. For absolute URLs, objects in links are also evaluated down to the link level specified. Choices range from one time only to any arbitrary time interval.

Default Refresh Time: Specify the time of day or the time interval when pinned objects will be refreshed. To specify an interval, you must first select Timed Interval from the Default Refresh Frequency drop-down list.

The Pin List

The pin list lets you specify URL patterns for identifying objects on the Web. You configure each URL pattern in the list with specific handling instructions as explained below.

URL Mask: Specify complete or partial URL patterns. A single URL mask might apply to a large set of URLs, or it might be so specific that only a single file on the Web matches it. For more information, see [“Pin List Examples” on page 274](#).

Pin Type: Specify whether and how the appliance will cache objects that match the URL mask.

- ◆ *Normal:* Excelerator handles objects matching the mask in the same way it handles any other requested objects. In other words, objects are cached but not pinned.
- ◆ *Cache:* Excelerator keeps the pinned objects in cache as long as possible, although they might be written to the appliance’s hard disk.
- ◆ *Memory:* Excelerator keeps the pinned objects in memory as long as possible, writes them to disk when memory gets too full, and places them back in memory as soon as they are requested by a user of the cache.
- ◆ *Bypass:* Excelerator does not cache the objects. In other words, you can use this option to prevent objects from being cached.

Pin Links: Specify how many link levels Excelerator will follow the pin type rule you’ve established.

Pin Images: Pin image files that reside on a different host than the page requested.

Refresh Frequency/Time: Specify a refresh frequency and time for the URL that is different from the default values shown above the pin list. After inserting the URL, click Modify to see the dialog box that lets you change these fields. Functionality is the same as described for the default refresh frequency fields.

Do Not Purge Pinned Content: Start the purge with this option selected to purge everything from the appliance’s cache except for the content which you have pinned.

Enable Dynamic Bypass

The Dynamic Bypass feature lets you configure the appliance so that specific errors from Web sites are explicitly not cached and subsequent requests to the Web sites are simply passed through for a specific time period. For more information, see [Chapter 32, “Dynamic Bypass,” on page 231](#).

IMPORTANT: The Dynamic Bypass feature only works in transparent proxy mode.

Enable Dynamic Bypass (Requires Router to Be Enabled): Checking this option enables the Dynamic Bypass feature.

Dynamic Bypass Duration: The period of time in hours and minutes that the URL remains in the Dynamic Bypass list, causing Excelerator to transparently pass matching requests through to the origin Web server.

Dynamic Bypass Error Codes to Enable: This section lists all valid HTTP 1.1 errors. You can check and uncheck all the errors using the All and None buttons, or you can check only the errors you want enabled for dynamic bypass.

Enable Dynamic Bypass Logging: Checking this option enables logging of all dynamic bypass transactions. Log entries include the time of the log entry, the word Bypassed, and the URL added to the Dynamic Bypass list.

Click Log Options to set log file rollover options and specify the handling of the oldest log files. For help with setting these options, see [“The Appliance Offers Two Log Rollover Options” on page 239](#). For information relevant to the dynamic bypass feature, refer to [“Configuration Step 3:](#)

Specifying Rollover Options” on page 243, and “Configuration Step 4: Specifying Handling of Older Files” on page 244.

Caching Based On URL Content

The two drop-down lists below the pin list let you specify object caching based on the following:

- ♦ Whether URLs contain a question mark
- ♦ Whether URLs have /cgi in the path

The Cache option lets you specify whether cacheable objects that meet the criteria are always cached.

The Cache option is sometimes misinterpreted to imply that objects meeting the criteria are always cached. That is not the case. The Excelerator appliance will not cache objects that Web server administrators have marked non-cacheable.

IMPORTANT: Even if either of the Cache options is selected, objects with a question mark or /cgi in the path will not be cached unless you have enabled the Cache Objects That Have no Validator or Expiration Date option in Tuning dialog box for the services you want affected. For more information, see “Advanced Options (Tuning) Dialog Box” on page 341.

By the same token, the Do Not Cache option is sometimes misinterpreted to imply that objects meeting the criteria are never cached. Objects containing question marks and/or /cgi in the path might meet other criteria that cause them to be cached. This option actually only causes Excelerator to ignore question marks and /cgi in determining whether to cache objects.

IMPORTANT: Even if the Do Not Cache option is selected, objects with a question mark or /cgi in the path may be cached if they have a Validator and/or an Expiration in the header.

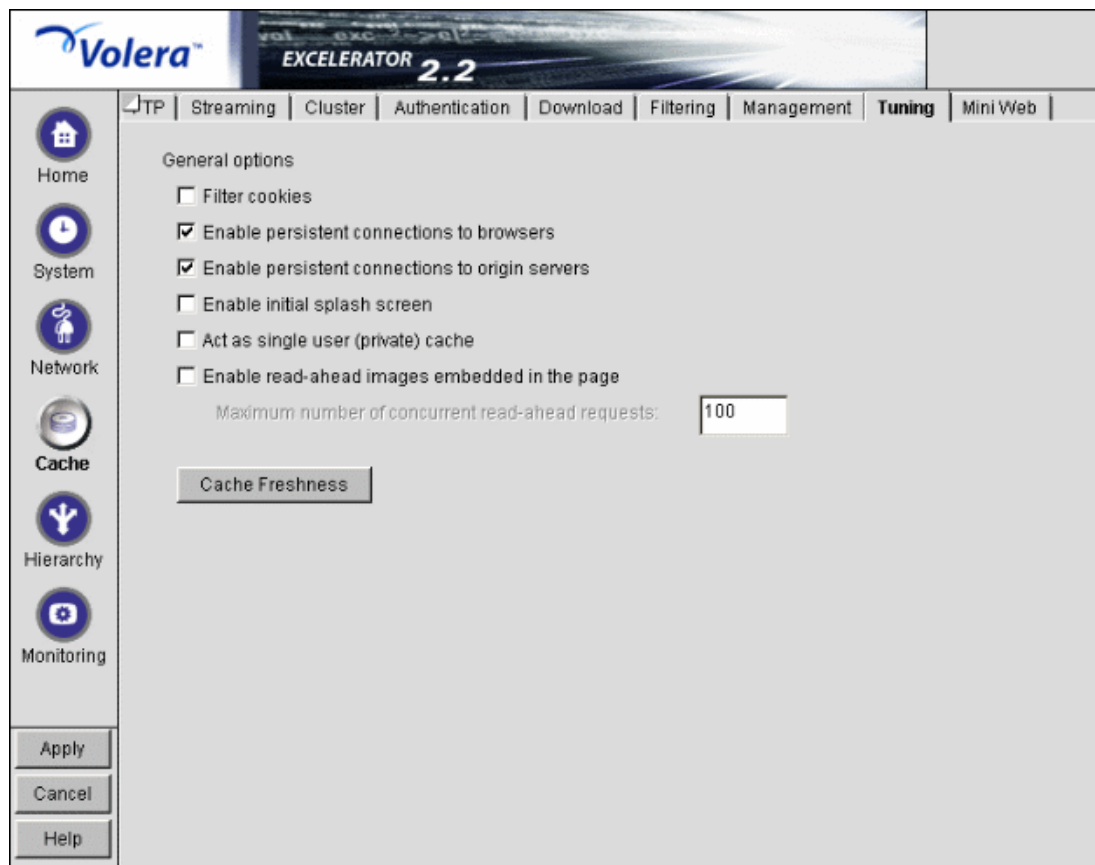
The Reset Button

Click the Reset button to return the two drop-down lists to their default (do not cache) settings.

Tuning Tab

Path: Cache > Tuning

Figure 149



The Tuning tab lets you restrict and enable functionality that affects all appliance operations. The implications for each option are explained below.

Filter Cookies: Excelerator removes all cookie HTTP request headers from requests forwarded to Web servers. It also removes all set-cookie HTTP reply headers from replies coming from Web servers.

Enable Persistent Connections to Browsers: All connections from browsers to the appliance remain open. This makes the response time between the appliance and browsers faster.

Enable Persistent Connections to Origin Servers: All connections from browsers to Web servers (through the appliance) remain open. This can cause some Web servers and their networks to crash, depending on the number of simultaneous connections they support.

Enable Initial Splash Screen: Browsers receive first-time and periodic notification that their requests are being processed by the appliance. The splash screen is customizable so that ISPs, for example, can advertise the fact that they are providing accelerated Web service. For information on customizing the splash screen, see [“Customizing the Appliance Splash Screen with FTP” on page 264](#). The splash screen is disabled (turned off) by default.

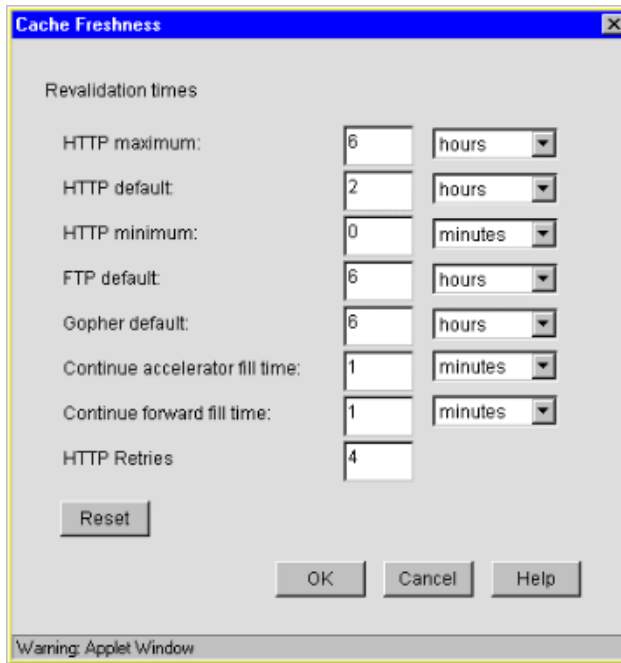
Act as Single User (Private) Cache: Excelerator caches objects that have been flagged for private caches only.

Enable Read-Ahead Images Embedded in the Page: Excelerator retrieves and caches objects that have been flagged Read-Ahead. You specify the maximum number of Read-Ahead objects Excelerator will retrieve in the Maximum Number of Concurrent Read-Ahead Requests field.

Cache Freshness Dialog Box

Path: Cache > Tuning > Cache Freshness

Figure 150



The Cache Freshness dialog box lets you set time values governing when Excelerator revalidates requested cached objects against those on their respective origin Web servers. If requested objects have changed, Excelerator re-caches them. Default field values are shown in [Figure 150](#).

Excelerator does not automatically re-cache objects when they expire. Expired objects are revalidated (and re-cached if they have changed) only when requested by browsers and in accordance with the time values set in the Cache Freshness dialog box. For more information on the appliance's cache freshness features, see [“Cache Freshness” on page 185](#).

HTTP Maximum: The maximum number of hours or days Excelerator will serve HTTP data from cache before revalidating it against content on the origin Web server. If an object has not been revalidated when this value expires, the object cannot be served from cache.

This overrides a freshness or Time to Expire header value specified by the Webmaster if he or she specified a longer time.

You can use this value to reduce the maximum time Excelerator waits before checking to see if requested objects need to be refreshed.

HTTP Default: The value Excelerator uses to determine when to revalidate requested objects for which Webmasters have not specified a freshness or Time to Expire header value.

HTTP Minimum: The minimum number of hours or minutes Excelerator will serve HTTP data from cache before revalidating it against content on the origin Web server. No requested object will be revalidated sooner than specified by this value.

This overrides the freshness or Time to Expire header value specified by the Webmaster if he or she specified a shorter time.

You can use this value to increase the minimum time Excelsator waits before checking to see if requested objects need to be refreshed. This parameter does not override No Cache or Must Revalidate directives from the origin Web server.

FTP Default: The number of hours or days FTP data remains in cache before it expires. Excelsator does not revalidate expired FTP objects. It simply re-caches them when they are requested.

Gopher Default: The number of hours or days Gopher data remains in cache before it expires. Excelsator does not revalidate expired Gopher objects. It simply re-caches them when they are requested.

Continue Accelerator Fill Time: The number of minutes or hours that the appliance's Web acceleration services ignore browser request cancellations and continue downloading objects from the target Web server until the download is complete.

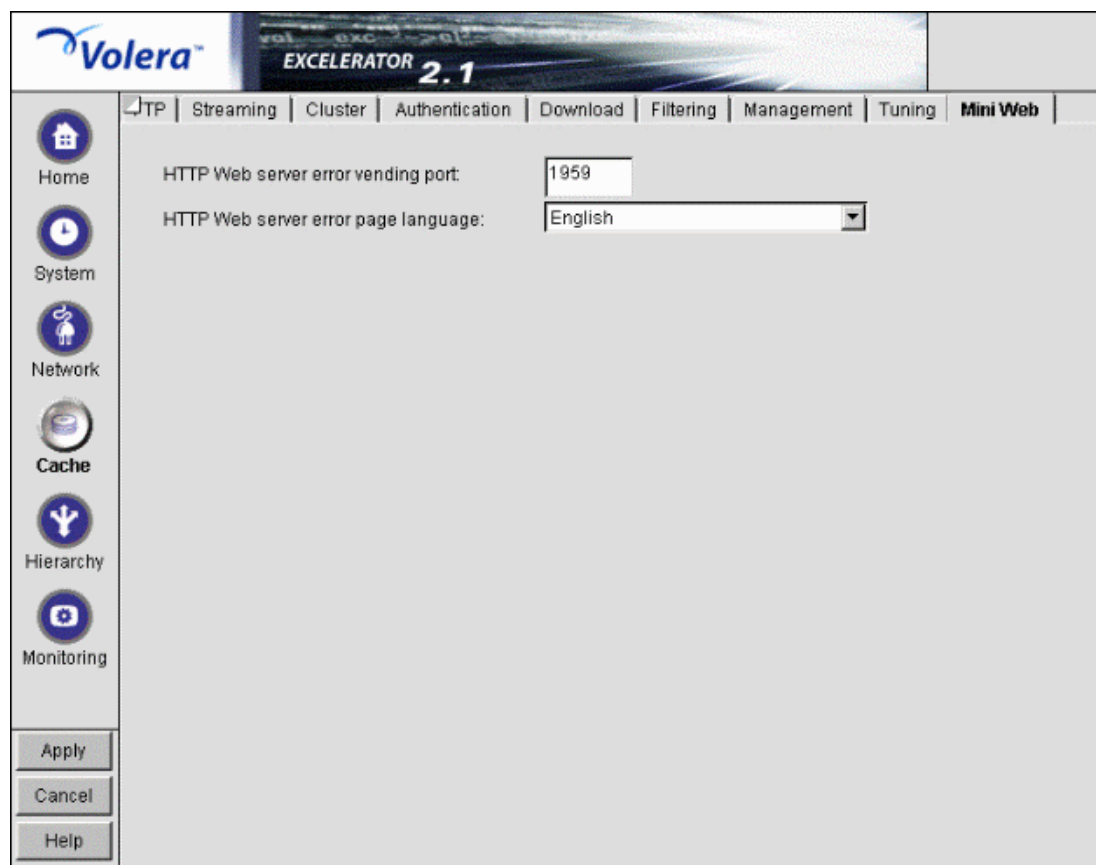
Continue Forward Fill Time: The number of minutes or hours that the appliance's HTTP forward and transparent proxy services ignore browser request cancellations and continue downloading objects until the download is complete.

HTTP Retries: The number of retry requests that will be issued.

Mini Web Tab

Path: Cache > Mini Web (to see this tab, click the upper-right corner of the Tuning tab.)

Figure 151



The Mini Web tab lets you configure how appliance-generated error pages are vended to browsers.

HTTP Web Server Error Vending Port: The port the browser will use when requesting objects that are part of the error pages. Changing this value does not affect the port for appliance administration, which is fixed at 1959.

HTTP Web Server Error Page Language: The language that will be used for appliance-generated error messages sent to browsers.

47

Using the Hierarchy Panel

The Hierarchy panel lets you configure the appliance to participate in ICP and CERN hierarchies.

ICP caches cooperate in finding URLs, thereby generating packet traffic on the network. CERN uses static routing without queries, offering less flexibility and power but generating much less network traffic.

If you define both ICP and CERN hierarchical relationships, the ICP hierarchy takes precedence.

The caching system does not support CARP.

For a brief explanation of how hierarchies function, see [“Overview” on page 107](#).

For tips and guidelines on setting up hierarchies, see [“CERN Hierarchy Setup” on page 109](#) and [“ICP Hierarchy Setup” on page 111](#).

For more information on caching hierarchies, see one of the TCP/IP references available at any bookstore carrying computer reference manuals.

ICP/CERN Configuration Tab

Path: Hierarchy > ICP/CERN Configuration

Figure 152

Volera™ EXCELERATOR 2.2

ICP/CERN Configuration | Bypass

☐ Must only forward through hierarchy

☐ Enable ICP/CERN client

ICP/CERN parents and peers

Host address/name	Type	Proxy port	ICP port	Priority	Domains
-------------------	------	------------	----------	----------	---------

ICP cache hierarchy timeout: seconds

☐ Enable ICP server

☒ Enable source round trip time

Listening port:

☐ Use ICP multicast

ICP multicast address	Proxy port	ICP port
-----------------------	------------	----------

The ICP/CERN Configuration tab lets you configure the appliance to participate in ICP and CERN hierarchies.

Must Only Forward through Hierarchy: Forces the appliance to always forward requests through the hierarchy. This is required in cases where the appliance would normally bypass the hierarchy. For example, the appliance would not normally access the hierarchy to get a non-cachable object. However, if its only path through the firewall is through the hierarchy, this box must be checked.

Additionally, you can make the appliance more secure behind a firewall by configuring it with a CERN parent that does DNS resolution and by checking this option. Otherwise, the appliance always attempts DNS resolution.

IMPORTANT: Enabling this option can potentially disrupt appliance-based filtering. For more information, see [“Critical Information about Filtering in CERN and ICP Hierarchies” on page 207](#).

Enable ICP/CERN Client: Enables the appliance to function as a child to both CERN and ICP parents and as a peer to other ICP servers in the hierarchy. This box must be checked for the appliance to request and accept data from hierarchies.

ICP/CERN Parents and Peers: A list of other proxy servers you define as peers or parents of the appliance using the [“ICP Parent Dialog Box” on page 408](#), the [“ICP Peer Dialog Box” on page 408](#), and the [“CERN Parent Dialog Box” on page 409](#).

By creating this list, you link the appliance into hierarchical relationships with other members of the hierarchy. You create hierarchical peer and parent relationships between the appliance and other members of the hierarchy.

The appliance's ICP client supports both ICP and CERN access and is dynamically configured, depending on the types of peers and parents (ICP or CERN) defined for the appliance.

NOTE: As you create the list, remember that in ICP you can configure another proxy server as either a peer or as a parent to the appliance, but not both. Only one relationship in the hierarchy is allowed.

ICP Cache Hierarchy Timeout: Specifies how long the appliance waits for a response from ICP servers (parents and peers) before directly requesting that the ICP parent or origin server fill the request. Valid field values range from 0 through 3600 seconds.

Enable ICP Server: Enables the appliance to be accessed as an ICP parent or ICP peer by other appliances in an ICP hierarchy. (A CERN parent is simply a forward proxy serving a second forward proxy [its child] and requires no special configuration.)

IMPORTANT: ICP and CERN servers must also be configured as forward proxy (HTTP) servers. The ICP service locates a URL; the forward proxy service returns the data. For this reason, an appliance that is an ICP server must have forward proxy services activated and IP addresses checked in the Proxy IP Addresses list. (See "[Client Accelerator Tab](#)" on page 333.)

Enable Source Round Trip Time: Enables the appliance to measure entire ICP hierarchy aggregate times for the routes it uses and to determine which routes to pursue first when seeking data. If this box is unchecked, the appliance chooses routes based strictly on local ICP hierarchy round-trip times sorted by priority.

Source round-trip time is an ICP function and is not available in CERN hierarchies.

Listening Port: The port that the appliance listens on for ICP traffic. The industry standard port number for ICP traffic is 3130.

Valid port numbers are 0 through 65535.

Use ICP Multicast: Enables each multicast-enabled appliance in a multicast group to accept multicast requests transmitted to the group's multicast address. A multicast group is a group of ICP peers that communicate with each other about caching information by using a single multicast address you designate.

You can create a multicast group composed of ICP peers by completing the following steps for each member of the group:

- 1** Check both the Enable ICP Client and Enable ICP Server options.
- 2** Check Use ICP Multicast > specify the multicast address you have assigned to the group > specify the ports on which the multicast group handles HTTP proxy and ICP traffic.

Multicast addresses are Class D addresses (the first decimal number in the dotted decimal notation is in the range of 224 to 239, inclusive).

This setting must be identical for each member of the group.
- 3** Click Apply.
- 4** Repeat the process on each appliance that is a member of the multicast group.

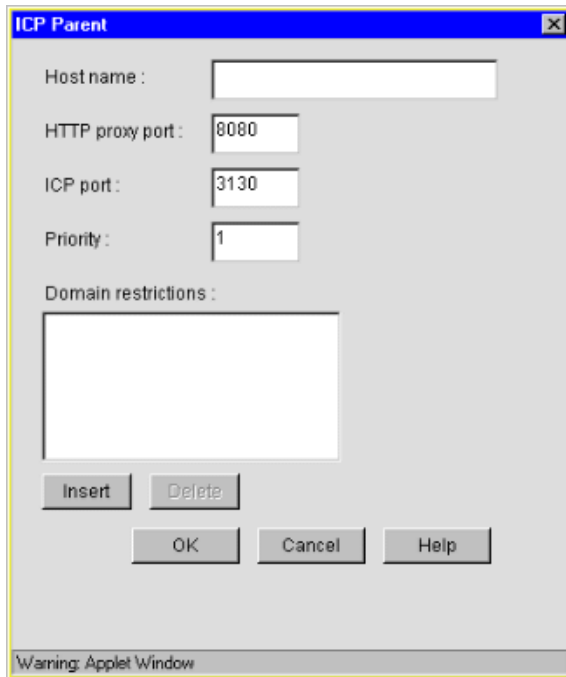
For more information on multicasting, see [the Web \(http://www.erg.abdn.ac.uk/users/gorry/course/intro-pages/uni-b-mcast.html\)](http://www.erg.abdn.ac.uk/users/gorry/course/intro-pages/uni-b-mcast.html).

ICP Access Control: Clicking this button opens the ICP Access Control dialog box. See "[ICP Access Control Dialog Box](#)" on page 410.

ICP Parent Dialog Box

Path: Hierarchy > ICP/CERN Configuration > Enable ICP/CERN Client > ICP Parent

Figure 153

The image shows a Java applet window titled "ICP Parent". It contains several input fields: "Host name:" with an empty text box, "HTTP proxy port:" with a text box containing "8080", "ICP port:" with a text box containing "3130", and "Priority:" with a text box containing "1". Below these is a "Domain restrictions:" label followed by a large empty text area. At the bottom left of the text area are "Insert" and "Delete" buttons. At the bottom right are "OK", "Cancel", and "Help" buttons. A status bar at the very bottom reads "Warning: Applet Window".

The ICP Parent dialog box lets you define an ICP parent for the appliance.

Hostname: The IP address or DNS name of the ICP parent.

HTTP Proxy Port: The port on which the parent services HTTP request.

ICP Port: The port on which the parent services ICP requests.

Priority: The priority the appliance should follow when evaluating which ICP parent to access after the initial request has timed out. Parents with lower priority numbers are accessed first, assuming their source round-trip times are relatively equal.

Domain Restrictions: Determines whether the appliance requests data through this ICP parent. If the request is to one of the domains in this list, the client considers this parent when selecting a server. If the requested domain is not listed, this parent is not considered.

If the list is empty, the client considers the parent for requests to all domains.

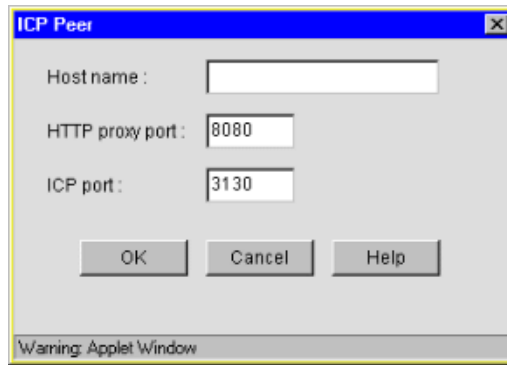
A domain can be a fully qualified domain name (FQDN). If it is, the client considers using the parent only for requests to a specific host. Domain restrictions are used to create virtual hierarchies for expediting request resolution.

Virtual hierarchies should contain parents without domain restrictions (for default request resolution) in addition to the more restrictive hierarchies defined by domain restrictions.

ICP Peer Dialog Box

Path: Hierarchy > ICP/CERN Configuration > Enable ICP/CERN Client > ICP Peer

Figure 154

The ICP Peer dialog box is a standard Java applet window with a blue title bar. It contains three text input fields: 'Host name' (empty), 'HTTP proxy port' (containing '8080'), and 'ICP port' (containing '3130'). Below these fields are three buttons: 'OK', 'Cancel', and 'Help'. At the bottom of the window, there is a warning message: 'Warning: Applet Window'.

The ICP Peer dialog box lets you define an ICP peer for the appliance.

Hostname: The IP address or DNS name of the ICP peer.

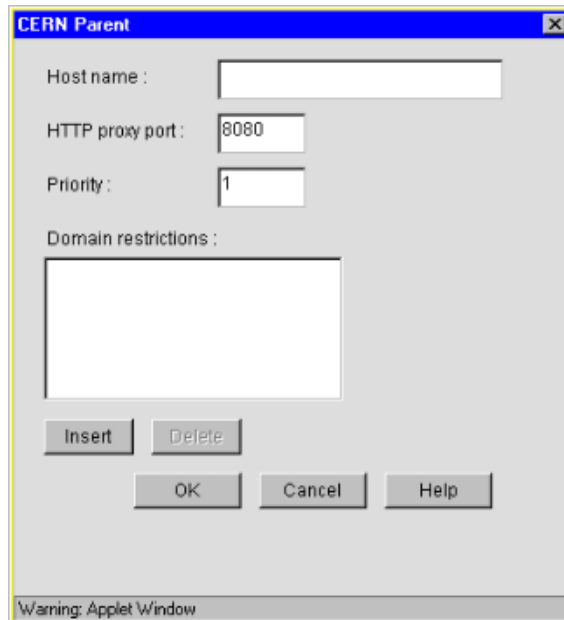
HTTP Proxy Port: The port on which the peer services HTTP requests.

ICP Port: The port on which the peer services ICP requests.

CERN Parent Dialog Box

Path: Hierarchy > ICP/CERN Configuration > Enable ICP/CERN Client > CERN Parent

Figure 155

The CERN Parent dialog box is a standard Java applet window with a blue title bar. It contains four text input fields: 'Host name' (empty), 'HTTP proxy port' (containing '8080'), 'Priority' (containing '1'), and 'Domain restrictions' (empty). Below the 'Domain restrictions' field are two buttons: 'Insert' and 'Delete'. At the bottom of the window are three buttons: 'OK', 'Cancel', and 'Help'. At the very bottom, there is a warning message: 'Warning: Applet Window'.

The CERN Parent dialog box lets you define a CERN parent for the appliance.

Hostname: The IP address or DNS name of the CERN parent.

HTTP Proxy Port: The port on which the parent transmits HTTP requests.

Priority: The priority that the appliance should follow when evaluating which CERN parent to access. Parents with lower numbers are accessed first as the primary route.

If you are creating hierarchical routes using domain restrictions and you define the same route (domain restrictions) for multiple CERN parents, you must ensure that one of these parents is the primary parent for these restrictions. The primary parent has the lowest (meaning the first) priority number.

Domain Restrictions: Determines whether the appliance requests data through this parent. If the request is to one of the domains in this list, the client considers the parent when selecting a server. If the requested domain is not listed, the parent is not considered.

If the list is empty, the client considers the parent for requests to all domains.

A domain can be a fully qualified domain name (FQDN). If it is, the client considers using the parent only for requests to a specific host. Domain restrictions are used to create virtual hierarchies for expediting request resolution.

Hierarchy routes consist of one hierarchy route for parents without domain restrictions and a route for each restricted domain.

Virtual hierarchies should contain parents without domain restrictions (for default request resolution) in addition to the more restrictive hierarchies defined by domain restrictions.

ICP Access Control Dialog Box

Path: Hierarchy > ICP/CERN Configuration > Access Control

Figure 156



The ICP Access Control dialog box lets you define the ICP proxy servers with which the appliance will communicate.

ICP Server Replies to Trusted List Only: If this box is checked, the appliance will allow only those child and peer proxies listed in the ICP Trusted List to access and use its forward proxy services.

If the box is not checked, the appliance will allow access by all ICP neighbors.

ICP Client Accepts Replies from Trusted List Only: If this box is checked, the ICP client will only accept responses from the appliances listed in the ICP Trusted list.

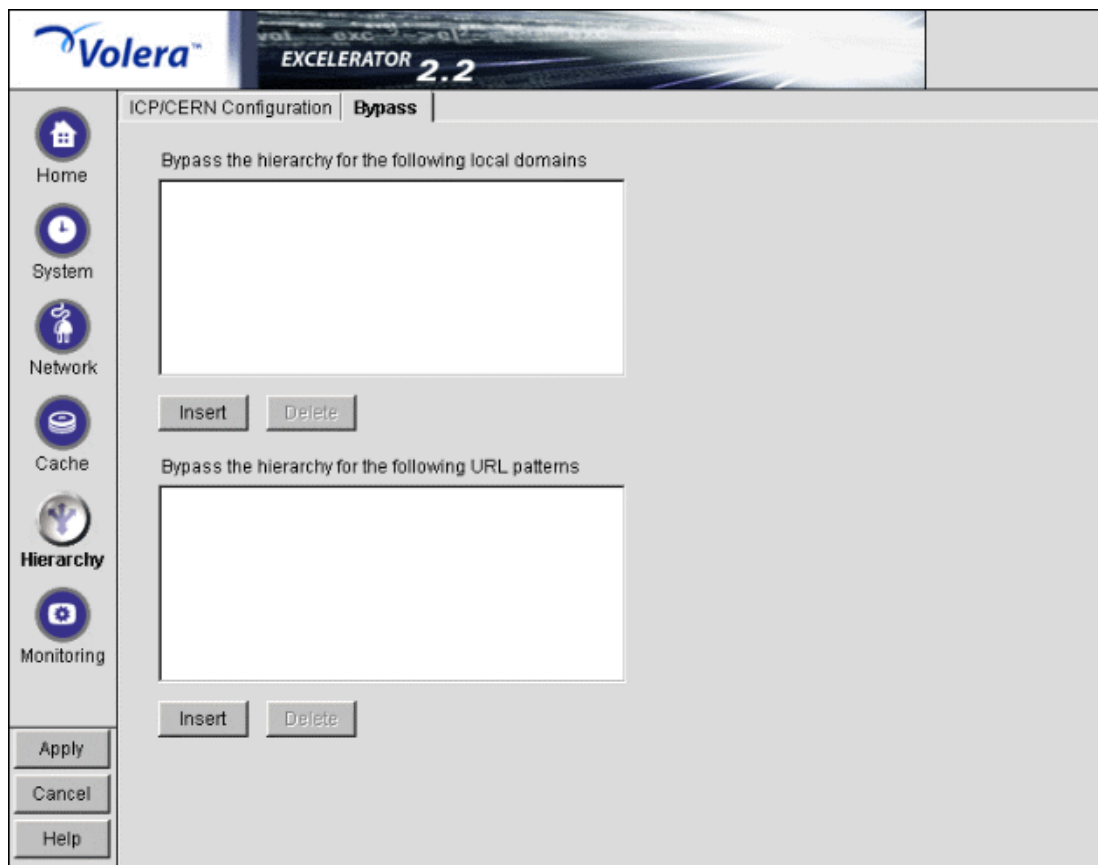
If the box is not checked, the client will accept responses from any of its ICP neighbors.

ICP Trusted List: This list contains the unicast addresses of ICP neighbors and ICP multicast neighbors.

Bypass Tab (Hierarchy)

Path: Hierarchy > Bypass

Figure 157



The Bypass tab lets you provide a mechanism for bypassing the hierarchy when requests match specified domain names or URL path substrings. If the object requested is not available in the appliance's cache and a list entry match is found, the appliance requests the object directly from the Web server.

The appliance uses the local domains and URL patterns lists in relation to both ICP and CERN hierarchies.

Bypass the Hierarchy for the Following Local Domains: This list should contain only domain (DNS) names and subdomain names. The domain names can be complete or partial domain names. Hosts matching the listed domain names are contacted directly rather than through the hierarchy.

Excelerator always processes the local domains list, whether or not the Must Only Forward through Hierarchy option in the **ICP/CERN Configuration Tab** is checked.

IMPORTANT: Entries in the local domains list must not contain wildcard characters or scheme (protocol) identifiers such as http://.

Entries such as http://www.foo.gov, *.foo.gov, and *.gov are not valid.

When a request comes through Excelerator, the domain or hostname specified in the URL is matched against the list to find trailing matches.

For example, the entry www.foo.gov matches only the www.foo.gov host, while the entry gov matches all hosts whose DNS names end with .gov.

Use this list to specify domains that are in good topology relationships with the appliance as well as domains that are within a firewall with the appliance and are therefore directly reachable by the appliance.

Bypass the Hierarchy for the Following URL Patterns: This list contains paths or path substrings for Excelerator to match against URL path statements.

IMPORTANT: Entries in the URL patterns list must not contain wildcard characters, scheme (protocol) identifiers such as http://, or DNS names such as www.foo.gov.

The entry http://www.foo.gov/html is not valid because it contains a scheme identifier and a DNS name. The entry *gov/html is not valid because it contains a wildcard.

Excelerator matches the entries in the URL patterns list against the paths of requested URLs.

For example, the entry gov/html would match all of the following URLs:

- ◆ http://www.foo.gov/gov/html
- ◆ http://www.goo.gov/fedgov/html/page
- ◆ http://www.hoo.gov/pages/yourgov/htmlpage

It would not, however, match the URL http://www.foo.gov/gov/htm because the gov/html substring is not found.

If you have URLs that don't work properly with hierarchies, you can avoid problems by listing their path patterns in the URL patterns list.

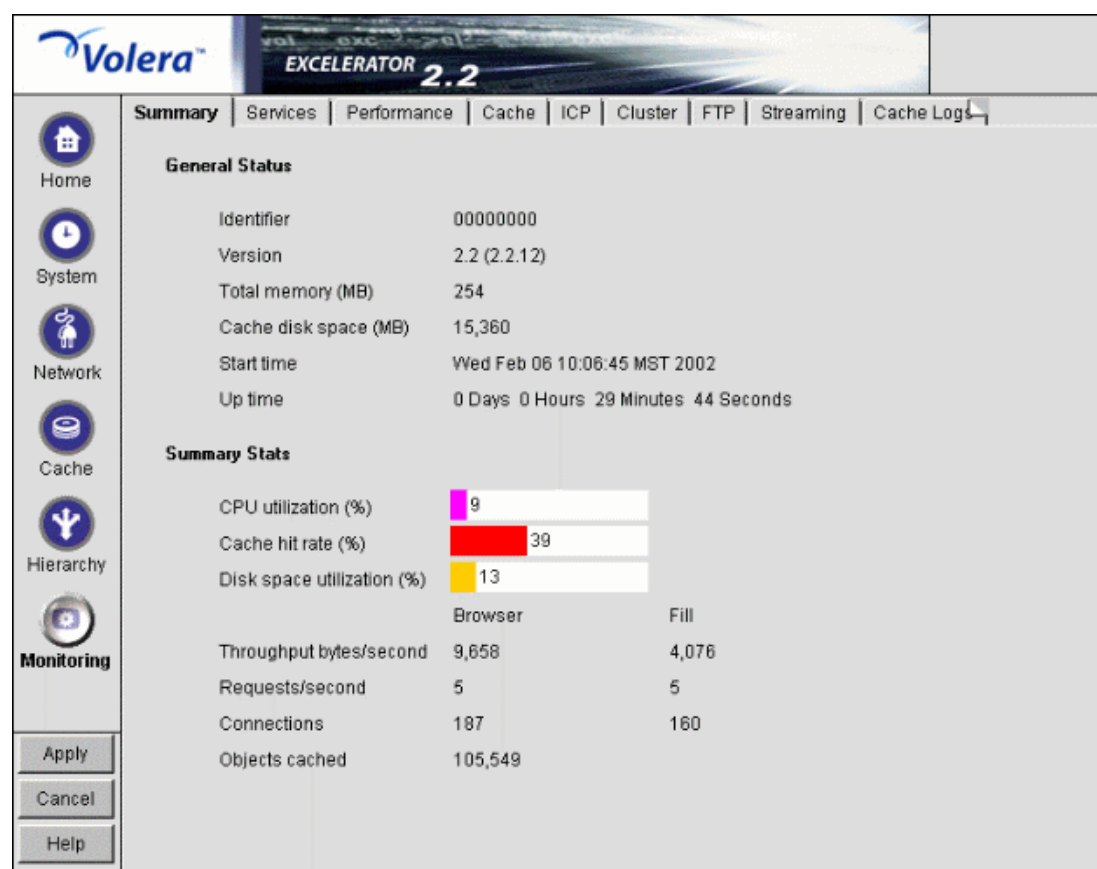
48 Using the Monitoring Panel

You can monitor various appliance activities and statistics in the browser-based tool.

Summary Tab

Path: Monitoring > Summary

Figure 158



The Summary tab shows key appliance statistics at a glance. Statistics are refreshed every second.

Identifier: The make, model, and serial number of the appliance.

Version: The current system software version.

Total Memory (MB): Total available memory.

Cache Disk Space (MB): Total disk space available for caching. The amount shown is smaller than the total appliance disk space because it doesn't include the operating system and log partitions. Check this field to verify whether Excelerator has detected all disks installed on the appliance.

Start Time: The last time the appliance was started.

Up Time: Total time the appliance has been running since last started.

CPU Utilization (%): The current CPU utilization rate. Use this chart for capacity planning.

Cache Hit Rate (%): The current cache hit rate. A high cache hit rate indicates that the caching system is off-loading significant request processing from Web servers whose objects have been cached. Use this chart for capacity planning.

Disk Space Utilization (%): The percentage of caching disk space currently in use.

Throughput Bytes/Second: Current throughput.

Requests/Second: The rate at which browser clients are requesting Web objects.

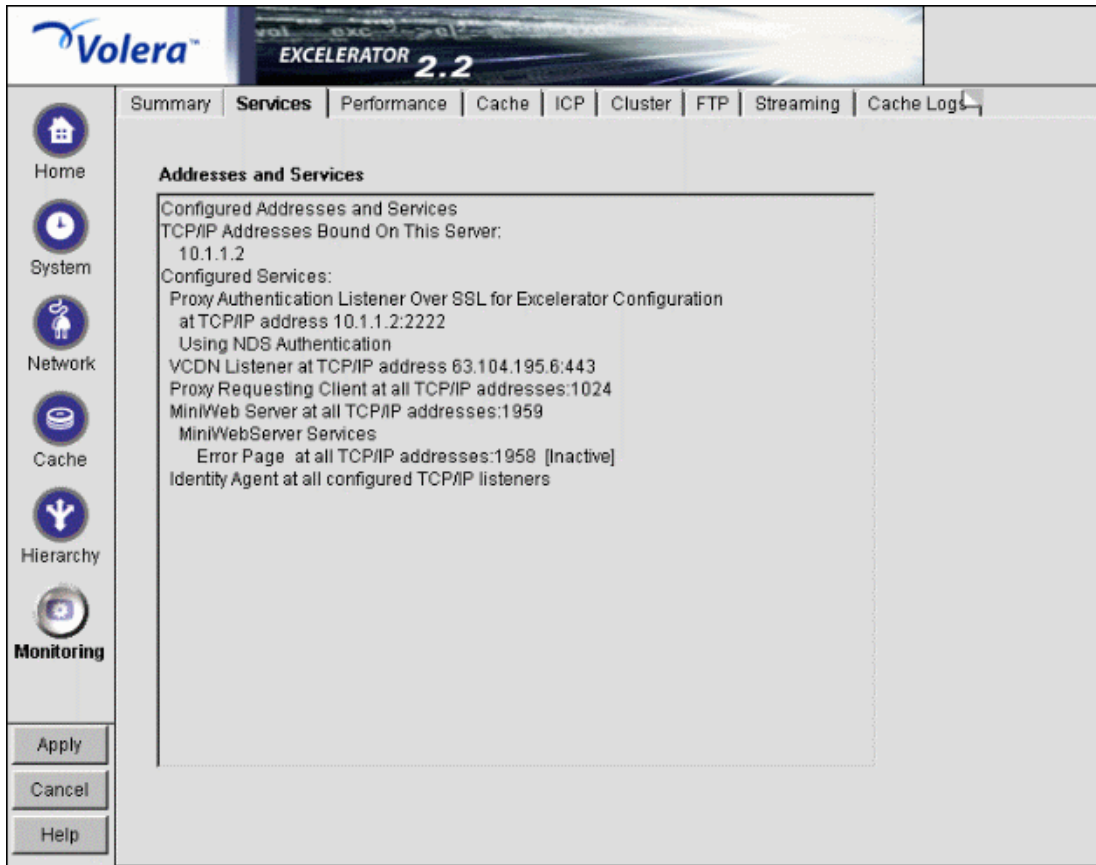
Connections: The total number of TCP connections that are active, idle, or closing.

Objects Cached: The total number of Web objects that have been cached.

Services Tab

Path: Monitoring > Services

Figure 159



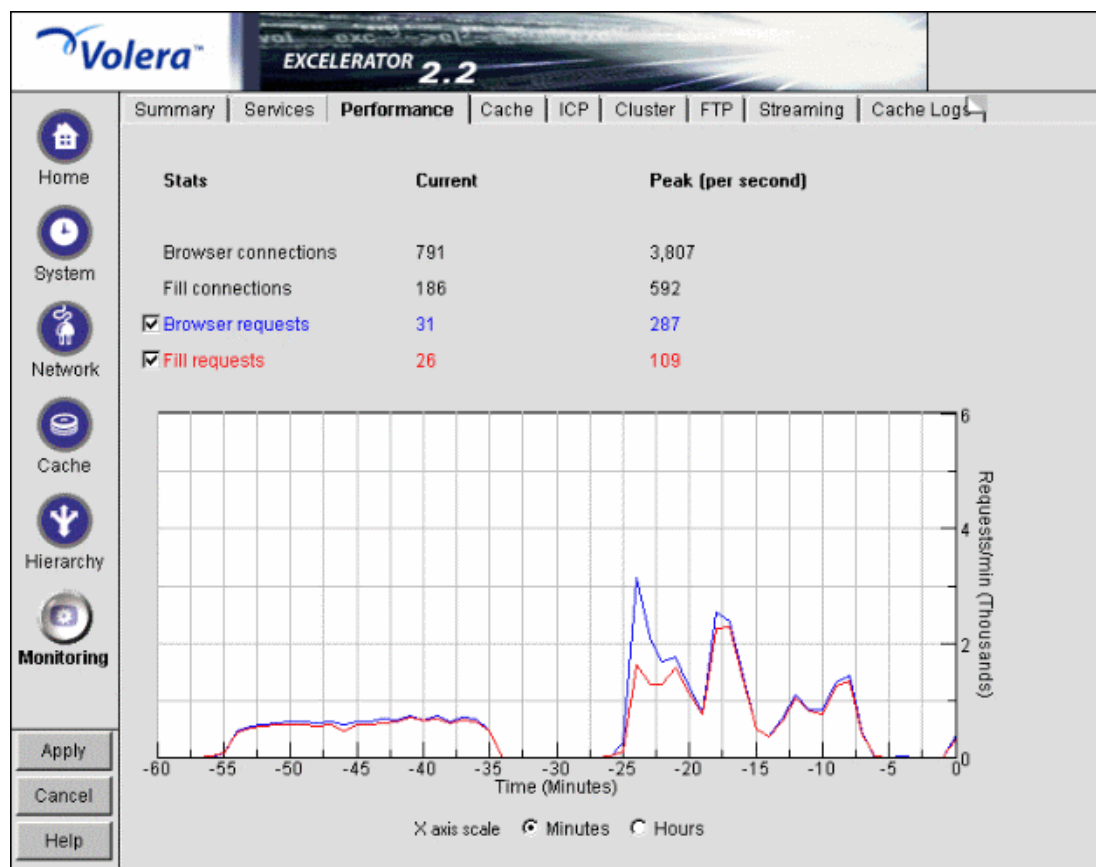
The Services tab shows you the IP addresses that are bound to appliance network cards and the services that are active. This information is refreshed every minute.

Addresses and Services: Displays active services along with the IP addresses and ports that they are running on. When the appliance detects errors, it displays appropriate error messages next to the services. Use this list for troubleshooting problems with configured services.

Performance Tab

Path: Monitoring > Performance

Figure 160



The Performance tab shows current and peak levels of usage in terms of TCP connections and HTTP requests. The tab also displays a graph of HTTP requests from browsers to the appliance and from the appliance to origin Web servers.

Statistics are updated every ten seconds. The graph is updated once a minute.

Browser Connections: The current and peak numbers of browser connections to the appliance.

Fill Connections: The current and peak numbers of connections that the appliance has opened to origin Web servers.

Browser Requests: The current and peak numbers of browser HTTP requests per second made to the appliance. Check this box to enable graphing of browser requests.

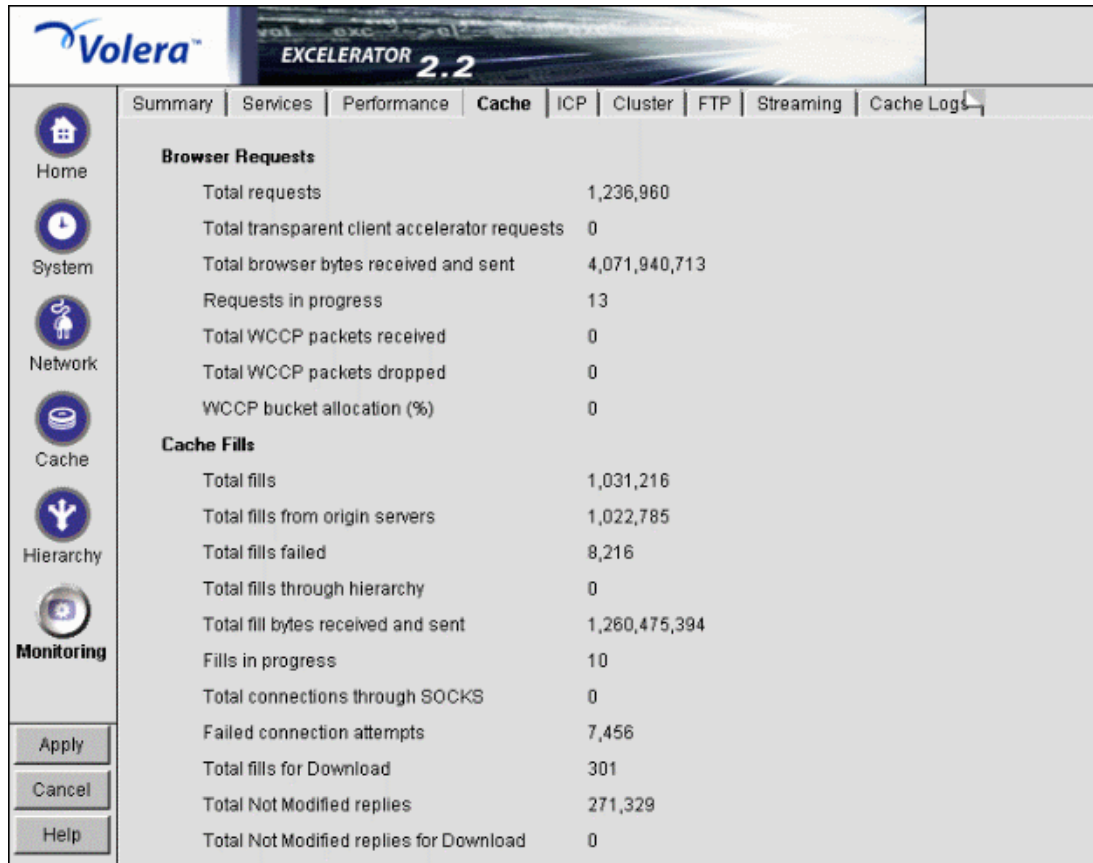
Fill Requests: The current and peak numbers of appliance requests per second to origin Web servers. Check this box to enable graphing of requests to origin Web servers.

Requests Graph: The HTTP browser requests to the appliance per minute (blue line) and HTTP fill requests to origin Web servers per minute (red line). Click Minutes or Hours to select the scale of the X axis. Minutes displays a one-hour history; Hours shows a 24-hour view.

Cache Tab

Path: Monitoring > Cache

Figure 161



The Cache tab shows statistics for browser requests to the appliance and for appliance requests to origin Web servers. Statistics are refreshed every ten seconds.

Total Requests: The total number of requests that browser clients have made since the appliance was started.

Total Client Accelerator Requests: The total number of client accelerator browser requests that came directly to the appliance or were routed to it through an L4 switch.

Total Browser Bytes Received and Sent: The total bytes that browser clients have sent to and received from the appliance.

Requests in Progress: The number of active browser requests that are currently being processed by the appliance.

Total WCCP Packets Received: The total number of packets that have been redirected to the appliance by a WCCP-capable router.

Total WCCP Packets Dropped: The total number of packets routed to the appliance from a WCCP-capable router that were dropped by Excelsator.

Packets are dropped for one of two reasons: either Excelsator did not expect the packets to be redirected or the packets were malformed. If the transparent proxy service has an Exception IP addresses list, Excelsator does not expect the router to redirect packets bound for addresses in the list. Use this field to troubleshoot operational problems with WCCP.

WCCP Bucket Allocation (%): The percentage of hash buckets allocated for this appliance. The higher the percentage, the more requests the router redirects to this appliance. If the percentage is 0, the router does not redirect requests to the appliance. Use this field to troubleshoot operational problems with WCCP.

Total Fills: The total number of fill requests the appliance has made to origin Web servers and to neighbors in its cache hierarchy.

Total Fills from Origin Servers: The total number of fill requests the appliance has made to origin Web servers.

Total Fills Failed: The total number of failed fill requests.

Total Fills through Hierarchy: The total number of fill requests the appliance has made to neighbors in its cache hierarchy.

Total Fill Bytes Received and Sent: The total bytes the appliance has sent and received in order to fill its Web object cache.

Fills in Progress: The number of active fill requests that the appliance is waiting for.

Total Connections through SOCKS: The total number of connections the appliance has made through a firewall in order to fill its Web object cache.

Failed Connection Attempts: The total number of failed connection attempts the appliance has made while attempting to fill its Web object cache.

Total Fills for Download: The total number of fill requests sent to origin servers that originated from system read-ahead functionality and batch downloads.

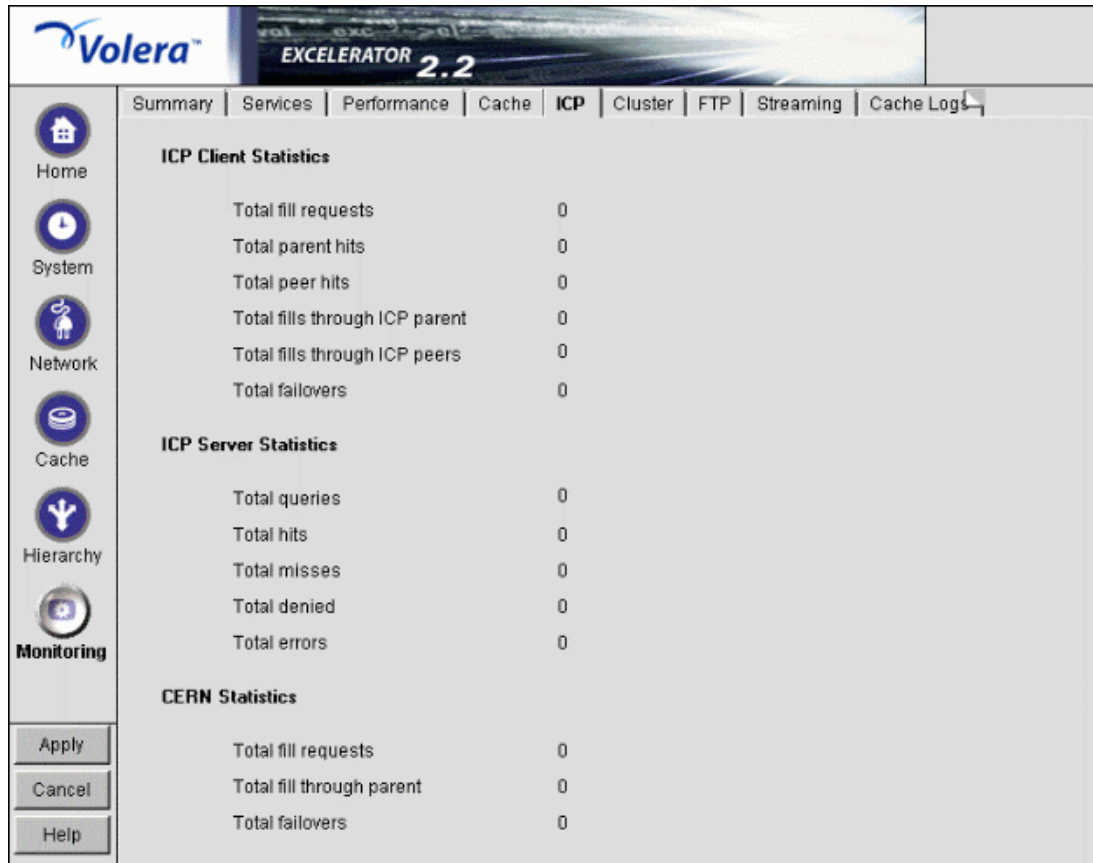
Total Not Modified Replies: The total number of 304 Not Modified replies received for all fill requests to origin servers.

Total Not Modified Replies for Download: The total number of 304 Not Modified replies received for fill requests to origin servers that originated from system read-ahead functionality and batch downloads.

ICP Tab

Path: Monitoring > ICP

Figure 162



The ICP tab shows you statistics relating to Web cache hierarchies. Statistics are refreshed every ten seconds.

ICP Client Statistics

Total Fill Requests: The total number of requests the appliance has made to fill its Web object cache through an ICP hierarchy since it was started.

Total Parent Hits: The total number of times that parents reported cache hits.

Total Peer Hits: The total number of times that peers reported cache hits. Because peers are usually in close proximity, a high number of peer hits indicates reduced time and bandwidth costs for filling the Web object cache.

Total Fills through ICP Parent: The total number of times the appliance chose one of its ICP parents to retrieve Web objects after its initial request timed out.

Total Fills through ICP Peers: The total number of times the appliance chose one of its ICP peers to retrieve Web objects after its initial request timed out. Because peers are usually in close proximity, a high number of fills through peers indicates reduced time and bandwidth costs for filling the Web object cache.

Total Failovers: The total number of times the appliance had to choose another neighbor to fetch Web objects because a direct request to a neighbor failed. A number of failovers could indicate unreliable network connections, faulty hierarchy configuration, or misbehaving neighbors.

ICP Server Statistics

Total Queries: The total number of ICP queries that the appliance has received from its ICP neighbors.

Total Hits: The total number of times that ICP-requested Web objects were found in the appliance's cache.

Total Misses: The total number of times that ICP-requested Web objects were not found in the appliance's cache.

Total Denied: The total number of times that ICP queries from neighbors were rejected because their IP addresses were not listed in the Access Control List. This might indicate a configuration problem; see [“ICP/CERN Configuration Tab” on page 405](#).

Total Errors: The total number of times the appliance received ICP queries which were malformed (included wrong op-codes).

CERN Statistics

Total Fill Requests: The total number of requests the appliance has made to fill its Web cache through a CERN hierarchy since it was started.

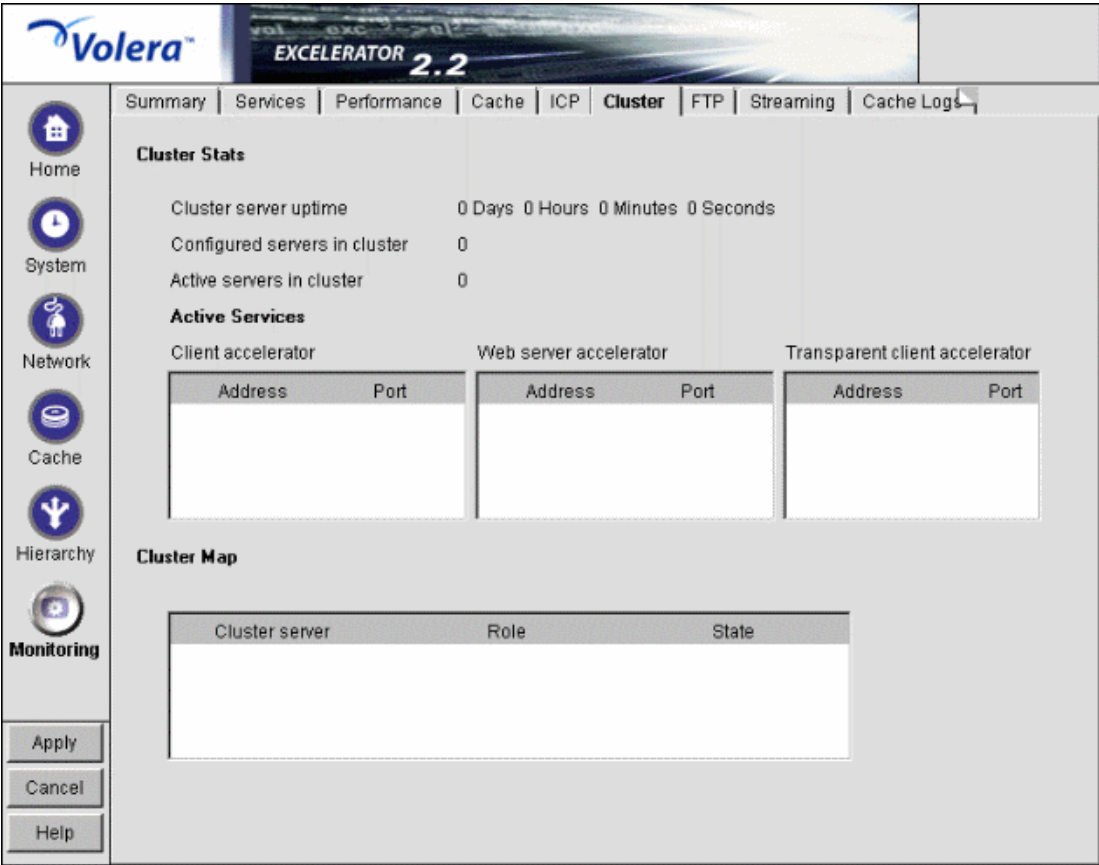
Total Fill through Parent: The total number of times the appliance chose one of its CERN parents to retrieve Web objects.

Total Failovers: The total number of times the appliance had to choose another route to fill a request because a CERN parent failed to return the objects requested. A number of failovers could indicate unreliable network connections, faulty hierarchy configuration, or malfunctioning CERN parents.

Cluster Tab

Path: Monitoring > Cluster

Figure 163



The Cluster tab shows status and statistics for the servers and services running in a cluster. Cluster information is refreshed every minute.

Cluster Server Uptime: The up-time of clustered services on this appliance.

Configured Servers in Cluster: The number of servers configured for the cluster.

Active Servers in Cluster: The number of servers currently active in the cluster.

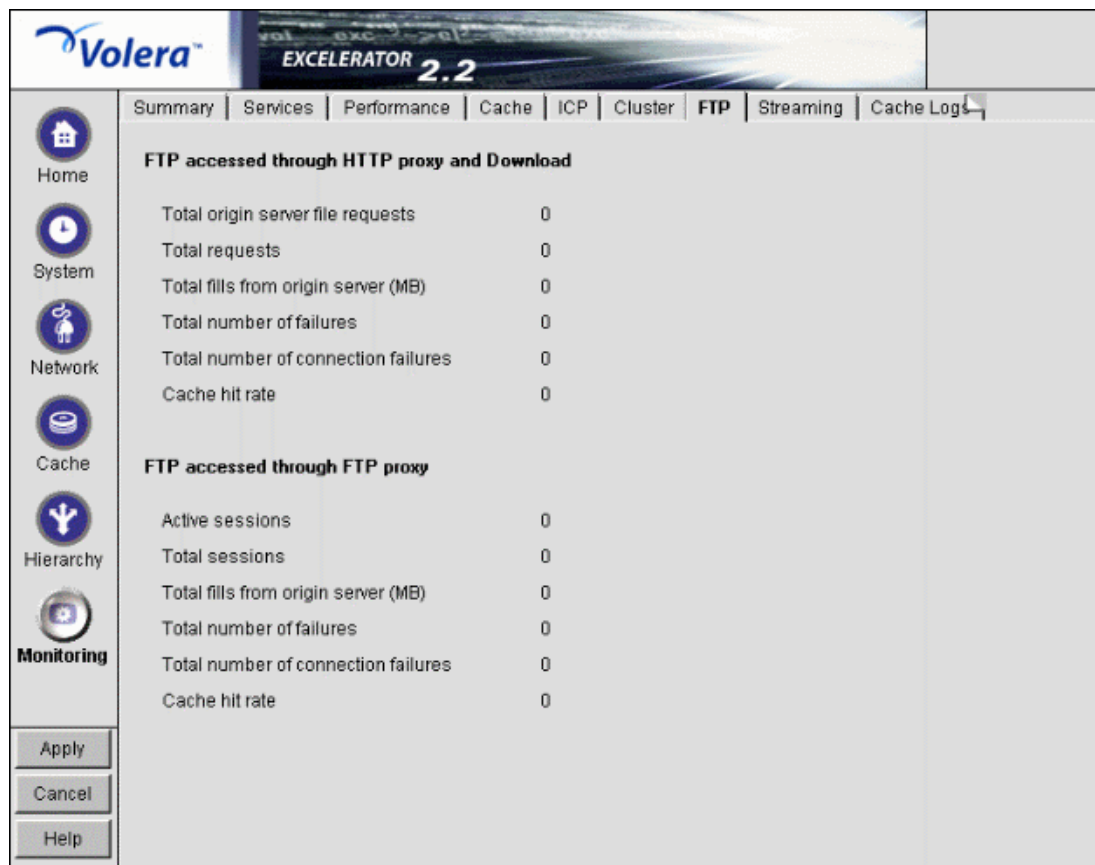
Active Services: The addresses and ports on which this appliance is hosting client accelerator, Web server accelerator, and transparent client accelerator.

Cluster Map: The role and state of each configured server in the cluster. Possible roles include Active, Standby, and Offline. Possible states are ServerUp, ServerDown, ServerJoining, and ServerLeaving.

FTP Tab

Path: Monitoring > FTP

Figure 164



The FTP tab shows key statistics for all FTP transactions. FTP requests through HTTP proxy are displayed separately from those accessed through FTP proxy. Statistics are refreshed every five seconds.

FTP Accessed through HTTP Proxy and Download: These statistics reflect FTP requests from browsers using the HTTP protocol.

- ◆ *Total Origin Server File Requests:* Number of files requested from browsers.
- ◆ *Total Requests:* Total number of FTP requests since the appliance was started.
- ◆ *Total Fills from Origin Server (MB):* Total number of megabytes transferred to cache.
- ◆ *Total Number of Failures:* Total number of errors, including memory allocation errors, connection establishment failures, connection aborts, and login failures.
- ◆ *Total Number of Connection Failures:* Total number of connection establishment failures.
- ◆ *Cache Hit Rate:* Total number of times FTP requests were satisfied from cache to the total number of FTP requests received, as a percentage.

FTP Accessed through FTP Proxy: These statistics reflect FTP requests that used the native FTP protocol.

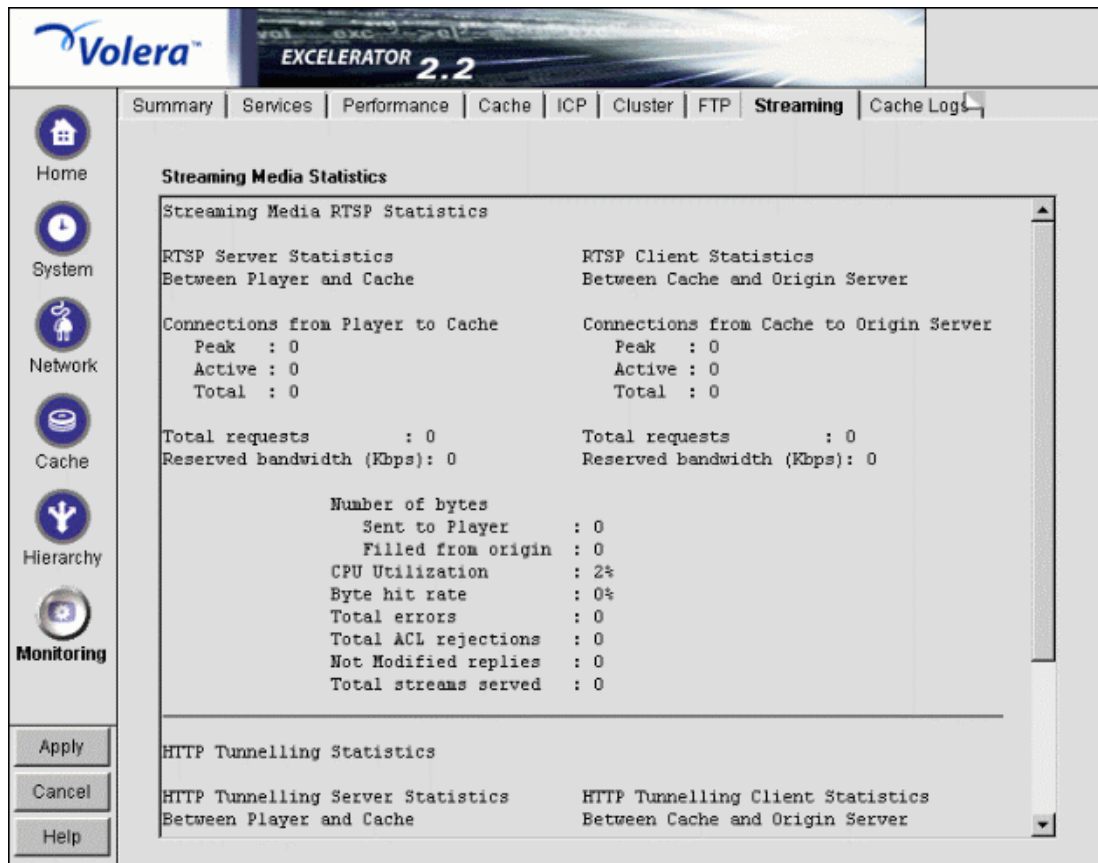
- ◆ *Active Sessions:* Total number of FTP sessions active on the system
- ◆ *Total Sessions:* Total number of sessions using FTP
- ◆ *Total Fills from Origin Server:* Total number of fills in megabytes transferred to cache

- ♦ *Total Number of Failures*: Total number of errors, including memory allocation errors, connection establishment failures, connection aborts, and login failures
- ♦ *Total Number of Connection Failures*: Total number of connection establishment failures
- ♦ *Cache Hit Rate*: Total number of times FTP requests were satisfied from cache to the total number of FTP requests received, as a percentage

Streaming Tab

Path: Monitoring > Streaming

Figure 165



The Streaming Media tab shows streaming media usage on the appliance. It can be used to monitor usage, verify that bandwidth limits are set appropriately, and determine what percentage of bytes are being delivered from cache.

Streaming RTSP Statistics

RTSP Server Statistics

Connections from Player to Cache: These statistics are for the RTSP/RTP connections from clients (players) to the appliance. If QuickTime HTTP Tunnel Caching is enabled on the **Streaming Tab**, HTTP tunneled statistics are included.

- ♦ *Peak*: The peak number of active connections from players.
- ♦ *Active*: The current number of active connections from players.
- ♦ *Total*: The total number of connections from players since the appliance was started.

Total Requests: The total number of player requests received since the appliance was started.

Active Bandwidth (Kbps): The downstream bandwidth currently being used.

RTSP Client Statistics

Connections from Cache to Origin Server: These statistics are for the RTSP/RTP connections from the appliance to the origin streaming servers.

NOTE: The appliance fills audio and video tracks separately, so a single connection from a player will often result in two (and sometimes more) connections to the origin streaming server.

- ♦ *Peak*: The peak number of active connections to origin streaming servers.
- ♦ *Active*: The current number of active connections to origin streaming servers.
- ♦ *Total*: The total number of connections to origin servers since the appliance was started.

Total Requests: The total number of appliance requests to origin streaming servers since the appliance was started.

Active Bandwidth (Kbps): The upstream bandwidth currently being used.

General RTSP Statistics

Number of Bytes: The total number of streaming bytes processed by the appliance traffic since it was started.

- ♦ *Sent to Player*: The total number of bytes sent to players since the appliance was started.
- ♦ *Filled from Origin*: The total number of bytes filled from origin servers since the appliance was started.

CPU Utilization: The system CPU utilization.

Byte Hit Rate: The percentage of bytes which were already in cache when requested. This is more meaningful than overall hit rate, because streaming objects can be very large and might be partially cached.

Total Errors: The total number of errors since the appliance was started.

Total ACL Rejections: The total number of connections rejected due to IP ACL controls since the appliance was started.

Not Modified Replies: The total number of not-modified messages since the appliance was started.

Total Streams Served: The total number of streams served since the appliance was started.

HTTP Tunneling Statistics (Pass Through Only)

These statistics reflect content which passes uncached through the appliance.

This occurs under the following conditions:

- ♦ The QuickTime HTTP Tunnel Enable checkbox is not checked.

- ♦ A compatible MediaCache service is not available to cache the content.

HTTP Tunneling Server Statistics

Total Connections: Total number of HTTP tunneling pass-through connections from players since boot time.

HTTP Tunneling Client Statistics

Total Connections: Total number of HTTP tunneling pass-through connections to origin servers since boot time.

General HTTP Tunneling Statistics

Number of Bytes: The total number of bytes that have passed through the appliance since it was started.

- ♦ *Sent to Player:* The total number of bytes sent to players using HTTP tunneling pass-through mode since the appliance was started.
- ♦ *Filled from Origin:* Total number of bytes received from origin streaming servers using HTTP tunneling pass-through mode since the appliance was started.

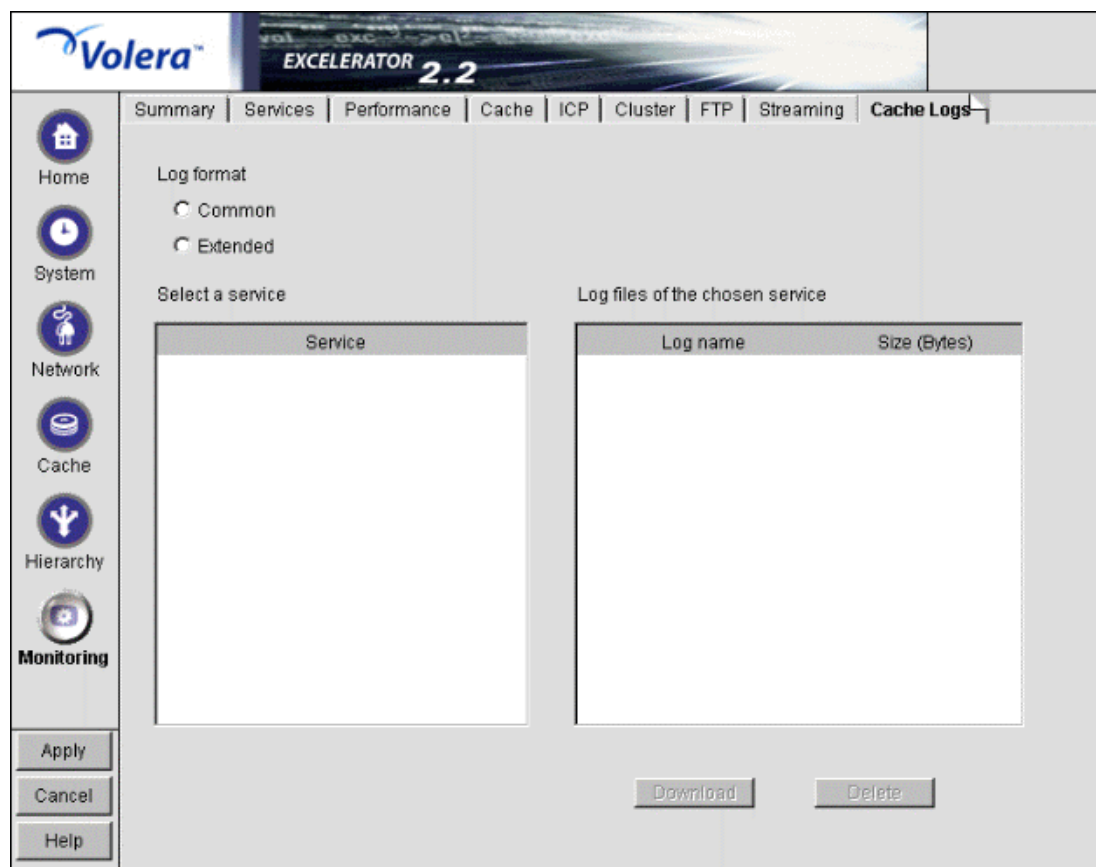
CPU Utilization: System CPU utilization.

Byte Hit Rate: This will always be zero because HTTP tunneling pass-through content is never cached.

Cache Logs Tab

Path: Monitoring > Cache Logs

Figure 166



The Cache Logs tab provides access to logs by format and service.

Log Format: These options let you choose the format of the logs you want to download and view.

Select a Service: Clicking a service name displays the associated logs in the Log Files of the Chosen Service list.

Log Files of the Chosen Service: Contains a list of log files matching the format and service options you have selected.

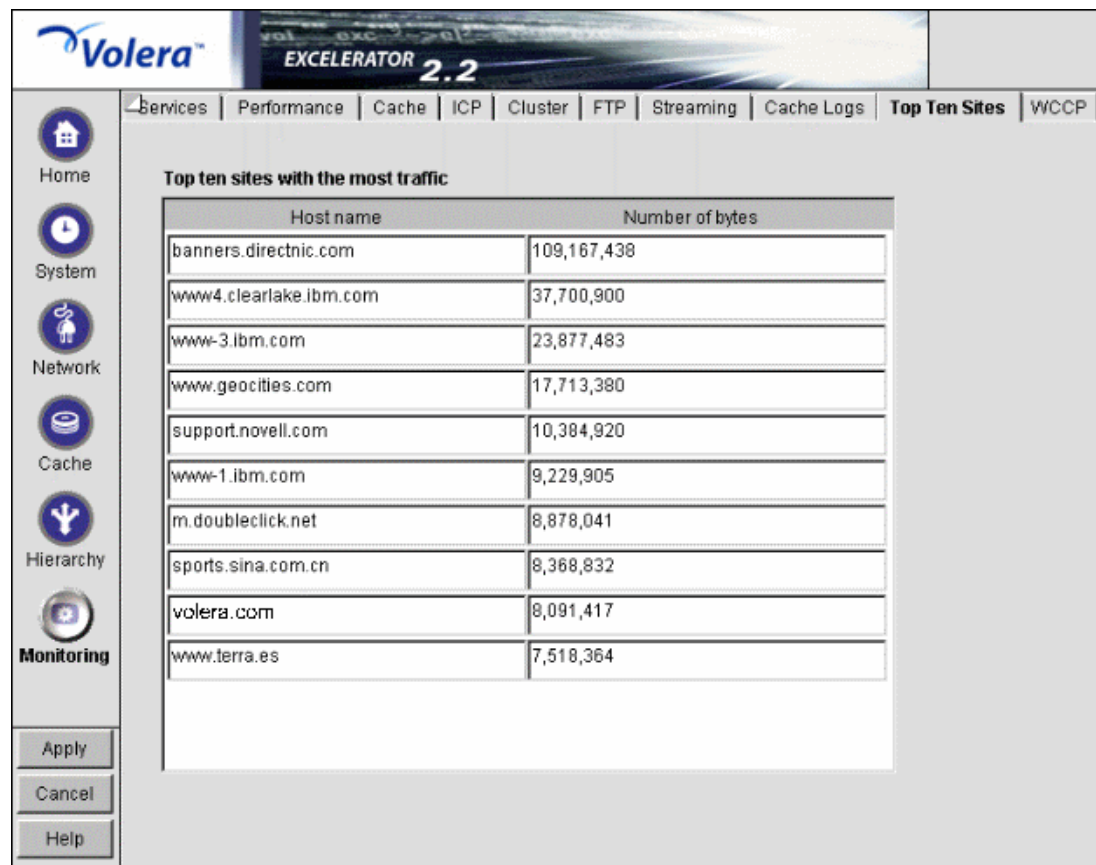
Download: Loads the log file into a separate browser window.

Delete: Removes the log file from the appliance.

Top Ten Sites Tab

Path: Monitoring > Top Ten Sites

Figure 167



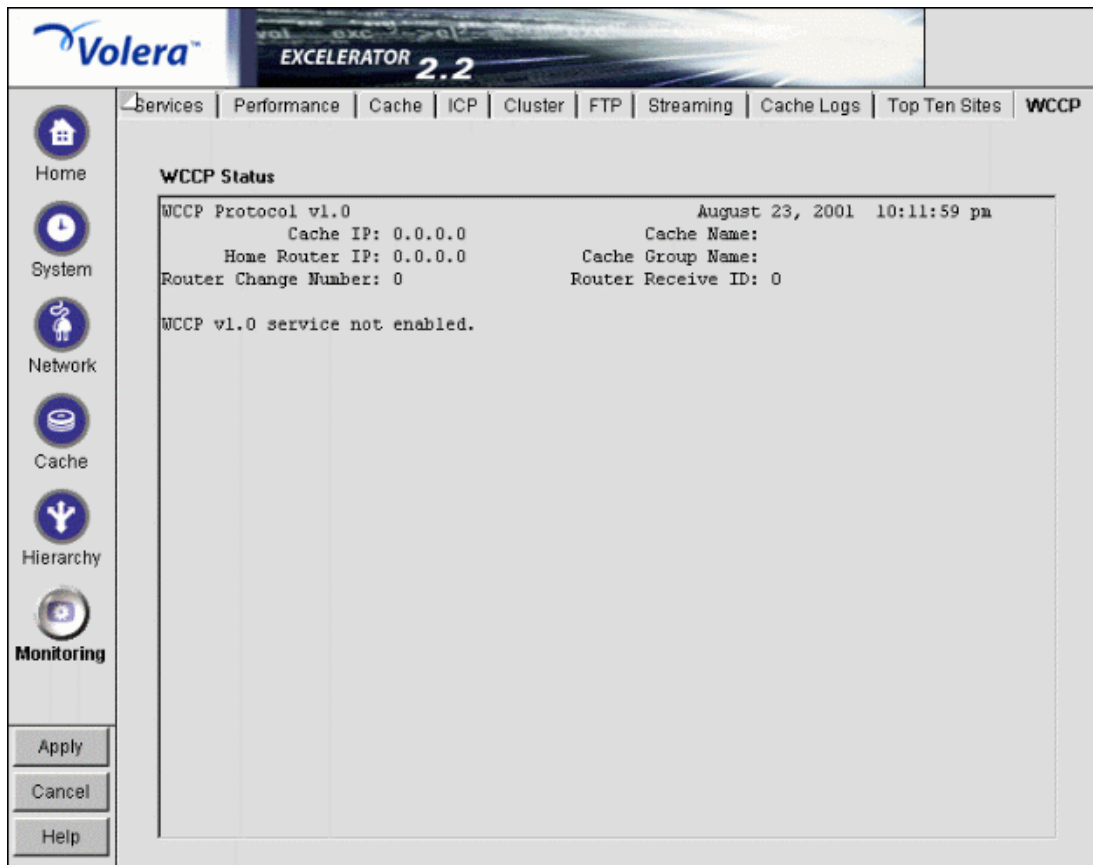
The Top Ten Sites tab displays a list of origin Web servers with more than 0 bytes cached on the appliance. The ten sites with the most total bytes cached are sorted in descending order.

NOTE: Only HTTP entries are listed on this tab.

WCCP Tab

Path: Monitoring > WCCP

Figure 168



The WCCP tab displays WCCP general status. The fields are updated continually to reflect the current WCCP status.

VIII

Reference Guides

The following table summarizes the tasks covered in this section.

To	See
Learn about the appliance's command line interface	Appendix A, "Command Line Reference," on page 431
Set up a Telnet or null-modem connection with the appliance	Appendix B, "Connecting Through Telnet," on page 433
Learn how to upgrade the appliance	Appendix C, "Upgrading the Appliance," on page 441

A

Command Line Reference

If you're working in a Telnet session or from the command line, you have 23 appliance commands available, several of which can be used with various parameters and values. These commands are listed in the table below.

Exceclerator includes help for all commands and their associated arguments and parameters.

To see a list of commands, enter **help** at the command line.

To see a list of arguments for a command, enter **help command**.

To get help for a specific command/argument combination, enter **help command argument**.

Command	Function	Requires Arguments?
add	Adds the new value to the current value.	Yes
apply	Applies the changes made at the command line. Some changes require a system restart, some merely suspend proxying and then restart, and others do not interrupt any process.	No
cancel	Discards all changes that are pending since the last apply command.	No
clear	Removes all items from a list or all settings from an argument.	Yes
clearscreen	Clears the current screen.	No
export	Exports the named file.	Yes
factorysettings	Restores the appliance to original factory settings.	No
get	Displays current settings.	Yes
Health	Displays the appliance health status.	No
Help	Displays a list of available commands.	No
identity	Displays the appliance manufacturer, serial number, and hardware configuration.	No
import	Imports the named file.	Yes
ping	Sends a ping request to the addresses specified. Ports are optional.	Yes

Command	Function	Requires Arguments?
purgecache	Purges the cache buffers.	No
purgednscache	Purges the DNS cache buffers. DNS requests for hostnames that fail to resolve are cached for a short time as negative lookups. IP Addresses used in place of hostnames are also cached.	No
remove	Deletes the specified value.	Yes
reset learnedroutes	Resets the internal router table when the appliance is acting as a router.	No
restart	Restarts the appliance. All proxying ceases until the system restarts.	No
restorefromclones	Replaces the current appliance image with the clone image.	No
scan	Rescans disk drives on the appliance.	No
set	Sets an option by executing a clear command followed by an add command. (Existing settings are cleared when the set command is used.)	Yes
shutdown	Shuts down the appliance. All functionality ceases.	No
updateclones	Replaces the clone image with the current appliance image.	No
version	Displays the current version.	No

Troubleshooting the Command Line

Commands entered return an error

- ☐ Make sure you use the equal (=) sign when setting or adding, as in the following example:

```
set forward enable=yes
```

- ☐ Try the command again. Sometimes a command will fail the first time it's entered.

I made several changes that don't appear when I use the GET command to display them

- ☐ You must conclude with the apply command to make the values take effect.

B

Connecting Through Telnet

You can manage the Excelerator appliance using commands from a workstation that has a Telnet connection to the appliance.

IMPORTANT: Telnet access is not secure unless a password is set. We strongly recommend you set system passwords as part of the appliance initialization process. For more information, see [“Change Password Dialog Box” on page 305](#).


Starting a Telnet Session

This section assumes you are using a Windows* 95/98 or Windows NT* 4 workstation. If you are using another type of workstation, use the following steps as general guidelines to configure your Telnet software to work with the appliance.

IMPORTANT: The appliance Telnet connection supports only the VT-100 terminal type.


- 1** Ensure that you have a network connection between the workstation that you are running Telnet on and the appliance.
- 2** Click Start > run TELNET.EXE.
- 3** From the Telnet screen, click Terminal > Preferences.

Figure 169



- 4** Under Terminal Options, check VT 100 Arrows.
- 5** Under Emulation, check VT-100/ANSI.
- 6** Set the Buffer Size to 25.
- 7** (Optional) Set any of the other preferences that you desire.
- 8** Click OK.
- 9** Click Connect > Remote System.

Figure 170



- 10 Enter the appliance IP address in the Host Name field > select Telnet for the port and VT100 for the terminal.

The uppercase VT100 option usually works better with the appliance than the lowercase vt100 option.

- 11 Click Connect.
- 12 Type the Config user password.

All appliance passwords are case-sensitive.

You can now use all console commands to manage the appliance.

Setting Up an Appliance Using Telnet

The *Getting Started* guide explains how to initialize an appliance and configure forward proxy services using the browser-based tool. The following procedures explain how to complete this task from the command line:

- 1 Start a Telnet session on the client machine.

For help starting a Telnet session, see [“Starting a Telnet Session” on page 433](#).

The starting address for Telnet should be 10.1.1.1, which is the address of eth0 on the appliance.

IMPORTANT: Versions later than 1.0 have no default password for Telnet access. Telnet is not secure unless a password is set for the Config user. (Telnet doesn't provide View user access.)

We strongly recommend you set system passwords as part of the initialization process. For more information, see [“Change Password Dialog Box” on page 305](#).

Telnet always prompts for a password. If you have not set a password for the Config user, enter a null password by pressing Enter.

After logging in to Telnet, you see the following prompt:

```
Internet Caching System
```

```
>
```

- 2 At the Internet Caching System prompt, enter the following:

```
set eth1 address=iii.iii.iii.iii, mask=mmm.mmm.mmm.mmm
set dns server=ddd.ddd.ddd.ddd
set dns domain=x
set gateway nexthop=ggg.ggg.ggg.ggg, metric=t
apply
```

The variables i = the IP address, m = the subnet mask, d = the DNS server IP address, x = your domain name, g = the gateway IP address, and t = the number of hops to the next hop.

3 (Optional) Configure the appliance to provide forward proxy service.

At the Internet Caching System prompt, enter the following:

```
add forward address=iii.iii.iii.iii
apply
```

The variable i = the IP address you entered in [Step 2](#).

The appliance is now configured to begin providing forward proxy service. To configure client browsers to use the forward proxy service, see the browser vendor's instructions.

Additional Information

If the appliance has a monitor attached, you will notice that commands issued through a Telnet connection are echoed on the appliance monitor.

If you get a message asking whether you want the X-session displayed on a display other than the default, you have selected the wrong terminal type. Click Connect > click Disconnect > repeat the connection procedure starting with [Step 9 on page 433](#) and ensuring you have selected the VT100 terminal type in [Step 10 on page 434](#).

Disabling Telnet Access

In Excelerator 2.3, you can disable and re-enable Telnet access to a cache device.

The action is immediately effective and any active Telnet connections are immediately terminated.

To disable Telnet access to a cache device, enter the following commands at the System prompt:

```
set listener telnet enable=no
apply
```

To re-enable Telnet access, enter the following commands:

```
set listener telnet enable=yes
apply
```

Establishing a Null-Modem Connection

This section assumes you are using a Windows 95/98 or NT 4 workstation. If you are using another type of workstation, use the following steps as general guidelines to configure your terminal emulation software to work with the appliance.

IMPORTANT: The appliance null-modem connection supports only the ANSI terminal type.

To establish a null-modem connection with the appliance, do the following:

- 1** Connect a null-modem cable to the serial port on the workstation and the appliance.
- 2** Click Start > run hypertrm.exe.

Figure 171



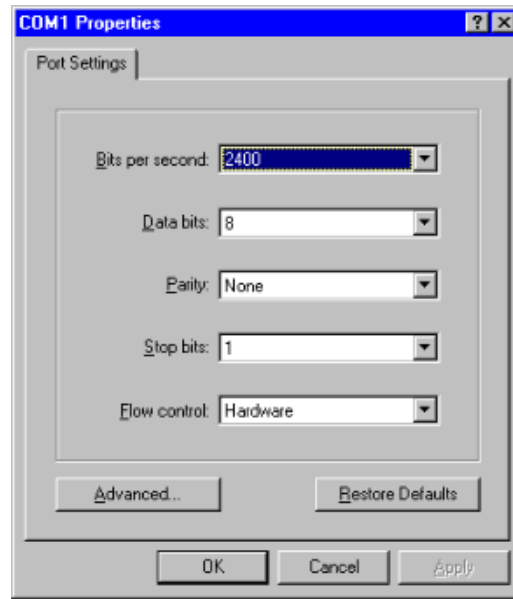
- 3** Enter a name for the connection > select an icon as instructed > click OK.

Figure 172



- 4** Click the Connect Using drop-down list > select a Direct to Com option corresponding to the serial port connection on your workstation > click OK.

Figure 173



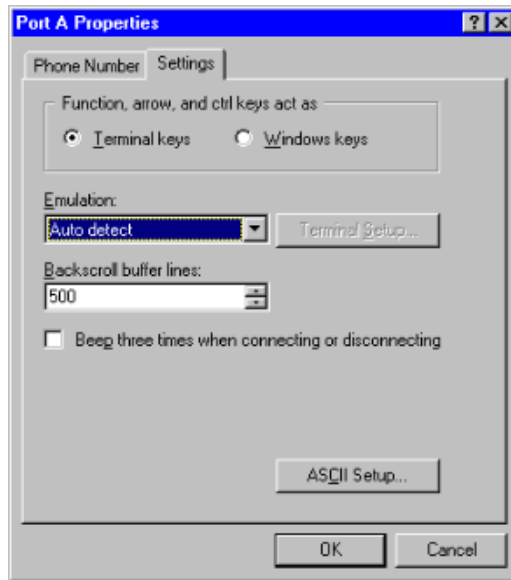
- 5** Set the properties according to the following table, and then click OK.

Property	Value
Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

- 6** Click File > Properties > Settings.

A dialog box similar to the following appears:

Figure 174



- 7** Click Terminal Keys.
- 8** Click the Emulation drop-down list > select ANSI.
- 9** (Optional) Specify cursor behavior by clicking Terminal Setup.
- 10** Click ASCII Setup > ensure that only Wrap Lines That Exceed Terminal Width is checked.
- 11** Click OK > OK.
- 12** Press Enter.

A command line prompt appears in Hyper Terminal.

You can now use all console commands to manage the appliance.

Additional Information

If the appliance has a monitor attached, commands issued through a null-modem connection are not echoed on the appliance monitor.

The following commands are available when using a null-modem connection:

- ♦ **cls** clears the screen
- ♦ **_info** displays the Com port settings for the appliance

When accessing the appliance through a null-modem connection, arrow keys function in the following ways:

- ♦ The Back arrow acts like the backspace or erase key.
- ♦ The Forward arrow acts like the space bar.
- ♦ The Up arrow displays a history of previously executed commands (beginning with the most recent command).
- ♦ The Down arrow scrolls forward through the command history and ends with a blank line.

Troubleshooting Telnet

Telnet never starts

- ☐ To establish a Telnet connection, you must be able to ping the server.

Telnet starts after a long time

- ☐ If you previously ran a session and turned on Transparent Proxy, Telnet might be very slow in starting.

Commands at the bottom of the screen look strange or don't make sense

- ☐ The display might not update correctly. Enter **clearscreen** to get a new screen and start at the top.



Upgrading the Appliance

You can apply over-the-wire upgrades to the appliance as the upgrades become available. Check with your appliance vendor for more information about these upgrades.

Upgrading

As system upgrades become available, you can access them quickly and apply them at predetermined times. For more information, see [“Upgrade Tab” on page 310](#).

Critical Information

Upgrading Version 1.0 Purges the Cache

If you are upgrading from version 1.0, the existing cache is automatically purged during the upgrade because of enhancements to the Cache Object Store in versions 1.2 and later. All configuration settings are retained.

Making Sure You Update the Clone Image Before and After Upgrading

To avoid having the upgraded system overwritten by the older clone image both before and after an upgrade or support pack installation, you must update the appliance's clone image. For more information, see [Step 8 on page 442](#).

Disabling Forward Proxy Authentication

Forward proxy authentication must be disabled in order to upgrade the appliance. Otherwise, the upgrade will fail.

Re-Imaging the Appliance from the Vendor CD

After re-imaging the appliance from the vendor CD, the Excelerator appliance must be fully shut down and restarted without the CD-ROM in the drive before applying an upgrade.

Preserving Configuration Settings

As mentioned previously, the system upgrade retains all appliance configuration settings. As a precaution, we recommend you also export your appliance's configuration settings to a DOS-formatted floppy diskette prior to the upgrade.

For assistance, refer to [“Backing Up the Appliance Configuration” on page 198](#).

If you need to import the configuration file after the upgrade is completed, refer to the table entry **“Apply a named configuration file on a floppy” on page 198.**

Upgrading through a Firewall

In most cases upgrading through a firewall is not a problem. If your environment allows HTTP access to the Web, the appliance should be able to retrieve the upgrade files as easily as a browser downloads Web pages.

If normal HTTP access is restricted within your firewall, the appliance attempts to retrieve upgrade packages through firewalls in one of the following ways:

1. First, the over-the-wire upgrade determines if the appliance can use an ICP or CERN parent. If so, the appliance uses the parent to download the upgrade package.
2. If an ICP or CERN parent is not available, the over-the-wire upgrade determines if the appliance is configured as a forward proxy with access through the firewall. If it is, the appliance tries two methods, in the following order:
 - a. If the firewall acts as a SOCKS server, you must configure the appliance as a SOCKS client. It can then retrieve the upgrade package from the origin server.
 - b. If the firewall is not acting as a SOCKS server, you must create a hole through the firewall that allows the appliance to make HTTP connections to the origin server with the upgrade package.

Close the hole as soon as the upgrade is downloaded.

3. If neither of the previous two methods is available, the over-the-wire upgrade attempts to establish a direct connection with the origin server.

To enable this connection, you must create a hole through the firewall and close it as soon as the upgrade is downloaded.

Downloading and Installing the Upgrade

To upgrade your appliance, do the following:

- 1** In the browser-based management tool, click System > Actions > Update Clone > Update Clone.
- 2** When the update is complete, click the Upgrade tab.
- 3** Check Enable Download > type the URL for the upgrade.
The URL is available from your appliance vendor.
- 4** Click the Download Time drop-down list > select the time you want the download to occur.
- 5** Check Enable Install.
- 6** Click the Install Time drop-down list > select the time you want the upgrade to be installed.
- 7** Click Apply.
- 8** As soon as possible after the upgrade is installed, update the appliance's clone image by clicking System > Actions > Update Clones.

This is necessary to avoid the appliance automatically applying an earlier clone image which could make Excelsior unstable.

Uninstalling the Most Recent Upgrade

To uninstall the most recent upgrade, do the following:

- 1** In the browser-based management tool, click System > Upgrade.
- 2** Check Enable Uninstall.
- 3** Click Apply.

Only the most recently installed upgrade can be uninstalled.

