AUTHORIZED DOCUMENTATION

Administration Guide

Novell®
iChain®

2.3 SP6

February 11, 2009

www.novell.com

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.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. See the Novell International Trade Services Web page (http://www.novell.com/info/exports/) for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Novell, Inc., has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed on the Novell Legal Patents Web page (http://www.novell.com/company/legal/patents/) and one or more additional patents or pending patent applications in the U.S. and in other countries.

Novell, Inc.
404 Wyman Street, Suite 500
Waltham, MA 02451
U.S.A.
www.novell.com

Online Documentation: To access the latest online documentation for this and other Novell products, see the Novell Documentation Web page (http://www.novell.com/documentation).

Novell Trademarks

For Novell trademarks, see the Novell Trademark and Service Mark list (http://www.novell.com/company/legal/trademarks/tmlist.html).

Third-Party Materials

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

Contents

About This Guide 15

1 Overview 17

1.1 iChain Benefits 17
1.1.1 iChain Controls Access to Web Servers and Web Applications 17
1.1.2 iChain Accelerates Access to Web Servers and Web Applications. 18
1.1.3 iChain Secures Data Sent from the Browser to Your Network 18
1.1.4 iChain Increases the Overall Security of a Web Server. 18
1.1.5 iChain Reduces Firewall Administration Costs 18
1.1.6 iChain Simplifies Management and Administrative Duties. 18
1.2 iChain Features 19
1.2.1 Reverse HTTP Proxy Cache 19
1.2.2 Authentication Support 19
1.2.3 URL-Based Access Privileges 20
1.2.4 Identity-Based Access Control. 20
1.2.5 User Convenience of Web Single Sign-on 21
1.2.6 Dynamic Data Encryption 21
1.2.7 Secure Access to Citrix Thin-Client Services 22

2 Installing iChain Components 23

2.1 Product Components 23
2.2 Installation Scenario. 23
2.3 System Requirements 24
2.3.1 iChain Proxy Server Requirements 24
2.3.2 iChain Authorization Server Requirements 24
2.3.3 Administrator Workstation Requirements 25
2.4 Installing iChain Services Software 25
2.4.1 Installing the iChain Proxy Services Software 25
2.4.2 Installing iChain Services Schema Extensions on the iChain Authorization Server . 27
2.4.3 Installing the iChain ConsoleOne Snap-Ins 28
2.5 Managing the iChain Proxy Server 29
2.5.1 The Proxy Administration Tool. 29
2.5.2 The Command Line Interface 29
2.6 Installing the iChain Proxy Server on Your Network 35
2.7 Preparing the Network 35
2.7.1 Basic Network Configuration Setup 35
2.7.2 Configuring the Client Workstation 36
2.7.3 Configuring the iChain Proxy Server 37
2.8 Troubleshooting iChain Proxy Server Issues 38
2.8.1 Custom Error Pages 38
2.8.2 Troubleshooting iChain Proxy Server Connectivity Issues 39
2.8.3 Troubleshooting iChain Proxy Server Authentication Issues. 41
2.8.4 Troubleshooting iChain Proxy Server Authorization Issues. 41

3 Configuring a Typical Accelerator 43

3.1 Setting Up an Accelerator 43
3.1.1 Configuring the Accelerator 43
3.2 Configuring the Authentication Server 49

3.2.1 Creating an iChain Service Object. 50
3.2.2 Setting Up a Protected Resource 50
3.3 Configuring iChain for NetIdentity Authentication 51
3.3.1 How NetIdentity Works with iChain Authentication 51
3.3.2 Configuring an iChain Accelerator to Use NetIdentity 52
3.3.3 Configuring An Accelerator to Use an Authentication Profile 53

4 Understanding the Web Server Accelerator 55

4.1 Accelerating Web Servers 55
4.1.1 Overview of Web Server Acceleration 55
4.1.2 How Origin Web Server Acceleration Works 55
4.1.3 Benefits of Origin Web Server Acceleration 56
4.2 Web Server Accelerator Setup. 56
4.2.1 Working with DNS 59
4.3 Accelerator/Web Server Page 59
4.3.1 Controls for Accelerator/Web Server. 61
4.4 Using Web-Based Distributed Authoring and Versioning (WebDAV). 62
4.4.1 Using WebDAV With NetStorage 62
4.5 Setting Up Authentication Using the Wireless Application Protocol (WAP) 63
4.6 Setting Up Secure Exchange 63
4.6.1 Creating a Trusted Root Container and Trusted Root Objects 64
4.6.2 Configuring iChain to Use the Trusted Root Container and Objects 65
4.7 Multi-Homing 65
4.7.1 What Is Multi-Homing? 65
4.7.2 Multi-Homing Web Server 66
4.7.3 Host-Based Multi-Homing 66
4.7.4 Domain-Based Multi-Homing. 67
4.7.5 Path-Based Multi-Homing 67
4.7.6 Path-Based Multi-Homing Example. 68
4.8 Custom Login Pages 69
4.8.1 Path or Host-Based Multi-Homing 70
4.8.2 Limitations 71
4.8.3 Coding of Login Pages 71
4.8.4 Custom Logout Page 72
4.8.5 Using the USERVOL for Custom Login and Logout Pages. 72
4.9 Cache Freshness 73
4.9.1 Managing Cache Freshness 73
4.10 GZIP Compression Support 75
4.11 Understanding HTTP 1.1 75
4.11.1 Troubleshooting HTTP 1.0/1.1. 76
4.11.2 Troubleshooting HTTP 1.1 and Recompression 76

5 Understanding Mutual Authentication 77

5.1 Setting Up Mutual SSL. 77
5.1.1 Mutual SSL Configuration 77
5.1.2 Using Certificate Mapping 80
5.2 Disabling Mutual SSL. 86
5.3 Using Third-Party Certificates 87
5.4 Using Multiple Certificate Authorities 87
5.4.1 Configuring Multi CAs 87
5.5 Configuring the Online Certificate Status Protocol. 88
5.5.1 New OCSP Setting in iChain 2.3 Support Pack 1 91

6 Understanding LDAP Authentication 93

6.1 Allowing Authentication Through the HTTP Authorization Header. 93
6.2 Enabling Authentication Through the HTTP Authorization Header 93

7 Using RADIUS Authentication 95

7.1 Using Token Authentication with iChain. 95
7.2 Installing NMAS, Novell RADIUS, and a Token Method 95
7.3 Configuring Novell RADIUS Components 96
7.3.1 Creating a Dial Access System (DAS) Object 96
7.3.2 Configuring the Login Policy Rules 97
7.3.3 Adding the iChain Proxy Server As a Client of the DAS Object 97
7.3.4 Creating a RADIUS Dial Access Profile (DAP) Object. 97
7.3.5 Adding an Attribute to the RADIUS DAP Object 98
7.3.6 Assigning the Token Method to Each User Object 98
7.3.7 Assigning the DAS Object to Each User Object 98
7.3.8 Starting Novell RADIUS Services on Your NMAS Server 98
7.4 Setting Up the iChain RADIUS Client 98
7.4.1 Adding a RADIUS Authentication Profile 99
7.4.2 Adding RADIUS Authentication to an Accelerator 99
7.4.3 Adding an LDAP Profile for Mapping a RADIUS User in the Authentication LDAP Directory 99

8 Accelerator Authentication Options 101

8.1 Accelerator Authentication Parameter Page 101
8.1.1 Controls for Accelerator Authentication Parameters 103
8.1.2 Advanced Authentication Options Dialog Box 103
8.1.3 RADIUS Authentication Load Balance and Failover 105
8.1.4 Add (Modify) Authentication Profile Dialog Box 105
8.1.5 Mutual Certificate Mapping Dialog Box 106
8.1.6 Controls for Authentication Profiles 108
8.1.7 Controls for Authentication Profile Options 110
8.1.8 New LDAP Authentication Server Dialog Box 111
8.1.9 Add LDAP Context Dialog Box 111
8.1.10 Controls for Add LDAP Context 112
8.1.11 Radius Options Dialog Box 112
8.1.12 Controls for RADIUS Options 114
8.2 Authentication Profiles and How They Are Used 114
8.2.1 The Authentication Process With Mutual Authentication 115
8.2.2 The Authentication Process With RADIUS Authentication 116
8.2.3 Using LDAP AND Mutual Over HTTP 116
8.2.4 Using LDAP AND Mutual Over HTTPS 116
8.2.5 Using LDAP OR Mutual, RADIUS OR Mutual 116
8.2.6 Using RADIUS AND Mutual. 116
8.2.7 Using LDAP AND RADIUS 116
8.2.8 Using LDAP OR RADIUS 117
8.2.9 Using LDAP AND RADIUS AND Mutual 117
8.2.10 Using LDAP OR RADIUS OR Mutual 117

9 Using Cross-Domain Authentication 119

9.1 CDA Scenario and Examples. 119
9.2 Selecting Accelerators as Members of CDA and Cross-Domain Brokers 120
9.3 Configuring the CDA 120
9.4 CDA and Session Broker 121

10 Configuring Session Broker 123

10.1 What is Session Broker? 123
10.2 Configuring Session Broker 124
10.3 Configuring Session Broker without a Floppy Drive 125
10.4 Troubleshooting Session Broker Issues 126
10.4.1 Troubleshooting Tools 126
10.4.2 Troubleshooting Steps 127

11 Using the iChain Service Object (ISO) 129

11.1 General Page 129
11.2 Protected Resource Page 130
11.2.1 Types of Protected Resources 131
11.3 Form Fill Policy Page 133
11.4 Session Broker Page 134
11.5 iChain Page 135

12 Using iChain Access Control Rules 137

12.1 Defining iChain Access Control Rules 137
12.1.1 iChain Access Control Object 139
12.1.2 ACL Exceptions 140
12.1.3 ACL Theory of Operations 141
12.2 Defining Dynamic Access Control Rules 145
12.2.1 Setting Up Dynamic Access Control Rules 146
12.3 Advanced Access Control Configuration 149
12.3.1 Enabling ACL Rule Checking for Community Objects 149
12.3.2 Enabling Debugging Messages for Access Control 149
12.3.3 Using ACLCHECK options 150

13 Using Object-Level Access Control 153

13.1 Setting Up Object-Level Access Control 153
13.2 Using OLAC Custom Header Variables 156
13.3 Customizing the Authorization Header 156
13.4 Using OLAC Caching 157
13.5 Using OLAC for Shared Secrets 157
13.5.1 How Shared Secrets Works 157
13.5.2 Defining OLAC for Shared Secrets 158
13.5.3 Configuring OLAC for Shared Secrets 158
13.5.4 Known Limitations With OLAC and Shared Secrets 159
13.6 The Effects of Disabling OLAC 159

14 Form Fill 161

14.1 Understanding an HTML Form 161
14.1.1 Analyzing the HTML Form 162
14.2 Designing a Basic Form Fill Policy 165
14.2.1 Basic Components of a Form Fill Policy 165
14.2.2 Sample Form Fill Policy 174
14.3 Creating a Form Fill Policy 175
14.4 Using Additional Form Fill Policy Options 176
14.4.1 Using Other Form Fill Tags 176
14.4.2 Injecting Static Values 178

14.4.3 Submitting Form Data to the Server 178
14.4.4 Using JavaScript in Forms 179
14.4.5 Using Login Pages in Multiple Language Character Sets 182
14.4.6 Specifying Form Fill Switches 182
14.5 Troubleshooting a Form Fill Policy 183
14.5.1 Tags 183
14.5.2 Tools and Documentation 185
14.6 Setting Up SecretStore 186
14.6.1 Installing SecretStore 187
14.6.2 Configuring SecretStore to Use Shared Secrets 189
14.6.3 Using the SecretStore Proxy Service 191

15 Using iChain to Manage Certificates 193

15.1 Naming Certificates 193
15.2 Creating Certificates Using the Appliance CA 194
15.3 Obtaining a Certificate from an External CA 195
15.3.1 Requesting the CSR 195
15.3.2 Sending the CSR 196
15.3.3 Storing the Certificate 197
15.3.4 Importing a CSR Signed by Intermediates 197
15.4 Viewing (Exporting) a Certificate’s CA 198
15.5 Modifying a Certificate 199
15.6 Deleting a Certificate 199
15.7 Backing Up a Certificate 199
15.8 Restoring a Certificate 200
15.9 Renewing a Third-Party Certificate 201

16 Using Advanced Accelerator Features 203

16.1 Using Secure Thin Client Services 203
16.1.1 Secure Access to Citrix Thin Clients 204
16.1.2 Configuring iChain to Accelerate Citrix MetaFrame Servers with Nfuse 204
16.1.3 Configuring iChain to Accelerate Citrix MetaFrame Servers without Nfuse 207
16.1.4 Using the Citrix Java Client 215
16.2 Using iChain With Novell Nsure Audit 216
16.2.1 Nsure Audit Overview 216
16.2.2 Auditing Background and Fundamentals 216
16.2.3 Accessing Novell Nsure Audit Documentation 217
16.2.4 Using iChain With Nsure Audit 217
16.3 Logging 219
16.3.1 Using Appliance Logging 219
16.3.2 Planning Your Logging Strategy 221
16.3.3 Configuring Logging Options 225
16.3.4 About the FTP Log Push Feature 226
16.3.5 Using FTP Push to Automatically Download and Delete Log Files 227
16.3.6 Manually Downloading and Deleting Log Files 228
16.3.7 About Extended Log Field Headers 231
16.3.8 Enabling and Viewing the ACL Rule Log File 233
16.4 Object Pinning 234
16.4.1 The Pin List 234
16.5 Using the Proxy Server to Record IP Addresses When Resolving URL Masks. 239
16.5.1 Router Capabilities 240
16.5.2 Using Appliance Routing 240
16.6 Concurrent Login Restriction 240

17 Rewriter Support 243

17.1 The Internal Rewriter . . . 243
17.1.1 Which URL References Are Rewritten? . . . 244
17.1.2 What Other Criteria Are Considered? . . . 244
17.1.3 Configuring the Internal Rewriter . . . 246
17.1.4 Sample Rewriter Scenario . . . 249
17.1.5 Disabling the Internal Rewriter . . . 251
17.1.6 Internal Rewriter Summary . . . 252
17.2 The Custom Rewriter . . . 252
17.2.1 Rewrite Filter Configuration File . . . 253
17.2.2 Enabling the Custom Rewriter . . . 255
17.2.3 Replacement Rules . . . 257

18 Using Start Up Options and the Command Line Interface 259

18.1 Customizing the Command Line Prompt . . . 259
18.2 Using the Enhanced Configuration Export . . . 260
18.2.1 File Format . . . 260
18.2.2 New Commands . . . 262
18.2.3 Using a Comment Specifier in NAS Files . . . 263
18.3 Managing Appliance Security Features . . . 263
18.3.1 Using the Console Lock Feature . . . 263
18.3.2 Accessing Proxy Internals . . . 264
18.3.3 Setting Up Enhanced Security Within the Authentication Cookie . . . 265

19 Using the Proxy Services Interface 267

19.1 Using the Proxy Administration Tool . . . 267
19.1.1 Prerequisites for Running the Administration Tool . . . 267
19.1.2 Starting the Administration Tool . . . 268
19.1.3 Applying and Cancelling Changes . . . 268
19.1.4 The Help Button . . . 269
19.1.5 Encryption . . . 269
19.2 The Home Panel Options . . . 269
19.2.1 Introduction Page . . . 269
19.2.2 Health Status Page . . . 270
19.2.3 Certificate Maintenance Page . . . 272
19.3 The System Panel Options . . . 273
19.3.1 Timezone Page . . . 273
19.3.2 Date/Time Page . . . 274
19.3.3 Actions Page . . . 275
19.3.4 SNMP Page . . . 279
19.3.5 Import/Export Page . . . 280
19.3.6 Upgrade Page . . . 281
19.3.7 Alerts Page . . . 283
19.3.8 Admin ACL Page . . . 284
19.4 The Network Panel Options . . . 285
19.4.1 IP Addresses Page . . . 286
19.4.2 DNS Page . . . 289
19.4.3 Gateway/Firewall Page . . . 291
19.5 The Configure Panel Options . . . 295
19.5.1 Web Server Accelerator Page . . . 296
19.5.2 Add Authentication Profiles Dialog Box . . . 305
19.5.3 Using Authentication Profiles . . . 307
19.5.4 LDAP Search Query Authentication . . . 312
19.5.5 Authentication Page . . . 314

19.5.6 Access Control Page 319
19.5.7 FTP Page. 322
19.5.8 Management Page 323
19.5.9 Tuning Page 325
19.5.10 Mini Web Page 327
19.6 The Monitoring Panel Options 328
19.6.1 Summary Page 328
19.6.2 Services Page 329
19.6.3 Performance Page. 330
19.6.4 Cache Page 331
19.6.5 Cache Logs Page 332
19.6.6 Top Ten Sites Page 332
19.7 Using Appliance Commands 333
19.7.1 Troubleshooting the Command Line 335
19.7.2 Connecting through Telnet 335
19.7.3 Establishing a Null-Modem Connection. 338
19.8 Performing Patch Upgrades. 341
19.8.1 Making Sure You Update the Clone Image before and after Upgrading 341
19.8.2 Preserving Configuration Settings 341
19.8.3 Upgrading through a Firewall. 341
19.8.4 Downloading and Installing the Upgrade 342

20 Using Logging Tools 343

20.1 Using Strong Cryptography 343
20.1.1 Cryptography Settings 344
20.1.2 Configuring Federal Information Processing Standards in iChain 344
20.2 Using Step-Up Cryptography 345
20.3 Automatic Configuration Mechanisms 345
20.3.1 About Appliance Configuration Files 346
20.3.2 System-Generated Configuration Files 346
20.3.3 Using Customized Configuration Files to Change the System Configuration 347
20.3.4 Managing Configuration Files 347
20.3.5 Creating Appliance Configuration Shortcuts 348
20.3.6 Restoring Factory Settings 349
20.3.7 Restoring the Appliance to the Clone Image 349
20.3.8 Reimaging and Restoring the Appliance System 350
20.4 DNS Name Resolution 351
20.4.1 How the Appliance Resolves DNS Names and Formulates Subsequent DNS Queries 351
20.5 Appliance Error Messages 353
20.5.1 Checking the Language Directories on Your Appliance 353
20.5.2 Customizing the Appliance Error Template and Message Files 353
20.6 FTP Services 357
20.6.1 Accessing Active and Passive FTP 358
20.6.2 Setting Up Appliance FTP Services. 358
20.6.3 Starting an FTP Session with the Appliance 359
20.6.4 Changing the FTP Working Directory 359
20.6.5 Managing Configuration Files with FTP. 360
20.6.6 Using FTP to Download a Configuration File to Your Workstation 360
20.6.7 Using FTP to Move a Configuration File to the appliance from a Workstation 360
20.6.8 Using FTP to Move a Configuration File to the Appliance and Execute It 360
20.6.9 Using FTP to Customize the Appliance Splash Screen. 360
20.7 Shutting Down and Restarting 361
20.7.1 Restarting from the Browser-Based Management Tool. 361
20.7.2 Shutting Down and Restarting from the Command Line 361
20.8 Using the SOCKS Client Service 362

20.9 Time Synchronization. 362
20.9.1 Synchronizing Time 362
20.9.2 Using the Browser-Based Management Tool to Set the Time. 362
20.9.3 Using the Command Line to Set NTP Time Sources 363

21 Upgrading Your iChain System 365

21.1 Upgrading from iChain 2.0, 2.1, and 2.2 365
21.1.1 1. Prepare the Current iChain Platform 365
21.1.2 2. Back Up the Existing iChain Configuration 365
21.1.3 3. Upgrade eDirectory with the iChain Schema Using the Install CD 366
21.1.4 4. Install ConsoleOne 1.34 and the iChain Snap-Ins. 367
21.1.5 5. Convert and Modify Existing ACL/ISO Definitions. 367
21.1.6 6. Upgrade the Proxy Server to iChain 2.3 367
21.1.7 7. Test the System. 368
21.1.8 8. Implement New Features. 368
21.1.9 Schema Differences Between 2.0, 2.1, 2.2, and 2.3 368
21.2 Upgrading from iChain 1.5 368
21.2.1 5. Converting and Modifying Existing ACL/ISO Definitions When Upgrading from iChain 1.5 369

A Known Issues 371

A.1 Installation and Upgrade Issues. 371
A.1.1 Over-the-Wire Upgrade (OTWUG) Requirements. 371
A.1.2 OTWUG vs. CD Upgrade 371
A.1.3 Upgrading From iChain 2.2 to iChain 2.3 372
A.1.4 OTWUG Might Have to Use IP Addresses 372
A.1.5 Additional RADIUS Configuration Required 372
A.1.6 Importing a .NAS File, Changing to the Factory Settings Can Cause a Reboot . . . 372
A.1.7 Password Management Information Lost When Importing .NAS File 372
A.1.8 The Servlet Directory Has Been Removed in iChain 2.3. 373
A.1.9 Accessing the SecretStore Client Utilities 373
A.2 Known Issues With iChain 2.3 and NetWare 6.5 373
A.2.1 Install the Latest Support Pack for NetWare 6.5 Before Testing iChain in Your Environment. 373
A.2.2 Using the iFolder Client Through iChain 373
A.2.3 The iPrint Client for Novell Open Enterprise Server No Longer Supports iChain . . 374
A.2.4 Accessing iFolder Through an Accelerator Requiring Authentication 374
A.2.5 Leaving the Challenge/Response Blank in RADIUS Causes Error 374
A.2.6 Issue With RADIUS Loading On a NetWare 6.5 Server 374
A.3 Known Issues With iChain Form Fill. 374
A.3.1 New Version of the SecretStore Plug-in 374
A.3.2 ConsoleOne Crashes When Trying to Pass Large or Numerous Form Fill Polices. 375
A.3.3 Form Fill Does Not Use Modified LDAP Options Without Being Restarted 375
A.3.4 Form Fill Fails to Remember Attribute When Spaces Exist in the <Name> Field . . 375
A.3.5 Using <JavaScript> to Keep Specified Functions Does Not Work 376
A.4 Known Authentication Issues 376
A.4.1 iChain Single Sign-on is Not Compatible With iManager 2.5. 376
A.4.2 Users Prompted to Authenticate Twice When Accessing Password Management Servlet 376
A.4.3 iChain Prompts for SSL Mutual Authentication Even After Disabling Authentication on an Accelerator 377
A.4.4 Time Restriction, Intruder Lockout, and Login Disabled are not Checked During Radius Token Authentication. 377
A.4.5 Certificate Revocation Checking (revocationcheckmethod =) 377
A.4.6 OSCP ‘Unknown' Response Lets Certificate Authenticate User 377

A.4.7 Certificate Authentication Problems When the CRL Is Invalid 378
A.4.8 Session Times Out During the 0 TTL State 378
A.4.9 NMAS RADIUS Update Required for iChain RADIUS Token Authentication 378
A.4.10 Use the Latest NetIdentity Client With iChain 2.3 378
A.4.11 Using an iChain Accelerator for WebDAV Connections to NetStorage 378
A.4.12 Do Not Use NetIdentity with LDAP/RADIUS ANDing 379
A.4.13 iChain is Unable to And LDAP with XTier Based Profiles 379
A.4.14 OLAC Authentication Profile Parameters Still Passed When OLAC Is Disabled . . . 379
A.5 Known Proxy Administration Tool Issues 379
A.5.1 Command Line and Proxy Administration Tool Sluggishness During Startup 379
A.5.2 Daylight Saving Time Adjustment 380
A.5.3 Speed and Duplex Settings Might Not Be Synchronized Between the Proxy Administration Tool and INETCFG 380
A.5.4 Proxy Administration Tool Slowness With Sun JRE 380
A.5.5 The PIN List Might Not Display Properly with SUN JRE 380
A.6 Known Multi-Homing and Path-Based Issues 380
A.6.1 Disabling a Multi-Homing Master Accelerator 381
A.6.2 Cannot Use Double-byte Path for Path Based Multi-Homing 381
A.6.3 Multi-Homing Display in the iChain GUI 381
A.6.4 Path-Based Multi-Homing Might Not Function Properly with JavaScript 381
A.7 Known Session Broker Issues 381
A.7.1 Configuration Via the ConsoleOne Wizard 381
A.7.2 Enabling Basic Authentication for LDAP and Disabling Secure Exchange Causes Session Broker to Fail 382
A.8 Known Certificate Issues 382
A.8.1 Creating a Certificate With an Unspecified Purpose Breaks XTier 382
A.8.2 Cannot Use UTF8 Encoded Server Certificates with iChain 2.3 382
A.8.3 Netscape Browsers and Certificate Database Passwords 382
A.9 Known Rewriter Issues 382
A.9.1 Changes to the Rewriter 382
A.9.2 Rewriting Fails on a Page with Numerous HREFs 383
A.10 Miscellaneous Issues 383
A.10.1 Sending the Via Header to the Web Server 384
A.10.2 Cannot Delete Email Messages in Outlook Web Access 384
A.10.3 License Agreement Does Not Appear in GUI When Logged In as the View User. . 385
A.10.4 Trustedroot Configured .nas File Fails on Import 385
A.10.5 Server Might Abend When Restart Server is Entered on the iChain Server 385
A.10.6 Creating Custom Login Pages without a UTF-8 Charset 385
A.10.7 Issue With Compressed Data Not Being Sent From Web Servers 385
A.10.8 Nsure Audit Logging Issue When Using iChain to Log Events 386
A.10.9 Cannot Add a Second Server to an Accelerator 386

B Using LDAP Server Load Balancing and Failover 387

C Documentation Updates 389

C.1 February 2009 (SP6) 389
C.2 March 2007 (SP5) 390
C.3 March 16, 2006 (SP4) 391
C.4 October 24, 2005 391
C.5 August 8, 2005 (SP3) 391
C.6 June 24, 2005 (SP3) 392
C.7 April 12, 2005 (SP2) 392
C.8 March 3, 2005 (SP2) 392
C.9 January 26, 2005 (SP2) 393

About This Guide

This guide describes how to install, configure, and administer the Novell® iChain® 2.3 product with the latest support pack. For the latest iChain 2.3 documentation, including updates to this guide, see the online documentation on the Novell iChain 2.3 Documentation Web page (http://www.novell.com/documentation/lg/ichain23/index.html).

This guide is divided into the following sections.

- Chapter 1, "Overview," on page 17
- Chapter 2, "Installing iChain Components," on page 23
- Chapter 3, "Configuring a Typical Accelerator," on page 43
- Chapter 4, "Understanding the Web Server Accelerator," on page 55
- Chapter 5, "Understanding Mutual Authentication," on page 77
- Chapter 6, "Understanding LDAP Authentication," on page 93
- Chapter 7, "Using RADIUS Authentication," on page 95
- Chapter 8, "Accelerator Authentication Options," on page 101
- Chapter 9, "Using Cross-Domain Authentication," on page 119
- Chapter 10, "Configuring Session Broker," on page 123
- Chapter 11, "Using the iChain Service Object (ISO)," on page 129
- Chapter 12, "Using iChain Access Control Rules," on page 137
- Chapter 13, "Using Object-Level Access Control," on page 153
- Chapter 14, "Form Fill," on page 161
- Chapter 15, "Using iChain to Manage Certificates," on page 193
- Chapter 16, "Using Advanced Accelerator Features," on page 203
- Chapter 17, "Rewriter Support," on page 243
- Chapter 18, "Using Start Up Options and the Command Line Interface," on page 259
- Chapter 19, "Using the Proxy Services Interface," on page 267
- Chapter 21, "Upgrading Your iChain System," on page 365

Audience

This guide is intended for network administrators who manage network resources such as Web servers and Web applications.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comment feature at the bottom of each page of the online documentation, or go to Novell Documenation Feedback (http://www.novell.com/documentation/feedback.html) and enter your comments there.

Documentation Updates

For the most recent version of the Novell iChain 2.3 Administration Guide, see Novell iChain 2.3 (http://www.novell.com/documentation/ichain23/index.html).

Addtional Documentation

For a Cool Solutions article that provides an overview of the major features found in iChain, see iChain 2.3 Technical White Paper (http://www.novell.com/coolsolutions/feature/2555.html).

For questions and support from other iChain users, see the iChain Support Forum (http://support.novell.com/forums/).

Documentation Conventions

In Novell documentation, a greater-than symbol (>) is used to separate actions within a step and items in a cross-reference path.

A trademark symbol (®, ™, etc.) denotes a Novell trademark. An asterisk (*) denotes a third-party trademark.

Overview 1

Novell® iChain® is an integrated security and access-management infrastructure that protects your network and safeguards sensitive business and identity data. iChain facilitates your Internet and remote-access initiatives by providing secure authentication and access management to portals, Web-based content, and Web applications. iChain provides premier stability and scalability for your enterprise by incorporating the award-winning Novell eDirectory™, recognized as the most scalable and widely used enterprise directory.

iChain is a caching reverse proxy that restricts access to Web-based content, portals, and Web applications based on authentication and access control policies. iChain also provides single sign-on to multiple Web servers and Web applications by securely providing authenticated users' credential information to the protected servers and applications. With iChain, you can simplify, secure, and accelerate your Internet business initiatives.

This section provides an overview of the benefits provided by iChain, and lists the features and services iChain delivers.

- iChain Benefits
- iChain Features

1.1 iChain Benefits

iChain secures, simplifies, and accelerates your Internet initiatives. It is an integrated security solution that offers identity management and access management services within a powerful infrastructure. Through iChain, all types of Internet and remote-access initiatives are more securely available than ever before.

This section discusses some specific iChain benefits:

- iChain Controls Access to Web Servers and Web Applications
- iChain Accelerates Access to Web Servers and Web Applications
- iChain Secures Data Sent from the Browser to Your Network
- iChain Increases the Overall Security of a Web Server
- iChain Reduces Firewall Administration Costs
- iChain Simplifies Management and Administrative Duties

1.1.1 iChain Controls Access to Web Servers and Web Applications

Web servers and Web applications are often highly vulnerable to attacks from hackers who want to crash, corrupt, or disrupt service. iChain sits as a reverse proxy between the user in the Internet space and your business Web services. With iChain, your Web servers and Web applications can be placed in a private area of your data center, and never be directly accessible from the Internet. iChain acts as a firewall so that users on the Internet can access your enterprise Web content only through iChain, allowing for tight control and protection of your Web servers and applications.

1.1.2 iChain Accelerates Access to Web Servers and Web Applications

iChain incorporates in-memory and disk-based caching of content coming from Web server and Web applications. Content caching from an HTTP proxy not only gets commonly used data to the client faster, it also reduces the total number of Web servers needed to serve up a site.

1.1.3 iChain Secures Data Sent from the Browser to Your Network

It is easy to eavesdrop on Internet connections, looking for unprotected data. Sensitive business data, confidential user information, and financial records are examples of data that should always be protected over the Internet. Requiring each Web server to protect data using SSL can significantly reduce the Web server's performance. iChain can offload the expensive SSL encryption process from the Web servers, while obscuring and protecting the Web servers from Internet attack.

1.1.4 iChain Increases the Overall Security of a Web Server

Using iChain as the gatekeeper (the only point of access) to enterprise Web applications increases the security of the Web server and associated identity information by preventing direct user access to the Web server itself. iChain significantly reduces the risks of hacker attacks on sensitive servers by only allowing HTTP requests to be sent to the Web server. iChain also ensures that only requests from registered DNS names, rather than anonymous IP addresses, are allowed to reach the Web server.

1.1.5 iChain Reduces Firewall Administration Costs

New company sites are continually adding new Web servers for internal employees, customers, or potential customers. Adding a new Web server that is accelerated by iChain offers the following unbeatable benefits:

- By using the path-based multihoming feature within iChain, a new Web server can be made publicly available without adding a new DNS name or IP address to a DNS server or firewall.
- With domain-based multihoming, host names such as www.novell.com and products.novell.com can share the same IP address.
- Although a Web server might have been developed with a private host name such as www.private.com, iChain can rewrite all private host name references to a public DNS host name such as www.public.com.

1.1.6 iChain Simplifies Management and Administrative Duties

Today, many companies manage user access to internal Web-based material on a server-by-server basis. These servers often run on different platforms, especially in large enterprises that have many divisions spread across a wide geographic area. A good example is a government agency with many separate departments. Each department employs its own set of standalone servers and Web applications. Something as common as modifying a user's access rights requires the IT staff to manually change all the involved systems, a time-consuming process that could necessitate a physical visit to each network server. If those servers are scattered across the entire country, the situation becomes expensive and impractical—either a single IT staff member is constantly traveling, or it becomes necessary to maintain a separate IT staff for each part of the network.

iChain solves this problem by centralizing all administrative tasks. Changes can be made through ConsoleOne®, a single utility that defines the access control policies to all resources protected by iChain, regardless of the platform or Web server used. Moreover, ConsoleOne can be run from any workstation in the network, thereby avoiding the costly upgrades and retrofits that would otherwise be needed to unify all your network resources.

iChain also delivers standard login pages for each secure Web site protected by the iChain Proxy Server. Using an HTML editor, these pages can be customized to reflect the standard look and feel of the organization or department's Web sites.

1.2 iChain Features

This section highlights iChain's core features:

- Reverse HTTP Proxy Cache
- Authentication Support
- URL-Based Access Privileges
- Identity-Based Access Control
- User Convenience of Web Single Sign-on
- Dynamic Data Encryption
- Secure Access to Citrix Thin-Client Services

1.2.1 Reverse HTTP Proxy Cache

iChain sits between your Web applications (servers, portals, etc.) and your users. iChain acts as an HTTP proxy with support for both HTTP 1.0 and HTTP 1.1 features. Data that is marked as cacheable is cached in an iChain Web cache, significantly reducing the load on the Web application server.

1.2.2 Authentication Support

To guard against unauthorized user access, iChain supports a number of authentication methods, including user identifiers (name, e-mail address, and other LDAP attributes), passwords, token-based authentication, and X.509 digital certificates. The following authentication mechanisms are supported:

LDAP Authentication: This is a way to log in a user based on criteria used to locate a unique user and the user's password. There are many options that refine how a unique user is found within an LDAP directory. iChain supports contextless logins, as well as administrator-defined searches.

Mutual Authentication: This applies when the user is issued a certificate from a trusted source. The certificate identifies the user in some way. To ensure the validity of X.509 certificates, iChain supports both Certificate Revocation Lists (CRLs) and Online Certificate Status Protocol (OCSP) methods of verification.

RADIUS Authentication: To accommodate secure, token-based authentication, iChain uses the Remote Authentication Dial-in User Service (RADIUS) protocol. RADIUS enables communication between remote access servers and a central server. Secure token authentication through RADIUS is possible in iChain, out of the box, because iChain includes Novell Modular Authentication Service (NMAS™) RADIUS software that can be run on an existing NetWare® server. iChain supports both PIN and challenge and response methods of token-based authentication. In other words, RADIUS represents token-based authentication methods used to authenticate a user, based on something the user possesses (for example, a token card). Token challenge-response is supported for two-step processes that are necessary to authenticate a user.

Multi-Factor Authentication: This method combines several authentication methods to produce an even higher level of security. For example, a company might require that a user present a valid User ID and password, as well as an X.509 certificate, before access is granted. Different levels of access can be required for different Web servers.

Secure Cookie: iChain requires at least 48 bits of random data to be matched for authentication. You set up the secure cookie in the command line interface. See Section 18.3.3, "Setting Up Enhanced Security Within the Authentication Cookie," on page 265.

NOTE: An old version of session broker can handle the new secure cookie. However, all iChain boxes that are connected to a session broker must use the same cookie version. There cannot be mixed mode support for old and new cookies, which is why the new cookie version is not the default. For more information about Session Broker, see Section 10.1, "What is Session Broker?," on page 123.

1.2.3 URL-Based Access Privileges

The iChain Proxy Server is the gatekeeper to your Web-based applications. If a user is not permitted to access a particular Web page or application, the Web server/application server will never receive the request. iChain manages access to services through "protected resources," which are defined in eDirectory. A protected resource can be defined as:

- Public: Access is granted with no security checks
- Restricted: Requires only user authentication
- Secure: Requires authentication and association to an access control object (see "Identity-Based Access Control" on page 20 for more information)

To add flexibility when defining these protected URLs, iChain offers wildcarding (*) and entire folder (?) options.

1.2.4 Identity-Based Access Control

iChain's formidable security infrastructure begins with the iChain Access Control object. This object contains a list of iChain protected resources or URLs and a list of users, groups, or containment (O, OU) objects that can access these resources. Multiple iChain Access Control objects can be configured to provide maximum flexibility to meet an organization's security policy. Without an association to an Access Control object, a user is denied access to any protected resource that is defined as secure.

In addition to a user being granted access based on his or her association to an iChain Access Control object, iChain also provides a dynamic access control process that looks at specified details of a user's identity (for example, jobTitle=manager) and grants access based on that information.

1.2.5 User Convenience of Web Single Sign-on

Whether the user of your Web application is an employee or a potential customer, the experience that person has with the Web application is often determined by convenience. To enhance each user's experience, iChain incorporates an innovative service called Web single sign-on, where users need to log in only once to gain access to multiple applications and platforms.

Single sign-on is possible because iChain authenticates the user from a centralized eDirectory profile. When the user requests access to a specific server, iChain retrieves the appropriate user credentials and transparently submits them to the Web server, usually in the form of username and password. The user does not see a login request, but is immediately granted or denied access based on the single sign-on policy.

iChain also offers users a convenient form-fill authentication feature that simplifies access to Web applications. With the form-fill feature, the user first authenticates to iChain before accessing the Web-application's authentication form. As the user enters his or her credentials, the information is automatically stored securely to the user's object in eDirectory using Novell SecretStore™ technology. From then on, when the user connects to that Web application, iChain automatically retrieves the user's credentials and completes the form on the user's behalf.

By making your services more readily available, you can strengthen customer loyalty and offer employees convenient access to business-critical information. Single sign-on also lowers the overhead costs associated with maintaining many different tables of usernames and passwords on numerous servers.

1.2.6 Dynamic Data Encryption

When organizations want to ensure the confidentiality of data as it crosses the Internet, they usually implement SSL services for the Web servers, which increases management (certificates must be installed on each Web server), can increase costs, and can reduce the performance of content delivery (Web server processing power is dramatically reduced when it needs to encrypt data).

To address these issues, iChain provides Secure Exchange, which can dynamically encrypt the data channel between the browser and the iChain Proxy Server. The content between the iChain Proxy Server and the Web server can be either HTTP (non-encrypted) or HTTPS (encrypted), depending on specific requirements.

Secure Exchange provides a single place to manage SSL certificates (iChain proxy), and allows Web servers to do what they are designed to do: deliver content as quickly as possible. When combined with the caching technology on the iChain proxy, the speed of the overall service is greatly increased.

This feature not only performs on-the-fly encryption of data, but it also rewrites the HTML links from HTTP to HTTPS, meaning that there is no need for an administrator to change HTML content, a task that must be performed when SSL is implemented at the Web server.

The combination of iChain's multi-homing and Secure Exchange features means that content from many Web servers can be encrypted over the network using a single SSL certificate, further reducing management and deployment costs.

1.2.7 Secure Access to Citrix Thin-Client Services

iChain provides secure access to Citrix* thin-client services that use the Citrix ICA protocol. This includes single sign-on to Citrix Nfuse* and iChain-authenticated access to Citrix MetaFrame* servers.

Before a session can be established through iChain to a protected Citrix MetaFrame server, the user must have an iChain authentication token. The only way for the user to receive this token is to be authenticated to iChain at the time the Citrix client attempts to connect to the MetaFrame server.

iChain's secure access for the Citrix thin client also provides the additional benefit of supporting Citrix client connections over standard port:80 (client to iChain proxy). iChain relies on the Citrix client's own encryption capabilities.

Installing iChain Components

2

This section provides instructions for installing Novell® iChain® services software and contains the following topics:

- Section 2.1, “Product Components,” on page 23
- Section 2.2, “Installation Scenario,” on page 23
- Section 2.3, “System Requirements,” on page 24
- Section 2.4, “Installing iChain Services Software,” on page 25
- Section 2.5, “Managing the iChain Proxy Server,” on page 29
- “Installing the iChain ConsoleOne Snap-Ins” on page 28
- Section 2.7, “Preparing the Network,” on page 35
- Section 2.8, “Troubleshooting iChain Proxy Server Issues,” on page 38

2.1 Product Components

Your iChain installation includes the following components:

- iChain Proxy Server
- iChain Authorization Server
- Novell Modular Authentication Service (NMAS™) RADIUS Server

2.2 Installation Scenario

iChain is a flexible software solution that can be implemented in a variety of configurations depending on the needs of your network. Because of the possible variations in installation scenarios, procedures in this section describe the installation of a basic or standard infrastructure consisting of the following:

- One iChain Proxy Server
- One iChain Authorization Server
- An iChain extended schema on the Novell eDirectory™ tree where the iChain Authorization Server resides
- An iChain Service Object (ISO)
- An administrator workstation with ConsoleOne® iChain snap-ins installed

For increased security, we recommend installing the iChain Authorization Server in a tree separate from your corporate file/print tree. DirXML® can be used to synchronize user account information between trees. For more information about synchronizing user account information, see Password Synchronization across Connected Systems (http://www.novell.com/documentation/dirxml20/admin/data/an4bz0u.html).

2.3 System Requirements

Review the following system requirements to ensure that your server and client environments meet installation prerequisites:

- "iChain Proxy Server Requirements" on page 24
- "iChain Authorization Server Requirements" on page 24
- "Administrator Workstation Requirements" on page 25

2.3.1 iChain Proxy Server Requirements

The iChain Proxy Server is a self-contained install and does not require licensed hardware to run. However, we recommend that you perform a test install before the hardware is purchased. If the iChain Proxy Server installs and the Mini Web Server can be configured and accessed correctly, then the hardware should be fully compatible. For detailed hardware requirements and limitations, along with information about tested hardware, see Novell Software Downloads (http://www.novell.com/products/ichain/sysreqs.html).

Hardware issues that have been logged have almost always related to disk, CD, or LAN adapter drivers. If no matching drivers are found, multiple matching drivers are found, or other manual parameter input is needed during driver configuration, the install will hang (infinite dots). Included drivers for disk arrays are limited.

Resources

You will need to increase your resources as you increase your iChain configurations. Memory size and processor speed are more important than disk size. We recommend that most iChain installations have processor speeds of at least 1 gigahertz and memory sizes of at least 1 gigabyte. Using multiple LDAP servers for authentication and access control has also proven to increase performance.

2.3.2 iChain Authorization Server Requirements

The iChain Authorization Server can be installed on the following Novell eDirectory version 8.7 platforms:

- NetWare® 5.1, 6.0, and 6.5
- Linux*
- Windows* NT* 4.0 and Windows 2000
- Solaris*

The iChain Authorization Server will run on eDirectory 8.5 and above. However, because of key fixes in eDirectory, we recommend that users upgrade to at least eDirectory 8.6.2.

For additional information on the supported platforms and full system requirements for Novell eDirectory 8.6.2, refer to the *Novell eDirectory 8.6.2 Quick Start* available at the Novell Documentation site (http://www.novell.com/documentation/lg/ndsedir86/index.html).

Novell eDirectory can be downloaded at Novell Software Downloads (http://download.novell.com).

NOTE: For increased security, we recommend that you install the iChain Authorization Server in a tree that is separate from your corporate file/print tree. DirXML can be used to synchronize user account information between trees if needed. DirXML 1.0 requires a master replica; however, in higher versions of DirXML, read/write replicas can be utilized. iChain has no such limitations. As long as you use eDirectory 8.6*x* or later, both master replicas and read/write replicas can be used.

2.3.3 Administrator Workstation Requirements

The administrator workstation requirements are as follows:

- ❑ Pentium* 233 MHz processor or higher
- ❑ Minimum 45 MB of free disk space
- ❑ Minimum 128 MB of RAM
- ❑ One LAN card
- ❑ Windows NT, Windows 2000, or Windows XP
- ❑ Current Service Pack for Windows
- ❑ Current Novell Client™
- ❑ ConsoleOne 1.3.4 or later
- ❑ IP Connectivity between the client, the iChain Authorization Server, and the iChain Proxy Server

The latest Novell Client and ConsoleOne can be downloaded at the Novell Download Web site (http://download.novell.com).

2.4 Installing iChain Services Software

To install a basic iChain infrastructure, complete the following procedures:

- "Installing the iChain Proxy Services Software" on page 25
- "Installing iChain Services Schema Extensions on the iChain Authorization Server" on page 27
- "Installing the iChain ConsoleOne Snap-Ins" on page 28

2.4.1 Installing the iChain Proxy Services Software

The iChain Proxy Server should only be installed on compatible hardware (see "iChain Proxy Server Requirements" on page 24). To install the proxy server software:

1 Insert the *iChain Proxy Server* CD in the CD drive of the appliance or machine.

2 At the license page, type `YES` if you accept the agreement, then press Enter.

During the proxy installation process, the server reboots twice. Do not remove the CD until the proxy prompt is visible, indicating the installation is complete.

2a After the system reboots the first time, you will hear a series of beeps and the installation will prompt you about whether you want to select custom drivers. If you click Yes, the installation stops in HDetect.nlm and allows you to select the correct drivers for the system in the same manner as the NetWare 6 installation. Because of the iChain imaging process, you will need to do this twice during the installation.

If you click No (or if no selection is made within 30 seconds from the time of the prompt), iChain automatically detects the drivers as it does in earlier versions of iChain.

IMPORTANT: When installing iChain 2.3 with custom drivers, remove the CD immediately after the drivers are copied. Otherwise, the installation might hang when the system reboots.

If you opt to select custom drivers and the wrong drivers are selected, the iChain 2.3 Proxy Server software installation fails. We recommend that you attempt an automatic installation first, and only attempt to select your own drivers if the automatic installation fails.

2b If the installation does not complete, remove the CD, reboot the server, and delete the C:/rdw file.

3 Make sure the LAN adapter IP address is configured correctly.

After installation, the first LAN adapter on the iChain Proxy Server is preconfigured with the IP address 172.16.0.1 and subnet mask 255.255.255.0. In order to administrate the server using the Proxy Administration Tool, you either need to have a client workstation with an IP address on the same subnet (such as 172.16.0.2) or you need to use the iChain command line interface to set the IP address on the iChain Proxy Server.

The following commands from the iChain proxy server console configure the first LAN adapter with an IP address of 123.45.67.89 and a subnet mask of 255.255.252.0:

```
>unlock
```

At the Password prompt, press Enter (no password exists yet).

```
>set eth0 address = 123.45.67.89/255.255.252.0
>apply
```

After resetting the eth0 address, remove the CD, then type *restart* to restart the server.

If you are going to configure the iChain Proxy Server from a different segment than the one the iChain Proxy Server is on, you also need to use the following commands to configure the gateway:

```
>set gateway nexthop = 123.45.69.254
>apply
```

After installation, your iChain Proxy Server requires some basic setup to support your iChain implementation, and might require FTP to be enabled. The basic steps are detailed in Chapter 3, “Configuring a Typical Accelerator” in the online *Novell iChain 2.3 Administration Guide*.

To enable FTP, use the following commands:

```
>set miniftpserver address = 123.45.67.89
>apply
```

NOTE: Because FTP is an insecure protocol, enabling FTP can be a security risk on your network. We recommend that you enable the FTP server on an IP address that is only accessible from a private network such as an isolated hub or cross-over cable.

2.4.2 Installing iChain Services Schema Extensions on the iChain Authorization Server

The iChain Authorization server is the access point that iChain Proxy Services uses to retrieve authentication, access privileges, user, and group information for your iChain implementation from the eDirectory database. To make your eDirectory server platform into an iChain Authorization Server, install the iChain schema extensions onto the eDirectory tree for that server.

To install iChain schema extensions on the iChain Authorization Server:

1. If you have not already done so, install eDirectory on the machine that will be your iChain Authorization Server.
2. Insert the iChain authorization CD into the CD drive of a Windows client machine with IP connectivity to the iChain Authorization Server.

   If this is a Windows 2000 or Windows NT machine, you need administrator-level access to the client. The installation program launches automatically.
3. Click Install iChain Schema.
4. On the Welcome page, click Next.
5. Read the license agreement. If you accept the terms of the agreement, click Yes.
6. Enter the administrator user name in comma-delimited LDAP format (for example, cn=admin, o=novell).
7. Enter the administrator password.
8. Enter the IP address (and port, if necessary) for the server where you want to extend the schema.
9. Click Next.

The installation program notifies you whether the schema extension was successful. If an error occurs, look at the log file to determine what LDAP errors occurred. If a bind error occurs, the installation was not able to log in to the LDAP server.

Common Bind Errors

Some of the most common bind errors are:

ldap_simple_bind failed: 49(Invalid credentials), dn: cn=admin,o=novell: Usually denotes an incorrect password. Check the password and try again.

ldap_simple_bind failed: 32(No such object), dn: cn=adm,o=novell: The specified administrator does not exist. Verify the username and try again.

ldap_simple_bind_failed: 13(Confidentiality required), dn: cn=admin,o=novell: You need to enable the Allow Clear Text Passwords option on the LDAP Group object. Open the LDAP Group object in ConsoleOne and make sure the check box labeled Allow Clear Text Passwords is selected.

ldap_simple_bind failed 81(Can't contact LDAP server), dn: cn=admin,o=novell: Either the IP address/port combination is incorrect or the LDAP server is not running. Verify the IP address and LDAP port, make sure the server is running, and try again.

Common Log File Errors

Sometimes the LDAP bind succeeds but there are other errors in the log file. In these cases, there are usually multiple instances of the same error. Some common non-bind-related errors are:

The LBURP extension is not available on the server. Using standard LDAP calls: This generally means the LDAP server is out of date. You should verify that the latest LDAP server (included with eDirectory) is installed on the server to ensure that the schema is completely extended.

Record1: LBURP operation failed: 50(Insufficient access), dn:cn=schema: This error means that the specified administrator does not have sufficient rights to extend the schema.

Record1: LBURP operation failed: 20(Type or value exists), dn:cn=schema: This error is expected if the server has already been extended with a previous version of iChain with this attribute or class.

If you are unable to resolve an error, refer to the Knowledgebase on the Novell Support Web site (http://support.novell.com). This site includes information for resolving a number of LBURP operation failure issues.

2.4.3 Installing the iChain ConsoleOne Snap-Ins

You must install the iChain ConsoleOne snap-in files in order to administer the iChain eDirectory objects such as the iChain Service Object. You can install the snap-in files to be used with ConsoleOne running from the iChain Authorization Server, another server in the tree, or from an administrator workstation.

NOTE: iChain 2.3 requires ConsoleOne 1.3.4 or later for all of the snap-ins to function correctly.

To install the iChain ConsoleOne snap-ins to a server or an administrator workstation:

1. If the server or workstation does not already have ConsoleOne installed, install ConsoleOne.

   After ConsoleOne is installed, make sure you close it before starting to install the snap-ins.
2. Insert the iChain authorization CD into the CD drive of the server or the administrator workstation.

   The installation program launches automatically.
3. Click Install ConsoleOne Snapins for iChain.
4. On the Welcome page, click Next.
5. Read the license agreement. If you accept the terms of the agreement, click Yes.
6. Select the target drive where you want to copy the snap-in files.
7. Click Next to start copying the files.
8. Click Finish.

After completing the full installation, you need to use ConsoleOne to create the iChain Service Object, along with the access control list (ACL) rule objects, and make any other configuration adjustments.

2.5 Managing the iChain Proxy Server

The proxy server can be configured and managed in the following ways:

- Through the Proxy Administration Tool.
- From the command line interface through a Telnet or null-modem connection. (You can also use an attached keyboard and monitor if your proxy server has the required connections.)

2.5.1 The Proxy Administration Tool

The Proxy Administration tool is unlike other management utilities because its interface appears in your browser as an HTML page originating from the proxy server. The only programs running on your workstation are a Java-compatible Web browser and the Java* components required by the HTML page. To access the Proxy Administration tool, open a browser and enter http://*ipaddress*:1959/appliance/config.html in the address bar. The *ipaddress* is the IP address of your iChain server.

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

2.5.2 The Command Line Interface

Although it is possible to configure and monitor an iChain Proxy Server using only the iChain command line interface, we strongly recommend that you use the Proxy Administration Tool for all administrative tasks whenever possible.

The Proxy Administration Tool includes extensive cross-checking, helpful messages, and other program features to ensure that the iChain Proxy Server 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 the iChain Proxy Server from the command line interface.

For more information about using the command line interface, see Chapter 18, “Using Start Up Options and the Command Line Interface,” on page 259.

The evaluation version of iChain expires and does not work after 90 days. When the evaluation version expires, the iChain Proxy Server does not function, and you need to re-image your machine. You must purchase an iChain license and apply a Product Activation Credential in order for iChain to continue working after 90 days.

You can activate iChain by using one of two methods. The first method involves the following tasks:

- Purchasing an iChain License
- Activating iChain Using a Generic Credential
- Installing the Product Activation Credential Received from Novell

The second method involves the following tasks:

- Purchasing an iChain License
- Generating a Product Activation Request

- Submitting an Activation Request
- Installing the Product Activation Credential Received from Novell

This section also contains the following topics:

- "Viewing Product Activations for iChain" on page 34
- "Troubleshooting iChain Activation Problems" on page 34

Purchasing an iChain License

To purchase an iChain license, see the iChain How to Buy Web page (http://www.novell.com/products/ichain/howtobuy.html)

After you purchase an iChain license, Novell sends you a Customer ID via e-mail. The e-mail also contains a URL to the Novell site where you can obtain a generic credential. If you do not remember or do not receive your Customer ID, please call the Novell Activation Center at 1-800-418-8373 in the U.S. In all other locations, call 1-801-861-8373. (You will be charged for calls made using the 801 area code.)

Activating iChain Using a Generic Credential

1 After purchasing a license, you will receive an e-mail from Novell with your Customer ID. The e-mail also contains a link under the Order Detail section to the site where you can obtain your generic credential. Click the link to go to the site. See Figure 2-1.

Figure 2-1 *Order Detail Section of the E-Mail*

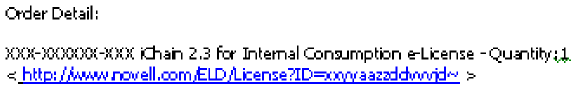

Order Detail:

XXX-XXXXXX-XXX iChain 2.3 for Internal Consumption e-License - Quantity: 1
< http://www.novell.com/ELD/License?ID=xxyyaazzddwwjd~ >

IMPORTANT: Only three differing e-mail addresses can be used to access the link where you can obtain the generic credential. If you try to access the link with more than three e-mail addresses, it is considered as a security risk and you are denied access. Additionally, only the e-mail address designated as the owner/contract for the Customer ID will receive the e-mail containing the Order Detail section with the information on obtaining the generic license. If your response e-mail does not contain the Order Detail section, you need to contact the Customer ID person within your organization to obtain the generic credential.

After clicking the link, you should see a page similar to Figure 2-2:

Figure 2-2 Electronic License Distribution Page

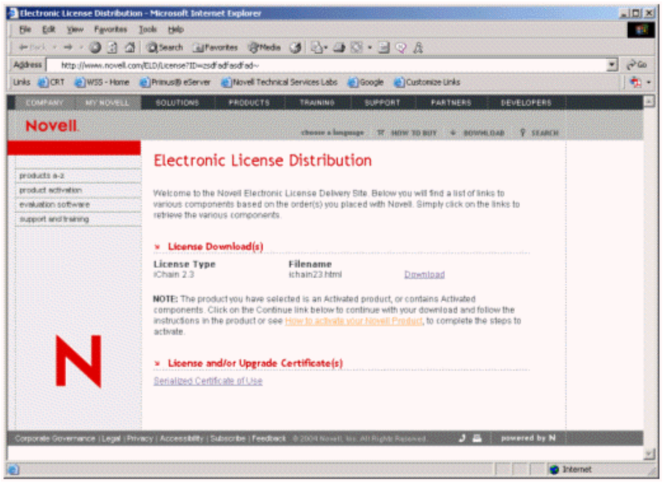
Electronic License Distribution - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://www.novell.com/ELD/License?ID=zsdfadfasdfad~
Links CRT WSS - Home Primus® eServer Novell Technical Services Labs Google Customize Links
COMPANY MY NOVELL SOLUTIONS PRODUCTS TRAINING SUPPORT PARTNERS DEVELOPERS
Novell
choose a language HOW TO BUY DOWNLOAD SEARCH
products a-z
product activation
evaluation software
support and training
Corporate Governance | Legal | Privacy | Accessibility | Subscribe | Feedback
powered by N
Internet

Electronic License Distribution

Welcome to the Novell Electronic License Delivery Site. Below you will find a list of links to various components based on the order(s) you placed with Novell. Simply click on the links to retrieve the various components.

License Download(s)

License Type	Filename	
iChain 2.3	ichain23.html	Download

NOTE: The product you have selected is an Activated product, or contains Activated components. Click on the Continue link below to continue with your download and follow the instructions in the product or see How to activate your Novell Product, to complete the steps to activate.

License and/or Upgrade Certificate(s)

Serialized Certificate of Use

2 Click the license download link and either save (download) or open the .html file.

After the file is opened, its content should be similar to the content shown in Figure 2-3:

Figure 2-3 Activation Credential

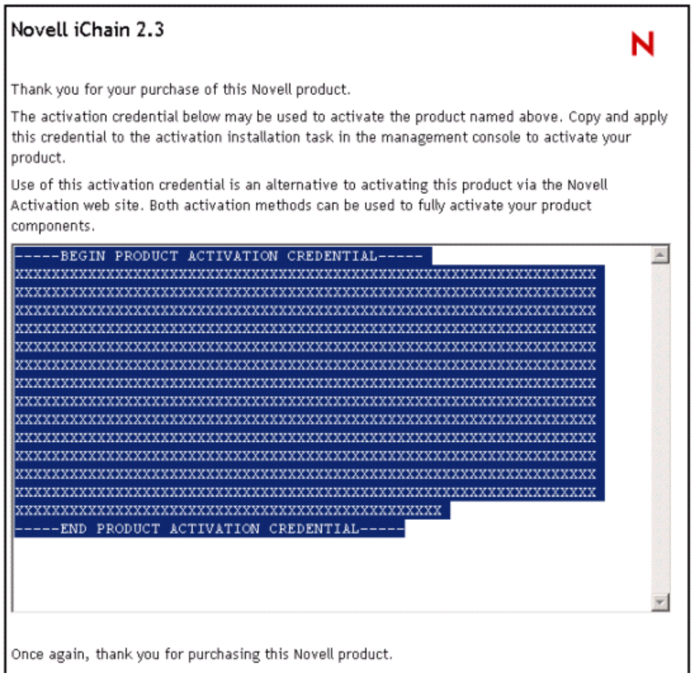

Novell iChain 2.3

Thank you for your purchase of this Novell product.

The activation credential below may be used to activate the product named above. Copy and apply this credential to the activation installation task in the management console to activate your product.

Use of this activation credential is an alternative to activating this product via the Novell Activation web site. Both activation methods can be used to fully activate your product components.

```
-----BEGIN PRODUCT ACTIVATION CREDENTIAL-----
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
-----END PRODUCT ACTIVATION CREDENTIAL-----
```

Once again, thank you for purchasing this Novell product.

3 Proceed to "Installing the Product Activation Credential Received from Novell" on page 33 for instructions on how to activate iChain.

Activating iChain Using a Product Activation Request

You use your Customer ID to create a Product Activation Request. After you create a Product Activation Request, you submit this request to Novell at the Novell Product Activator Web site (http://www.novell.com/activator). After you submit the Product Activation Request, Novell sends you an e-mail containing a Product Activation Credential that you use to activate iChain.

Generating a Product Activation Request

You should not attempt to generate a Product Activation Request until you have configured your Access Control page in the Proxy Administration Tool and have enabled FTP.

If you are using a generic credential, you do not need to generate a Product Activation Request. See "Activating iChain Using a Generic Credential" on page 30.

You use your Customer ID to generate a Product Activation Request in ConsoleOne. You must have an iChain Service Object (ISO) already set up to generate a Product Activation Request to activate iChain. See Chapter 11, "Using the iChain Service Object" in the online iChain 2.3 documentation for details. The iChain Proxy Server must be able to access the ISO via LDAP in order to generate the activation request.

You need to generate only one Product Activation Request for each ISO where iChain is installed for each iChain license that you have acquired.

1. Open ConsoleOne.

   You must use the version of ConsoleOne that you installed with iChain 2.3.
2. Click Wizards > Create an iChain Activation Request.
3. Select the iChain Service Object (ISO) where you want iChain to be activated, then click Next.
4. Type your Novell Customer ID, then click Next.
5. A dialog box appears, asking whether you want to automatically or manually notify the proxy server of the activation request. If you have FTP enabled on the proxy (see "Installing the iChain Proxy Services Software" on page 25), you can click Yes and enter the IP address and config password of any iChain Proxy Server that has been configured to read this ISO.

   Otherwise, click No and enter the command:

   ```
   getactivationrequest
   ```

   from the iChain command line interface.
6. The wizard then waits for the iChain Proxy Server to generate a Product Activation Request. Every 30 seconds you are prompted to ask whether you want to continue to wait for a response. If no response occurs, ensure that there is LDAP communication between the iChain Proxy Server and the ISO.
7. Do one of these steps:

   Specify the name of the Activation Request file and where you want the file written to, then click Next.

   or

   Copy the Activation Request in the text area to the Clipboard. You paste the contents of this file in a text area at the Novell Product Activator Web site.

   IMPORTANT: Do not edit the contents of the Product Activation Request.

   or

   If you are working on a machine without a browser, save the file to a diskette to upload at a machine that has a browser.
8. Click Launch to launch the Novell Product Activator Web site.
9. Continue with Submitting an Activation Request.

Submitting an Activation Request

You submit the Product Activation Request that you generated in ConsoleOne to the Novell Product Activator Web site (http://www.novell.com/activator).

1 Log in at the Product Activator Web site (http://www.novell.com/activator).

  You must have an eLogin account to access the Product Activator Web site. If you do not already have an eLogin account, you must create one when you visit the Product Activator site.

2 Click Browse to specify the path to the Product Activation Request file, or paste the text of the Product Activation Request into the text area.

  If you copied the Product Activation Request to a diskette, make sure you put the request on the computer you are working on.

  IMPORTANT: Do not edit the contents of the Product Activation Request.

3 Click Submit.

4 Mark the product you are activating.

  You need to activate each line item.

5 Click Submit.

  Novell generates a Product Activation Credential based on the Product Activation Request you submitted and sends that credential to you via e-mail.

Installing the Product Activation Credential Received from Novell

You activate iChain by installing the Product Activation Credential you received from Novell. (The process is the same, whether you have a Product Activation Credential from Novell or you are using a generic credential.)You install this file via ConsoleOne. If a tree has multiple iChain Service Objects (ISOs), the Product Activation Credential must be installed for each ISO.

NOTE: Make sure you install the Product Activation Credential on the same tree where you generated the Product Activation Request.

1 Open the e-mail that contains the Product Activation Credential from Novell.

  You can also use the generic credential on your product page for use with iChain. All of the following steps are identical, and you don't need to submit an activation request in this case.

2 Do one of these steps:

  Save the Product Activation Credential file.

  or

  Open the Product Activation Credential file, then copy the contents of the Product Activation Credential file to your clipboard.

  IMPORTANT: Do not edit the contents of the Product Activation Request.

3 Open ConsoleOne.

4 Click Wizards, then click Install an iChain Activation.

5 Select the iChain Service Object (ISO), then click Next.

6 Do one of these steps:

Specify where you saved the iChain Activation Credential, then click Next.

or

If you copied the contents to a clipboard, paste the contents of the iChain Activation Credential in the text area, then click Next.

7 A dialog box appears, asking whether you want to automatically or manually notify the proxy server of the new credential. If you have FTP enabled on the proxy (see "Installing the iChain Proxy Services Software" on page 25), you can click Yes and enter the IP address and config password of any iChain Proxy Server that has been configured to read this ISO.

   Otherwise, click No and enter the command:

   ```
   refreshcredentials
   ```

   from the iChain command line interface.

   If you have multiple iChain Proxy Servers reading the same ISO, such as when the Session Broker is used for load balancing and failover, you can only automatically notify one of them of the new credential. You need to enter the command `refreshcredentials` on the console of all other iChain Proxy Servers in order for them to immediately recognize the new credential.

8 Click Finish.

Viewing Product Activations for iChain

For each of your purchases of an iChain license, you can see the product activations you have installed for those licenses for iChain on the Proxy Administration Tool on the Home > Introduction panel.

Troubleshooting iChain Activation Problems

You must configure the LDAP/Authorization Server information in the Proxy Administration Tool. This information is entered by going to Configure, then clicking the Access Control tab. Make sure you enter the following information: the iChain Service Object LDAP Name (this requires that the ISO object has been configured in eDirectory), the LDAP Server Address, LDAP Port, LDAP Proxy User and Password. Also, the Proxy User must have sufficient rights (if you aren't sure, use *admin*). For more information about setting up proxy users, see the Novell Technical Information Document, "LDAP Proxy User Minimum Rights for iChain" (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10084506.htm).

Most problems with generating activation requests and installing credentials are caused by faulty LDAP communication between the iChain Proxy Server and the ISO, and are also caused by the LDAP user having only browse rights to the LDAP server. To troubleshoot these problems, enter No when prompted to notify the iChain Proxy Server automatically in the wizard, and enter `getactivationrequest` from the iChain Proxy Server console if you are generating an activation request, or `refreshcredentials` if you are installing a new credential.

If an error occurs, make a note of the error code and contact the Novell Activation Center at 1-800-418-8373 in the U.S. In all other locations, call 1-801-861-8373. (You will be charged for calls made using the 801 area code.)

In addition, make sure that:

- The eDirectory schema has been correctly extended for iChain. (You can check the local ice.log file, created by default in the c:\winnt directory, to make sure the schema updates were successful.)
- You are not trying to install the Activation Request file as an Activation Credential.
- The credential you are using was created from an activation request generated on the same tree where your ISO is located. The credential functions only on a specific tree.
- The LDAP user is not limited to having only browse rights to the LDAP server.

When you go into debug mode (unlock / debug) on the proxy command line interface, select Proxy menu, option 19, option 4 to look at the ISO object information. It indicates whether a problem exists.

2.6 Installing the iChain Proxy Server on Your Network

The iChain Proxy Server is an integral part of iChain. This section is included to help you understand the concepts behind accelerating a Web server.

This section contains the following topics:

- Section 2.7, "Preparing the Network," on page 35
- Section 2.8, "Troubleshooting iChain Proxy Server Issues," on page 38

2.7 Preparing the Network

After you complete the initial proxy server installation and activation, review this section to ensure that all your network components are properly configured.

2.7.1 Basic Network Configuration Setup

Figure 2-4 on page 36 provides a visual map for the information in this section.

NOTE: The letters in Figure 2-4 on page 36 are referenced in the tables that follow. The addresses shown are for illustration purposes only. You need to substitute actual addresses for your network.

Figure 2-4 Basic Network Configuration Setup

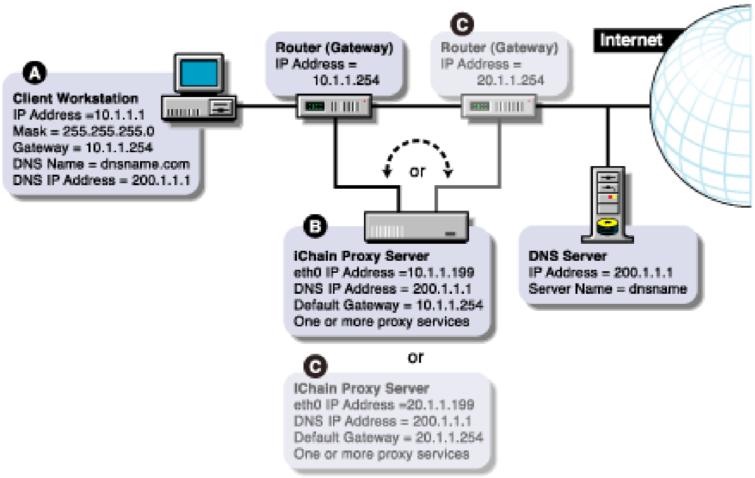

2.7.2 Configuring the Client Workstation

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 the iChain Proxy Server must be configured with the IP address information listed in the following table.

IP Address Information	Must Be on the Same Subnet (Y/N)
A numeric IP address on the subnet	Y
The subnet mask	Y
The numeric IP address of the default gateway for the subnet	Y
The numeric IP address of the DNS server the browser will use to resolve DNS names	N
The domain name for the DNS server the client will use (optional)	N

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 The subnet mask The numeric IP address of the default gateway for the subnet The numeric IP address of the DNS server the browser uses to resolve DNS names The domain name for the DNS server the client will use (optional)	Refer to the setup instructions for the system. The procedure is different for each platform. On a Windows NT workstation, for example, right-click the Network Neighborhood icon on the desktop.	See **A** in Figure 2-4 on page 36. The IP address, subnet mask, and gateway address must all be on the same subnet.

2.7.3 Configuring the iChain Proxy Server

IMPORTANT: When possible, connect the network cable to the network card on the iChain Proxy Server before assigning an IP address to the card. If this is not possible, you might need to restart the server after the cable is attached for the IP address assignment to take effect.

Configure the iChain Proxy Server with the following information:

To Configure	Do This	Notes
IP addresses and subnet masks for the network connections (eth0, eth1, etc.) that handle proxy services	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, then select the appropriate adapter. 3. Click Apply.	See **B** and **C** in Figure 2-4 on page 36. The proxy server need not be on the same subnet as the browser. If not, its IP address reflects a different subnet. Eth0, eth1, etc., can be on different subnets.
At least one DNS server IP address	1. In the browser-based tool, click Network > DNS. 2. Enter the addresses in the appropriate fields. 3. Click Apply.	

To Configure	Do This	Notes
The numeric IP address for a gateway (router) on the same subnet as the proxy server	1. In the browser-based tool, click Network > Gateway/ Firewall. 2. Specify the address in the Default Gateway IP Address field. 3. Click Apply.	See **B** in Figure 2-4 on page 36. If the iChain Proxy Server is on the same subnet as the client workstation, the proxy server and the workstation have the same gateway address. If the proxy server is on a different subnet than the browser, its gateway address is the IP address of the router on the other subnet. See **C** in Figure 2-4 on page 36.
Passwords for Config and View users	1. In the browser-based tool, click System > Actions > Password. 2. Click the User drop-down list, then select Config. 3. Enter the password information, then click Change. 4. Repeat Step 2 and Step 3 for the View user. 5. Click Apply.	**NOTE:** Telnet is not secure unless a password is set. We strongly recommend you set system passwords as part of the initialization process.
One or more proxy services	See the sections that follow.	

IMPORTANT: If you are reinitializing the system, you should remove the CD, shut down iChain Proxy Server, turn the proxy server off, and restart it again.

2.8 Troubleshooting iChain Proxy Server Issues

This section contains information on how to troubleshoot problems you might encounter with your iChain Proxy Server. The following topics are covered:

- Section 2.8.1, "Custom Error Pages," on page 38
- "Troubleshooting iChain Proxy Server Connectivity Issues" on page 39
- "Troubleshooting iChain Proxy Server Authentication Issues" on page 41
- "Troubleshooting iChain Proxy Server Authorization Issues" on page 41

2.8.1 Custom Error Pages

By default, iChain provides only English error pages. However, iChain can plug in translated error pages. This section explains how to determine the error page languages.

1 Open the iChain Proxy Server Administration tool.

2 Click Configure, then click the Mini Web tab.

3 Select the HTTP Web server error page language.

4 At the iChain prompt, create a messages.cfg file for the selected language.

For example, sys:etc\proxy\data\errpage\nls\english\messages.cfg.

5 Type the following commands:

```
cd etc/proxy/data/errpage/nls/english
cp messages.cfg ../spanish/message.cfg
edit
sys:etc\proxy\data\errpage\nls\spanish\messages.cfg
```

6 Follow the instructions in the messages.cfg file to localize the error page.

IMPORTANT: When customizing the messages, do not use a semi-colon (;) in the message. If you do, the message truncates at the semi-colon and only the portion of the message before the semi-colon is displayed.

2.8.2 Troubleshooting iChain Proxy Server Connectivity Issues

This section contains information on how to troubleshoot connectivity problems you might encounter with your iChain Proxy Server. The following topics are covered:

- "My proxy server isn't working" on page 39
- "I can't ping the proxy server from my client" on page 39
- "All the numbers are correct and it still won't ping" on page 40
- "My browser can't find the application" on page 40
- "Nothing ever comes up on my browser" on page 40
- "None of the changes I made in the browser application are taking effect" on page 40

My proxy server 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 proxy server from my client

The IP address for the client must be 172.16.0.1 and the subnet mask must be 255.255.255.0. Its gateway must be the address of the proxy server; in the original configuration that address is 172.16.0.1. DNS on the client must also be set to the IP address of the proxy server. If the ping fails with these addresses, dump the arp table on the browser (for example, using arp -a on Windows) and verify that an entry exists for the proxy server IP address 172.16.0.1. If there is no entry, the problem is likely a hardware issue. Check the cables and confirm that there is a physical connection between both IP hosts.

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

- Some Ethernet cards under Windows NT do not allow a cross-over cable. If that is the case, 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 iChain Proxy Server 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) on the Web.

My browser can't find the application

- The correct URL is http://172.16.0.1:1959/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

- We recommend that you use Internet Explorer 5.5 (or higher) with the proxy server. Also, the product release notes might contain more information regarding browser compatibility.
- You must have a JVM* (Java Virtual Machine) installed.
- Try exiting from your client OS and restarting.
- Try the SHUTDOWN command from a Telnet or command line session on the proxy server. Then turn the proxy server off and on again and wait for it to come up.
- If you just started the proxy server, you might be trying to start before the server is up. When the proxy server starts, you hear the startup beep pattern (two longs and four shorts) repeated four times. When the beeping stops, the proxy server is ready.
- If the browser appears to hang, check the URL to see if it is trying to go to a URL with port 2222 in it. If this is the case, then there is probably an issue with the default certificate used for authentication to the iChain GUI. Go to the iChain Proxy Server console and do the following:

1. Unlock the console.
2. At the iChain command line interface, enter an_kill. This causes all iChain modules to be unloaded.
3. When you have control of the iChain Proxy Server console, enter the following:

   ```
   unload certappstart.ncf
   ```

4. After loading TCPCON at the iChain Proxy Server console, confirm that TCP port 2222 is listening on that server. You can do this by going to Protocol Information > TCP > TCP Connections, then confirming that an entry exists for TCP port 2222.

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.

2.8.3 Troubleshooting iChain Proxy Server Authentication Issues

For information on troubleshooting iChain Proxy Server authentication issues, see the Novell Developer Web site (http://developer.novell.com/research/appnotes/2002/septembe/01/a020901.htm).

2.8.4 Troubleshooting iChain Proxy Server Authorization Issues

For information on troubleshooting iChain Proxy Server authorization issues, see the Novell Developer Web site (http://developer.novell.com/research/appnotes/2002/october/02/a021002.htm).

Configuring a Typical Accelerator 3

This section explains how to configure a typical Novell® iChain® accelerator and includes the following topics:

- Section 3.1, "Setting Up an Accelerator," on page 43
- Section 3.2, "Configuring the Authentication Server," on page 49
- Section 3.3, "Configuring iChain for NetIdentity Authentication," on page 51

3.1 Setting Up an Accelerator

The example setup in this section uses the following servers: an origin Web server, an LDAP authentication server, and an iChain server.

Origin Web Server Details

```
DNS name: originserver.ichain.net
IP address: 192.168.10.1
Web server port: 80
```

LDAP Authentication Server Details

IP address: 172.16.10.2

iChain Server Details

Single network card with an IP address of 172. 16.10.1

The DNS name of the accelerator is accelerator.ichain.net

accelerator.ichain.net resolves to 172.16.10.1

3.1.1 Configuring the Accelerator

The iChain Proxy Server functions as the primary access point into your iChain infrastructure. This section provides a brief introduction to the basic steps needed to set up the iChain Proxy Server.

Using the Network Configuration Page

To set up the iChain Proxy Server for an iChain implementation:

1. Access the URL of the proxy server where you installed the iChain Proxy Services software to launch the proxy server browser-based administration tool.

   For this example, this would be http://172.16.10.1:1959/appliance/config.html.

   NOTE: If the iChain Proxy Server is located behind a firewall and you are accessing the proxy server browser-based administration utility from a browser outside that firewall, you must open ports 1959, 2222, and 51100 on the firewall to administer the proxy server.

2. Accept the default username (do not enter a password), then click OK.

3 Click System > Actions > Password, then set a password for the proxy server.

4 Click Home > Introduction, then verify that the iChain Proxy Server is installed and is running on the server.

This is shown as a bitmap that indicates if you are running version 2.3.

5 Click Network, then the IP Addresses tab.

6 Configure, accept, or verify the Eth0 adapter setting (172.16.10.1).

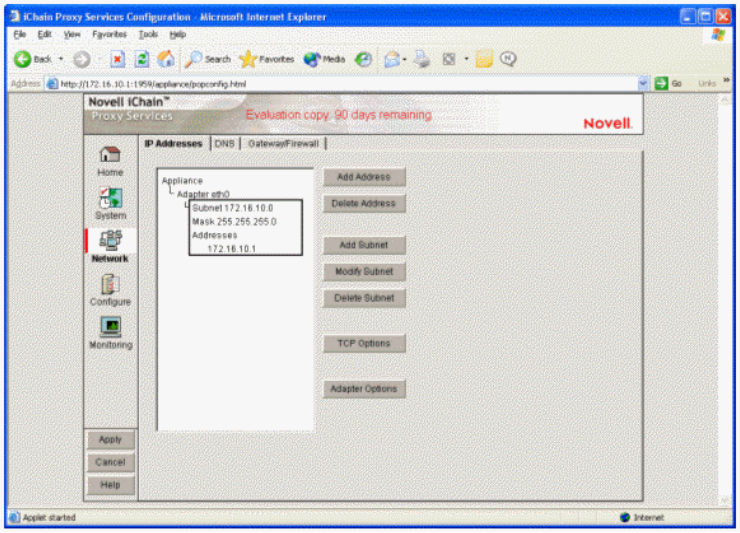
iChain Proxy Services Configuration - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://172.16.10.1:1959/appliance/popconfig.html
Go Links
Novell iChain™ Proxy Services
Evaluation copy 90 days remaining
Novell.
Home
System
Network
Configure
Monitoring
Apply
Cancel
Help
IP Addresses | DNS | Gateway/Firewall
Add Address
Delete Address
Add Subnet
Modify Subnet
Delete Subnet
TCP Options
Adapter Options
Applet started
Internet

Appliance
Adapter eth0
Subnet 172.16.10.0
Mask 255.255.255.0
Addresses
172.16.10.1

7 Click the Gateway-Firewall tab, then set the iChain Proxy Services default gateway to the gateway necessary to access your public IP address.

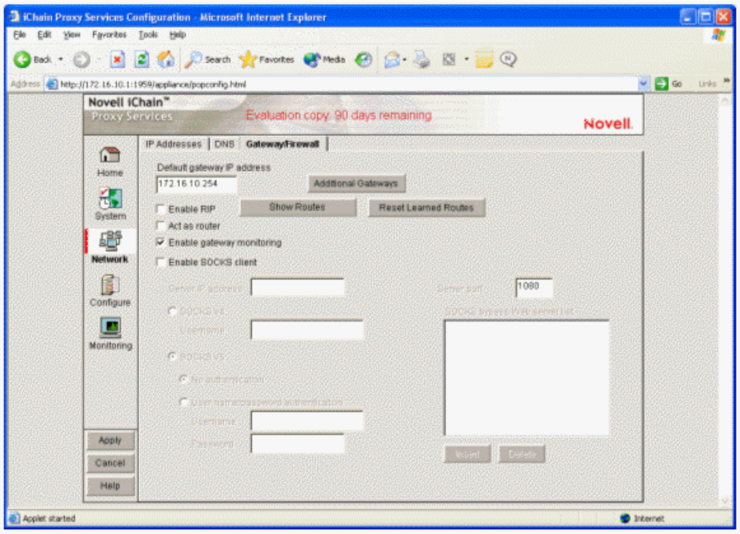
iChain Proxy Services Configuration - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://172.16.10.1:1959/appliance/popconfig.html
Go Links
Novell iChain™ Proxy Services
Evaluation copy 90 days remaining
Novell.
Home
System
Network
Configure
Monitoring
Apply
Cancel
Help
IP Addresses | DNS | Gateway/Firewall
Additional Gateways
Show Routes
Reset Learned Routes
Server IP address
Server port
SOCKS V4
Username
SOCKS V5
No authentication
User name/password authentication
Username
Password
SOCKS bypass Web server list
Insert
Delete
Applet started
Internet

Default gateway IP address: 172.16.10.254

☐ Enable RIP
☐ Act as router
☑ Enable gateway monitoring
☐ Enable SOCKS client

Server port: 1080

8 Click Network, then the DNS tab.

9 Specify the DNS domain name (for example, novell.com), the IP address of the DNS server, and the Appliance domain name or alias.

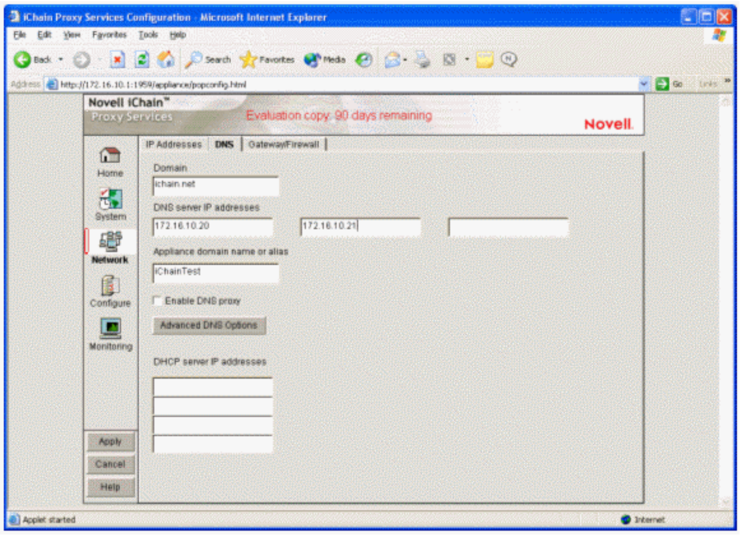
iChain Proxy Services Configuration - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://172.16.10.1:1959/appliance/popconfig.html
Go Links
Novell iChain™
Proxy Services
Evaluation copy: 90 days remaining
Novell.
IP Addresses | DNS | Gateway/Firewall
Home
System
Network
Configure
Monitoring
Advanced DNS Options
Apply
Cancel
Help
Applet started
Internet

Domain: ichain.net

DNS server IP addresses: 172.16.10.20 | 172.16.10.21

Appliance domain name or alias: iChainTest

☐ Enable DNS proxy

DHCP server IP addresses

10 Click Apply for the new settings to take effect.

11 Click System > Actions, then verify the internal and external connections to your network by pinging the origin server you will be accelerating within your internal network and an external Host on the Internet.

The following figure shows an example of iChain unsuccessfully pinging the origin server. If this is your experience, you must resolve this issue before proceeding.

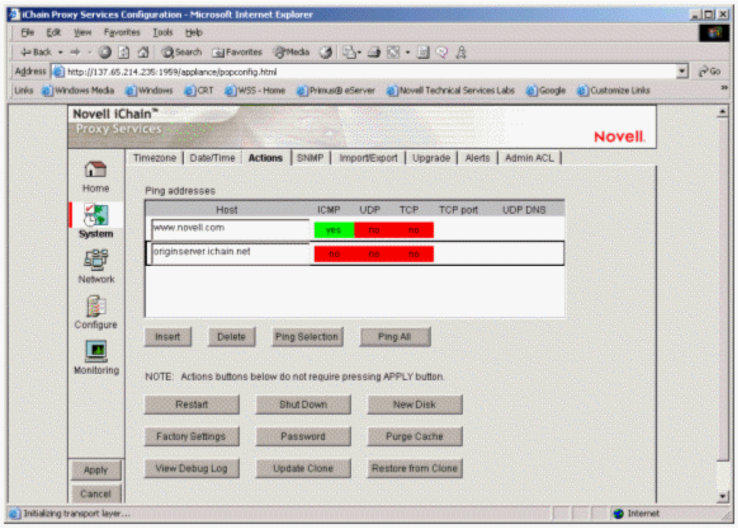
iChain Proxy Services Configuration - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://137.65.214.235:1959/appliance/popconfig.html
Go
Links Windows Media Windows CRT WSS - Home Primus® eServer Novell Technical Services Labs Google Customize Links
Novell iChain™
Proxy Services
Novell.
Timezone | Date/Time | Actions | SNMP | Import/Export | Upgrade | Alerts | Admin ACL
Home
System
Network
Configure
Monitoring
Insert
Delete
Ping Selection
Ping All
Restart
Shut Down
New Disk
Factory Settings
Password
Purge Cache
View Debug Log
Update Clone
Restore from Clone
Apply
Cancel
Initializing transport layer....
Internet

Ping addresses

Host	ICMP	UDP	TCP	TCP port	UDP DNS
www.novell.com	yes	no	no		
originserver.ichain.net	no	no	no		

NOTE: Actions buttons below do not require pressing APPLY button.

Configuring Authentication

To set up access to the iChain Authorization Server for the authentication function, you need to create an authentication profile. Follow these steps to create an LDAP profile that authenticates users to your iChain Authorization Server:

1 In the proxy server administration tool, click Configure, then click Authentication.

2 Insert a new profile, name the profile, select LDAP Authentication, then click LDAP Options.

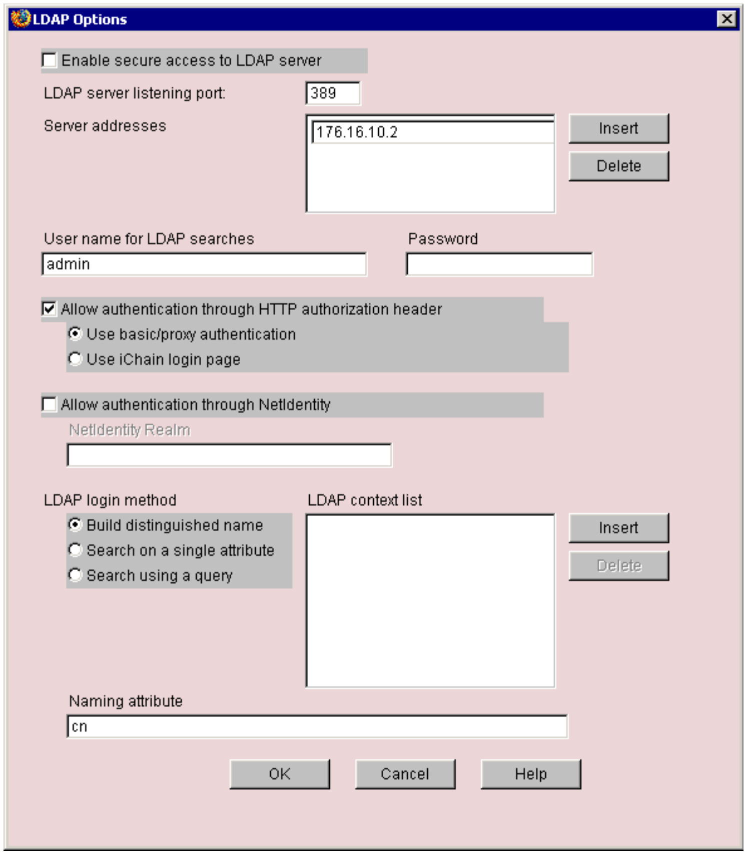

☐ Enable secure access to LDAP server

LDAP server listening port: 389

Server addresses: 176.16.10.2

User name for LDAP searches: admin

Password:

☑ Allow authentication through HTTP authorization header

● Use basic/proxy authentication

○ Use iChain login page

☐ Allow authentication through NetIdentity

NetIdentity Realm:

LDAP login method

● Build distinguished name

○ Search on a single attribute

○ Search using a query

LDAP context list

Naming attribute: cn

3 Specify 389 as the LDAP server listening port for non-secured LDAP.

4 Click Insert next to Server Addresses and set the server IP address to the iChain Authorization Server address.

5 Specify a username and password for LDAP access.

For the initial setup, try using the Admin user to avoid rights issues. If this is not possible, create an LDAP proxy user with rights set up as follows:

5a Make the LDAP Proxy User a trustee of the user's container (or ROOT) and give it a specific assignment of Read, Compare, and Write rights to the Object Class property.

5b The LDAP Proxy User also needs Write rights to the ISO object for license activation.

5c The LDAP user also must have Read and Write rights to all users (in the user's containers).

For more information about setting up proxy users, see the Novell Technical Information Document, “LDAP Proxy User Minimum Rights for iChain” (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10084506.htm).

6 (Optional) Select Allow authentication through HTTP authorization header.

You can use basic/proxy or the iChain login page for this authorization.

7 (Optional) Select Allow authentication through NetIdentity.

8 Specify the NetIdentity Realm.

The NetIdentity Realm is the NetIdentity tree name.

9. Select one of the following LDAP login methods:
    - Build distinguished name.
    - Search on a single attribute.
    - Search using a query.
10. Click Insert, then specify an LDAP context. For example, ou=test,o=mycompany.
11. Specify the Naming Attribute.

    This name uniquely identifies each entry in the Directory Information Tree.
12. Repeat Step 10 for each context users will authenticate from.
13. Click OK twice, then click Apply.

    To set login and search timeouts for this profile, see “Setting Timeouts and Pool Limits for LDAP Profiles” on page 310.

Configuring Authorization (Access Control)

To set up access to the authorization server for access control functions:

1. In the proxy server administration tool, click Configure, then click Access Control.
2. Specify the fully distinguished name of the ISO object name for the iChain service.

    You must use commas as delimiters. For example, cn=myISO,o=novell.
3. Specify the following LDAP profile settings:

    LDAP server addresses for the iChain LDAP access control servers
    LDAP port on the iChain LDAP access control servers
    LDAP proxy user
    Password

    The LDAP user name and password must have the following rights:
    - Write rights to the ISO object for license activation.
    - Read and write rights to all users (in the user's containers).
    - Make the LDAP Proxy User a trustee of the user's container (or ROOT) and give it a specific assignment of Read, Compare and Write rights to the Object Class property.
4. Click Apply.
5. Click Refresh ACLCHECK.

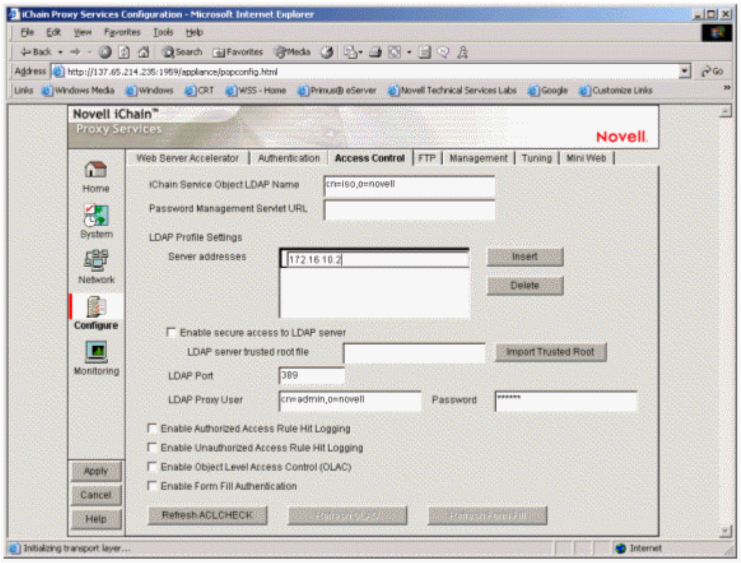

iChain Proxy Services Configuration - Microsoft Internet Explorer

Address: http://137.65.214.235:1959/appliance/popconfig.html

iChain Service Object LDAP Name: cn=iso,o=novell

Password Management Servlet URL:

LDAP Profile Settings

Server addresses: 172.16.10.2

□ Enable secure access to LDAP server

LDAP server trusted root file:

LDAP Port: 389

LDAP Proxy User: cn=admin,o=novell Password: ******

□ Enable Authorized Access Rule Hit Logging

□ Enable Unauthorized Access Rule Hit Logging

□ Enable Object Level Access Control (OLAC)

□ Enable Form Fill Authentication

Initializing transport layer...

Configuring FTP

Enable FTP on the administration IP address. In this example, it would be 172.16.10.1.

Figure 3-1 *Configuring FTP*

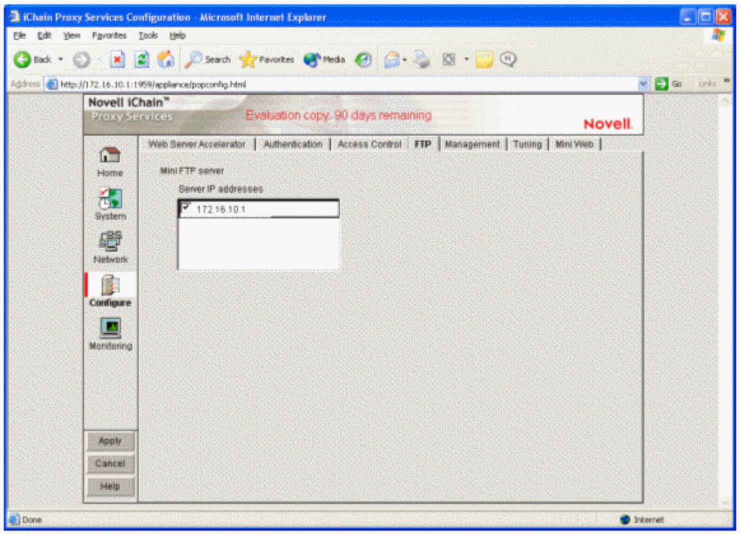

iChain Proxy Services Configuration - Microsoft Internet Explorer

Address: http://172.16.10.1:1959/appliance/popconfig.html

Evaluation copy: 90 days remaining

Mini FTP server

Server IP addresses

☑ 172.16.10.1

Configuring the Accelerator

To set up a Web Server accelerator:

1. In the proxy server administration tool, click Configure, click Web Server Accelerator, then click Insert.
2. In the Web Server Accelerator dialog box, specify a Name for the accelerator using a maximum of 8 characters.

   The name must be unique for each Web Server accelerator. In this example, it is named TEST.
3. Specify a DNS name for the accelerator.

This is the DNS name by which users access the resource. It should resolve to the public IP address of the iChain Proxy Server. In this example, the name is accelerator.ichain.net.

4 Specify the Alternate host name.

This is the DNS name of the origin server. In this example, it is originserver.ichain.net.

5 In the Web server addresses field, click Insert, then specify the IP addresses of the origin Web server that contains the desired content.

This will usually be on your private network. In this example, it is 192.168.10.1. Clients should not be able to access this server directly or the iChain infrastructure will be bypassed.

6 In the Accelerator IP addresses field, check the public IP address or address that the DNS name specified in Step 3 resolves to.

In this example, it is 172.16.10.1.

7 Check Enable authentication.

8 Click Authentication Options, select an existing profile from the list, then click Add to set the profile as the Service Profile.

9 Check Enable Secure Exchange.

10 Click OK twice, then click Apply.

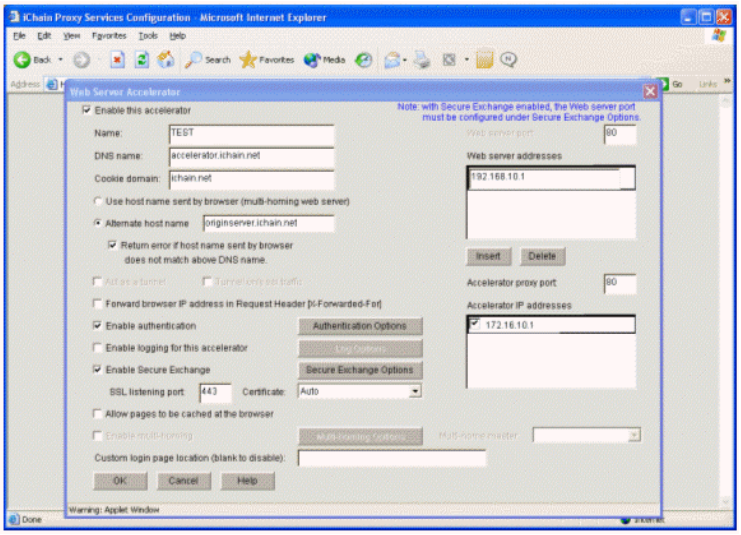

iChain Proxy Services Configuration - Microsoft Internet Explorer

Web Server Accelerator

Note: with Secure Exchange enabled, the Web server port must be configured under Secure Exchange Options.

☑ Enable this accelerator

Name: TEST

DNS name: accelerator.ichain.net

Cookie domain: ichain.net

Web server port: 80

Web server addresses: 192.168.10.1

○ Use host name sent by browser (multi-homing web server)

● Alternate host name: originserver.ichain.net

☑ Return error if host name sent by browser does not match above DNS name.

□ Act as a tunnel

□ Tunnel only ssl traffic

Accelerator proxy port: 80

□ Forward browser IP address in Request Header [X-Forwarded-For]

Accelerator IP addresses: ☑ 172.16.10.1

☑ Enable authentication

□ Enable logging for this accelerator

☑ Enable Secure Exchange

SSL listening port: 443

Certificate: Auto

□ Allow pages to be cached at the browser

□ Enable multi-homing

Multi-home master:

Custom login page location (blank to disable):

Warning: Applet Window

Proceed with Section 3.2, “Configuring the Authentication Server,” on page 49.

3.2 Configuring the Authentication Server

To configure the Authentication Server, you must first create an iChain Service Object (ISO). An iChain service is a logical entity that defines an iChain domain and the resources of that domain. Configuration for your iChain service is contained in an ISO. To set up a basic iChain implementation, you must create an ISO and set up the service parameters for the object. In addition, to set up your basic implementation, you must also set up access to Web-based application resources.

3.2.1 Creating an iChain Service Object

1 From ConsoleOne®, select an OU in which to create your ISO, then select File > New > iChain Object.

or

Click the New iChain object icon (to the right of the New Object icon).

2 Select iChain Service Object, define a name for the service or domain (for example, ISO), then click OK.

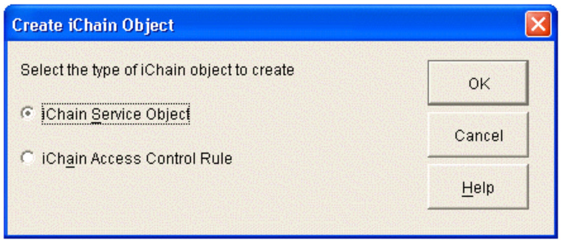

Create iChain Object

Select the type of iChain object to create

● iChain Service Object

○ iChain Access Control Rule

3.2.2 Setting Up a Protected Resource

To integrate and allow access to Web-based application resources, you must set the appropriate parameters in the ISO.

To set up a protected resource for an iChain service:

1 From ConsoleOne, click the Protected Resources tab on the ISO object you created for this configuration.

2 Click Add (the icon with the plus [+] sign).

3 Specify a name for the resource and the URL for the resource in this format: http://www.resource.com/*, where www.resource.com is the DNS name specified when you created the Web Server Accelerator.

In this example, it is accelerator.ichain.net/*.

4 Set the Access to this resource as Restricted.

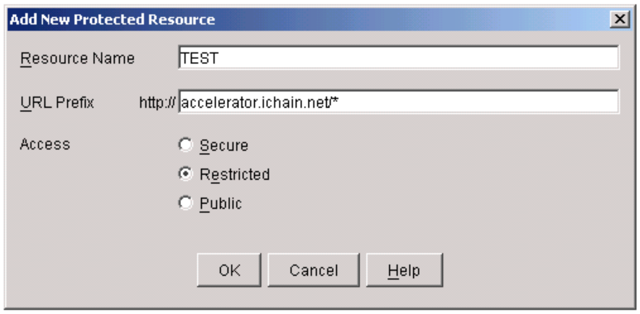

Add New Protected Resource

Resource Name: TEST

URL Prefix: http:// accelerator.ichain.net/*

Access:
○ Secure
● Restricted
○ Public

5 Click OK, then click Apply to save the resource.

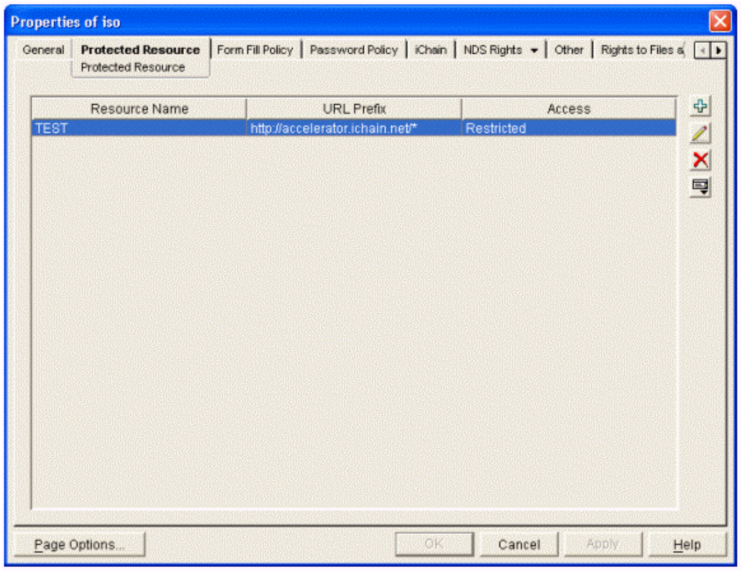

Resource Name	URL Prefix	Access
TEST	http://accelerator.ichain.net/*	Restricted

6 If you configured or enabled FTP in the Web-based administration utility, you are prompted to refresh the iChain Proxy Server. Otherwise, you need to go to the Web-based administration utility and click Configure, click Access Control, then click Refresh ACLCheck to read the new protected resource.

NOTE: Novell Technical Support does not support the password management servlets used in the iChain configuration. These servlets are not developed by the iChain development team; therefore, no bugs related to these servlets will be fixed.

The ConsoleOne Password Policy TAB contains options for writing to the ISO object. This means that the password management servlet from forge.novell.com or a custom developed servlet can read password properties from an iChain object. However, Novell Technical Services does not support servlets from forge.novell.com. For more information, see TID 407040 (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10097123.htm).

3.3 Configuring iChain for NetIdentity Authentication

iChain lets you use the NetIdentity protocol for proxy authentication. This also allows for single sign-on to back-end applications like NetStorage and Zen for Desktops.

NetIdentity authentication requires the NetIdentity client 1.2.1 or later or Zen for Desktops Agent plugins that are included in the Zen for Desktop 401 patch to be installed on the workstation.

- Section 3.3.1, “How NetIdentity Works with iChain Authentication,” on page 51
- Section 3.3.2, “Configuring an iChain Accelerator to Use NetIdentity,” on page 52

3.3.1 How NetIdentity Works with iChain Authentication

IMPORTANT: The NetIdentity client is compatible only with Internet Explorer*.

For NetIdentity to work with iChain authentication, assume you have an accelerator that is using an authentication profile with the option Allow authentication through NetIdentity enabled. A workstation with the NetIdentity client needs to access a Web server through the accelerator. The browser sends a GET request to iChain, which does not include a valid iChain cookie because the workstation is not yet authenticated to iChain. Because the workstation has NetIdentity installed, the GET request includes a header with the value NovINet: v1.2. When iChain receives the GET, it returns a 302 Found packet and redirects the browser to the ICSLogin URL. The browser then sends a GET request to that URL, again with the NovINet header.

iChain replies with a 401 Unauthorized packet. This packet includes NetIdentity specific information such as NetIdentity Realm name and other information used by the NetIdentity protocol. When the NetIdentity client receives this information, if Strict Trust is enabled (default), NetIdentity verifies that the server and Certificate Authority (CA) are trusted (according to the list of trusted authorities in Internet Explorer). If the server and CA are trusted, NetIdentity checks its existing credential store (i.e. “wallet”) to see if authentication credentials already exist for that Realm name. If they do exist, they are sent back to iChain in attempt to authenticate. If the wallet credentials do not exist for that realm name, and the NetIdentity registry setting called Try Local Credentials is enabled, the user's desktop login credentials are sent. If no entry exists in the wallet for this realm and Try Local Credentials is disabled or fails, the NetIdentity client provides a pop-up dialog prompting the user for login credentials. The NetIdentity pop-up dialog color scheme differs from that of the browser's pop-up dialog.

After NetIdentity successfully completes the login, the credentials for that Realm are stored in the wallet.

3.3.2 Configuring an iChain Accelerator to Use NetIdentity

This section explains how to set up an Authentication Profile with NetIdentity.

1 In the proxy server administration tool, click Configure, then click Authentication.

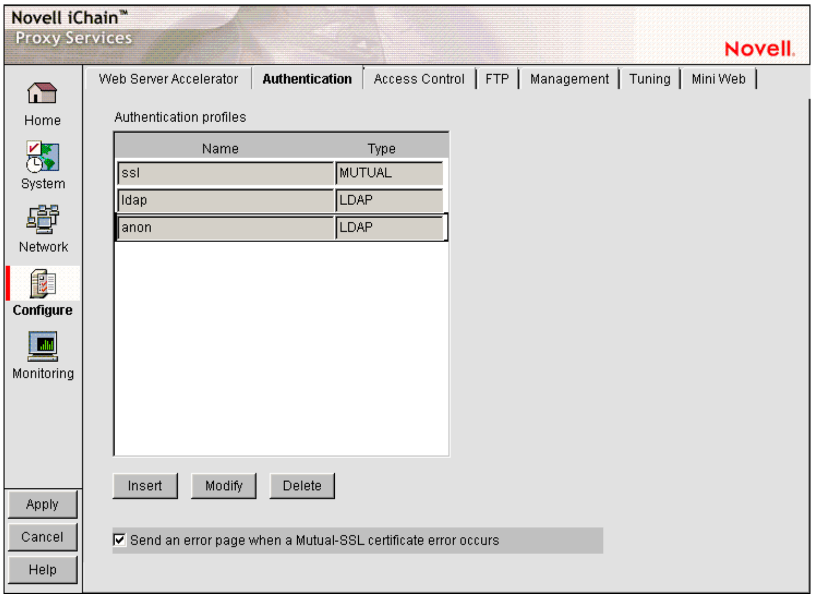

Authentication profiles

Name	Type
ssl	MUTUAL
ldap	LDAP
anon	LDAP

☑ Send an error page when a Mutual-SSL certificate error occurs

2 Do one of the following:

- Select the desired LDAP authentication profile, then click Modify.
- Click Insert to create a new authentication profile.

  If you create a new authentication profile, type the name in the Authentication Profile name field, then go to Step 3.

3 Click the LDAP Authentication radio button, then click LDAP Options.

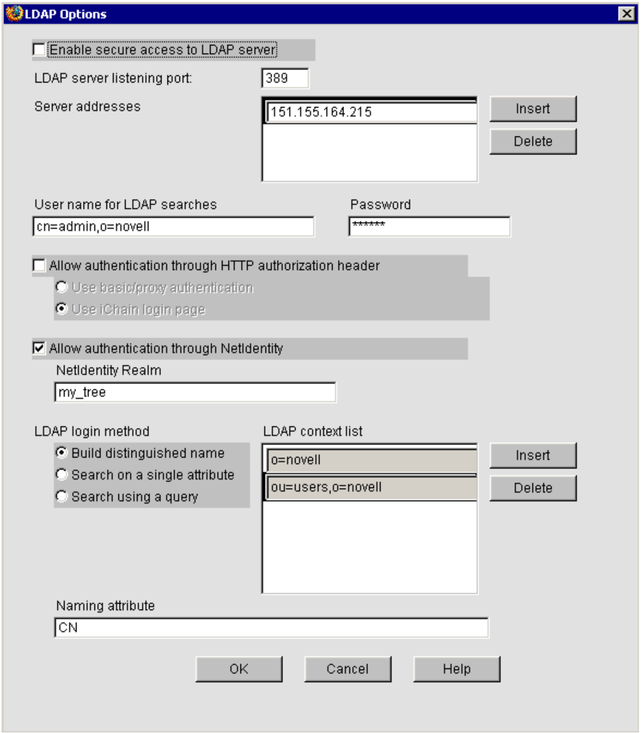

LDAP Options

☐ Enable secure access to LDAP server

LDAP server listening port: 389

Server addresses: 151.155.164.215

User name for LDAP searches: cn=admin,o=novell

Password: ******

☐ Allow authentication through HTTP authorization header

○ Use basic/proxy authentication

● Use iChain login page

☑ Allow authentication through NetIdentity

NetIdentity Realm: my_tree

LDAP login method

● Build distinguished name

○ Search on a single attribute

○ Search using a query

LDAP context list: o=novell
ou=users,o=novell

Naming attribute: CN

4 Check the Allow authentication through NetIdentity check box.

5 Type a name you want in the NetIdentity Realm Name field (case-sensitive).

IMPORTANT: The Realm Name is especially important with accelerating NetIdentity aware applications, such as NetStorage. The Realm name needs to exactly match the name of the realm used by the application. Typically, the Realm name of a NetIdentity aware application is the same as the eDirectory tree name.

6 Click OK twice, then click Apply.

After you set up an authentication profile with NetIdentity, you need to configure the accelerator to use that profile, see Section 3.3.3, "Configuring An Accelerator to Use an Authentication Profile," on page 53.

3.3.3 Configuring An Accelerator to Use an Authentication Profile

1 In the proxy server administration tool, click Configure, then click Web Server Accelerator.

2 Select the desired accelerator, then click Modify.

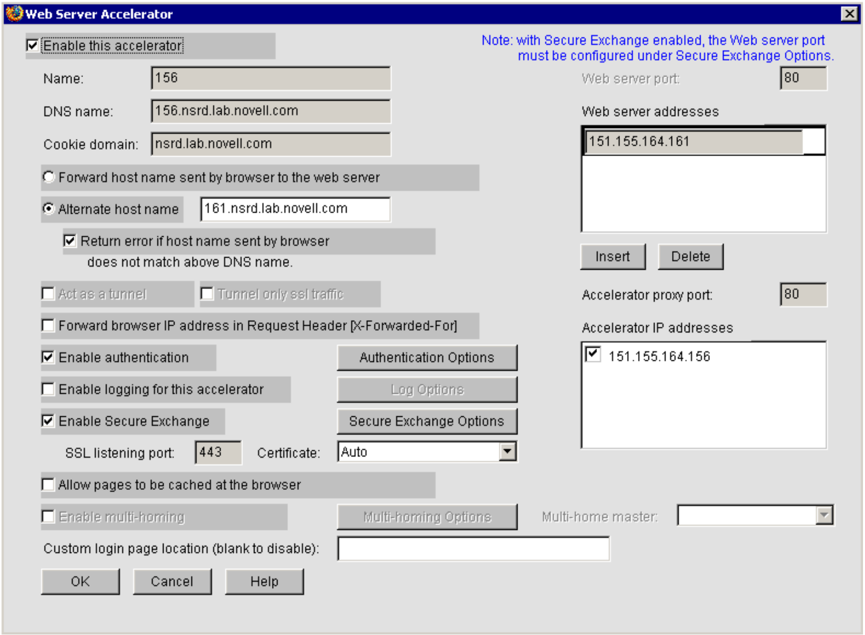

Web Server Accelerator

☑ Enable this accelerator

Note: with Secure Exchange enabled, the Web server port must be configured under Secure Exchange Options.

Name: 156

DNS name: 156.nsrd.lab.novell.com

Cookie domain: nsrd.lab.novell.com

○ Forward host name sent by browser to the web server

● Alternate host name: 161.nsrd.lab.novell.com

☑ Return error if host name sent by browser does not match above DNS name.

□ Act as a tunnel □ Tunnel only ssl traffic

□ Forward browser IP address in Request Header [X-Forwarded-For]

☑ Enable authentication

□ Enable logging for this accelerator

☑ Enable Secure Exchange

SSL listening port: 443 Certificate: Auto

□ Allow pages to be cached at the browser

□ Enable multi-homing

Custom login page location (blank to disable):

Web server port: 80

Web server addresses: 151.155.164.161

Accelerator proxy port: 80

Accelerator IP addresses: ☑ 151.155.164.156

Multi-home master:

3 Click Authentication Options.

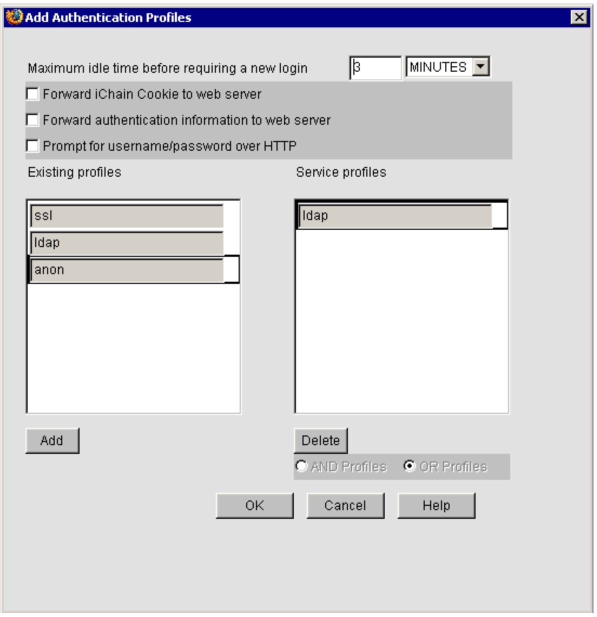

Add Authentication Profiles

Maximum idle time before requiring a new login: 3 MINUTES

□ Forward iChain Cookie to web server

□ Forward authentication information to web server

□ Prompt for username/password over HTTP

Existing profiles: ssl, ldap, anon

Service profiles: ldap

○ AND Profiles ● OR Profiles

The authentication profile displays in the Service Profiles column. If the profile you want is not in that column, select the profile in the existing profiles column, then click Add.

4 Click Ok twice, then click Apply.

Understanding the Web Server Accelerator

4

This section describes the following components of the Web server accelerator:

- Section 4.1, “Accelerating Web Servers,” on page 55
- Section 4.2, “Web Server Accelerator Setup,” on page 56
- Section 4.3, “Accelerator/Web Server Page,” on page 59
- Section 4.4, “Using Web-Based Distributed Authoring and Versioning (WebDAV),” on page 62
- Section 4.5, “Setting Up Authentication Using the Wireless Application Protocol (WAP),” on page 63
- Section 4.6, “Setting Up Secure Exchange,” on page 63
- Section 4.7, “Multi-Homing,” on page 65
- Section 4.8, “Custom Login Pages,” on page 69
- Section 4.9, “Cache Freshness,” on page 73
- Section 4.10, “GZIP Compression Support,” on page 75
- Section 4.11, “Understanding HTTP 1.1,” on page 75

4.1 Accelerating Web Servers

This section covers the following topics:

- Overview of Web Server Acceleration
- How Origin Web Server Acceleration Works
- Benefits of Origin Web Server Acceleration
- Working with DNS

4.1.1 Overview of Web Server Acceleration

The proxy server's origin Web server accelerator relies on DNS causing 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.

4.1.2 How Origin Web Server Acceleration Works

The mechanism for routing browser requests meant for Web servers to the Web server accelerator instead 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 Novell® iChain® Proxy Server Web server accelerator (reverse proxy) service.

Figure 4-1 *Web Server Acceleration*

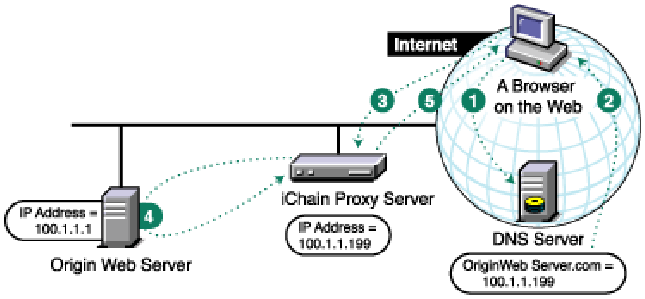

1. A browser on the Web requests an origin Web server Web page. This generates a request to DNS for the numeric IP address of the Web server.
2. Instead of returning the origin Web server's numeric IP address, DNS returns the numeric IP address, of the accelerator service on the proxy server.
3. The 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.

4.1.3 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 proxy server 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 Section 4.2, “Web Server Accelerator Setup,” on page 56.

The procedure for configuring DNS to work with Web server accelerators is explained in “Working with DNS” on page 59.

4.2 Web Server Accelerator Setup

Figure 4-2 provides a visual map for the information in this section.

NOTE: The letters in Figure 4-2 are referenced in the table that follows. The addresses shown are for illustration purposes only. You need to substitute actual addresses for your network.

Figure 4-2 *Web Server Accelerator Setup*

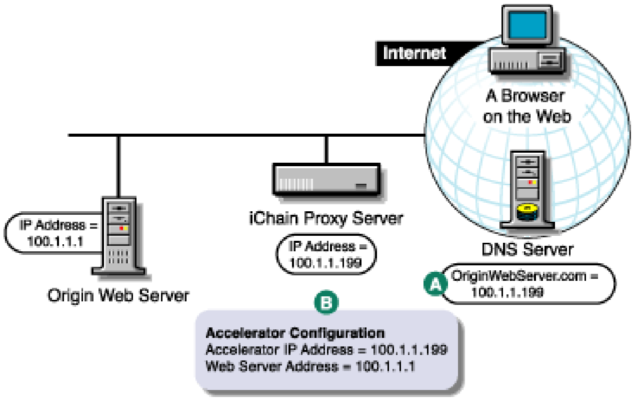
Internet
A Browser on the Web
iChain Proxy Server
IP Address = 100.1.1.199
DNS Server
A
OriginWebServer.com = 100.1.1.199
IP Address = 100.1.1.1
Origin Web Server
B
Accelerator Configuration
Accelerator IP Address = 100.1.1.199
Web Server Address = 100.1.1.1

As you set up your Web server, use the information in the following table to determine what tasks you need to complete and how.

To	Do This	Notes
Ensure that your basic network configuration is complete for each proxy server	1. See "Configuring a Typical Accelerator" on page 43.	
Ensure that DNS resolves browser requests to the proxy server IP addresses configured for the Web server accelerator services	1. See "Working with DNS" on page 59.	See **A** in Figure 4-2.

To	Do This	Notes
Set up one or more Web server accelerator services	1. In the browser-based tool, click Configure, click Web Server Accelerator, then click Insert. 2. Select a name for the Web server accelerator for your tracking purposes. 3. Specify a DNS name. 4. In the Accelerator Proxy Port field, specify the port that the Web server accelerator will receive requests and vend data on. 5. In the Accelerator IP Addresses list, select one or more addresses that the Web server accelerator will receive requests and vend data on. (DNS resolves requests to these addresses.) 6. In the Web Server Port field, specify the port that the proxy server and origin Web server will communicate on. 7. In the Web Server Addresses list, insert one or more IP addresses (or DNS names) that the Web server accelerator will fill its cache from. (The proxy server must be able to fill all requests through any of these names or addresses.) 8. To activate the Web server accelerator, select Enable This Accelerator. 9. Click OK, then click Apply.	See **B** in Figure 4-2 on page 57. If server persistence is enabled in the Web Server Accelerator tab, the proxy server uses the same Web server to fill browser requests during a session. This setting affects all accelerators on the proxy server and saves e-business users from logging in multiple times. If logging is enabled, accelerator log files for the Web server accelerator has the same name as the Web server accelerator. The DNS name is required when: ◆ You are accelerating multiple Web servers on the same IP address. (When multiple accelerator services use the same IP address.) ◆ You are accelerating a single Web site using path-based multi-homing. If you enter DNS names in the Web Server Addresses list, make sure they are not the names that now resolve to proxy server numeric IP addresses. That would create an endless loop.

IMPORTANT: Do not create more than 250 accelerators. If you create more than 250 accelarators, the proxy.nlm can cause a CPU Hog Abend.

4.2.1 Working with DNS

The steps you take to have DNS resolve requests to the proxy server rather than to the origin server depend on whether the proxy server and the origin Web server are on the same subnet. The following sections explain each alternative.

If the proxy server and the origin Web server are on the same subnet, you can swap IP addresses as shown below.

Figure 4-3 *Working with DNS*

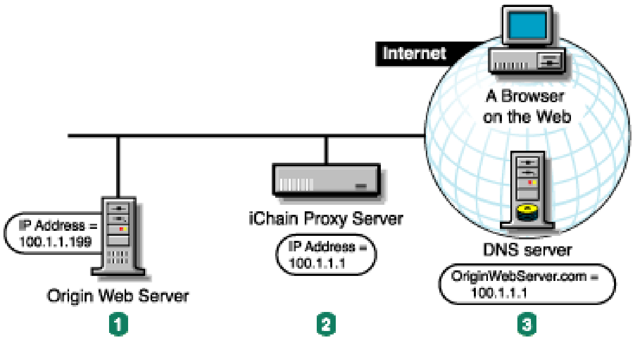

1. The origin Web server's IP address was 100.1.1.1. You change it to 100.1.1.199.
2. You assign the proxy server the IP address 100.1.1.1.
3. DNS is unchanged, but now sends browsers to the appliance instead of the origin Web server.

If the origin Web server is on a remote network, you need to alter DNS as shown below.

Figure 4-4 *Altering DNS*

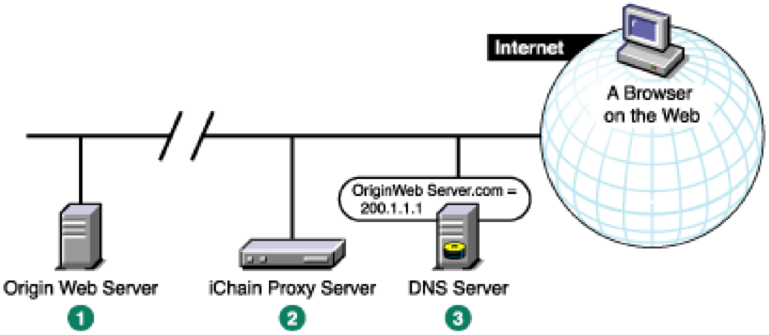

1. The origin Web server's IP address is 200.1.1.1.
2. You assign the proxy server the IP address 100.1.1.1.
3. DNS had 200.1.1.1 as the IP address for OriginWebServer.com. You change the DNS so that OriginWebServer.com now resolves to 100.1.1.1.

4.3 Accelerator/Web Server Page

On the fourth page of the wizard in ConsoleOne, the user can enable the accelerator and specify the Web servers to be accelerated, as well as which ports they are to be accelerated through and which proxy (accelerator) address and port is to be used, as shown below.

Figure 4-5 *Accelerator/Web Server Page*

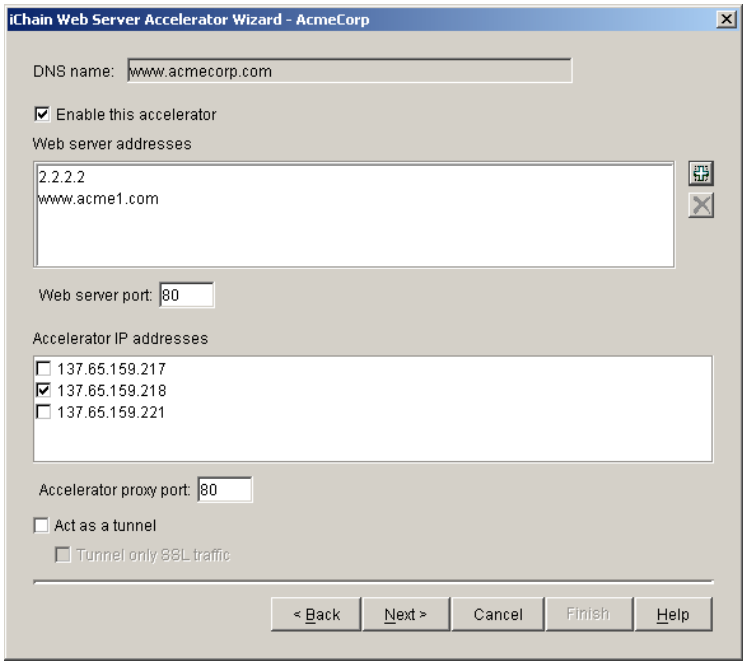

iChain Web Server Accelerator Wizard - AcmeCorp

DNS name: www.acmecorp.com

☑ Enable this accelerator

Web server addresses

2.2.2.2
www.acme1.com

Web server port: 80

Accelerator IP addresses

☐ 137.65.159.217
☑ 137.65.159.218
☐ 137.65.159.221

Accelerator proxy port: 80

☐ Act as a tunnel

☐ Tunnel only SSL traffic

The following table describes the fields on this page:

Field Name	Description	Status
DNS name	Displays the DNS name of the accelerator currently under construction or modification. The accelerator name is shown in the caption following the hyphen after the page title. In Figure 4-5, the accelerator name is AcmeCorp. The DNS Name field is non-editable and exists as an information source for the user.	Non-editable
Enable this accelerator	Selecting this check box enables the accelerator so that the data entered can be used in accelerating the specified Web servers. When the box is checked, all of the fields are enabled unless the accelerator is a path-based multi-homing child (in which case only the Web server addresses and Web server port fields are enabled). If deselected, the fields are not editable.	Optional
Web server addresses	Lists all the Web servers being accelerated by this accelerator. The Web server address can be in one of two formats: IP address or DNS name. Entries in the list can be added or deleted by using the buttons to the right of the field on the interface.	Required when the accelerator is enabled
Web server port	Specifies the port on the Web server by which the proxy server will communicate with the Web server. The default value is 80.	Required when the accelerator is enabled.

Field Name	Description	Status
Accelerator IP addresses	Displays which IP addresses does the actual acceleration. This list is populated by the proxy server. That is, it is not populated by the user, but shows the IP addresses on the proxy server that are available for accelerating. Selecting an entry in the table toggles the check box on the line on or off. A checked entry participates in accelerating the Web server.	Required when the accelerator is enabled
Accelerator proxy port	Specifies the port on the proxy server by which the proxy server will communicate with the Web server. The default value is 80.	Required when the accelerator is enabled.
Act as a Tunnel	The Act as a Tunnel option 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. When this option is selected, the accelerator sets up a tunnel for all incoming traffic.	Optional
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. If this option is checked the service then verifies that the address and port being accessed are actually an SSL Web site. If verification fails, the service tears down the connection. **NOTE:** The SSL port number for the SSL tunnel is specified via the Accelerator Proxy Port and not the SSL listening port.	Optional

4.3.1 Controls for Accelerator/Web Server

This section describes the following buttons:

- Add
- Delete

Add

When the Add button is selected, a dialog box appears where the user can enter an IP address or DNS name for a Web server, as shown in Figure 4-6:

Figure 4-6 *Add Web Server Dialog Box*

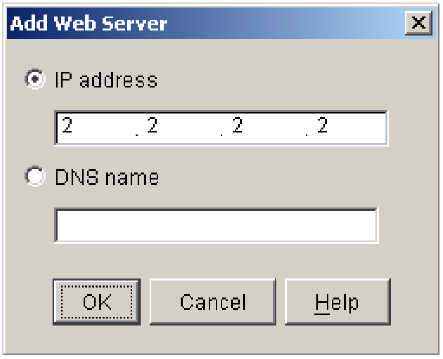

● IP address

2 . 2 . 2 . 2

○ DNS name

The following table describes the fields in this dialog box:

Field Name	Description
IP address	If selected, an IP address is used to specify the Web server address. The address is entered in the field directly below this button, as shown in Figure 4-6.
DNS name	If selected, a DNS name is used to specify the Web server address. The DNS name is entered in the field directly below this button on the interface, as shown in Figure 4-6.

Delete

The Delete button allows the user to delete an entry in the table. The actual deletion of the entry on the proxy server takes place only when the user clicks Finish at the end of the wizard session.

4.4 Using Web-Based Distributed Authoring and Versioning (WebDAV)

Web-based Distributed Authoring and Versioning (WebDAV) is a set of extensions to the HTTP protocol. This allows you to edit and manage files on remote Web servers.

iChain supports the use of WebDAV; however, there are limitations. iChain cannot rewrite hostname references in binary data. Use the option “Forward host name sent by the browser to the Web server”.

If the client to the WebDAV application is not a browser, you must enable the settings: Allow Authentication through HTTP Authorization Header, and Use Basic/Proxy Authentication in the Authentication Profile used by the accelerator. Otherwise, the application will not work.

4.4.1 Using WebDAV With NetStorage

If you are using an iChain accelerator for WebDAV connections to NetStorage, the Allow Authentication through HTTP Authorization Header and the Use Basic/Proxy Authentication options must be enabled in the Authentication Profile used by the accelerator. Users are required to provide login credentials both for iChain and for NetStorage authentications; however, if the workstation has the NetIdentity client installed and configured to trust the NetStorage CA, the user is not prompted for the additional login to NetStorage.

The Authentication profile option, Allow Authentication through NetIdentity, might give unexpected results if you enable it in addition to basic/proxy authentication. If you enable the Allow Authentication through NetIdentity option, the workstations with the NetIdentity client installed are unable to make WebDAV connections through the accelerator. This problem occurs because the WebDAV OPTIONS header is a non-redirectable request. In this case, iChain returns a 409 Conflict error, resulting in a failed authentication and connection.

4.5 Setting Up Authentication Using the Wireless Application Protocol (WAP)

iChain looks for Wireless Application Protocol (WAP) device information in the HTTP headers. When it sees WAP information, it uses the .wml templates (or smaller HTML templates if the device is expecting HTML) instead of the full-sized HTML templates. These can be found in the directory with all of the other login page templates. If there are issues, the templates can be altered so that they work with the WAP device.

There are inconsistencies with how different devices support .wml tags. During the SSL handshake, the WAP devices (phones or PDAs) need to be able to validate the server certificate being returned from iChain. This implies that the WAP devices need the trusted roots of the iChain server certificates built into them. With a browser, it is easy to import these trusted root certificates but with WAP phones, it is not an option. If the WAP device you are using does not have a built-in trusted root certificate for the iChain server certificate, it fails. Verisign* and Thawte* trusted root certificates are built in to almost all devices and these work fine; Novell trusted root certificates are not and therefore WAP devices fail if the auto-certificates are enabled for the iChain accelerator. One workaround to this problem is to authenticate over HTTP using the Prompt Username/ Password over HTTP option. However, if you do this, the authentication data is posted (POST) in clear text over the WAP network, which is a security concern. For more information, see the Novell Technical Information Document (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10085481.htm).

4.6 Setting Up Secure Exchange

Secure Exchange is typically used when a Web server does not provide Secure Sockets Layer (SSL) functionality but you still want to access Web pages securely over the Internet. If Secure Exchange is enabled, then all of the HTTP requests coming to the iChain Proxy Server are redirected to HTTPS, causing the data exchanged between the browser and the server to be encrypted using SSL.

You must enable Secure Exchange for Basic Authentication to function correctly (see Section 6.2, "Enabling Authentication Through the HTTP Authorization Header," on page 93).

You can set up Secure Exchange between the client browser and the iChain proxy and also between the iChain proxy and the origin Web server. If you want Secure Exchange between the proxy and the Web server, you must first have it enabled between the client and the proxy.

To set up Secure Exchange:

1. In the proxy server administration tool, click Configure, then click Web Server Accelerator.
2. Select the desired accelerator where Secure Exchange is being configured, then click Modify.
3. Check the Enable Secure Exchange check box.
4. Specify the SSL listening port to use for Secure Exchange traffic.
5. In the Certificate drop-down box, select the certificate to use for SSL.

   If you also want to set up Secure Exchange between the iChain proxy and the origin Web server, go to Step 6; otherwise, go to Step 9.
6. Click Secure Exchange Options.

   This page lets you enable secure access between the iChain Proxy and the origin Web server.

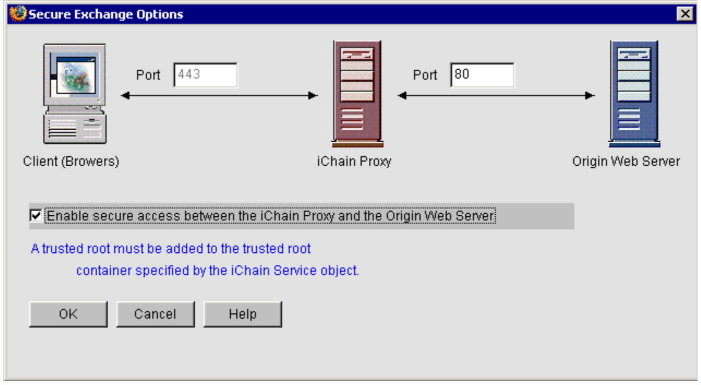

Port 443

Port 80

☑ Enable secure access between the iChain Proxy and the Origin Web Server

A trusted root must be added to the trusted root container specified by the iChain Service object.

7 Enter the SSL port in the port field between the iChain Proxy and the Origin Web server.

8 Check the Enable secure access between the iChain Proxy and the Origin Web Server check box.

With Secure Access enabled between the iChain proxy and the origin Web server, iChain needs to trust the Certificate Authority of the certificate used by the origin Web server. This means you must add a trusted root object into the trusted root container specified in the iChain Service Object configuration. For information about adding a trusted root container, see Section 4.6.1, "Creating a Trusted Root Container and Trusted Root Objects," on page 64.

9 Click OK, then click Apply.

4.6.1 Creating a Trusted Root Container and Trusted Root Objects

1 Export a base-64 trusted root file of the Certificate Authority used by the SSL certificate of the origin Web server.

2 From ConsoleOne, select the Security object located at the root of your LDAP tree.

3 Select File, select New, then select New Object

or

Click the New Object icon.

4 Select NDSPKI:Trusted Root, then click OK.

5 Specify a name for the trusted root container (for example, iChain Roots), then click OK.

6 Select the object you just created (for example, the iChain Roots object).

7 Select File, select New, then select New Object.

or

Click the New Object icon.

8 Select NDSPKI:Trusted Root Object, then click OK.

9 Define a name for the trusted root object (for example, Baltimore CA), then click OK.

10 Click the Read from File button, browse your system for the trusted root certificate, then import it into the dialog box

or

Paste your trusted root certificate into the dialog box.

To use this option, you must first open the trusted root certificate in a text editor or some other program and copy the contents to the clipboard. Then click inside the box and paste the certificate contents.

11 Click Finish.

If you want to add more trusted root objects, repeat Step 6 through Step 10 for each certificate.

After you create a trusted root container object, you need to configure the iChain Service Object (ISO) with the location of that container object. See Section 4.6.2, “Configuring iChain to Use the Trusted Root Container and Objects,” on page 65.

4.6.2 Configuring iChain to Use the Trusted Root Container and Objects

1 From ConsoleOne, click General on the iChain Security object (ISO).

2 Using the Browse button, browse to the trusted root container previously created in Section 4.6.1, “Creating a Trusted Root Container and Trusted Root Objects,” on page 64.

3 Click OK.

4.7 Multi-Homing

This section contains information about the following topics:

- “What Is Multi-Homing?” on page 65
- “Multi-Homing Web Server” on page 66
- “Host-Based Multi-Homing” on page 66
- “Domain-Based Multi-Homing” on page 67
- “Path-Based Multi-Homing” on page 67
- “Path-Based Multi-Homing Example” on page 68

4.7.1 What Is Multi-Homing?

Multi-homing is the ability to read from multiple origin Web servers over the same IP address and IP port.

Within a large company or corporation, IP addresses (for example, 10.1.1.1) are valued resources that are carefully managed. Opening additional ports can be time-consuming and political because of firewall and security issues.

For example, imagine using iChain to accelerate 100 back-end origin Web servers. Each of these Web servers has its own IP address. In order to grant access to these Web servers through iChain without multi-homing, a DNS server must map each Web server name to a separate IP address that iChain is listening on. A simple mapping would mean that 100 IP addresses would be needed by iChain to map to 100 backend Web servers.

With multi-homing, it is possible for iChain to provide access to all 100 back-end origin Web servers through one IP address and port. Most installations, however, will organize the origin Web servers into a handful of groups where each group is a group of multi-homing accelerators listening on one IP address.

There are four ways that iChain can multi-home origin Web servers on a single IP address and port. Before these features are explained, some terms must be defined:

Host name: In the URL http://www.novell.com/products/iChain, the host name is www.novell.com. The host name is used to name a Web server.

DNS name: A DNS name is a host name that has been mapped to an IP address using a DNS server or DNS mapping table. A browser may use the hosts file on the local drive as a DNS mapping table.

Cookie Domain: The cookie domain represents a group of host names that have in common two or more of the right-end pieces of the host name. For example, novell.com is a cookie domain for both www.novell.com and download.novell.com.

Origin Web Server: The origin Web server is a Web server that iChain accelerates through defining a Web Server Accelerator. A browser does not have direct access to an origin Web server and must obtain data from the Web server through iChain. The host name of an origin Web server might be different than the DNS name used by the browser. In fact, most iChain installations have a different DNS name than the host name of the origin Web server.

4.7.2 Multi-Homing Web Server

A multi-homing Web server is a Web server that listens on a single IP address and directs incoming requests to multiple logical Web servers running on the same machine. This is an easy way of having multiple alias names representing one host name. At the Web Server Accelerator page, you should select the Use Host Name Sent by Browser (Multi-homing Web Server) option. If a request is made on the accelerator IP address and port, the host name sent by the browser is not checked against the DNS name of the accelerator.

4.7.3 Host-Based Multi-Homing

Host-based multi-homing is not one of the standard multi-homing options for an accelerator. Host-based multi-homing is a way to listen for requests to a number of different host names on a single IP address and port. For example, www.a.com, www.b.com, and www.c.com can have the same DNS table entry of 100.1.1.1. There would be three accelerators defined, where each of the accelerators reads from three origin Web servers. See the table below:

DNS Name	Accelerator IP Address	Accelerator Proxy Port	Alternate Host Name	Web Server Address	Web Server Port
www.a.com	100.1.1.1	80	a.internal.com	127.12.12.1	80
www.b.com	100.1.1.1	80	b.internal.com	127.12.12.2	80
www.c.com	100.1.1.1	80	c.internal.com	127.12.12.3	80

If either authentication is enabled or Secure Exchange is enabled for two or more of the above accelerators, you must specify a unique SSL listening port. When authentication or Secure Exchange is enabled, an SSL listening port is issued to the accelerator. Two or more accelerators cannot listen on the same SSL port unless multi-homing is enabled. This is because a certificate must be used to establish an SSL connection. The certificate has the host name or a wildcard host name (*.novell.com). A single certificate cannot be served to secure both www.a.com and www.b.com. Usually each accelerator has its own unique SSL listening port for an IP address. See the table below.

DNS Name	Accelerator IP Address	Accelerator Proxy Port	SSL Listening Port
www.a.com	100.1.1.1	80	443
www.b.com	100.1.1.1	80	444
www.c.com	100.1.1.1	80	445

4.7.4 Domain-Based Multi-Homing

Domain-based multi-homing is selectable when Enable Multi-Homing is checked. To be more specific, this feature could be named Cookie-Domain-Based Multi-Homing because all of the DNS names for the accelerators in the multi-homing have the same cookie domain (see acme.com in the table below).

There must be a master accelerator already defined. A master accelerator is not part of a host-based multi-homing group and is not a child of a multi-homing group. The master defines the cookie domain, accelerator IP address, accelerator proxy port, SSL listening port, and certificate.

DNS Name	Accelerator IP Address	Type	Alternate Host Name	Web Server Address
a.acme.com	100.1.1.1	Master	a.internal.com	127.12.12.1
b.acme.com	100.1.1.1	Child	b.internal.com	127.12.12.2
c.acme.com	100.1.1.1	Child	c.internal.com	127.12.12.3

The same SSL port can be used for a group of domain-based accelerators if the SSL certificate supports wild card host names. In the table above, the SSL certificate would be *.acme.com.

NOTE: There are some browser versions that give a warning that the host name doesn't match the certificate name when a wildcard certificate is used. If the AUTO certificate is used, iChain creates a wildcard certificate using the cookie domain value.

4.7.5 Path-Based Multi-Homing

Path-based multi-homing is available when you select Enable Multi-Homing. Path-based multi-homing could be considered the most secure option because an SSL certificate containing a host name is used. The belief exists that wildcard certificates are less secure than a certificate with just a host name. All accelerators that participate in a path-based group have the same DNS name. All children must have a unique starting path.

NOTE: JavaScript* might obscure the URL data that iChain needs to access and modify in order to direct requests to backend servers. Because an absolute path reference can occur anywhere in JavaScript, the iChain word parse does not know how to assemble or parse through JavaScript.

We recommend that you do not use JavaScript applications with the iChain path-based multi-homing services. All other types of iChain accelerators should function correctly.

The same rules apply as they do for a domain-based multi-homing master: There must be a master accelerator already defined. A master accelerator is not part of a host-based multi-homing group and is not a child or a multi-homing group. The master defines the cookie domain, accelerator IP address, accelerator proxy port, SSL listening port, and certificate.

DNS Name	Multi-homing Path	Accelerator IP Address	Type	Alternate Host Name	Web Server Address
a.acme.com		100.1.1.1	Master	a.internal.com	127.12.12.1
a.acme.com	/Bstuff	100.1.1.1	Child	b.internal.com	127.12.12.2
a.acme.com	/Cstuff	100.1.1.1	Child	c.internal.com	127.12.12.3

There are two types of options for a path-based multi-homing child. The next two tables outline the URL requested from the browser and the URL that will be submitted to the origin Web server.

Sub-Path Match String = /Bstuff, Remove Sub-Path from URL

URL from browser	URL to Web server
http://a.acme.com/Bstuff/index.html	http://b.internal.com/index.html

This first example is the most common use of path-based multi-homing. A base path tells iChain that a specific accelerator is to be used to service the request. The base path is stripped when requesting the page from the Web server. Within the HTML pages that come from the Web server, all absolute and some relative references are rewritten by iChain so that the sub-path is present. For example, a reference in the HTML page might have href=/index.html. The rewriter rewrites this reference as href=/Bstuff/index.html so that the browser sets the correct GET request and iChain maps the request to the correct accelerator.

iChain can easily locate absolute path references because only the data coming from a path-based multi-homed accelerator is considered for this kind of rewriting. The word parser in iChain then locates all of the valid URL tags (href, src, action, etc.) and adds the subpath if the reference is an absolute-path reference. Absolute-path references that are outside of these tags are not considered for rewriting because there are too many variations of what the sub-path information could be.

Sub-Path Match String = /Bstuff

URL from browser	URL to web server
http://a.acme.com/Bstuff/index.html	http://b.internal.com/Bstuff/index.html

This second example expects the sub-path match string to be present as a valid path on the origin Web server. The rewriter doesn't need to update the absolute references as it did in the previous table.

4.7.6 Path-Based Multi-Homing Example

The information in the following example uses the Multi-homing Options dialog box.

Example

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

The administrators sets up two accelerators for www.zxy.org on the same IP address and port number, and configures them with path-based multi-homing 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 and have all requests starting with www.zxy.org/sales be directed to the sales Web server farm, and all requests starting with xxx.zxy.org/support be 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 using the check box Remove Sub-Path from URL option.

4.8 Custom Login Pages

iChain 2.1 and later provides the ability to create a custom login page per accelerator. For example, iChain could be fronting three different sites: novell.com, ctp.com, and silverstream.com. With the custom login page feature, each page could have its own unique login page.

To help you implement the custom login page feature, a brief explanation of iChain login, logout, and error pages is useful.

When you set up an accelerator, you have the option of enabling authentication. When this is done, you must specify an authentication profile. If the profile is an LDAP authentication profile, it has three options for login name format:

- User's e-mail
- Distinguished name
- Field name

Each login name format is presented to the user via a designated login page.

As an HTTP request comes in to be serviced by the accelerator, the proxy first checks the servicing accelerator to see if it has authentication enabled. If so, the proxy first presents a designated login page, contingent on the type of authentication profile specified for the servicing accelerator.

For example, assume there is an accelerator with an LDAP authentication profile where a distinguished name is the specified type of login name format. If a new HTTP request comes in to be serviced by this accelerator where no prior connection has been established, the proxy first presents the login page to the user:

```
sys:\etc\proxy\data\calogldp.htm
```

If the login name format was the user's e-mail, the login page that is presented is:

```
sys:\etc\proxy\data\caloglma.htm
```

In order to allow an administrator to specify a custom login page, in the iChain 2.3/iChain Proxy Server Web Administration tool, a field in the setup window exists to specify a subdirectory where the login page for that accelerator can be found. The actual file name continues to be predetermined by the profile and login name format.

For example, if the user specifies the subdirectory nike and specifies an LDAP authentication profile where the login name format is user's e-mail, the proxy attempts to find:

```
sys:\etc\proxy\data\nike\caloglma.htm
```

Currently the designated login pages are:

- Sys:\etc\proxy\data\calogldp.htm: Login page for LDAP profile based on a login page with a login name format of distinguished name or field name.
- Sys:etc\proxy\data\caloglma.htm: Login page for LDAP profile based on a login name format of user's e-mail.
- Sys:etc\proxy\data\calograd.htm: Login page for RADIUS authentication profile.

IMPORTANT: When you create or modify a custom login page, you must save the page in the UTF-8 format. You can also use <FORM.... ACCEPT-CHARSET=”UTF-8”> as the form statement on the login page. This is important because the decoder for the iChain login pages now assumes that the post data coming back from the browser is UTF-8. The HTTP/HTML specifications have no other mechanism in place to ensure that this is the case. This is necessary to support any usernames not falling into the 7-bit ASCII set.

For more information about this issue, see TID 10099466 (http://support.novell.com/cgi-bin/search/searchtid.cgi?10099466.htm).

To modify a login page specific to an accelerator:

1. Add the subdirectory to sys:\etc\proxy\data\.
2. Copy in the appropriate HTML and graphics files.
3. In the accelerator setup field, specify the name of the directory.
4. Apply the changes.

The designated error pages are the same as the ones above, but different text is placed in certain fields in these files to indicate an error has occurred. Since proxy.nlm currently hard-codes strings into designated HTML pages, it is best to allow for the specification of a unique error login page per accelerator.

4.8.1 Path or Host-Based Multi-Homing

The same mechanism that is described in Section 4.8, “Custom Login Pages,” on page 69 is provided for child accelerators that are part of either path-based or host-based multi-homing.

4.8.2 Limitations

Because the login pages are serviced from memory, this limits the types of graphics that can be supported. All streaming types of graphical/sound widgets (that is, AVI, WAV, MPEG, QuickTime*, etc.) are not supported; however, BMP, JPG, GIF, and other types of clip-art graphics are supported. Login pages and links referenced by the login pages should follow the 8.3 naming convention. Failure to do so results in broken links and possible authentication problems.

4.8.3 Coding of Login Pages

Default login pages are contingent on the type of authentication profiles that are being employed for the accelerator. Coding of specific login pages is most readily accomplished via modifying copies of these files. The following information describes the specific login pages:

- Sys:\etc\proxy\data\calogldp.htm: Used for accelerators that have an LDAP profile with login name formats of distinguished with LDAP contexts.
- Sys:etc\proxy\data\caloglnc.htm: Used for accelerators that have an LDAP profile with login name formats of distinguished without LDAP contexts (fully distinguished).
- Sys:etc\proxy\data\caloglfn.htm: Used for accelerators that have an LDAP profile with login name formats of field name.
- Sys:etc\proxy\data\caloglma.htm: Used for accelerators that have an LDAP profile with login name formats of e-mail.
- Sys:etc\proxy\data\caloglma.htm: Used for RADIUS-based authentication profiles.

Sys:\etc\proxy\data\calogldp.htm

The requirements for an LDAP profile login page for login name formats of Distinguished with LDAP contexts can be found in calogdp.htm:

- The attribute name *context* needs to be present.
- The attribute name *username* needs to be present and be of Type TEXT.
- The attribute name *password* needs to be present and be of Type PASSWORD.
- The attribute name *url* needs to be present and be of Type TEXT.

Sys:etc\proxy\data\caloglnc.htm

The same attributes as explained for calogldp.htm apply to accelerators using profiles with login name formats of Distinguished without LDAP contexts (fully distinguished). The difference is that there is no context and the username should contain a fully distinguished LDAP name for the value.

Sys:etc\proxy\data\caloglfn.htm

The same attributes as explained for calogldp.htm apply to accelerators using profiles with login name formats of Field Name. The difference is that there is no context and the username should contain the field name value for the value.

Sys:etc\proxy\data\caloglma.htm

The same attributes as explained for calogldp.htm apply to accelerators using profiles with login name formats of e-mail address. The difference is that there is no context and the username should contain an e-mail address for the value.

Sys:etc\proxy\data\calograd.htm

The same attributes as explained for calogldp.htm apply to accelerators using RADIUS profiles. The difference is that there is no context and the username should contain a RADIUS username for the value.

4.8.4 Custom Logout Page

If you want to set up a custom logout page, you need to provide a link in your respective HTML/XML page that reads as follows:

```
href="/cmd/BM-Logout"
```

or

```
href="/cmd/ICSLogout"
```

When this is completed, the following logout page is presented to the user:

```
sys:etc\proxy\data\calogout.htm
```

Custom logout pages are similar to custom login pages.

When the user specifies a subdirectory for login/logout pages, either through the GUI or at the iChain Proxy Server command line, as follows:

```
set accelerator accelerator name loginpage=subdir from etc proxy data
```

(for example: set accelerator nike loginpage=nike)

then the proxy looks for:

```
sys:etc\proxy\data\nike\calogout.htm
```

4.8.5 Using the USERVOL for Custom Login and Logout Pages

Because of the free space issues that can arise on the sys: volume, you can move your custom login/logout pages to the uservol volume if needed. (If you do not have issues with space, we recommend that you leave your custom login/logout pages on the sys: volume.)

To move the pages to the uservol volume:

1 Copy all of the files (including the subdirectories) from sys:\etc\proxy\data to uservol:\etc\proxy\data.

For example, if you are using toolbox's commands, you would do the following:

1a Load toolbox at the iChain server console.

1b Change the directory to the sys:\etc\proxy\data directory (where all of your custom pages are located).

1c Copy all of the customized data to the uservol volume by using the following syntax:

```
copy *.* uservol:\etc\proxy\data /s
```

1d Change the directory to the USERVOL and confirm that the files you copied now exist there. You can verify this by using the following command:

```
cd uservol:\etc\proxy\data
```

NOTE: If you do not want to copy the full list of customized pages, you can use the copy command to copy specific directories. The directory format must remain the same as that on the sys: volume. For example, the custom files in the \etc\proxy\data *custom* directory under USERVOL.

2 Administrators with existing custom directories on the sys: volume (under the etc\proxy\data directory) must rename these custom directories to other names. For example, if the original custom directory is called nike, you could rename it to nike.sav.

The directories must be renamed because the proxy.nlm reads the sys: volume first, and if it finds that a directory with custom pages exists, it will not read the custom pages from the similarly named directory in the uservol: volume.

3 Reboot the server. This action is required in order for the proxy to be able to read the custom login pages from the uservol: volume instead of sys: volume.

4.9 Cache Freshness

When first introduced to Web content caching, many network administrators assume that the object cache on an proxy server 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 iChain Proxy Services 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. The proxy server honors all flags that affect cache freshness, including Time to Expire, Don't Cache, and Must Revalidate directives.

In addition, the proxy server 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 objects for freshness that do not include Time to Expire headers.

This is administered at Configure > Tuning > Cache Freshness. For more information on configuring the iChain Proxy Services for cache freshness, see Section 4.9.1, “Managing Cache Freshness,” on page 73.

4.9.1 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 iChain Proxy Server 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 iChain Proxy Server 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 vends 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, the iChain Proxy Server retrieves the updated object into cache and serves it to the requesting browser before restarting the timer.

If the object has not changed, the iChain Proxy Server 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, the iChain Proxy Server honors the browser request, retrieves and caches the object regardless of whether it has changed, and restarts the timer.

How a Proxy Server Keeps the Oldest Cached Objects Fresh

More than 80% of all Web objects have either no Time to Expire directives or they 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, administered at Configure > Tuning > Cache Freshness.

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 iChain Proxy Services does not vend an object that has been in cache longer than six hours without first checking whether it should be refreshed.

HTTP Default: The iChain Proxy Server applies this timer to objects that don't have Time to Expire settings.

The default timer value is two hours. This means that the iChain Proxy Server does not vend an object that has no Time to Expire setting that has been in cache longer than two hours without first checking whether it should be refreshed.

How the iChain Proxy Server 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 proxy server 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 the proxy server serves HTTP data from cache before revalidating it against content on the origin Web server. No requested object is revalidated sooner than specified by this value.

The default value for this timer is 0, meaning that the proxy server 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 need the default settings to be 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 iChain Proxy Services 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 refreshing more often than is necessary. Otherwise you could negate the bandwidth and response-time benefits of having the appliance on your network.

4.10 GZIP Compression Support

iChain supports GZIP-compressed data between the browser and the proxy server. This additional support improves the overall performance of content delivery when the Web application has GZIP enabled.

The GZIP feature relies on the Web server to send GZIP-compressed data.

NOTE: Some Web servers, including Microsoft* Internet Information Services (IIS), do not respond with compressed (GZIP) data if a Via HTTP header is sent from a proxy to the Web server.

4.11 Understanding HTTP 1.1

iChain allows Web servers to return content according to HTTP 1.1 specifications, mainly chunked entity data and compressed entity data. When compressed entity data comes into iChain, at times it might need to be rewritten.

For example, HTML data such as href=http://privateside.com/index.html would be rewritten to https://securepublic.com/index.html. iChain rewrites the appropriate HTTP references, then recompresses the data back to the client.

4.11.1 Troubleshooting HTTP 1.0/1.1

If your Web servers are experiencing issues with having HTTP 1.1 requests sent to them, you can using the following troubleshooting command that enables an HTTP 1.1 request from a browser to be translated to an HTTP 1.0 request so that the Web server will respond correctly. The following is an example of how you would use this command:

```
SET ACCELERATOR name ForceHTTP10ToOrigin=Yes
```

Replalce *name* with the name of the accelerator for which you want to translate HTTP 1.1 requests to HTTP 1.0. Purge the cache afterwards, then HTTP 1.0 requests can be sent to the origin server. This action is permanent upon reboot and is exported to the .nas file.

4.11.2 Troubleshooting HTTP 1.1 and Recompression

HTTP 1.1 has the ability to deal with compressed data in either a Deflate or GZIP format. This reduces the size of data being sent across the wire. Because HTML pages are just text, they typically compress very well.

In previous versions, iChain accepted compressed content from the origin server, decompressed it for rewriting, and then sent the uncompressed page to the browser. iChain 2.3 has the ability to recompress the content before sending it to the browser. This is the default behavior.

To determine if the compression algorithm employed by iChain has a problem, you can turn it off. With recompression turned off, pages which had problems loading might now load properly. To turn off the compression algorithm, add the following command to the `load proxy` line of the `appstart.ncf` file and then restart the system:

```
load proxy -gzip 0
```

This causes iChain to accept compressed data from the origin server and to send uncompressed data to the browser.

To return the proxy to the default behavior, add the following command to the `load proxy` line of the `appstart.ncf` file and then restart the system:

```
load proxy -gzip 1
```

This command causes iChain to accept compressed data from the origin server and to send recompressed data to the browser.

Understanding Mutual Authentication

5

Mutual authentication is used when a user is issued a certificate from a trusted source. The certificate identifies the user in some way. To ensure the validity of X.509 certificates, Novell® iChain® supports both Certificate Revocation Lists (CRLs) and Online Certificate Status Protocol (OCSP) methods of verification.

This section describes the following:

- Setting Up Mutual SSL
- Disabling Mutual SSL
- Using Third-Party Certificates
- Using Multiple Certificate Authorities
- Configuring the Online Certificate Status Protocol

5.1 Setting Up Mutual SSL

SSL provides:

- Authentication and nonrepudiation of the server, using digital signatures
- Data confidentiality through the use of encryption
- Data integrity through the use of authentication codes

Mutual SSL provides the same things as SSL, with the addition of authentication and nonrepudiation of the client, using digital signatures.

IMPORTANT: There is a security guideline that you should be aware of when setting up iChain in a production environment. For more information about this guideline, see TID 10096315.

5.1.1 Mutual SSL Configuration

There are many different certificate authority vendors and varying methods to configure Mutual SSL. Although it is not possible to cover all the possibilities, the following is an example using the Novell Certificate Authority:

iChain Server Certificate Setup: Certificate Signing Request

1. In the proxy server administration tool, click Home, click Certificate Maintenance, then click Create.
2. Specify an appropriate name for the certificate and subject name.
3. Click the Signature Algorithm drop-down list, then select the algorithm you want to use (SHA-1 or MD-5).
4. Click the RSA Key Size drop-down list, then 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.

5 Click Use External Certificate Authority.

6 If desired, specify a name for your organizational unit or division.

This is commonly referred to as the Organizational Unit and is used to differentiate between organizational divisions or to describe departments or divisions.

7 If desired, specify a name for your organization.

8 Specify the city or town where your organization does business.

9 Specify the non-abbreviated name of the state or province where the organization does business.

This is commonly referred to as the state.

10 Specify the International Standards Organization 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 country code.

11 Click OK.

Examine 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.

12 Click Apply.

If any errors occur during the certificate request process, they are displayed in the Error field on a red background.

If an error occurs:

12a Click Modify.

12b In the Modify Certificate dialog box, make the changes necessary to resolve the errors, then click OK.

12c Click Apply.

Repeat the modification process until the Status field displays the words CSR in Progress on a yellow background.

Extracting the CSR from the iChain Proxy Server to Send the CSR

1 Click View CSR to open a new browser window that displays the CSR contents.

2 Select and copy the complete CSR text into your computer's Clipboard.

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 varies 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 is similar to the following:

```
----- BEGIN NEW CERTIFICATE REQUEST -----
```

The trailer line is 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.

Using Novell as the External Certificate Authority

1. In ConsoleOne, click Tools > Issue Certificate, then paste the CSR and follow the prompts to sign the certificate.
2. Click Finish and save the file in a .b64 format.
3. In ConsoleOne, go to the Organization CA object's properties page (in the Security container). Go to the Certificates > Self Signed Certificate page, then export the Self Signed Certificate in .b64 format.

Storing the Certificate in the iChain Proxy Server

After the external CA responds with the certificate:

1. In the proxy server administration tool, click Home, click Certificate Maintenance, then click the name of the certificate you want to store, then click Store Certificate.
2. In the Store Certificates dialog box, paste the CA certificate into the CA Certificate Contents box. If you are using Novell CA, this is where the Self Signed Certificate should be placed.

   If the CA Certificate Contents and the Server Certificate Contents are in the same Base-64 encoded file, select the No Trusted Root Certificate Available check box. This will dim the CA Certificate Contents box and allow the single Base-64 encoded file containing the entire certificate chain to be pasted into the Server Certificate Contents box.
3. Paste your newly issued certificate in the Server Certificate Contents box.
4. Click Create.

   Examine 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.
5. Click Apply.

   If no error occurs during the certificate creation process, the status changes to Active.

   If an error occurs during the certificate creation process, it is displayed in the Error field on a red background.

   If an error occurs:

   5a Click Store Certificate.

   5b In the Store Certificate dialog box, verify that the correct certificates are pasted in the boxes, then click OK.

   5c Click Apply. Repeat this process until the Status field displays the words Active on a green background.

Create a New Authentication Profile Via the Proxy Server Administration Tool

1. In the proxy server administration tool, click Configure, click Authentication, then click Insert, then specify a name for the profile (for example, “mutual”).
2. Select SSL Certificate Mutual Authentication, click OK, then click Apply.

   NOTE: Before you click OK, you can select Mutual Options to configure certificate mapping, if it is required. See “Using Certificate Mapping” on page 80.
3. Add the new profile to the Web Server Accelerator by clicking Configure, then click Web Server Accelerator.

4 Highlight the appropriate accelerator, then click Modify.

5 Click enable Authentication, click Authentication Options, then select the newly created profile.

6 Click Add, then click OK.

7 On the Accelerator page, in the Certificate drop-down selection box, select the name of the certificate you created, click OK, then click Apply.

iChain Server: Secure Exchange Between the Browser and iChain

NOTE: When using Mutual Authentication, it is recommended that Secure Exchange also be enabled between the browser and iChain as follows:

1 In the proxy server administration tool, click Configure, then click Web Server Accelerator.

2 Select the appropriate Web accelerator, click Modify, then click Enable Secure Exchange.

Leave the SSL Listening Port as the default (443).

NOTE: Secure Exchange can also be set up between iChain and the Web server. See Section 4.6, “Setting Up Secure Exchange,” on page 63.

Creating a User Certificate From a Novell Certificate Authority

1 In ConsoleOne as an administrator, go to the user object's properties on the Security page, click Certificates, then click Create.

2 Select the default options (for example, with the private key). Change only what you need to change (for example, the expiration).

Do not change the subject name if it is shown in reverse (for example, o=novell,ou=stress,cn=user1020).

3 Save the file in .pfx format with a password.

You must be logged in as the user (not as an admin) to save this file in .pfx format.

Importing the .pfx Certificate to the Browser

The following instructions assume you are using Internet Explorer as your browser.

1 In the browser, click Tools, then click Internet Options.

2 Click Content, then click Certificates.

3 Import a Personal Certificate that has been signed by the Certificate Authority.

Follow the prompts to import the certificate. You will be prompted to enter a password.

5.1.2 Using Certificate Mapping

When using SSL Mutual Authentication, there must be a user in the iChain LDAP Authentication tree that corresponds with the user certificate. Certificate Mapping gives four different ways to map the user certificate to a user in the iChain LDAP Authentication tree. The four mapping types are Directory Name Mapping, Email Mapping, Subject Name Mapping, and Serial Number & Issuer

Name Mapping. The proxy server can be configured to use any combination of the four mapping types. When searching for a user with the configured mappings, the first user found is the user that is used for authentication and access control, even if the other users maps to the same certificate.

Configuring a Certificate Mapping Search Base

The search base is the location in the iChain LDAP Authentication tree to search for user objects that the certificate can map to. More than one search base can be configured. The search looks for matches starting at the search base. All containers below the search base are included in the search.

At least one search base needs to be configured for Certificate Mapping. The certificate that authenticates the user must be mapped to a user in the LDAP directory. To do this, create an LDAP authentication profile named ldapcert. Configure the LDAP profile to search on a single attribute, and insert one or more search bases.

Configuring Certificate Mapping Types

Certificate Mapping gives four different ways to map the user certificate to a user in the iChain LDAP Authentication tree. The four mapping types are Directory Name Mapping, Email Mapping, Subject Name Mapping, and Serial Number & Issuer Name Mapping. The proxy server can be configured to use any combination of the four mapping types. The certificate mapping types are configured from the iChain Proxy Server utility.

1 In the proxy server administration tool, click Configure, then click Authentication.

2 Select an authentication profile of type Mutual.

3 Click Modify, then click Mutual Options.

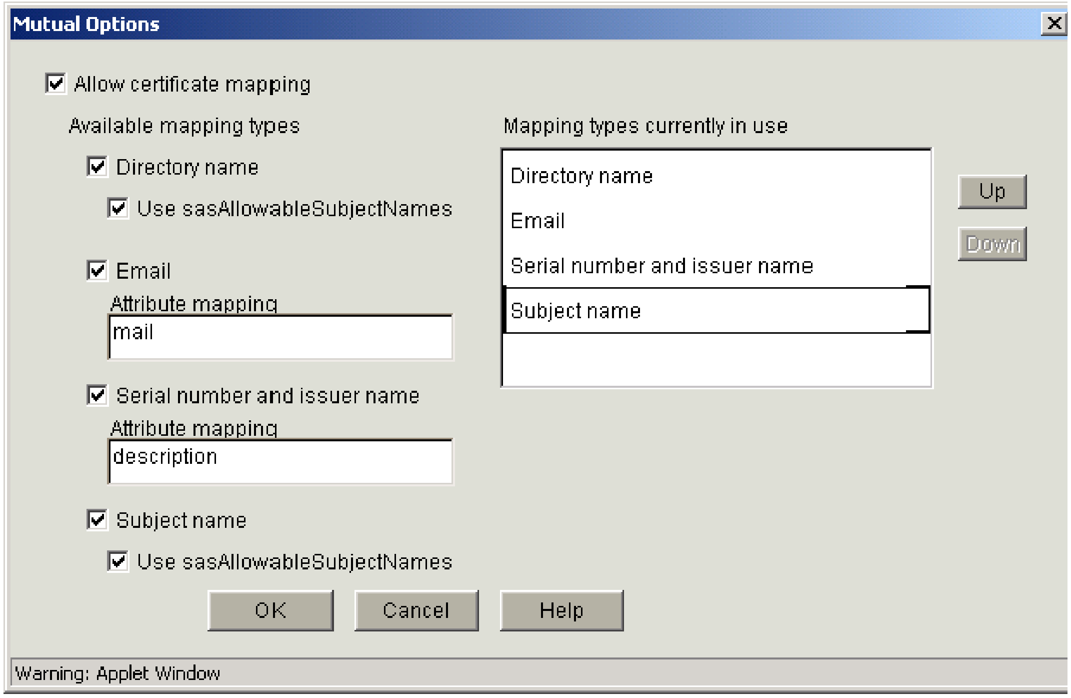

Mutual Options

☑ Allow certificate mapping

Available mapping types

☑ Directory name

☑ Use sasAllowableSubjectNames

☑ Email

Attribute mapping: mail

☑ Serial number and issuer name

Attribute mapping: description

☑ Subject name

☑ Use sasAllowableSubjectNames

Mapping types currently in use

Directory name
Email
Serial number and issuer name
Subject name

Warning: Applet Window

Directory Name Mapping

With directory name mapping, the Subject Alternative Name field in the user certificate, with a name type of Directory Name, is used to identify the certificate portion of the user (see Figure 5-1). The name in the certificate can be from root to leaf or from leaf to root.

***Figure 5-1** Subject Alternative Name*

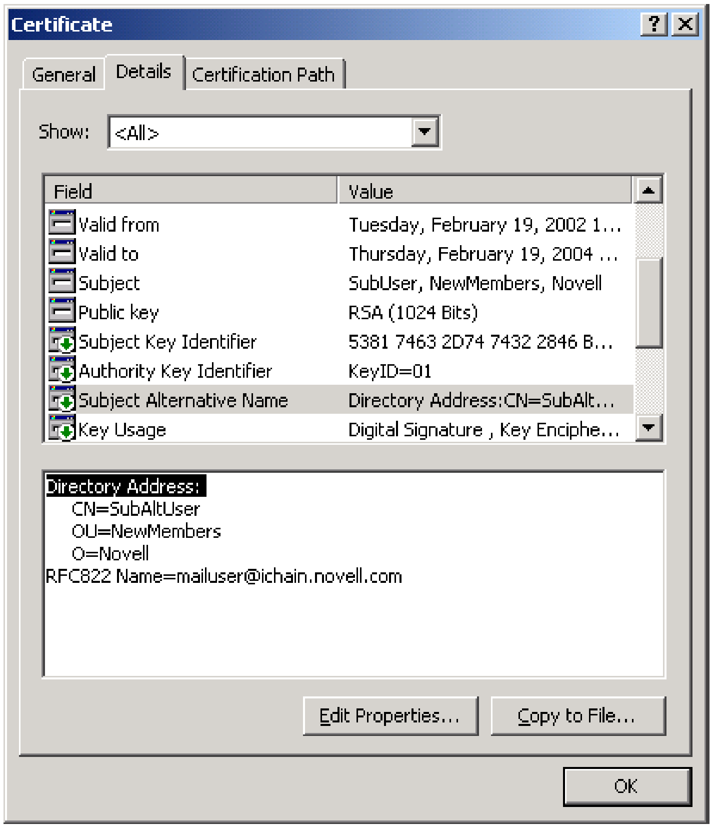

Show: <All>

Field	Value
Valid from	Tuesday, February 19, 2002 1...
Valid to	Thursday, February 19, 2004 ...
Subject	SubUser, NewMembers, Novell
Public key	RSA (1024 Bits)
Subject Key Identifier	5381 7463 2D74 7432 2846 B...
Authority Key Identifier	KeyID=01
Subject Alternative Name	Directory Address:CN=SubAlt...
Key Usage	Digital Signature , Key Enciphe...

```
Directory Address:
     CN=SubAltUser
     OU=NewMembers
     O=Novell
RFC822 Name=mailuser@ichain.novell.com
```

A user in the LDAP Authentication Tree matching the Directory Name in the Subject Alternative Name field of the certificate is checked first. If a user is not found and Use sasAllowableSubjectName is also enabled for directory mapping, the LDAP Authentication Tree is searched for a user containing an sasAllowableSubjectName attribute matching the Directory Name in the Subject Alternative Name field of the certificate. If sasAllowableSubjectName is enabled, the LDAP Authentication tree should be configured so that there is no duplication of allowed names between users in the sasAllowableSubjectName attribute.

The sasAllowableSubjectName attribute is the same attribute currently used by NMAS for certificate mapping. The ConsoleOne snap-ins and schema updates are part of the NMAS installation on the Authorization Server CD. Figure 5-2 shows the sasAllowableSubjectName attribute in ConsoleOne.

Figure 5-2 *sasAllowableSubjectName Attribute*

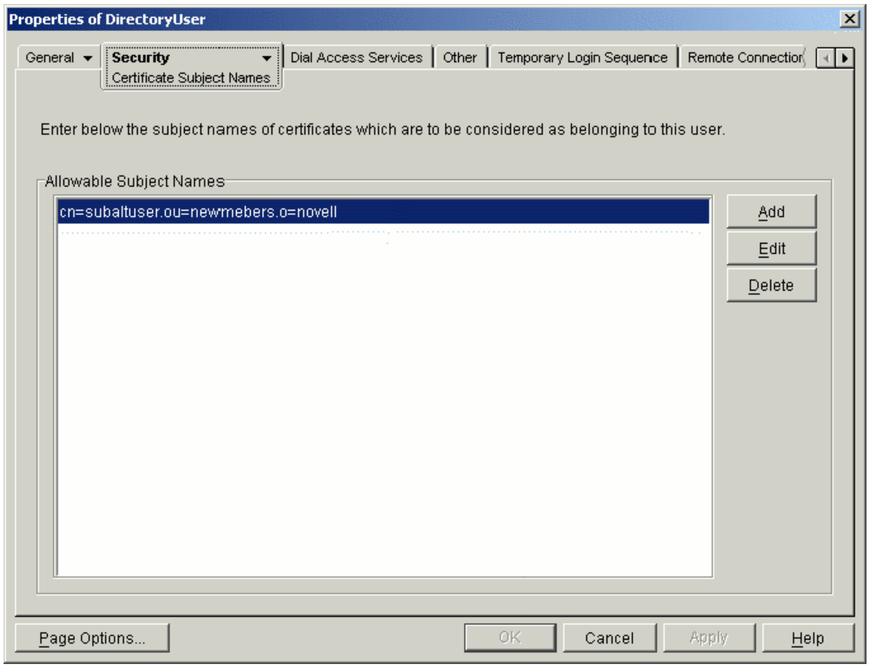

Enter below the subject names of certificates which are to be considered as belonging to this user.

cn=subaltuser.ou=newmebers.o=novell

Email Mapping

With Email mapping, there are two possible fields in the user certificate that can be used to identify the certificate portion of the user. The first is the Subject Alternative Name field in the user certificate, with a name type of RFC822 (see Figure 5-1). The second is when an e-mail name is embedded in the Subject field of the certificate (see Figure 5-3). If both the Subject field and the Subject Alternative Name field contain an e-mail address, the Subject Alternative Name is the only field used.

Figure 5-3 *Email Name Embedded in Certificate Subject Field*

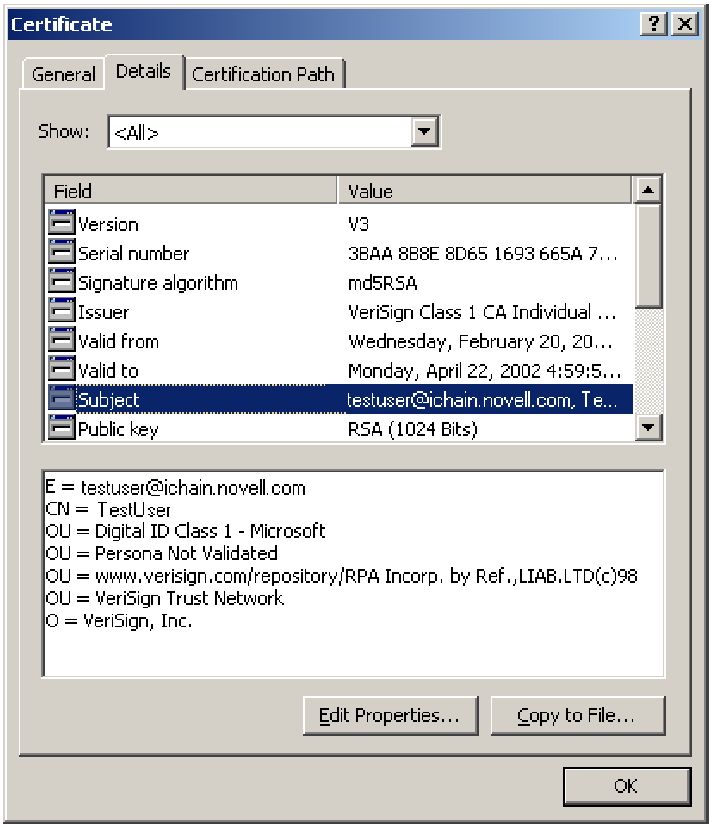

Show: <All>

Field	Value
Version	V3
Serial number	3BAA 8B8E 8D65 1693 665A 7...
Signature algorithm	md5RSA
Issuer	VeriSign Class 1 CA Individual ...
Valid from	Wednesday, February 20, 20...
Valid to	Monday, April 22, 2002 4:59:5...
Subject	testuser@ichain.novell.com, Te...
Public key	RSA (1024 Bits)

E = testuser@ichain.novell.com
CN = TestUser
OU = Digital ID Class 1 - Microsoft
OU = Persona Not Validated
OU = www.verisign.com/repository/RPA Incorp. by Ref.,LIAB.LTD(c)98
OU = VeriSign Trust Network
O = VeriSign, Inc.

The LDAP attribute configured in the Email attribute mapping is used to match the Email address from the certificate when searching for a user in the LDAP Authentication tree. The default LDAP attribute is mail, which is the attribute currently used by GroupWise and Novell Certificate Server. The LDAP Authentication tree should be configured so that there is no duplication of Email addresses between users in the configured e-mail attribute mapping.

Subject Name Mapping

With directory name mapping, the Subject field in the user certificate is used to identify the certificate portion of the user (see Figure 5-4). The Subject name in the certificate can be from root to leaf or from leaf to root.

***Figure 5-4** Subject Field in the User Certificate*

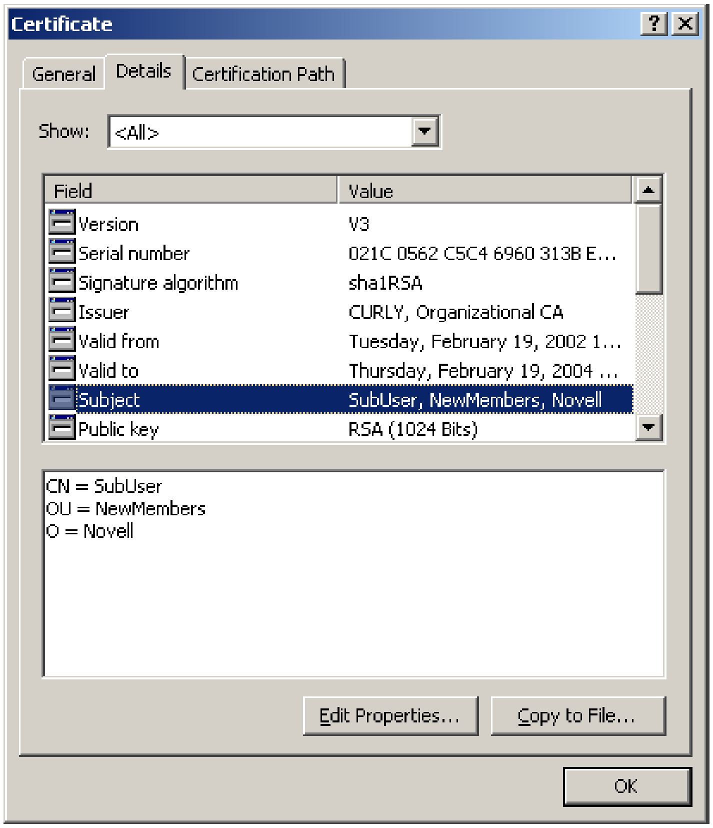

Show: <All>

Field	Value
Version	V3
Serial number	021C 0562 C5C4 6960 313B E...
Signature algorithm	sha1RSA
Issuer	CURLY, Organizational CA
Valid from	Tuesday, February 19, 2002 1...
Valid to	Thursday, February 19, 2004 ...
Subject	SubUser, NewMembers, Novell
Public key	RSA (1024 Bits)

CN = SubUser
OU = NewMembers
O = Novell

A user in the LDAP Authentication tree matching the Subject Name field of the certificate is checked first. If a user is not found and the Use sasAllowableSubjectName is also enabled for directory name mapping, the LDAP Authentication tree is searched for a user containing a sasAllowableSubjectName attribute matching the Subject Name field of the certificate. If the sasAllowableSubjectName attribute is enabled, the LDAP Authentication tree should be configured so that there is no duplication of allowed names between users in the sasAllowableSubjectName attribute.

The sasAllowableSubjectName is the same attribute currently used by NMAS for certificate mapping. The ConsoleOne snap-ins and schema updates are part of the NMAS installation on the Authorization Server CD. Figure 5-2 shows the sasAllowableSubjectName in ConsoleOne.

Serial Number & Issuer Name Mapping

With Serial Number & Issuer Name mapping, both the serial number and the issuer name fields from the certificate is used together to identify the certificate portion of the user (see Figure 5-5).

Figure 5-5 *Serial Number & Issue Name Mapping*

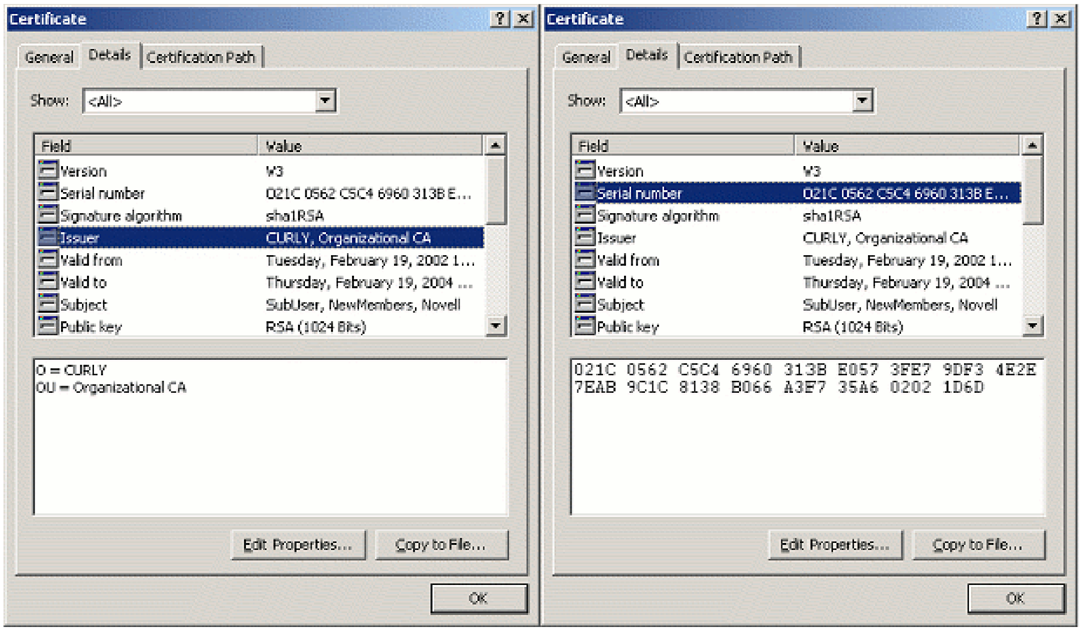

Show: <All>

Field	Value
Version	V3
Serial number	021C 0562 C5C4 6960 313B E...
Signature algorithm	sha1RSA
Issuer	CURLY, Organizational CA
Valid from	Tuesday, February 19, 2002 1...
Valid to	Thursday, February 19, 2004 ...
Subject	SubUser, NewMembers, Novell
Public key	RSA (1024 Bits)

```
O = CURLY
OU = Organizational CA
```

```
021C 0562 C5C4 6960 313B E057 3FE7 9DF3 4E2E
7EAB 9C1C 8138 B066 A3F7 35A6 0202 1D6D
```

Both the issuer name and the serial number need to be put into the same LDAP attribute of the user. The LDAP attribute that is used is specified in the Serial number and issuer name Attribute mapping field of the iChain Proxy Server utility. The LDAP attribute can be any Case Ignore List or Case Ignore String attribute of the user. If you are configuring your own attribute, make sure the attribute is added to the Person class.

When using a Case Ignore List attribute, both the issuer name and the serial number need to be in the same list. The issuer name needs to be the first item in the list, with the serial number being the second and last item in the list.

When using a Case Ignore String attribute, both the issuer name and the serial number need to be in the same attribute separated by a dollar sign ($) character. The issuer name needs to be in front of the $ character, with the serial number following the $ character. Do not use any spaces in front of or behind the $ character. (For example, O=CURLY.OU=Organization CA$021C0562C5C4... could be used for the certificate displayed in Figure 5-5).

The issuer name can be from root to leaf or from leaf to root. The issuer name is dot-delimited without a preceding dot. (For example, O=CURLY.OU=Organization CA or OU=Organization CA.O=CURLY could be used for the certificate displayed in Figure 5-5.

NOTE: The certificate number is displayed in Internet Explorer with a space after every fourth digit. The certificate number needs to be entered without spaces. For example, the certificate number displayed in Figure 5-5 is shown with spaces, but should be entered as: 021C0562C5C46960313BE0573FE79DF34E2E7EAB9C1C8138B066A3F735A602021D6D.

5.2 Disabling Mutual SSL

After you have enabled mutual SSL on an accelerator, complete the following steps to disable it:

1 In the proxy server administration tool, click Configure, then click Web Server Accelerator.

2 Select the accelerator, then click Modify.

3 Click Authentication Options.

4 Select the profile that is performing mutual authentication, then click Delete.

5 Click Apply.

6 Click System, click Actions, then click Purge Cache.

7 Select Purge all of cache, then click Start Purge.

Users are denied access to the site until cache is purged because the users are still prompted for the certificate.

5.3 Using Third-Party Certificates

Novell iChain includes Novell Public Key Infrastructure Services (PKIS 2.0) to provide cryptography and enable certificate services in your iChain infrastructure. A Novell server certificate is installed and configured automatically when you install Novell iChain; however, you might want to use other third-party certificates, such as VeriSign* certificates, in your infrastructure. In order to use third-party certificates in your iChain infrastructure, you must request a certificate from a Certificate Authority (CA), have the CA sign the certificate, collect and export the certificate and its trusted root, and then import the certificate and its trusted root to the iChain Proxy Server. For more information, see Chapter 15, "Using iChain to Manage Certificates," on page 193.

5.4 Using Multiple Certificate Authorities

The Multi CA feature enhances authentication to support alternate Certificate Authorities (CAs) during mutual SSL authentication. The Multi CA feature allows the iChain proxy to accept user certificates that are signed by a different CA than the CA that signed the iChain server certificate.

For example, if your iChain server certificate is signed by a VeriSign CA, then using the Multi CA feature could allow users with certificates signed by a Baltimore CA or an Entrust* CA to access your system (the Baltimore or Entrust certificates would need to be installed into your LDAP server tree).

5.4.1 Configuring Multi CAs

To configure Multi CAs, you need to place the alternate CA certificates into your LDAP tree, then configure the iChain proxy to use a specified trusted root container, as described below in "Placing Alternate CA Certificates Into Your LDAP Tree" on page 87, and "Configuring the iChain Proxy Server to Use a Specified Trusted Root" on page 88.

Placing Alternate CA Certificates Into Your LDAP Tree

1 From ConsoleOne, select the Security object located at the root of your LDAP tree.

2 Select File > New > New Object

or

Click the New Object icon.

3 Select NDSPKI:Trusted Root, then click OK.

4 Define a name for the trusted root container (for example, iChain Roots), then click OK.

5 Select the object you just created (for example, the iChain Roots object).

6 Select File > New > New Object

or

Click the New Object icon.

7 Select NDSPKI:Trusted Root Object, then click OK.

8 Define a name for the trusted root object (for example, Baltimore CA), then click OK.

9 Click the Read from File button, browse your system for the trusted root certificate, then import it into the dialog box

or

Paste your trusted root certificate into the dialog box.

To use this option, you must first open the trusted root certificate in a text editor or some other program and copy the contents to the clipboard. Then paste the contents in the box.

10 Click Finish.

If you want to add more trusted root certificates, repeat Step 5 through Step 10 for each certificate.

Configuring the iChain Proxy Server to Use a Specified Trusted Root

1 From ConsoleOne, click the Trusted Root Container tab on the iChain Security object (ISO) you previously created for this configuration.

2 Using the Browse button, browse to the trusted root container previously created (see "Placing Alternate CA Certificates Into Your LDAP Tree" on page 87), then click OK.

or

Specify the complete name of the previously created trusted root container (for example, iChain Roots.Security).

3 Click OK.

Define the Location of the Trusted Root Container for all Trusted Roots

This option lets you establish the location of the Trusted Root container where all the trusted roots are stored. You normally configure this setting at the ISO object level, but this set command provides the same configuration option from the iChain CLI. If you get an error when setting up mutual authentication or client authentication, you can use this option to resolve the issue.

1 Open the Command Line interface.

2 Set the following parameter to the correct context:

```
set authentication mutual mutual trustedrootfile = <trusted_root_container>
```

3 Apply the option.

5.5 Configuring the Online Certificate Status Protocol

Online Certificate Status Protocol (OCSP) is a protocol that can be used to determine the revocation status or validity of certificates. This protocol allows a service to provide revocation information rather than having Novell iChain determine it. Having a service provide the revocation information can also provide more up-to-date information about the revocation status of certificates.

As with Certificate Revocation Lists (CRLs), an extension in the certificate specifies the Uniform Resource Identifier (URI) to be used. For OCSP, the extension is the Authority Information Access (AIA) extension, which specifies the location of the OCSP service.

In addition, the OCSP standard allows for a local configuration option that is supported by iChain. This option allows the administrator to configure the location of the OCSP service. When the local configuration option is used, it overrides any information specified in the certificate. Also, when the local configuration option is used, it applies to all certificates, whether or not they contain the AIA extension.

During the certificate validation processes, iChain uses OCSP (if it is configured or specified) followed by CRLs (if specified) to determine the revocation status of the certificates. During this process, if OCSP is able to determine the revocation status of a certificate, then CRLs are not used. Thus, OCSP enables iChain to bypass the CRL downloads, which can save processing time when the CRLs have expired and need to be downloaded.

NOTE: Because iChain caches CRLs, any time saving would only apply where valid CRLs were not previously cached. In general, OCSP does not provide better processing time than CRLs. Rather, cached CRLs provide better performance than OCSP.

In the OCSP process, there is an OCSP client (iChain) and an OCSP responder (a third-party OCSP server). The client (iChain) makes a request to the OCSP responder and then waits for a response. The OCSP server should respond with the status of the certificate in question as valid, revoked, or unknown. It can also respond that the request is invalid for one of multiple reasons. If the server responds that the request is invalid, no indication is given as to the status of the certificate in question.

In iChain, you configure OCSP at the command line. The following are samples of the commands that relate directly and indirectly to OCSP and certificate validity, along with an explanation of each. This configuration is done through the SSL profile. In each of these examples, *ssl* refers to the name of your SSL mutual authentication profile.

```
authentication ssl mutual useocspconfiguredsource = No
```

Specifies whether to enable the local configuration option. The default is No. Change it to Yes if you want to configure iChain to always use a specific OCSP responder (for example, a locally determined OCSP server) for revocation checking.

NOTE: In order to use the local configuration option, you must also use the command `authentication ssl mutual url` to specify which OCSP server to use.

```
authentication ssl mutual url = (currently no values assigned)
```

Specifies the locally configured OCSP responder location. This value overrides any information stored in the certificate.

NOTE: In order to use the local configuration option, you must also use the command `authentication ssl mutual useocspconfiguredsource = Yes` to turn on the local configuration option.

Set this value to the URL of the OCSP server. The http protocol must be specified as part of the URL (http:// for non-secure or non-SSL, or https:// for secure or SSL). For example, http://ocsp.openvalidation.org.

```
authentication ssl mutual signedrequestcert = (currently no values assigned)
```

Specifies the server certificate to use to sign the OCSP request. You specify the certificate name exactly as it as viewed in the Certificate Maintenance page in the proxy server interface after you have created the certificate. See Chapter 15, “Using iChain to Manage Certificates,” on page 193 for more information. Request signing is an optional OCSP feature, and should only be used if required by the OCSP responder. Normally this command is not used. If it is used, the server certificate needs to be one that is accepted by the OCSP responder. (Contact the administrator of the OCSP service for details and requirement specifications.)

NOTE: To create this server certificate, follow the normal iChain procedure for creating a server certificate. See Chapter 15, “Using iChain to Manage Certificates,” on page 193 for more information.

The following settings are not directly specific to OCSP, but can influence it indirectly.

```
authentication ssl mutual verifyrootca = Yes
```

Specifies whether to provide revocation checking on the root Certificate Authority (CA). The default is Yes. If you change this setting to No, iChain does not check the revocation status of the root certificate (even if the certificate specifies a method for doing so).

WARNING: Changing this setting to No can reduce the security of your system.

```
authentication ssl mutual disablerevocationchecks = No
```

Specifies whether to disable revocation checking on all certificates. The default and recommended setting is No.

WARNING: Do not disable revocation checking in a production environment. Changing this setting to Yes can reduce the security of your system.

```
authentication ssl mutual mapx500crltoldap = (currently no values assigned)
```

Specifies the IP or DNS address of the LDAP server to use when mapping X.500 CRL distribution points using the LDAP protocol. In other words, the X.500 name is mapped to an LDAP name and then the LDAP protocol is used to download the CRL. iChain only connects over the default 389 port (you cannot change the port). You cannot specify the username and/or password.

NOTE: This option is available because X.500 does not support the concept of a multiple tree environment, but LDAP does.

```
set authentication mutual_auth_profile_name mutual revocationcheckmethod = OCSP.
```

Certificate revocation is checked in the URL of the certificate. The certificate revocation method is set using the following setting:

revocationcheckmethod=*method*

The *method* is the type of certificate revocation checking performed during mutual authentication. The following Certificate Revocation checks are available depending on the value of the revocationcheckmethod parameter:

- OCSP only. This method checks only the online certificate status protocol (OCSP). The CRL is checked in the URL and overrides the client certificate. The certificate does not pass if it is revoked or there is a miscommunication.
- CRL only. This checks only the Certificate Revocation list.
- OCSP- CRL. This method checks the OCSP first, then checks the CRL. The CRL protocols include HTTP, LDAP, and X.500 directory name (provides only the directory name).

The Certificate Revocation List (CRL) is checked when the following conditions exist:

- The client certificate contains a CRL distribution point (CDP).
- The client certificate contains a CRL CDP but not an Authority Information Access (AIA) extension for OSCP.
- The OCSP configured sources is disabled.

NOTE: If you have your certificate revocation method set to OCSP-CRL, the certificate is allowed to authenticate the user if an unknown response occurs and the client certificate does not contain a CRL distribution point. If you do not want the certificate to authenticate the user under these circumstances, you need to set your revocation method to OCSP only.

5.5.1 New OCSP Setting in iChain 2.3 Support Pack 1

In iChain 2.3 SP1 Support Pack 1, a new OCSP setting is introduced:

```
set authentication <ssl auth profile name> mutual
ocspconfiguredcerts = <trusted_root_container>
```

```
set authentication ssl auth profile name mutual ocspconfiguredcerts = trusted root
container
```

This setting is needed only when the OCSP server's signing certificate's CA is not the same CA as the certificate that iChain tries to validate through the OCSP protocol. Note that if the OCSP server's signing certificate's CA is the same CA as the certificate that iChain tries to validate through the OCSP protocol, this setting is not needed.

For this setting, the *trusted_root_container* is a fully qualified name of a trusted root container object in the LDAP authorization server tree. It must contain the trusted root of the OCSP server's signing certificate.

You create this object and the trusted root object the same way you create the trusted root container configured in the ISO object. It could be the same one that is in the ISO object, or it could be another trusted root container.

This new setting can only be set from the command line. The fully qualified name of the trusted root container needs to be semicolon (;)-delimited, and does not include the tree name.

For example, if you enter the following command:

```
get authentication ssl auth profile name
```

you see the following on your server for this setting:

```
authentication ssl mutual ocspconfiguredcerts = cn=OCSPTRContainer;cn=Security
```

Understanding LDAP Authentication

6

The previous sections in this document described how to install and set up the basic implementation described in Section 2.2, "Installation Scenario," on page 23. To meet your company's networking needs, you might need to augment or alter this implementation and use some of the more advanced features of Novell® iChain® services. This section describes the following Novell iChain configuration procedures:

- Section 6.1, "Allowing Authentication Through the HTTP Authorization Header," on page 93
- Section 6.2, "Enabling Authentication Through the HTTP Authorization Header," on page 93

6.1 Allowing Authentication Through the HTTP Authorization Header

The Allow authentication through HTTP authorization header check box on the LDAP Authentication options screen allows Basic (401) authentication as either an alternative or a substitute for the iChain login form/page.

This feature allows iChain to process a request, log in the user (if necessary), and return the response without having a programmer deal with login redirects or parsing login pages and forms. The iChain cookie is returned in the response for possible use in subsequent requests. If authorization headers are optional, a user who is not authenticated is redirected to the standard iChain login page. If the headers are mandatory, a 401 status is returned. The browser then requests the user's credentials, and the request is resubmitted along with the user's credentials. In this mode, the CDA features are disabled.

NOTE: We do not recommend Basic Authentication for use with users/browsers because of security issues relating to lack of control of the credentials on the wire. The primary use is anticipated to be programming-related, where the credentials can be passed in an authorization header along with a request. That way, a programmer retains control over the exposure of the credentials.

6.2 Enabling Authentication Through the HTTP Authorization Header

1. In the proxy server administration tool, click Configure, then click Authentication. Select the LDAP authentication profile on which you want to enable Basic Authentication.
2. Click Modify, then click LDAP Options.
3. Select the Allow Authentication through HTTP Authorization Header check box.
4. Select User iChain Login Page (Basic Authorization Headers are Optional),

   or

   Use Basic/Proxy Authentication (Basic Authorization Headers Are Mandatory).
5. Click OK twice, then click Apply.

IMPORTANT: If you enable Basic Authentication for LDAP, secure exchange must be enabled (see Section 4.6, “Setting Up Secure Exchange,” on page 63).

Using RADIUS Authentication

7

The Novell® Remote Authentication Dial-In User Service (RADIUS) is a Novell authentication service that enables remote users to securely dial in to NetWare® networks (version 5.1 or later) and access network information and resources. When installed on a server, RADIUS can be configured as an integral part of the Novell Modular Authentication Services™. This section describes the following topics:

- Section 7.1, "Using Token Authentication with iChain," on page 95
- Section 7.2, "Installing NMAS, Novell RADIUS, and a Token Method," on page 95
- Section 7.3, "Configuring Novell RADIUS Components," on page 96
- Section 7.4, "Setting Up the iChain RADIUS Client," on page 98

7.1 Using Token Authentication with iChain

You can configure Novell iChain® to leverage the authentication service provided by Novell Modular Authentication Services (NMAS™) Enterprise Edition. You can set up your iChain users so they are required to authenticate to eDirectory™ using a token device. This adds a higher level of protection to your information by ensuring that only those who have the proper token code can have access to your information.

There are three steps you must perform to set up token authentication with iChain:

1. Section 7.2, "Installing NMAS, Novell RADIUS, and a Token Method," on page 95.
2. Section 7.3, "Configuring Novell RADIUS Components," on page 96.
3. Section 7.4, "Setting Up the iChain RADIUS Client," on page 98.

IMPORTANT: Time restriction, intruder lockout, and login disabled are checked only if you are using the Novell NMAS Radius server in the same tree as the iChain Authorization tree. Also, the LDAP server does the checks only if the Radius Token authentication is ANDed with an LDAP authentication.

7.2 Installing NMAS, Novell RADIUS, and a Token Method

If you are not using a third-party RADIUS server, you must install NMAS into your eDirectory tree. NMAS is included on the *iChain Authorization Server* CD under the \nmas\nmasserver directory. Change to this directory and run install.exe.

As part of the NMAS installation, you can select and install the Novell RADIUS server components. Consider the following important information:

❑ You must install NICI 2.4.6 on the NMAS server. This version of NICI is included on the *iChain Authorization Server*

❑ You should also install the NICI on the workstation where you are running ConsoleOne. The NICI files for the workstation are located on the *iChain Authorization Server CD* at \nici\winclient. Run the WCNICIU0 application found in this folder to install the NICI snap-ins on the workstation.

❑ We recommend that your NMAS server reside in the same eDirectory tree as your iChain LDAP server that holds the Access Control List (ACL). Doing so allows the ACL to recognize users who are authenticating using NMAS and to allow the users access to the information they need.

❑ If you run ConsoleOne® from a different location than the NMAS server, you must install the NMAS ConsoleOne snap-ins to that location. To do this, change to the \nmas\nmasconsoleone directory on the *iChain Authorization Server* CD and run snapinInstall.exe. This allows you to install the NMAS ConsoleOne snap-ins to any location you choose. Extract the RadiusSnapin.zip file into the directory where you installed ConsoleOne.

After NMAS is installed, you can select, install, and set up a third-party token login method. The token login methods are available for download at the iChain Web site (http://www.novell.com/products/ichain). Documentation on how to install and use each token login method is provided by the partner who developed the login method.

For more information on installing and configuring NMAS and RADIUS, and for general information on installing and setting up a login method, see the NMAS product documentation on the Novell Product Documentation Web site (http://www.novell.com/documentation).

For specific information on installing and using a login method, see the documentation provided by the login method partner.

7.3 Configuring Novell RADIUS Components

After NMAS, Novell RADIUS, and the token login method have been installed, you must configure Novell RADIUS on your NMAS server.

Perform the following procedures in order:

“Creating a Dial Access System (DAS) Object” on page 96
“Configuring the Login Policy Rules” on page 97
“Adding the iChain Proxy Server As a Client of the DAS Object” on page 97
“Creating a RADIUS Dial Access Profile (DAP) Object” on page 97
“Adding an Attribute to the RADIUS DAP Object” on page 98
“Assigning the Token Method to Each User Object” on page 98
“Assigning the DAS Object to Each User Object” on page 98
“Starting Novell RADIUS Services on Your NMAS Server” on page 98

7.3.1 Creating a Dial Access System (DAS) Object

1. Start ConsoleOne.
2. Right-click an Organizational Unit container object, click New, click Object, then click RADIUS:Dial Access System.
3. Specify the object name.
4. Click OK.

5 Specify the password.

6 Click OK.

7.3.2 Configuring the Login Policy Rules

1 Start ConsoleOne.

2 From the Security Container, double-click the Login Policy object.

3 Click the Rules tab (if it isn't already open).

4 Click the plus sign (+) to add a login rule.

5 Click the browse button at the end of the Service Object field, then select the DAS object.

6 On the User list tab, click +, then select the user or container that you want the rule to apply to.

7 On the Sequences tab, click +, select the token method, then select Mandatory.

8 Click OK until you return to ConsoleOne.

7.3.3 Adding the iChain Proxy Server As a Client of the DAS Object

1 Start ConsoleOne.

2 Double-click the DAS object.

3 On the Clients page, click Add.

4 For Address, type the IP address of your iChain proxy server.

5 For Vendor Type, use the drop-down list to select Novell.

6 Type and confirm a secret for this client.

7 Click OK.

8 On the User Resolution page, click the Use Lookup Contexts List to Resolve User Name option if the users are not in the same context as the DAS object.

9 Click Add.

10 Browse and select the container where the User objects reside.

11 In the Object Name field, type a name for the object.

12 Click OK, then click OK again.

7.3.4 Creating a RADIUS Dial Access Profile (DAP) Object

1 Start ConsoleOne.

2 Right-click an Organizational Unit container object, click New, click Object, then click RADIUS:Profile.

3 Click OK.

4 Specify the object name.

5 Click OK.

7.3.5 Adding an Attribute to the RADIUS DAP Object

1. Start ConsoleOne.
2. Double-click the DAP object.
3. On the Attributes page, click Add.
4. Select the Novell eDirectory Name attribute.
5. Select the check box next to Novell eDirectory attribute.
6. Select FDN (Fully Distinguished Name).

   IMPORTANT: It is critical that you select FDN so that name resolution works properly. Otherwise, the users who use this profile will get a 403 User Name Mismatch error when they try to access Web pages.

7. Click OK twice.

7.3.6 Assigning the Token Method to Each User Object

1. Start ConsoleOne.
2. Double-click a User object.
3. Click the Login Methods tab, then select the Token method you previously installed.
4. Follow the partner's instructions for enabling this method.

7.3.7 Assigning the DAS Object to Each User Object

1. Start ConsoleOne.
2. Double-click a User object.
3. Click the Dial Access Services tab.
4. Select a Dial Access Control.
5. Browse and select the DAS object you want to assign to this user.
6. Click Add.
7. Browse and select the DAP object.
8. Click OK twice.

7.3.8 Starting Novell RADIUS Services on Your NMAS Server

From the NMAS server console, type `RADIUS` to start the RADIUS services.

7.4 Setting Up the iChain RADIUS Client

To set up the iChain RADIUS Client, complete these tasks:

- "Adding a RADIUS Authentication Profile" on page 99
- "Adding RADIUS Authentication to an Accelerator" on page 99
- "Adding an LDAP Profile for Mapping a RADIUS User in the Authentication LDAP Directory" on page 99

7.4.1 Adding a RADIUS Authentication Profile

1 In the proxy server administration tool, click Configure, click Authentication, then click Insert.

2 Specify a name for the Radius profile.

3 Click RADIUS Authentication, then click RADIUS Options.

4 Specify the RADIUS server's IP address.

5 Specify 1645 for the Novell NMAS RADIUS server's port number.

6 Specify the shared secret set up in "Adding the iChain Proxy Server As a Client of the DAS Object" on page 97.

7 Click OK twice, then click Apply.

7.4.2 Adding RADIUS Authentication to an Accelerator

1 In the proxy server administration tool, click Configure, then click Web Server Accelerator.

2 Select the accelerator you want to add the RADIUS authentication profile to.

3 Click Modify, click Enable Authentication, then click Authentication Options.

4 Select the Radius profile created in "Adding a RADIUS Authentication Profile" on page 99.

5 Click Add, click OK twice, then click Apply.

You are now ready to authenticate through RADIUS by using the token login method.

7.4.3 Adding an LDAP Profile for Mapping a RADIUS User in the Authentication LDAP Directory

If you are using RADIUS as your only authentication method for an accelerator, then the RADIUS user must be mapped to a user in the authentication LDAP directory. Create an LDAP authentication profile named ldaprad to be used for the mapping. Make sure you specify the following:

- Server address
- User name and password for searches
- LDAP login method

NOTE: If there is not an ldaprad profile to do the mapping, or if the configuration of the ldaprad profile does not locate a user, the authentication fails, even when the RADIUS server returns a success to the iChain server on the authentication.

Accelerator Authentication Options

8

This section provides information on the following:

- Section 8.1, "Accelerator Authentication Parameter Page," on page 101
- Section 8.2, "Authentication Profiles and How They Are Used," on page 114

8.1 Accelerator Authentication Parameter Page

In ConsoleOne, the fifth page of the iChain Web Server Accelerator Wizard is where you specify the accelerator authentication parameters. You can enable or disable authentication, enable or disable Secure Exchange, and create authentication profiles.

Figure 8-1 *Accelerator Authentication Parameter Page*

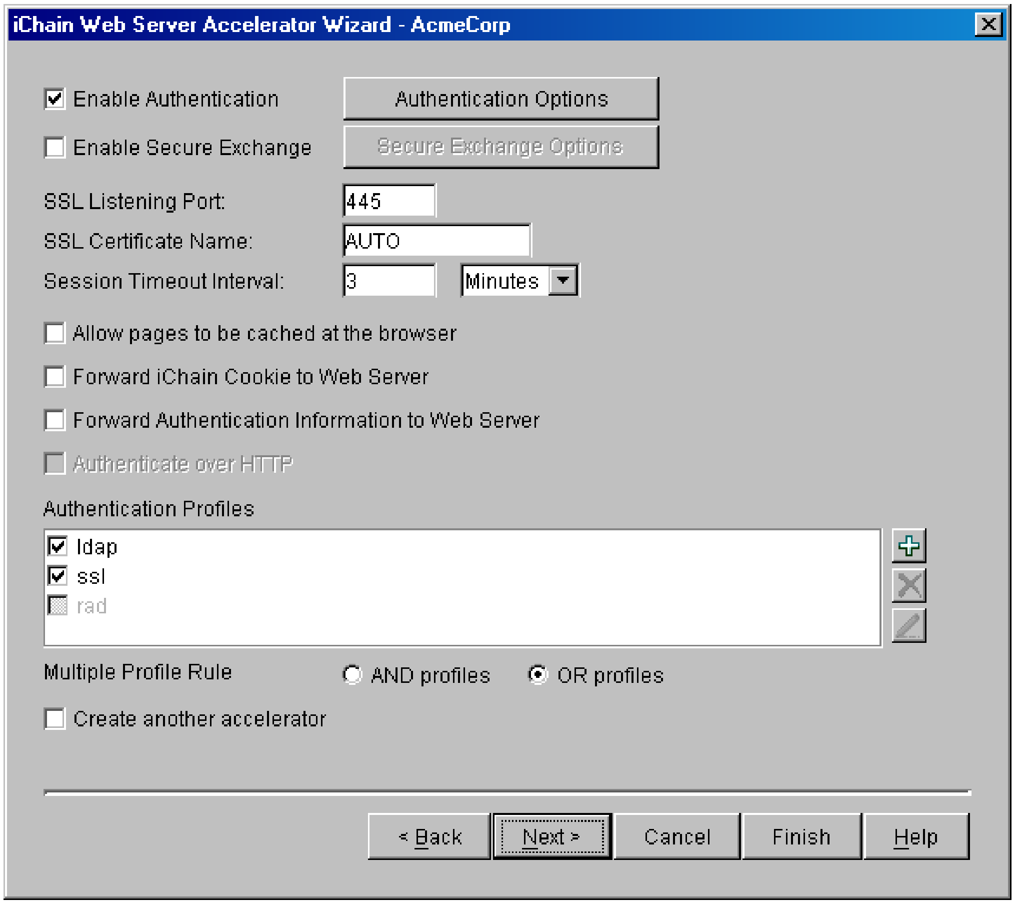

☑ Enable Authentication

☐ Enable Secure Exchange

SSL Listening Port: 445

SSL Certificate Name: AUTO

Session Timeout Interval: 3 Minutes

☐ Allow pages to be cached at the browser

☐ Forward iChain Cookie to Web Server

☐ Forward Authentication Information to Web Server

☐ Authenticate over HTTP

Authentication Profiles

☑ ldap
☑ ssl
☐ rad

Multiple Profile Rule ○ AND profiles ● OR profiles

☐ Create another accelerator

The following table describes the fields on this page:

Field Name	Description	Status
Enable Authentication	Selecting this option forces a user to authenticate to access this Web server	Optional
Enable Secure Exchange	Selecting this option enables Secure Exchange (formerly known as SSLizer). Advanced options for Secure Exchange are not currently available from the wizard, but can be set from the proxy server administration application.	Optional If you choose to enable this option, see Section 5.3, "Using Third-Party Certificates," on page 87 for instructions on how to import the trusted root. Required to use Basic Authentication for LDAP. See Section 6.2, "Enabling Authentication Through the HTTP Authorization Header," on page 93
SSL Listening Port	The SSL port that the user is redirected to for authentication if Secure Exchange is enabled.	Required if authentication or Secure Exchange is enabled
SSL Certificate Name	The certificate name for this accelerator. If the name does not appear in the drop-down list, it can be entered manually.	Required if Secure Exchange is enabled
Session Timeout Interval	The amount of time a connection can be inactive before re-authentication is required.	Required if authentication is enabled or Secure Exchange is enabled
Forward iChain Cookie to Web Server	Sends the Novell® iChain® cookie to the Web server along with the other data being sent.	Optional
Forward Authentication Information to Web Server	Sends username and/or password to the Web server	Optional
Authenticate over HTTP	Allows authentication over unencrypted HTTP instead of HTTPS. This feature is not compatible with RADIUS authentication profiles.	Optional
Authentication Profiles	Each existing profile is listed; those in use appear with a check box. At least one profile must be selected when authentication is enabled. When multiple profiles are in the list, more than one can be enabled. Currently, only Mutual SSL profiles may be used with LDAP or RADIUS profiles. LDAP and RADIUS profiles cannot be used together.	Required if authentication is enabled.

Field Name	Description	Status
Multiple Profile Rule	Only valid if multiple Authentication Profiles are checked. Selects whether only one profile is required (OR) or if all selected authentication methods need to be fulfilled before authentication is granted (AND). OR is the default when multiple profiles are checked.	
Create another accelerator	If this check box is selected when you select the Next button, the wizard returns to the Accelerator Specification Page where a new accelerator can be created. This saves you from needing to select the Next button followed by selecting the Back button three times to return to the Accelerator Specification Page.	Optional

8.1.1 Controls for Accelerator Authentication Parameters

Four buttons allow you to modify authentication parameters:

- **Advanced Options:** Launches the Advanced Authentication Options dialog box as shown in Figure 8-2.
- **Add:** Launches the Add Authentication Profile dialog box.
- **Delete:** Allows you to delete an existing Authentication Profile.
- **Edit:** Launches the Modify Authentication Profile dialog box.

8.1.2 Advanced Authentication Options Dialog Box

The Advanced Authentication Options dialog box allows you to specify advanced authentication options, including options that are set under special circumstances.

Figure 8-2 *Advanced Authentication Options Dialog Box*

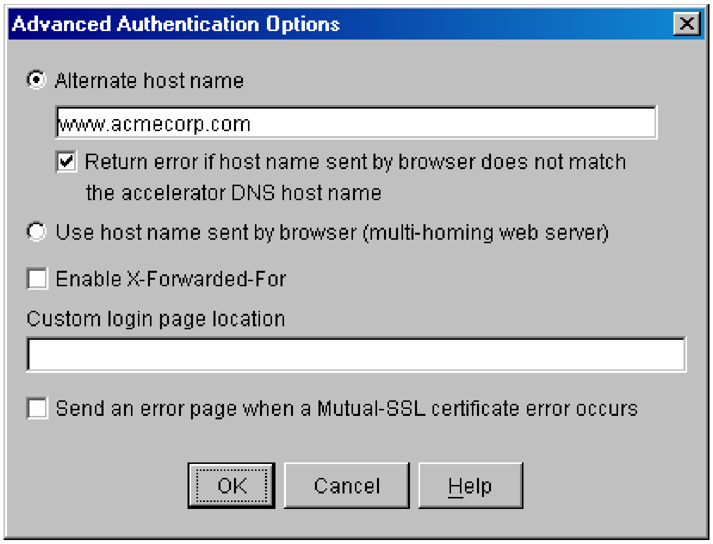
Advanced Authentication Options

● Alternate host name

www.acmecorp.com

☑ Return error if host name sent by browser does not match the accelerator DNS host name

○ Use host name sent by browser (multi-homing web server)

☐ Enable X-Forwarded-For

Custom login page location

☐ Send an error page when a Mutual-SSL certificate error occurs

The following table describes the fields in this dialog box:

Field Name	Description	Status
Enable X-Forwarded-For	Selecting 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.	Optional
Alternate Host Name	Selecting this option causes the specified string to be substituted for the host name in the HTTP header before the request is forwarded to the Web server.	Optional
Return Error if Host Name Sent by Browser Does Not Match the Accelerator DNS Host Name	Selecting this option causes iChain Proxy Services to match the host name 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, iChain Proxy Services returns an error to the requesting browser.	Optional
Use Host Name Sent by Browser	Selecting this option preserves the host name in the HTTP header exactly as it came in the browser request.	Optional

Field Name	Description	Status
Custom Login Page Location	Specifies the location of the login page for this accelerator. The login page must exist on the iChain Proxy server.	Optional
Send an Error Page When a Mutual-SSL Certificate Error Occurs	Select this option to send a specific error page when a Mutual-SSL certificate error occurs. Otherwise a “page not found” message is always given.	Optional

8.1.3 RADIUS Authentication Load Balance and Failover

Load balancing divides a computer's workload between two or more computers so more work can be accomplished in the same amount of time. For authentication, load balancing commonly distributes credential search requests in a fixed sequential order to the different servers.

This feature has a 60-second, non-configurable timeout period after which the authentication search cycle begins again.

When a configuration requires multiple servers, load balancing is often combined with failover. Failover is a backup operational mode where processes are shifted to another server if the primary server becomes unavailable. The failover process offloads tasks to a standby system component. Failover is different from load balancing. Instead of searching servers sequentially, it continues to search on one server until the information cannot be found or the server becomes inactive. Then it moves to the next server.

1. In the proxy server administration tool, click Configure, click the Authentication tab, select radius, then click Modify.
2. Select RADIUS Authentication, click RADIUS Options.
3. In the RADIUS Connect Mode drop-down menu, select FailOver or RoundRobin.
4. Click OK.

8.1.4 Add (Modify) Authentication Profile Dialog Box

The Add Authentication Profile dialog box allows you to name and create authentication profiles. The Modify Authentication Profile dialog box is exactly the same except for the dialog box title.

__Figure 8-3__ Add Authentication Profile Dialog Box

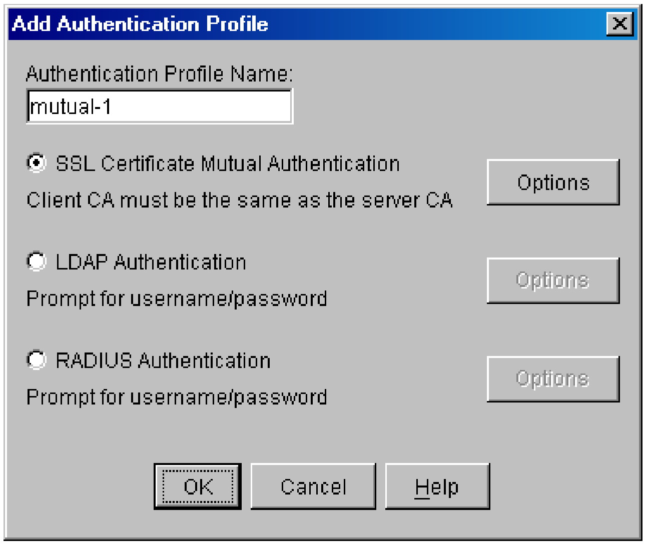

Add Authentication Profile

Authentication Profile Name: mutual-1

● SSL Certificate Mutual Authentication

Client CA must be the same as the server CA

○ LDAP Authentication

Prompt for username/password

○ RADIUS Authentication

Prompt for username/password

The following table describes the fields in this dialog box:

Field Name	Description	Status
Authentication Profile Name	The name of the authentication profile. This name must be unique and must be less than 8 characters with no special characters.	Required
SSL Certificate Mutual Authentication	Specifies a mutual authentication profile.	Optional
LDAP Authentication	Creates an LDAP profile. Selecting this button enables the corresponding options button.	Optional
RADIUS Authentication	Creates a RADIUS profile. Enables the Radius Options button.	Optional

8.1.5 Mutual Certificate Mapping Dialog Box

The Mutual Certificate Mapping dialog box allows you to configure certificate mapping types. See Figure 8-4.

Figure 8-4 *Mutual Certificate Mapping Dialog Box*

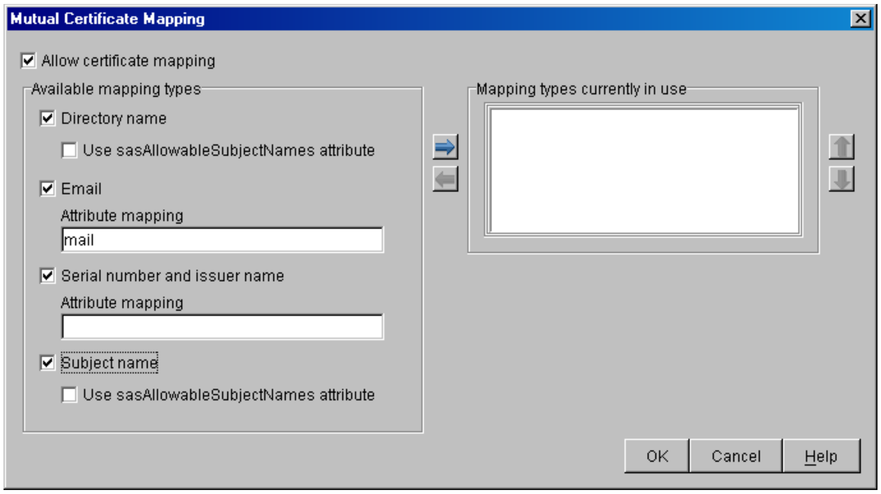
Mutual Certificate Mapping

☑ Allow certificate mapping

Available mapping types

☑ Directory name

□ Use sasAllowableSubjectNames attribute

☑ Email

Attribute mapping: mail

☑ Serial number and issuer name

Attribute mapping:

☑ Subject name

□ Use sasAllowableSubjectNames attribute

Mapping types currently in use

The following table describes the fields and buttons in this dialog box:

Field Name or Button	Description	Status
Directory Name	Enables certificate mapping, which gives four ways to map the user certificate to a user in the iChain LDAP Authentication tree.	Optional
Use sasAllowableSubject Names attribute	If a user is not found with Directory Name and Use sasAllowableSubjectNames is also enabled for directory mapping, the LDAP Authentication tree is searched for a user containing an sasAllowableSubjectName attribute matching the Directory Name in the Subject Alternative Name field of the certificate. If sasAllowableSubjectName is enabled, the LDAP Authentication tree should be configured so that there is no duplication of allowed names between users in the sasAllowableSubjectName attribute.	
Email Description	With Email mapping, there are two possible fields in the user certificate that can be used to identify the certificate portion of the user. The first is the Subject Alternative Name field in the user certificate, with a name type of RFC822. The second is when an e-mail name is embedded in the Subject field of the certificate. If both the Subject Field and the Subject Alternative Name field contain an e-mail address, the Subject Alternative Name is the only field used.	
Attribute Mapping	This attribute is used to match the Email address from the certificate when searching for a user in the LDAP Authentication tree. The default LDAP attribute is mail, which is the attribute currently used by GroupWise® and Novell Certificate Server™. The LDAP Authentication tree should be configured so that there is no duplication of Email addresses between users in the configured email attribute mapping.	
Serial Number and Issuer Name	With serial number and issuer name mapping, both the serial number and the issuer name fields from the certificate are used together to identify the certificate portion of the user.	

Field Name or Button	Description	Status
Attribute Mapping	Both the issuer name and the serial number need to be put into the same LDAP attribute of the user. The LDAP attribute that is used is specified in this field. The LDAP attribute can be any Case Ignore List or Cast Ignore String attribute of the user. If you are configuring your own attribute, make sure the attribute is added to the Person class.	
Subject Name	A user in the LDAP Authentication tree matching the Subject Name field of the certificate is checked first.	
Use sasAllowableSubjectNames Attribute	If a user is not found with Subject name and Use sasAllowableSubjectNames is also enabled for directory name mapping, the LDAP Authentication tree will be searched for a user containing an sasAllowableSubjectName attribute matching the Directory Name in the Subject Alternative Name field of the certificate. If sasAllowableSubjectName is enabled, the LDAP Authentication tree should be configured so that there is no duplication of allowed names between users in the sasAllowableSubjectName attribute.	
Add	The iChain Proxy Server can be configured to use any combination of the four mapping types. This button allows type to be added to the Mapping types currently in the use list.	
Remove	Allows a type to be removed from Mapping types currently in the use list.	
Order Up	Allows for a mapping type within the Mapping types currently in the use list to be moved up. **NOTE:** When searching for a user with the configured mappings, the first user found is the user that is used for authentication and access control, even if the other users map to the same certificate. See "Using Certificate Mapping" on page 80 for more information.	
Order Down	Allows for a mapping type within the Mapping types currently in the use list to be moved down. **NOTE:** When searching for a user with the configured mappings, the first user found is the user that is used for authentication and access control, even if the other users map to the same certificate. See "Using Certificate Mapping" on page 80 for more information.	

8.1.6 Controls for Authentication Profiles

- **LDAP Options:** Launches the LDAP Authentication Profile Options dialog box, which allows you to specify LDAP authentication parameters. It is functionally identical to the corresponding dialog box in the iChain Proxy Server administration application.
- **Radius Options:** Launches the RADIUS options dialog box.

Figure 8-5 *LDAP Authentication Profile Options Dialog Box*

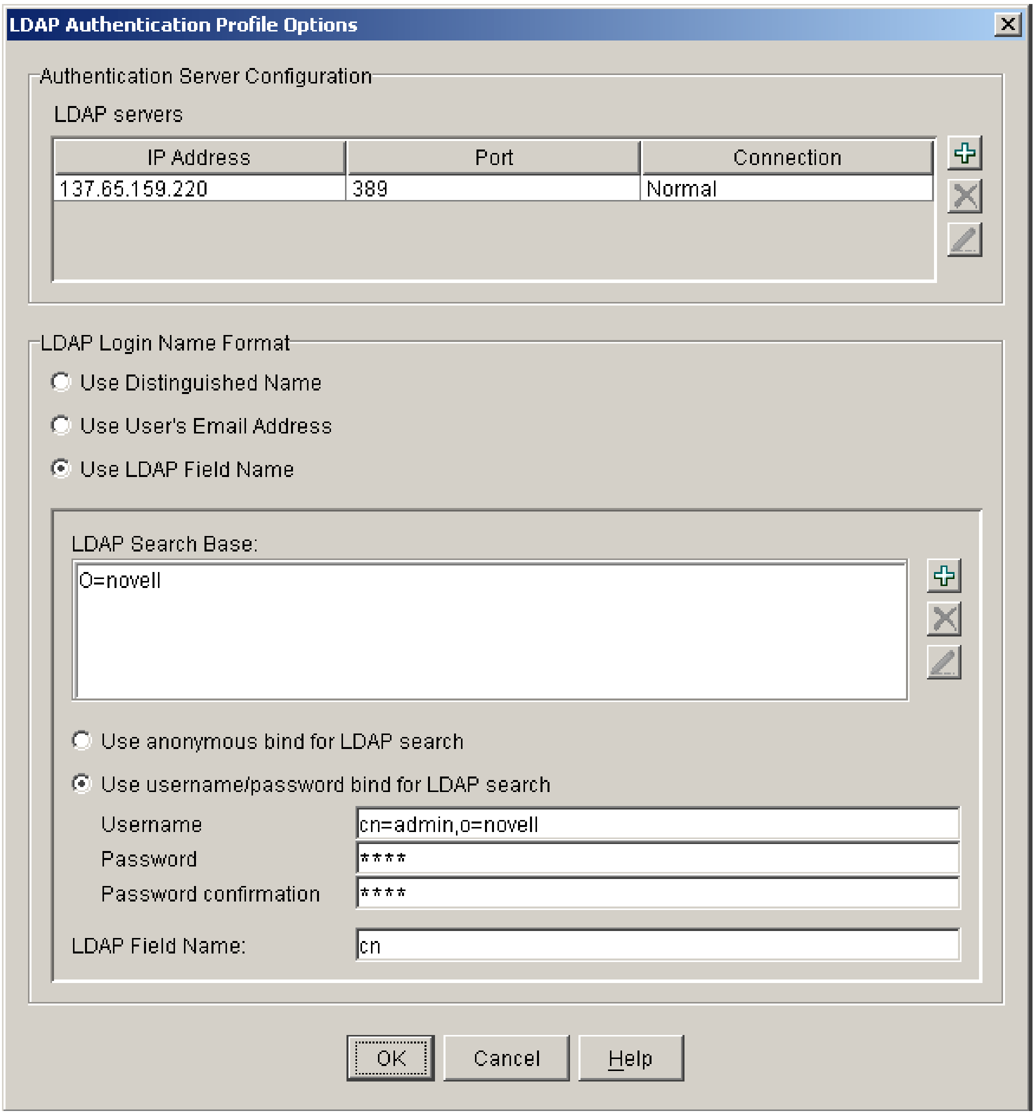

LDAP Authentication Profile Options

Authentication Server Configuration

LDAP servers

IP Address	Port	Connection
137.65.159.220	389	Normal

LDAP Login Name Format

○ Use Distinguished Name

○ Use User's Email Address

● Use LDAP Field Name

LDAP Search Base:

O=novell

○ Use anonymous bind for LDAP search

● Use username/password bind for LDAP search

Username: cn=admin,o=novell

Password: ****

Password confirmation: ****

LDAP Field Name: cn

The following table describes the fields in this dialog box:

Field Name	Description	Status
LDAP Servers	This table lists the IP address, port, and connection type for all the LDAP servers used for this profile. Currently, the port and connection type must be the same for all servers.	Required
Use Distinguished Name	Selecting this option requires users to log in using their DS names.	Optional
Use User's Email Address	Selecting this option requires users to log in using their e-mail addresses.	Optional

Field Name	Description	Status
Use LDAP Field Name	Selecting this option requires users to log in using some LDAP field.	Optional
LDAP Search Base (LDAP User Contexts)	This field displays as LDAP Search Base when either Use User's Email Address or Use LDAP Field Name is selected. It allows entry/deletion/modification of LDAP search bases or user contexts.	Required
Use Anonymous Bind for LDAP Search	Bind anonymously to search the LDAP directory.	Optional
Use Username/Password Bind for LDAP Search	Bind with a proxy server to search the LDAP directory.	Optional
Username	Proxy username in LDAP format.	Required when Use Username/Password Bind for LDAP Search is selected
Password	Proxy user password.	Required when Use Username/Password Bind for LDAP Search is selected
Password Confirmation	Proxy user password confirmation.	Required when Use Username/Password Bind for LDAP Search is selected
LDAP Field Name	LDAP field name to search for (only visible with Field Name).	Required when Use LDAP Field Name is selected.

8.1.7 Controls for Authentication Profile Options

Use the following buttons to control the authentication profile:

- **Add LDAP Server:** Allows you to launch the New LDAP Authentication Server dialog box.
- **Delete LDAP Server:** Allows you to delete an authentication server from the list.
- **Edit LDAP Server:** Allows you to launch the Modify LDAP Authentication server dialog box.
- **Add LDAP Context:** Allows you to launch the dialog box to add an LDAP Search Base/User Context (if DN is selected).
- **Delete LDAP Context:** Allows you to delete an LDAP Search Base/User Context from the list.
- **Edit LDAP Context:** Allows you to launch the dialog box to modify an LDAP Search Base/User Context (if DN is selected).

8.1.8 New LDAP Authentication Server Dialog Box

The New LDAP Authentication Server dialog box allows you to specify the parameters for new LDAP authentication servers. The Modify LDAP Authentication Server dialog box is exactly the same except for the dialog box title.

Figure 8-6 *New LDAP Authentication Server Dialog Box*

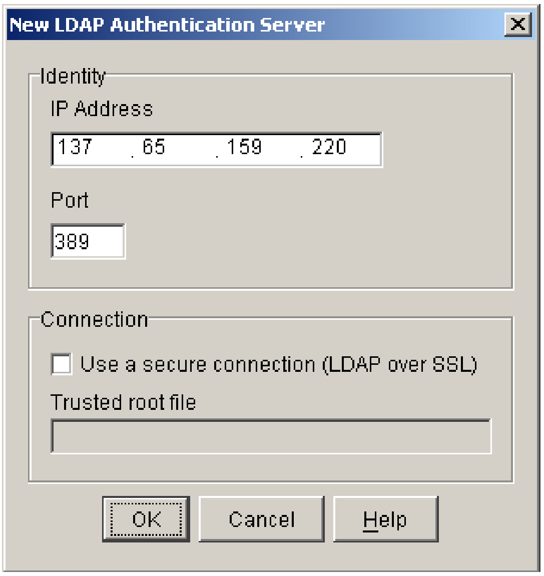

New LDAP Authentication Server

Identity

IP Address: 137 . 65 . 159 . 220

Port: 389

Connection

☐ Use a secure connection (LDAP over SSL)

Trusted root file:

The following table describes the fields in this dialog box:

Field Name	Description	Status
IP Address	The IP address of this LDAP server.	Required
Port	The LDAP port to communicate over. Currently, this is only modifiable for the first LDAP server in the list.	Required
Use a Secure Connection (LDAP over SSL)	If selected, authentication information is sent over LDAPS (encrypted). This is only modifiable for the first LDAP server.	Optional
Trusted Root File	Specifies the trusted root file to be used for secure communications. This is only modifiable for the first LDAP server.	Required when Use a Secure Connection is selected.

8.1.9 Add LDAP Context Dialog Box

The Add LDAP Context dialog box provides the input of LDAP search bases or user contexts. The Modify LDAP Context dialog box is exactly the same except for the dialog box title.

Figure 8-7 Add LDAP Context

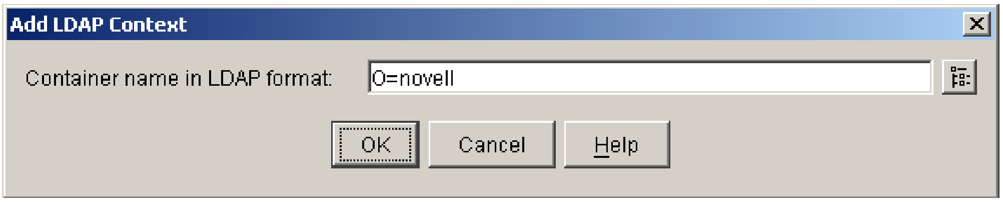

Container name in LDAP format: O=novell

The following table describes the field in this dialog box:

Field Name	Description	Status
Container name in LDAP format	The name of the container in LDAP (comma delimited) format	Required

8.1.10 Controls for Add LDAP Context

Use the Object Browser button to launch an object browser to select the desired container.

8.1.11 Radius Options Dialog Box

The Radius Options dialog box allows you to specify the parameters for RADIUS profiles. This dialog box is functionally identical to the corresponding iChain Proxy Server administration application dialog box.

Figure 8-8 *RADIUS Profile Options Dialog Box*

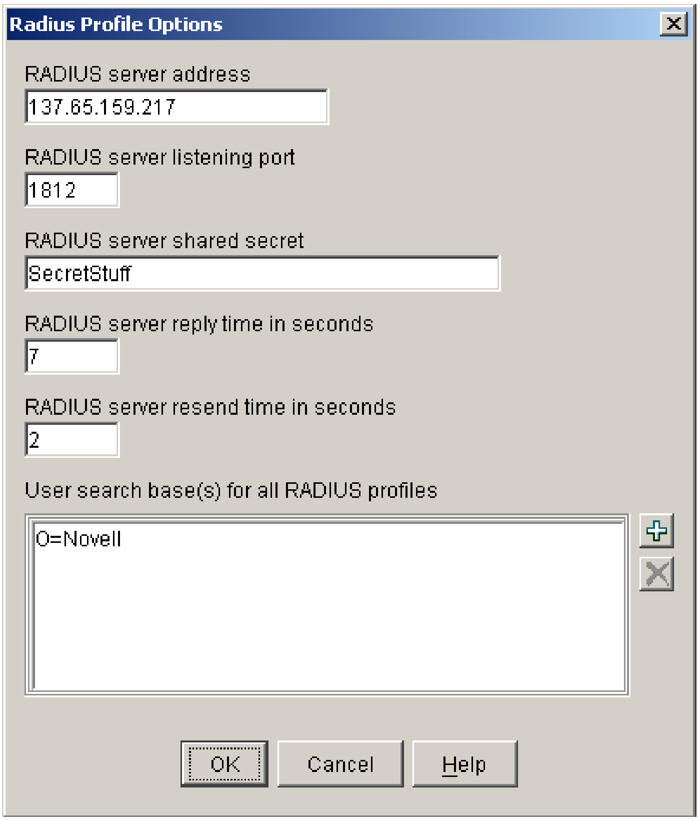

Radius Profile Options

RADIUS server address
137.65.159.217

RADIUS server listening port
1812

RADIUS server shared secret
SecretStuff

RADIUS server reply time in seconds
7

RADIUS server resend time in seconds
2

User search base(s) for all RADIUS profiles
O=Novell

The following table describes the fields in this dialog box:

Field Name	Description	Status
RADIUS Server Address	The IP address of the RADIUS server.	Required
RADIUS Server Listening Port	The port number on which the RADIUS server listens for incoming authentication.	Required
RADIUS Server Shared Secret	The string the RADIUS server uses to verify that the appliance can request authentication of users.	Required
RADIUS Server Reply Time in Seconds	The total time the appliance waits for a response from the RADIUS server before authentication fails. The default is 7 seconds.	Required

Field Name	Description	Status
RADIUS Server Resend Time in Seconds	The interval in seconds between appliance requests to the RADIUS server. The default is two seconds. This means that the appliance could send three requests before the 7-second default limit expires and the authentication request fails.	Required
User Search Base(s) for All RADIUS Profiles	Lists the contexts that the proxy server uses when searching for the user being authenticated when using non-Novell RADIUS authentication. This list applies to all RADIUS profiles, not just the current one being created or modified.	Optional

8.1.12 Controls for RADIUS Options

The following buttons allow you to add or delete search bases:

- **Add Search Base:** Allows you to launch an object browser to select the desired container.
- **Delete Search Base:** Allows you to delete a search base from the list.

8.2 Authentication Profiles and How They Are Used

There are three classes of authentication profiles:

- LDAP
- RADIUS
- Mutual

You cannot have two or more authentication profiles from the same class combined. If two or more authentication profiles are provided in an AND condition, then only one user can be represented from a successful login.

Mutual and RADIUS authentications first try to obtain a username or distinguished name (DN), or some unique information about a user. This data then is used to either search for the specific user or bind to the specific user using an LDAP authentication profile. The end result is that iChain has a full DN.

A user must be found so that OLAC variables can be obtained and ACLCheck can provide access control for the user. Only the ACLCheck authentication profile is used for OLAC and ACLCheck.

The following table provides information about the authentication profiles that are available:

Authentication Class	Sub-class	Description
Mutual	DN	The full DN is stored within the certificate. A potential issue is that the format of the DN might be in reverse order. No formatting is done to the DN.
Mutual	Certificate mapping	Map data in the certificate to a user in an LDAP profile.
RADIUS	Novell RADIUS	The username and token are used to authenticate to the RADIUS server. This can be configured to return the full DN of the user that authenticated. With this DN, a user can be found in an LDAP profile.
RADIUS	Normal	The username and token are used to authenticate to the RADIUS server. The username must be used in some fashion to match to a user in the LDAP profile. This limits the kinds of LDAP profiles that can be used.
LDAP	DN bind with DN input	This configuration lists no LDAP contexts. A full DN is required for input. This profile might not work in conjunction with RADIUS:normal.
LDAP	DN bind with search field	A DN is built from the search field and username values. A bind () is called for all contexts or a specific context with the DN and password. The default search field value is **cn**.
	Attribute search	A search is performed for equivalence on the defined attribute and the username value. The search is performed for each subtree listed until a match is found. Note that multiple matches are not allowed. The search is performed using the rights of the LDAP username defined in the LDAP profile. After a match is found, a bind () on the DN of the matched user with the input password is called to validate the password.

8.2.1 The Authentication Process With Mutual Authentication

The data is extracted from the certificate and used to locate a user in an LDAP tree. The LDAP tree that is selected comes from an LDAP authentication profile named ldapcert. If this authentication profile is not found, the LDAP server defined in the aclcheck profile is used. The LDAP authentication profile must have a user and password with rights to read and verify the information within the mutual certificate.

When a mutual authentication profile is only used to authenticate a user, a password is not provided. The username is extracted from searches that are performed with the certificate data and the mapping rules. This means that the password field value is empty if passed to a back-end Web server.

8.2.2 The Authentication Process With RADIUS Authentication

Only a username and token value are needed to log in to a RADIUS server. The LDAP tree that is selected comes from an LDAP authentication profile named ldaprad. If this authentication profile is not found, the LDAP server defined in the aclcheck profile is used. The LDAP authentication profile must have a user and password that has rights to read and verify the username provided in the form. With a properly configured Novell RADIUS server, the full DN is returned with a successful login. This full DN represents the user.

The LDAP password is optional. If a password is provided, the input username and LDAP password will be used in a bind () call against the selected LDAP authentication profile. If successful, the user is authenticated. If it is not successful, the user must try again. If the password is not provided, a search is made on the selected LDAP tree to locate the user.

8.2.3 Using LDAP AND Mutual Over HTTP

The user should be prompted for a certificate when a restricted or secure page is hit. The mutual screen is displayed first, followed by the LDAP login screen. Both username values must be the same or the authentication fails.

8.2.4 Using LDAP AND Mutual Over HTTPS

The user should be prompted for a certificate upon initial connection to an accelerator. If required, this could be changed when a restricted or secure page is hit. The LDAP login screen comes up when a restricted or secure page is hit. Both username values must be the same or the authentication will fail.

8.2.5 Using LDAP OR Mutual, RADIUS OR Mutual

This configuration raises issues when the user authenticates with mutual authentication and a password is not defined that can be passed to origin Web servers.

8.2.6 Using RADIUS AND Mutual

This configuration is similar to LDAP AND Mutual with the exception that the LDAP password is optional with a RADIUS profile. If a password is provided, then an LDAP bind takes place. The Mutual authentication can use the ldapcert authentication profile and the RADIUS authentication can use the ldaprad authentication profile.

8.2.7 Using LDAP AND RADIUS

This combination forces the user to input the LDAP password. Only one authentication page needs to be presented to the user. The current RADIUS login page lists the LDAP password as optional. The administrator is responsible for configuring it correctly.

8.2.8 Using LDAP OR RADIUS

This combination should give the user the RADIUS login page. The user can either input the username with the token or the username with the LDAP password (either will work). After a RADIUS authentication, the LDAP profile in the combination is used to locate the user within the tree. If the user is not found, the authentication fails.

8.2.9 Using LDAP AND RADIUS AND Mutual

This combination is similar to RADIUS AND Mutual, where once the user is located, the LDAP profile is used to validate that the user exists instead of the default aclcheck profile. The main difference is that the LDAP password is required and an LDAP bind takes place.

8.2.10 Using LDAP OR RADIUS OR Mutual

This combination is similar to RADIUS OR Mutual, where once the user is located, the LDAP profile is used to validate that the user exists instead of the default aclcheck profile. The main difference is that the LDAP password is required and an LDAP bind takes place.

Using Cross-Domain Authentication 9

Cross-Domain Authentication (CDA) provides a graded authentication feature that lets you build a trust relationship (CDA) among different domain types (for example, www.c.com, www.l.com, and www.lc.com). Inside this CDA, your users login only once when accessing accelerators with the same types of authentication methods.

For more information about graded authentication, see Using Graded Authentication (http://www.novell.com/documentation/nmas23/admin/data/a53visc.html) in the *Novell Modular Authentication Services Administration Guide.*

This section provides information on the following:

- Section 9.1, "CDA Scenario and Examples," on page 119
- Section 9.2, "Selecting Accelerators as Members of CDA and Cross-Domain Brokers," on page 120
- Section 9.3, "Configuring the CDA," on page 120
- Section 9.4, "CDA and Session Broker," on page 121

9.1 CDA Scenario and Examples

The following scenario clarifies how CDA works:

These accelerators are CDA-enabled and their authentication methods are as follows:

www.l.com — LDAP authentication

www.c.com — Certificate (Mutual) authentication

www.lc.com — Certificate (Mutual) and LDAP authentication

Single Sign-on Example (same grade with same authentication methods): If the authentication methods of all accelerators are the same (for example, all LDAP/certificate, or all certificate and LDAP), after your user logs in to a domain, he or she can access any other domains without logging in to them.

Graded Authentication Example 1 (from grade high to low): If your user accesses www.lc.com first, he or she is asked to log in twice, once for certificate and again for LDAP. After the user accesses www.lc.com, he or she can access www.l.com (or www.c.com) without any login.

Graded Authentication Example 2 (from grade low to high): If a user accesses www.l.com (or www.c.com) first, he or she is asked to log in using LDAP (or certificate). If the user wants to access www.lc.com, he or she is asked to log in using a certificate.

Graded Authentication Example 3 (same grade but with different authentication methods):

If a user accesses www.l.com first and then later accesses www.c.com, he or she is asked to log in twice; once for www.l.com with LDAP and again for www.c.com with certificate (then the user can access www.lc.com without logging in).

9.2 Selecting Accelerators as Members of CDA and Cross-Domain Brokers

In Cross-Domain Authentication, only one accelerator can be chosen as a Cross-Domain Broker (DB), while other accelerators are non-DBs. The DB works as a coordinator to see whether a successful authentication (login) has been performed. The CDA feature should be enabled only when there is more than one accelerator with authentication turned on and users want to use single sign-on and graded authentication among these accelerators.

Before choosing accelerators as members of CDA, consider the following criteria:

- **Security.** Security is the first concern. For example, if you have www.c.com and www.lc.com with certificate authentication for both accelerators, if the certificate for www.c.com is not trusted by www.lc.com, one (or both) of them should not be CDA-enabled. If both accelerators are CDA-enabled, a user can log in to one of the accelerators but will not be prompted to log in again when he or she accesses the other accelerator. Because CDA uses a single session cookie for all CDA-enabled accelerators, if a user logs out or times out from one of the accelerators, he or she will be logged out from all CDA-enabled accelerators.
- **Graded Authentication.** CDA provides a single sign-on feature by allowing accelerators with the same type of authentication to require log in only once.

  NOTE: For accelerators that use different authentication methods, we do not recommend that you use CDA unless one session cookie is important for these accelerators. (For example, www.lc.com uses Radius authentication, www.l.com uses LDAP, and www.c.com uses certificate. In this example, there are no common authentication methods among www.l.com, www.c.com, and www.lc.com.)

- **Performance.** CDA uses redirection to set and get the session cookie between DB and non-DB accelerators. The overhead for these additional redirections has little performance impact because it reduces the total number of logins that involve manual interaction. There is no extra redirection when accessing a DB-enabled accelerator.

  The selection of a DB is critical in CDA. You should never disable a DB-enabled accelerator, or disable its authentication. Also, because there is no extra redirection when accessing a DB-enabled accelerator, we recommend that you select the most frequently accessed accelerator to be the DB.

  NOTE: If these criteria are difficult to meet, you can select any CDA member to be the DB.

9.3 Configuring the CDA

1 At the Novell® iChain® Proxy Server GUI, click Configure, then click Domain.

The Cross-Domain Authentication Settings dialog box appears.

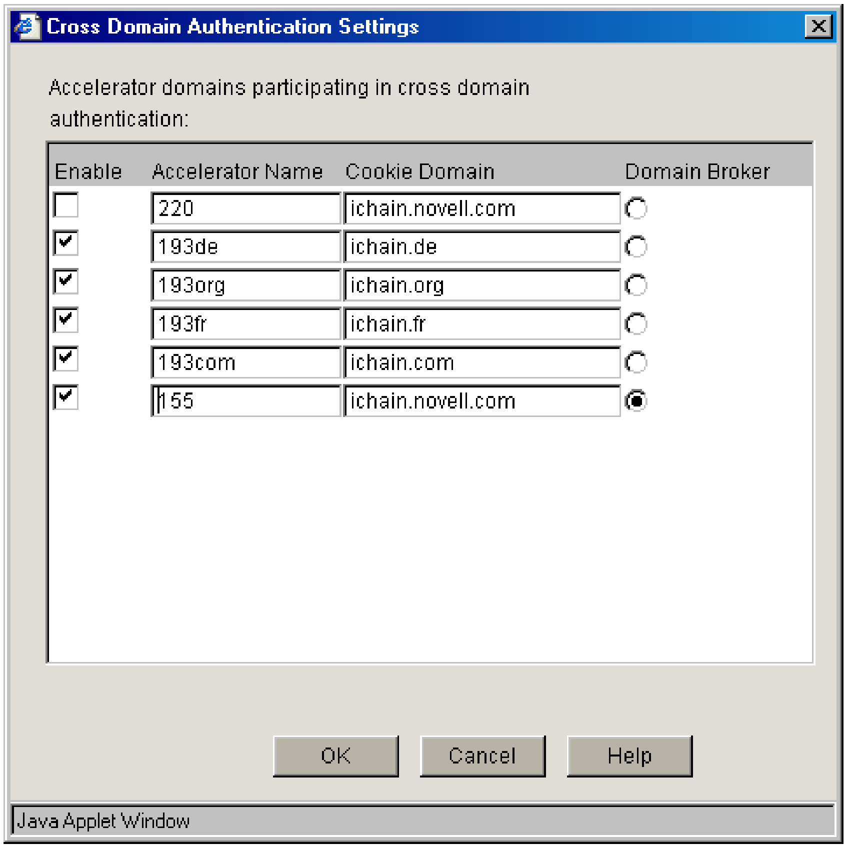

Cross Domain Authentication Settings

Accelerator domains participating in cross domain authentication:

Enable	Accelerator Name	Cookie Domain	Domain Broker
□	220	ichain.novell.com	○
☑	193de	ichain.de	○
☑	193org	ichain.org	○
☑	193fr	ichain.fr	○
☑	193com	ichain.com	○
☑	155	ichain.novell.com	●

2 In the Cross-Domain Authentication Settings dialog box, select the boxes of accelerators you want to be CDA members.

IMPORTANT: You must use unique cookie domain names to control and regulate authentication. If you have similar cookie domain names, users can authenticate even if an accelerator is not part of the CDA setup.

3 Select Domain Broker from the enabled CDA members by selecting the radio option among these members.

NOTE: You should use the same authentication profile for the same type of authentication (for example, you should use only the “ldapx” authentication profile for all LDAP authentication of all CDA accelerators).

9.4 CDA and Session Broker

When using Session Broker, all groups of accelerators participating in CDA should be configured identically. Session Broker relies on the fact that groups of accelerators are a single entity. Therefore, if an accelerator is participating in CDA on one machine, then it should be participating in an identical CDA on all other machines with that accelerator. Logouts are communicated to all iChain servers to ensure security.

Configuring Session Broker

10

This section contains information about the Session Broker feature. The following topics are discussed:

- Section 10.1, “What is Session Broker?,” on page 123
- Section 10.2, “Configuring Session Broker,” on page 124
- Section 10.4, “Troubleshooting Session Broker Issues,” on page 126

10.1 What is Session Broker?

If you need to have more than one Novell® iChain® box at your site, you might need Session Broker. Session Broker allows “sessions” (user authentication data) to be shared between multiple iChain boxes. This in turn lets a user authenticate only once when browsing across all of the boxes.

For example, an iChain site might contain multiple iChain servers, iChain A and iChain B, are accelerating back-end Web servers. These two iChain servers are sitting behind a layer-4 (L4) switch whose task is to load balance the http and https requests going to the iChain servers.

In this scenario, a user browses to a page on your Web site. The Web server's DNS entry resolves to the virtual IP address of the L4 switch. The L4 switch then transmits the request to the iChain A server. iChain asks the user to authenticate before granting access to the page. Once authenticated, iChain sets an iChain-specific cookie on the user's browser.

Suppose that while browsing, the user is directed by the L4 switch to a page protected by the iChain B server. Without Session Broker, the user is required to authenticate again. This is because the iChain B server has no way of knowing that the user was authenticated to the iChain A server. When a Session Broker server is running, each iChain server (acting as a client to the Session Broker server) relays information about their authenticated users if the incoming request has an iChain-specific cookie set. In the above scenario, when the user tries access a protected page through the iChain B server, the user is not already authenticated to this box, but the incoming request has an iChain cookie set. The iChain B server therefore asks the Session Broker server if the user is already authenticated to a different iChain box. If so, the user is granted access without needing to authenticate again. See Figure 10-1 for a visual diagram of this process.

Figure 10-1 *Session Broker: Visual Diagram*

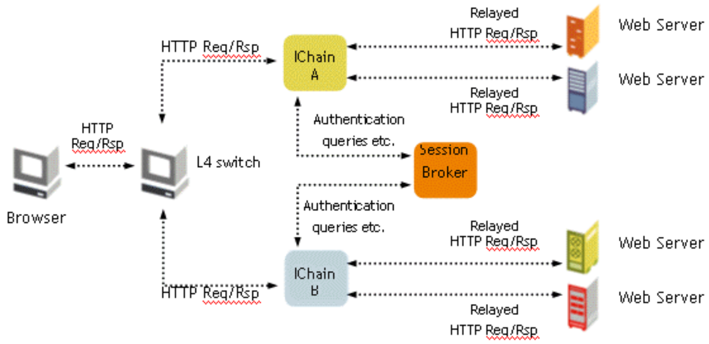

10.2 Configuring Session Broker

Before setting up Session Broker in your iChain configuration, you must install iChain on a standalone server that is dedicated as the primary Session Broker server. The primary Session Broker server should not be an iChain accelerator server. A typical setup would include two iChain accelerator servers and a third iChain server as a dedicated primary Session Broker server. One of the iChain servers could be designated as the secondary server. When users authenticate to either of the iChain servers, the session information would be updated to the primary Session Broker server. If either of the iChain servers were to go down in this setup, the user would not be prompted to re-authenticate.

IMPORTANT: Only authentication profiles with the same name on each iChain server are shared. If the second iChain server doesn't have an authentication profile with the same name as the authentication profile the user authenticated to on the first server, the user will be required to authenticate again.

Only the primary Session Broker server maintains a database of all authenticated sessions. The secondary server does not contain a duplicated/synchronized copy of the database. The secondary server is initialized only if the primary server fails. Only then does the secondary server begin building a new Session Broker database. For more information, see the following Technical Information Documents, "iChain Session Broker FAQ" (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10075532.htm), and "iChain Session Broker Operation Work Flow" (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10087576.htm).

You use ConsoleOne to configure Session Broker on an iChain server:

1 In ConsoleOne®, right-click the ISO object and select Properties.

2 Select the Session Broker tab.

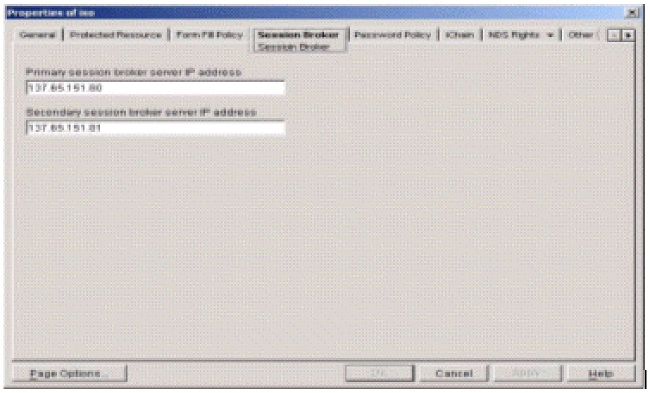

Primary session broker server IP address
137.65.151.80

Secondary session broker server IP address
137.65.151.81

3 Specify the two requested IP addresses: the Primary session broker server IP address and the Secondary session broker server IP address.

The Secondary session broker server IP address is optional and only becomes active if the Primary session broker server IP address is down; otherwise, it remains idle.

Place the primary Session Broker on an iChain server that has no other responsibilities. The secondary Session Broker can be configured on an iChain server with other duties since it is used only for short periods of time.

4 Do the following steps on the primary Session Broker:

4a Establish a shared secret between your iChain servers and the Session Broker(s) that can be used to encrypt data passed between them. To do this, enter the following command at any iChain console:

```
createsessionbrokerkey
```

This creates the Session Broker key on the primary Session Broker server and also copies the key to a floppy disk. This floppy disk is used to install the Session Broker keys on all iChain servers within the Session Broker setup, including the primary and secondary Session Broker servers.

It is possible to disable encryption of data passed between iChain and Session Broker. Do this only if you are certain that the messages passed between them are secure. To disable encryption of data, instead of entering the command above, enter the command, `createnullsessionbrokerkey`. This creates a null key, telling iChain and Session Broker that no encryption is desired.

4b When prompted, insert a floppy disk into the floppy drive and enter a password to encrypt the shared secret. The password you enter must be at least 6 characters in length.

4c When prompted, confirm the password.

5 Do the following on all machines participating in the Session Brokerage:

5a Insert the floppy disk containing the encryption key into the floppy drives of each of your iChain servers, including the primary and secondary Session Broker servers.

5b At the console of each Session Broker server, enter the following command:

```
installsessionbrokerkey
```

5c When prompted for the password, enter the password you gave when you created the encryption key (see Step 4b).

6 After creating or installing the encryption key, restart your proxy server in order for the server to read in the key and begin encrypting the Session Broker data.

7 At the iChain console of the servers you have designated to be the primary and secondary Session Broker servers, enter the following command:

```
set authentication sessionbrokerenable = yes
```

Set this parameter only on the designated primary and secondary iChain servers. All other iChain servers know they participate by reading the configuration information from the ISO object.

Session Broker should now be running. To confirm that it is running and is initialized on the primary Session Broker server, load tcpcon > protocol information > TCP > TCP listeners, then confirm that 5001 exists in the list.

10.3 Configuring Session Broker without a Floppy Drive

The current configuration of Session Broker requires a floppy drive because the encryption key is on a floppy disk. If you do not have a floppy drive, you can still configure Session Broker using an updated keyinstall.nlm.

1 Create a text file in sys:/etc/called keyinst.cfg on any server that uses Session Broker.

NOTE: If this file does not exist, Session Broker reads the file from the floppy drive.

2 In the keyinst.cfg file, add a single line with the path for your session.dat file.

For example:

sys:/etc/proxy/data/session.dat

3 Reboot the server to initialize the keyinst.cfg file.

4 On your primary Session Broker server, open the Command Line Interface (CLI) and initiate the following command:

```
createsessionbrokerkey
```

The session.dat file is specified in Step 2 during this procedure.

IMPORTANT: You need to get a copy of the session.dat file before you complete Step 5 because the file gets deleted for security reasons. You need to place the session.dat file on each server that uses Session Broker.

5 Open the Command Line Interface (CLI) and initiate the following command:

```
installsessionbrokerkey
```

6 Complete Step 5 on each server (including the primary server) that uses Session Broker.

7 On the primary and secondary servers, initiate the following set command:

```
set authentication sessionbrokerenable=yes
```

8 Reboot the primary session broker and let it come up. Then reboot all the other servers that participate in the session brokerage.

10.4 Troubleshooting Session Broker Issues

This section provides information on how to troubleshoot Session Broker issues, including providing information on specific areas, and providing information on steps you can take to troubleshoot Session Broker on your system.

10.4.1 Troubleshooting Tools

- TCPCON

  Go to Protocol Information, the click TCP > TCP Connections. Confirm that both TCP port 5001 and 5002 are listening.

- Generate unencrypted Session Broker sessions

  When creating the Session Broker keys, use the command `createnullsessionbrokerkey` instead of `createsessionbrokerkey`. Doing this allows trace information to be obtained where username, context, authentication profile, and other user authentication attributes are visible.

- PKTSCAN.NLM

  This allows administrators to take traces of traffic from the iChain servers going to the Session Broker servers. This can be used to verify that communication exists in both directions, and that they correct request/response data is being exchanged.

- SB command line parameters

Loading SB with the -d <n> option spawns a debug Session Broker screen with verbose information in it.

n=1 Broker Listens
n=2 Broker Activity
n=3 Client Connects
n=4 Packet detail
n=5 Debug (includes everything)

See Figure 10-2.

Figure 10-2 *Session Broker Screen*

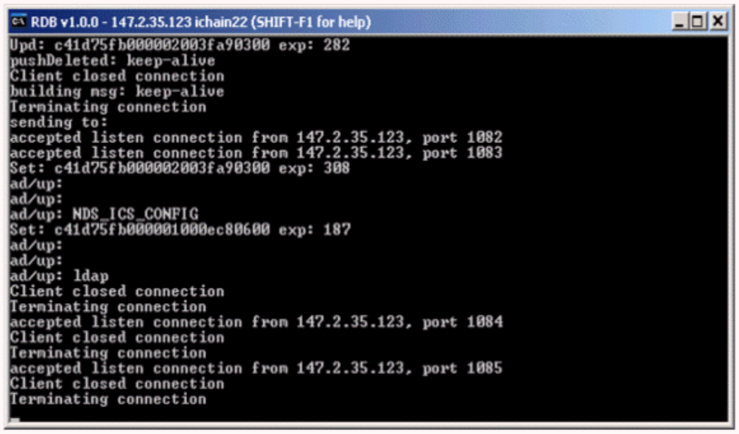

```
Upd: c41d75fb000002003fa90300 exp: 282
pushDeleted: keep-alive
Client closed connection
building msg: keep-alive
Terminating connection
sending to:
accepted listen connection from 147.2.35.123, port 1082
accepted listen connection from 147.2.35.123, port 1083
Set: c41d75fb000002003fa90300 exp: 308
ad/up:
ad/up:
ad/up: NDS_ICS_CONFIG
Set: c41d75fb000001000ec80600 exp: 187
ad/up:
ad/up:
ad/up: ldap
Client closed connection
Terminating connection
accepted listen connection from 147.2.35.123, port 1084
Client closed connection
Terminating connection
accepted listen connection from 147.2.35.123, port 1085
Client closed connection
Terminating connection
```

10.4.2 Troubleshooting Steps

The following are various steps you can take to troubleshoot your Session Broker issues:

- Verify that the latest patches are installed on your system.

  All the latest iChain 2.3 patches are available at the Novell® Support Web site (http://support.novell.com). They start with ic23*.exe. Check the patches to see if updates exist to either the Session Broker server (sb.nlm) or the Session Broker clients (proxy.nlm).

- Verify that the Session Broker keys exist and are installed.

  The error “Unable to initialize the Session Broker” is often triggered if the floppy disk used to copy the keys is not formatted or includes hidden files.

- Verify that “get authentication sessionbrokerenable” returns a yes.

- Verify that sb.nlm loads without errors on the primary and secondary (if one exists) Session Broker servers.

  If errors are reported on the SB debug screen, they are often related to the initialization of the Session Broker keys. After the keys are generated, the Session Broker server must be rebooted.

- Verify that no LDAP errors exist in the Proxy Console screen of the iChain servers.

  At startup, all iChain servers generate LDAP requests to the ISO object to get the Session Broker configuration attributes. If any LDAP communication issues exist, the iChain servers might not get the Session Broker information, and might result in issues.

  LAN traces with PKTSCAN can also confirm that LDAP information was received correctly.

- Verify that authentication with no SB works properly.

If the main proxy authentication engine has problems with authenticating users, the Session Broker services are also likely to fail.

- Verify that the issue being experienced is not load-related.

  We recommend that a dedicated Session Broker server be used. In cases where the Session Broker server is running on an iChain server with accelerators enabled, try to create a dedicated Session Broker server to see if your problem still exists.

Using the iChain Service Object (ISO)

11

The iChain Service object (ISO) is an object in eDirectory™ that functions as the main component of Novell® iChain® security and the single sign-on environment. You can create the ISO by using the iChain ConsoleOne snap-ins found on the *iChain Authorization Server Install* CD. Prior to creating the iChain Service object, you must perform an iChain schema extension to the eDirectory server (or tree).

The iChain Proxy looks at the ISO to determine the following:

- To determine activation details
- To determine whether URLs accessed by users are protected
- To determine OLAC parameters
- To determine the Form Fill policy
- To determine the trusted root container in eDirectory from which to copy the certificates to the proxy machine
- To determine the primary and secondary session broker addresses

When viewed in ConsoleOne, the iChain Service object in eDirectory has six important pages:

- Section 11.1, “General Page,” on page 129
- Section 11.2, “Protected Resource Page,” on page 130
- Section 11.3, “Form Fill Policy Page,” on page 133
- Section 11.4, “Session Broker Page,” on page 134
- Section 11.5, “iChain Page,” on page 135

11.1 General Page

This is where the location of the trusted root container is specified for the proxy to read and copy the certificates listed in this container. For example, if you are using LDAP to secure the authentication profile, the trusted root certificate (which was imported into this container) is copied to the proxy and is used to make the LDAP secure connection.

Figure 11-1 General Page

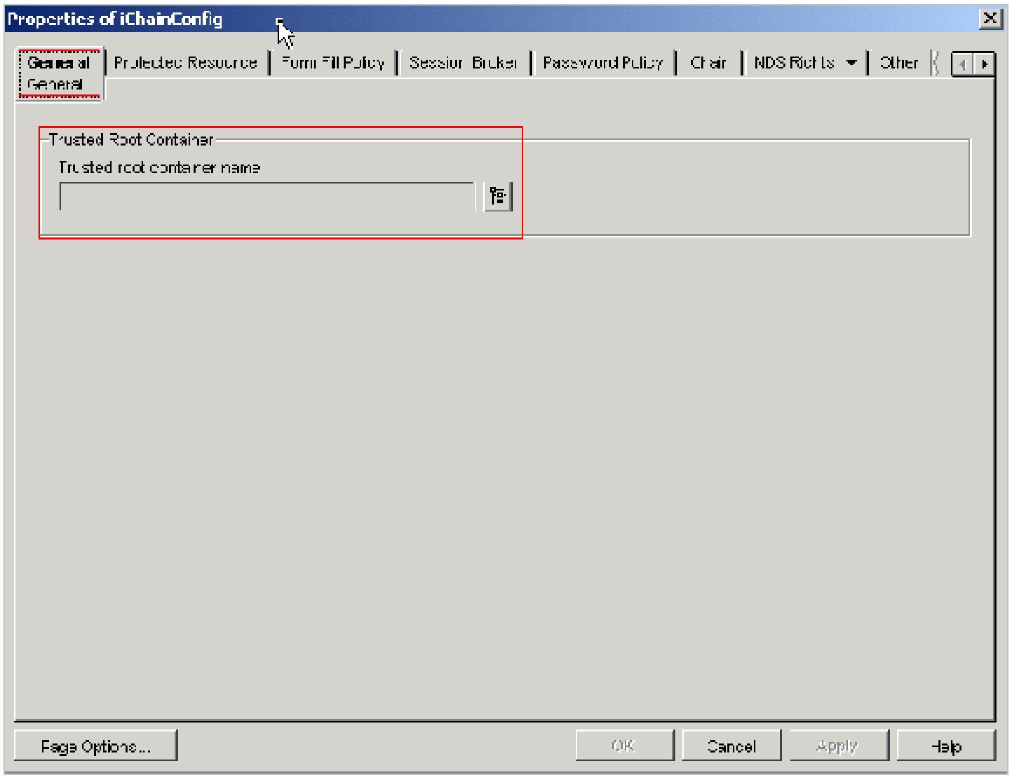

The Trusted Root Container is the eDirectory object you created and populated with certificates when setting up multiple certificate authorities. You must complete these steps before you can assign it to the proxy. For more information, see "Configuring Multi CAs" on page 87.

11.2 Protected Resource Page

The protected resource is the list of URLs defined by the iChain administrator.

Figure 11-2 Protected Resources Page

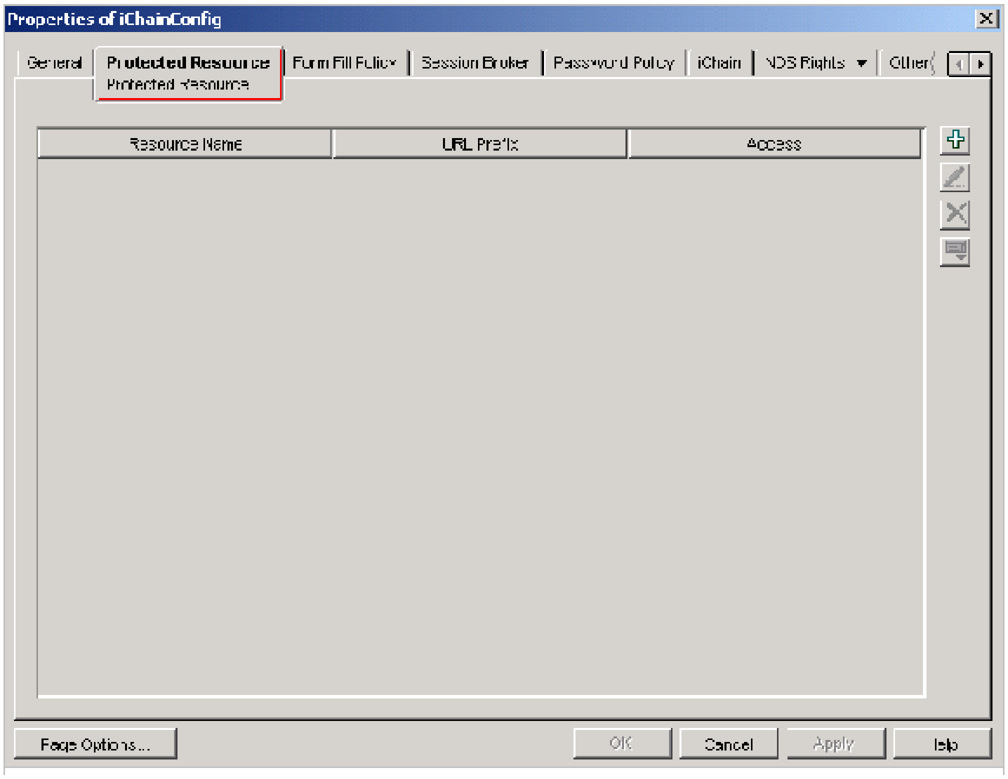

11.2.1 Types of Protected Resources

iChain provides three levels of security for protected resources:

- **Public:** No authentication or access control exists for the pages under this protected resource.
- **Restricted:** Authentication exists only for the pages under this protected resource.
- **Secure** : This is the most secured type of protected resource. To access the pages under this protected resource, the user needs to be authenticated through the proper authentication mechanism and also needs to pass through the access control. This is the default selection for the new protected resource.

In simple terms, if a user accesses a URL that has the domain name of an iChain proxy accelerator, the iChain Proxy Server checks with the ISO (found in eDirectory) to determine whether the URL is found in its URL list. Based on the type of access allowed, the proxy makes the decision of whether to grant access. For more details, see the scenarios below.

Scenarios:

Consider that the following is required:

1. http://ichain.novell.com/index.html should be accessed by all users. The index.html page is found on the Web server at the document root. The index.html page loads GIFs or bitmaps from the /images folder at the document root.

2. http://ichain.novell.com/restricted/index.html should be accessed by authenticated users. The index.html page loads Gifs or bitmaps from the /restricted/images/ folder at the document root.
3. http://ichain.novell.com/secure/index.html should be accessed by users who should be authenticated and also should have access control checking. The index.html page loads Gifs or bitmaps from the /secure/images/ folder at the document root.

Resource Name	URL Prefix	Access
Root Index	http://ichain.novell.com/index.html	Public
RootImages	http://ichain.novell.com/images/?	Public
RestrictFolder	http://ichain.novell.com/restrict/*	Restricted
SecureFolder	http://ichain.novell.com/secure/*	Secure

An accelerator is created with a DNS Domain Name of ichain.novell.com and is associated with an authentication profile (authentication is enabled). Consider that we are using an authentication profile that uses an LDAP distinguished name login mechanism. While matching the protected resource URLs, iChain looks for the most specific match in deciding URL access.

Case 1:

When the user accesses the URL, http://ichain.novell.com/index.html, he or she is able to access the page without authenticating. The images required by index.html are successfully loaded from the / images/ folder.

Case 2:

When the user accesses the URL, http://ichain.novell.com/restrict/index.html, he or she receives a login page and has to use his or her LDAP user name and password to view the page.

Case 3:

When the user accesses the URL, http://ichain.novell.com/secure/index.html, he or she receives a login page and must use his or her LDAP user name and password to view the page. If the user has an access control rule that allows him or her to access the page and its dependent resources (for example, gifs and bitmaps), then he or she can view the page. (See Chapter 12, "Using iChain Access Control Rules," on page 137 for more information.) If there is no access control rule or if the user is not allowed access to the page or its dependent resources, he or she sees a 403 Forbidden error message and cannot view the page or its dependent resources.

The above cases are summarized in the following table:

Cases	URL Accessed	Authentication Required	Access Control Required
Case I	http://ichain.novell.com/ index.html	No	No
Case II	http://ichain.novell.com/ restrict/index.html	Yes	Yes
Case III	http://ichain.novell.com/ secure/index.html	Yes	Yes

Case 4:

If the accelerator with the DNS domain as iChain.novell.com does not have any type of authentication enabled, all the pages accessed by the user would be seen without any authentication or access control:

URL Accessed	Authentication Required	Access Control Required
http://ichain.novell.com/ index.html	No	No
http://ichain.novell.com/restrict/ index.html	No	No
http://ichain.novell.com/secure/ index.html	No	No

IMPORTANT: Authentication must be enabled for the Public, Restricted, or Secure levels of security to function. If you have no authentication on the accelerator, you are using the iChain server as a caching appliance only.

As shown in the example above, iChain allows you to use wildcard characters when specifying the URL for a protected resource. If the protected resource's URL is absolute, ending with a trailing slash (/), iChain matches only the URL. However, if the protected resource's URL ends with a question mark (?), iChain matches all files in the specified folder. For example, http://ichain.novell.com/dir1/? matches all the files under the dir1 folder. If the protected resource's URL ends with an asterisk (*), iChain matches all the files under the specified folder and all the subfolders and their contents. For example, http://ichain.novell.com/dir1/* matches all files under the dir1 folder and any subfolders below dir1.

You can use wildcard (asterisk - *) folders to define protected resources. The wildcard can be used with one of more folders. The following are examples of how you could define URLs of a protected resource:

<DNS_Domain_Name>/*/public/*
<DNS_Domain_Name>/a*/public/
<DNS_Domain_Name>/a*c/public/

NOTE: Using wildcard folders could negatively impact the performance of your system. We recommend that you use this option only if it is required.

11.3 Form Fill Policy Page

This page shows the form fill policy information.

Figure 11-3 Form Fill Policy Page

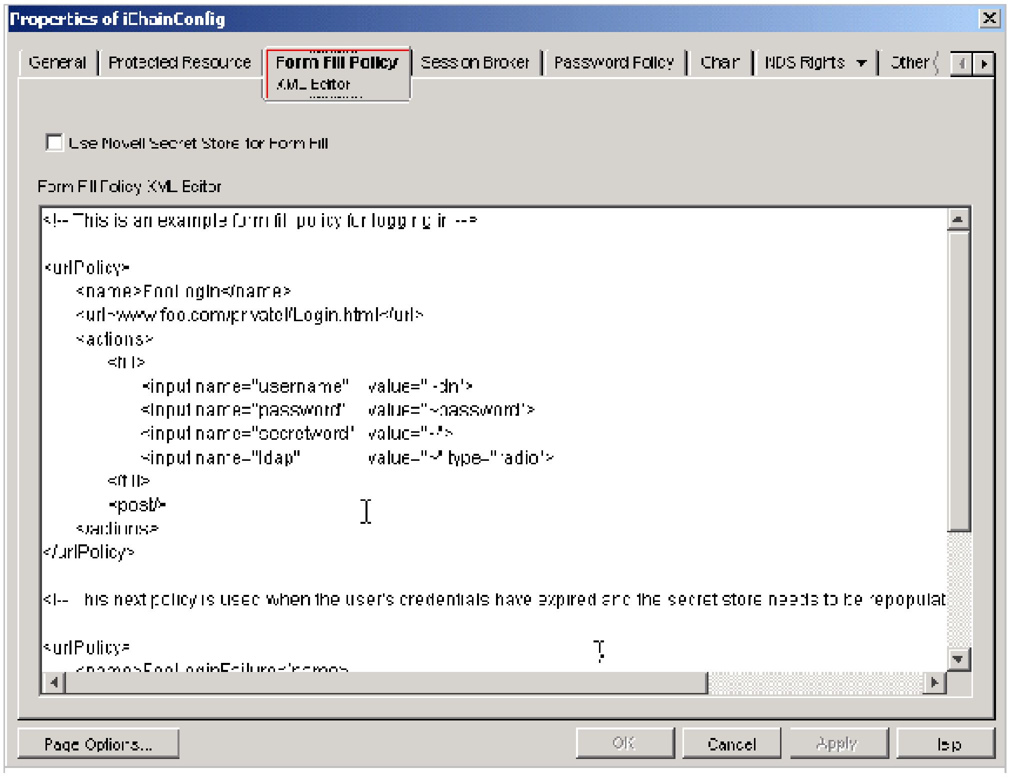

□ Use Novell Secret Store for Form Fill

Form Fill Policy XML Editor

```
<!-- This is an example form fill policy for logging in -->

<urlPolicy>
    <name>FooLogin</name>
    <url>www.foo.com/private/Login.html</url>
    <actions>
        <fill>
            <input name="username"   value="~dn">
            <input name="password"   value="~password">
            <input name="secretword" value="~">
            <input name="ldap"       value="~" type="radio">
        </fill>
        <post/>
    </actions>
</urlPolicy>

<!-- This next policy is used when the user's credentials have expired and the secret store needs to be repopulat

<urlPolicy>
    <name>FooLoginFailure</name>
```

11.4 Session Broker Page

This page shows the IP address of the primary and secondary session broker servers.

Figure 11-4 Session Broker Page

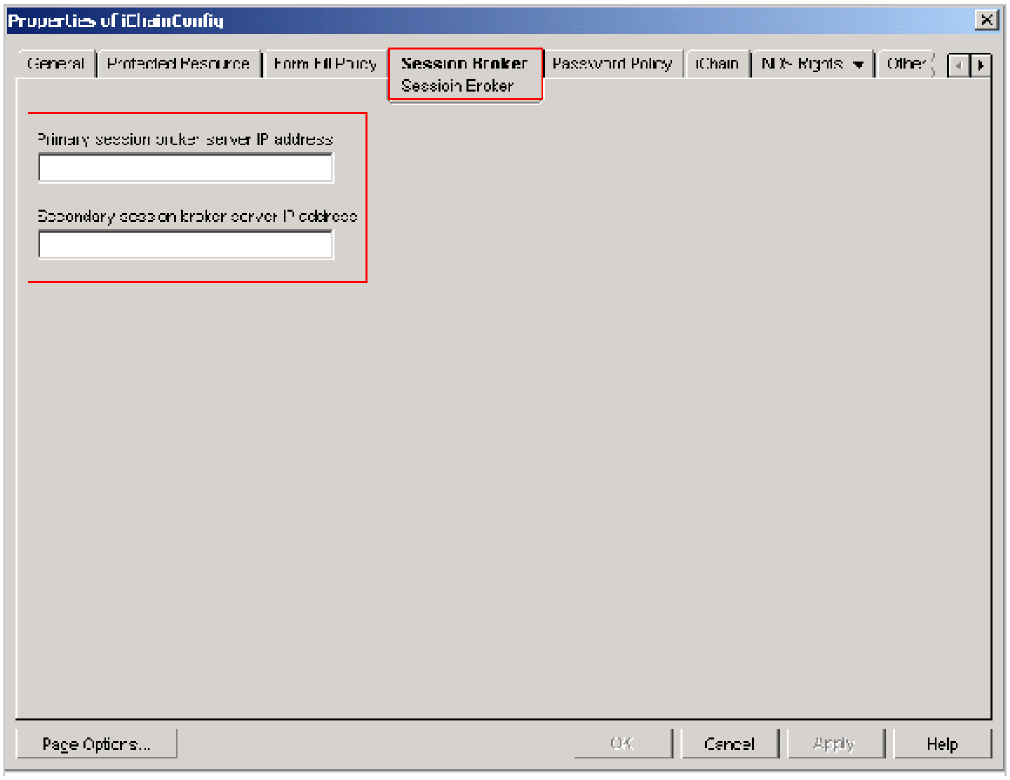

Primary session broker server IP address

Secondary session broker server IP address

11.5 iChain Page

The iChain page shows the activation information.

Figure 11-5 iChain Page

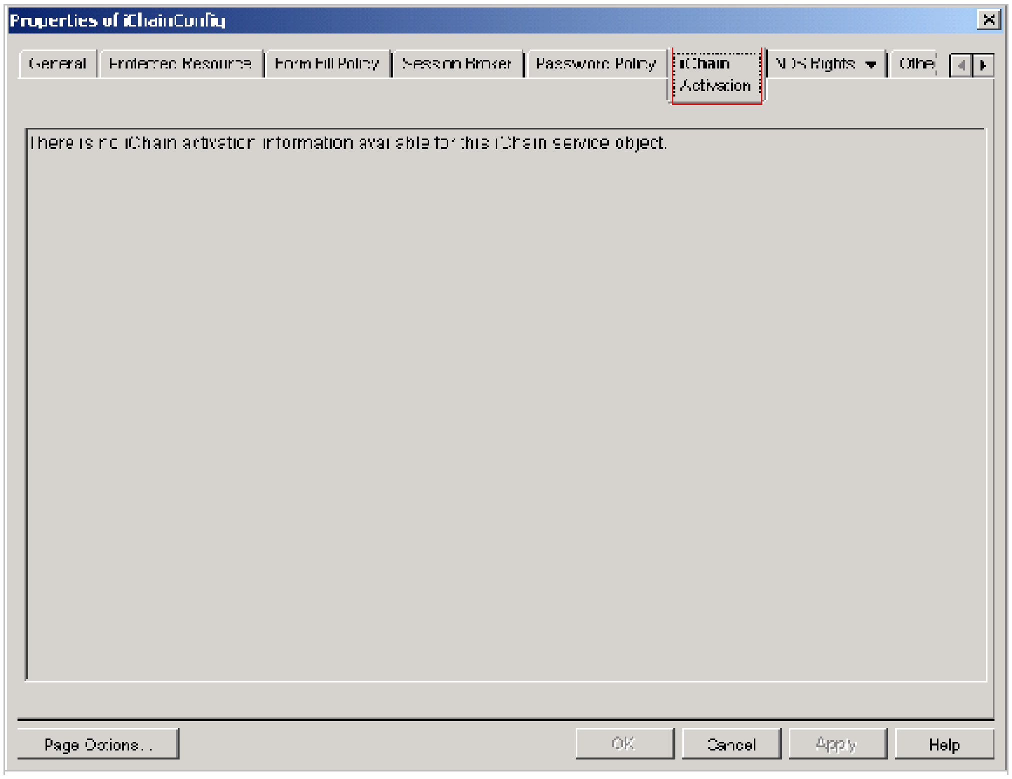

There is no iChain activation information available for this iChain service object.

Using iChain Access Control Rules

12

This section discusses the rules that control user access. The following topics are discussed:

- Section 12.1, “Defining iChain Access Control Rules,” on page 137
- Section 12.2, “Defining Dynamic Access Control Rules,” on page 145
- Section 12.3, “Advanced Access Control Configuration,” on page 149

12.1 Defining iChain Access Control Rules

After a user has logged in successfully, Access Control List (ACL) rules control what resources the user can access.

By default, the user has access to nothing. Selected users can access the resources that are explicitly listed in your ACL rules (as specified by the URL). The ACL Rule object can be applied to an organization (O), an organizational unit (OU), a Group object, or even to users listed in the Apply To list for the rule. Whenever possible, we recommend that you use the highest-level object in the list of allowed users, making it easier and faster to configure an ACL rule. Aclcheck.nlm is the module that performs the ACL checking, explained in detail throughout this chapter.

iChain® access control checks the ACL rules in the following sequence:

- Rules for all the user's containers, starting with the highest level container
- Rules for all communities of all the user's containers (if ACLCHECK was started using the /m option)
- Rules for the user's groups
- Rules for all communities of the user's groups (if ACLCHECK was started using the /m option)
- Rules for the user
- Rules for all communities of the user (if ACLCHECK was started using the /m option)

As explained in Chapter 11, “Using the iChain Service Object (ISO),” on page 129, when a user tries to access a protected resource that has been defined as Public, the user is immediately granted access. If the resource is defined as Restricted, the Novell® iChain® system checks the user's browser cookie address to see if he or she is a currently authenticated user and either lets the user access the resource if the user is authenticated or prompts the user for authentication. A current authenticated connection is all that is required. However, when a user attempts to access a URL that has been defined as Secure, the user must log in to eDirectory™ and provide a password. When the user is authenticated, the ACL rules are checked to see if the user is allowed to access the site.

ACL rules allow the use of an asterisk (*) or question mark (?) as wildcard characters when specifying URLs. The asterisk indicates that the user can have access to the folder contents and all subfolders. The question mark indicates the user can have access to the folder contents, but not the subfolders. Also, each ACL rule can be individually disabled or enabled, allowing you to turn on or off a particular rule for a time without losing its parameter settings.

ACL rules are stored in a cache that is updated periodically at a configurable interval. For performance reasons, the recommended cache refresh interval is three to six hours. If you make changes or additions to the ACL rules and want the cache to be updated immediately, use the Manual Refresh option available in the Configure > Access Control pages of the Proxy Administration Tool. If you have FTP enabled on the proxy, you can automatically refresh the iChain proxy when prompted by the snap-in.

When you create an entry in the URL list of an ACL rule, at least one of the two fields (Resource Name and URL) is required. If only the URL is specified, it must be given as an absolute URL (for example, http://www.novell.com/index.html, not /index.html). The URL can contain wildcards. The ACL rule matches any request for the URL (including wildcards). If only the Resource Name is specified, the ACL rule matches any request for the exact path of the Resource Name. For example, if the protected resource myserver has been defined as http://www.novell.com, and a URL list entry is created with myserver as the Resource Name and with no URL, then the ACL rule applies to the http://www.novell.com URL only.

If both the Resource Name and the URL are specified, the URL must be given as a relative URL (/index.html, not http://www.novell.com/index.html) and may include wildcards. The ACL rule will match requests for the combined Resource Name and URL, including wildcards. For example, if the Resource Name is myserver and the URL is /documentation/*, then the ACL rule applies to http://www.novell.com/documentation/*.

To create a new ACL rule for iChain:

1. In ConsoleOne®, select File > New > iChain Object.

   or

   Click the New iChain Object icon.
2. Select iChain Access Control Rule, then click OK.
3. Define a name for the rule, then click OK.
4. Select the rule you just created, then click Properties > Access Control.
5. Under the list of Allowed URLs, click Add. Define a name and URL for a resource that this rule will control access to.

   You can use an asterisk (*) or question mark (?) as a wildcard character when specifying URLs. The asterisk indicates that the user can have access to the folder contents and all subfolders. The question mark indicates the user can have access to the folder contents, but not the subfolders.
6. Under the Apply To List, click Add to browse to and select the Os, OUs, groups, and users to which this rule applies.

   The Os, OUs, groups, and users in the Apply to List are allowed access to the listed URLs.
7. Under the Exception List, click Add to browse to and select the Os, OUs, groups, and users that are exceptions to this rule.

   The Os, OUs, groups, and users in the Exceptions List are a subset of the Apply to List and are objects that are denied access to the listed URLs.

   WARNING: If you add the same object to both the Apply To List and the Exceptions List, results are unpredictable.

8 To enable the ACL rule, select the Enable Access Control check box on the General page.

9 To disable the ACL rule and save it for later use, deselect the Enable Access Control check box.

12.1.1 iChain Access Control Object

When viewed in ConsoleOne, the iChain Access Control Object in eDirectory has two important pages that are briefly explained below:

General Page

The General page displays the option to enable or disable a rule object. If a rule object is disabled, it is not used by ACLCHECK. This page also displays the option to enable or disable logging for this rule object.

Figure 12-1 *iChain Access Control Object: General Page*

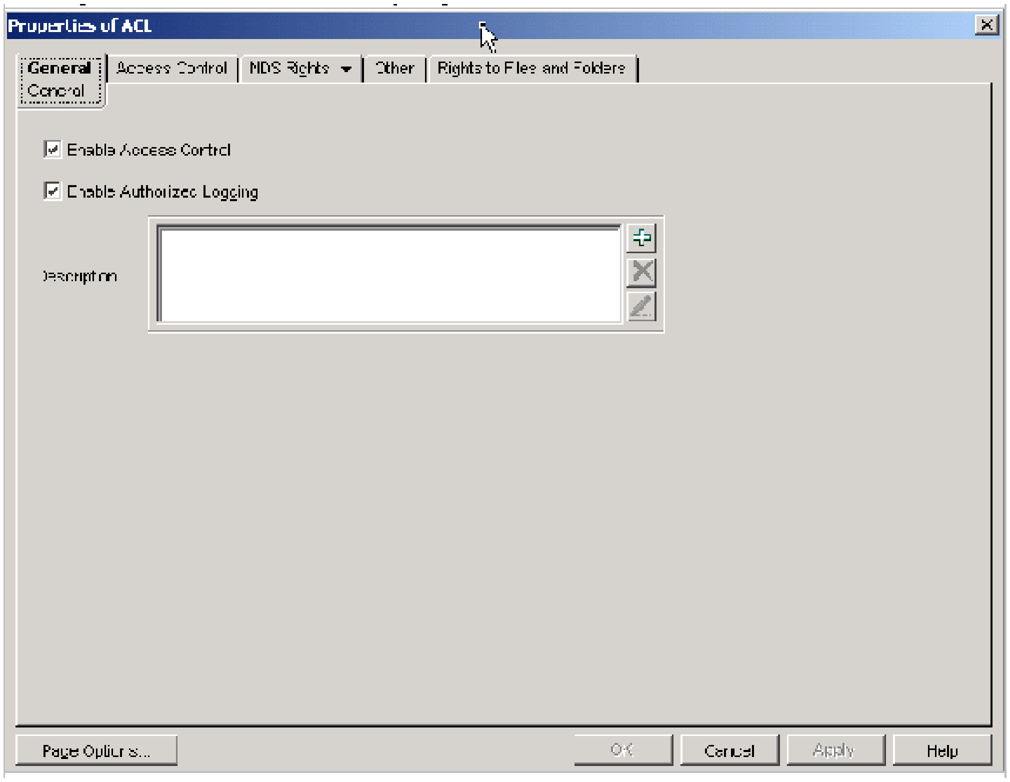

☑ Enable Access Control

☑ Enable Authorized Logging

Access Control Page

You configure the Access Control Policy on the Access Control page.

Figure 12-2 iChain Access Control Object: Access Control Page

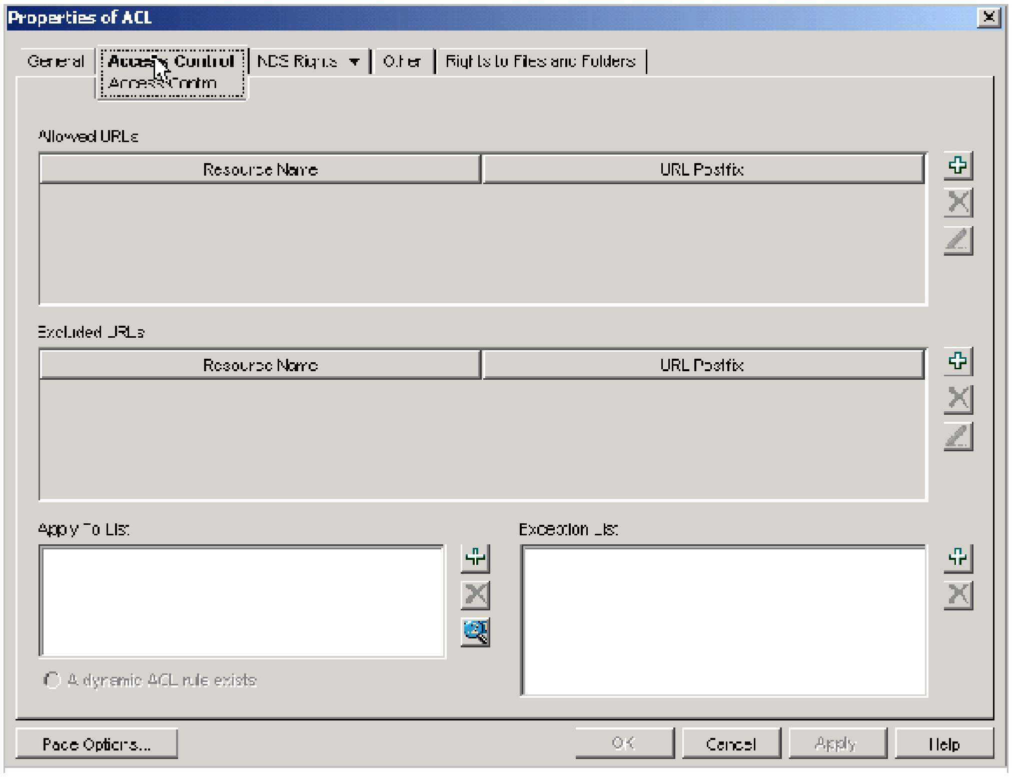

The following options are available:

- Allowed URLs: Shows the list of URLs that are allowed.
- Excluded URLs: Shows the list of URLs that are not allowed from the Allowed list.
- Apply To List: Shows the list of users who have access to the Allowed URLs, minus the Excluded URLs. You can also use this section to configure Dynamic Access Control. See Section 12.2, “Defining Dynamic Access Control Rules,” on page 145 for more information.
- Exception List: Shows the list of users and objects that are exceptions to this rule.

WARNING: If you add the same object to both the Apply To List and the Exceptions List, results are unpredictable.

12.1.2 ACL Exceptions

You can exclude certain users or group members listed in the Apply To List that you do not want to have access to the specified URLs. However, these exceptions are made on a per rule basis. So, although users might be excluded from one rule, they might still have access to the URL through other ACL rules. Double-check all ACL rules for the resource to be sure exceptions are as you expect.

You can also define a subset of the destination URL as an exception for an ACL rule. For example, an ACL rule could be set on http://ichain.novell.com/* for the users in the o=novell container. By using the URL exception feature, an administrator could define http://ichain.novell.com/private/* as a URL exception. iChain access control would then allow the users in the o=novell container to go to all the pages under http://ichain.novell.com/, except http://ichain.novell.com/private/.

12.1.3 ACL Theory of Operations

The following flow chart shows the basic operation of the ACL. (It is not meant to be an all-inclusive code translation.)

***Figure 12-3** Basic ACL Operation*

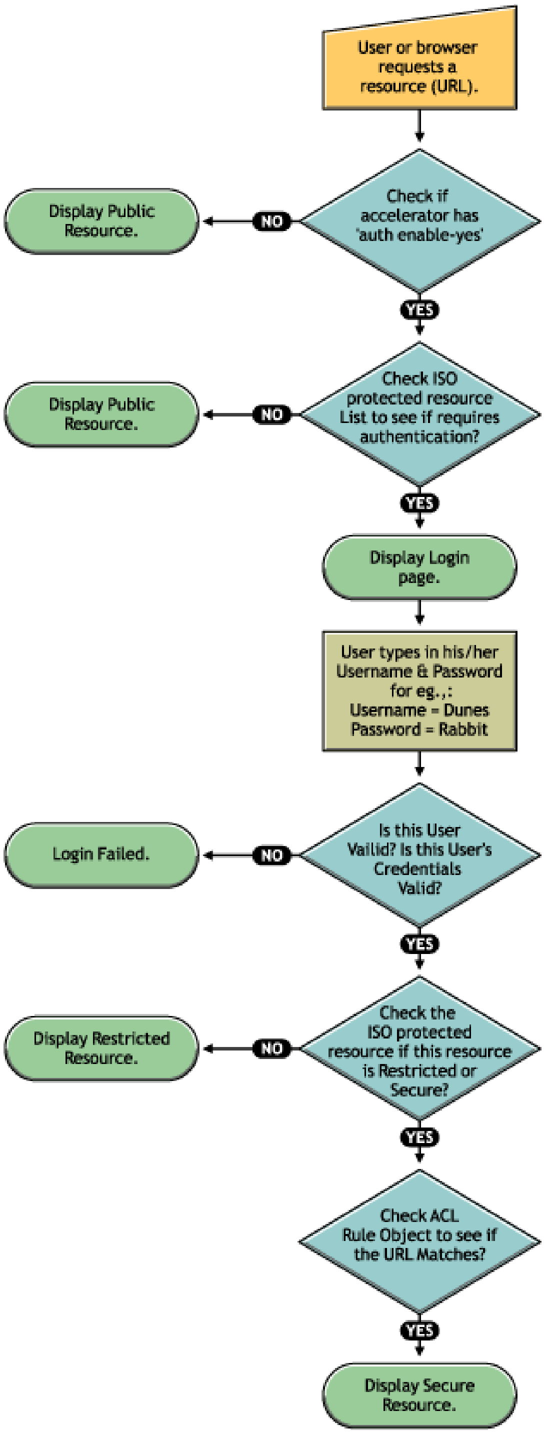
User or browser requests a resource (URL).
Check if accelerator has 'auth enable-yes'
NO
Display Public Resource.
YES
Check ISO protected resource List to see if requires authentication?
NO
Display Public Resource.
YES
Display Login page.
User types in his/her Username & Password for eg.,: Username = Dunes Password = Rabbit
Is this User Vailid? Is this User's Credentials Valid?
NO
Login Failed.
YES
Check the ISO protected resource if this resource is Restricted or Secure?
NO
Display Restricted Resource.
YES
Check ACL Rule Object to see if the URL Matches?
YES
Display Secure Resource.

NOTE: If the Authorization Server is using one LDAP source and the Access Control is using another, the ISO object (if it exists) in the Authentication tree is ignored and the ISO object in the Access Control tree is utilized.

The example below explains the process of ACLs. The example of the protected resource as explained in Chapter 11, “Using the iChain Service Object (ISO),” on page 129 is used in this example as follows:

SN	Resource Name	URL Prefix	Access
A	RootIndex	http://ichain.novell.com/index.html	Public
B	RootImages	http://ichain.novell.com/images/?	Public
C	RestrictFolder	http://ichain.novell.com/restrict/*	Secure
D	SecureFolder	http://ichain.novell.com/secure/*	Secure

An ACL Rule Object called ACL_Rule_One is created with the following details:

SN	Applied To URLs		Excluded URLs	
	Resource Name	**URL Prefix**	**Resource Name**	**URL Prefix**
E	SecureFolder	/index.html		
F	SecureFolder	/images/?		
G	SecureFolder	/folder1/*		
H			SecureFolder	/folder1/folder2/?
I	SecureFolder	http://ichain.novell.com/secure/folder1/folder2/folder3/*		

SN	Applied To	Exclusion List

J	OU=Permanent, OU=Users, O=Company
	CN=Group1, OU=Users, O=Company	

K		CN=Fuser456, OU=Permanent, OU=Users, O=Company

Use the flow chart and tables above to understand the following three cases:

Case 1

1. At the browser, Jack (who is a member of OU=Permanent, OU=Users, O=Company) enters the URL as http://www.novell.com/secure/index.html.
2. DNS resolves Jack's browser to the iChain machine.

3. The iChain box has a Web Accelerator with www.ichain.novell.com defined and the Enable Authentication switch is turned on. (If the switch was not turned on, iChain would simply cache the resource and would provide no security).
4. The ISO entries are compared to the URL request to determine if authentication is required. There is a match (D).
5. Jack is asked to enter his name and password for authentication. The username and password are checked for validity.
6. The URL Jack requested is verified against the ISO Protected Resource to identify whether it is Restricted or a Secure Resource. The Resource D matches the URL Jack requested, and it is secure.
7. ACL checking takes place. The ACL_Rule_One object's information is verified to see whether user=Jack has access for the URL he is requesting. The /index.html is found in the list in E. User=Jack is found in the container OU=Permanent,OU=Users,O=Comapny (J) and is not found in the exception list (K). Thus, user=Jack is given access to index.html. Index.html loads the images from the /secure/images/ folder. It is also checks to verify that the images from this location are allowed to be loaded for user=Jack. The images that are in the folder /secure/images/ are allowed to load (F). If there are any images that are referenced from /secure/images/folder1/, they are not loaded. A 403 Forbidden error would result for this type of request.

Jack is granted access as summarized in the following table:

NOTE: The most specific match, based on the URL, takes precedence. The ISO entry (D) http://ichain.novell.com/secure/* is essentially cut and pasted, then concatenated from the URL postfix of /index.html to form http://ichain.novell.com/secure/index.html.

User	URL Accessed	Status
Jack	http://ichain.novell.com/secure/index.html	Allowed
Jack	http://ichain.novell.com/index.html	Allowed
Jack	http://ichain.novell.com/restrict/index.html	Allowed
Jack	http://ichain.novell.com/secure/folder1/index.html	Allowed
Jack	http://ichain.novell.com/secure/folder1/folder2/index.html	Denied
Jack	http://ichain.novell.com/secure/folder1/folder2/folder3/index.html	Allowed

Case 2

User=Jane is a member of Group1 group object which is under OU=users,O=company; however, her user object resides under OU=TempUsers,O=Users,O=Company. As illustrated in Case 1, Jane would experience the following, depending on the URLs she attempts to access:

User	URL Accessed	Status
Jane	http://ichain.novell.com/secure/folder1/index.html	Allowed
Jane	http://ichain.novell.com/secure/folder1/folder2/index.html	Denied
Jane	http://ichain.novell.com/secure/folder1/folder2/folder3/index.html	Allowed

She is denied access to /secure/folder1/folder2/index.html because it is part of the Excluded URLs (H). She is allowed access to /secure/folder1/folder2/folder3/index.html because this is part of the Allowed URLs (I).

Case 3

As in Case 1, user=fuser456 (a member of OU=users,O=company) would experience the following, depending on which URLs she accesses:

User	URL Accessed	Status
Fuser456	http://ichain.novell.com/index.html	Allowed
Fuser456	http://ichain.novell.com/restrict/index.html	Allowed
Fuser456	http://ichain.novell.com/restrict/index.html	Allowed
Fuser456	http://ichain.novell.com/secure/index.html	Denied
Fuser456	http://ichain.novell.com/secure/folder1/folder2/index.html	Denied
Fuser456	http://ichain.novell.com/secure/folder1/folder2/folder3/ index.html	Denied

Fuser456 is denied access to all of the secure resources since this user is on the Exception List (K).

12.2 Defining Dynamic Access Control Rules

Dynamic Access Control Rules allow an administrator to set up an access control rule based on a query of user attributes. (If a user's attribute value satisfies a predefined value, the rule would be applied to that particular user.) This type of query can be based on a user's attributes (the user's location, salary, hobby, etc.). For example, an administrator could configure a rule that says “Apply this rule to all the users who are in San Jose and have a salary greater than (>) $50,000.”

Dynamic Access Control Rules are based on the following principles:

- A rule is applied to the users who satisfy a condition (query) provided by the administrator.
- Static applicability of the rules (object-based ACLs) will still exist in its current form. The DN exception and URL exception are also supported with the new type of rules. An ACL rule can contain both a dynamic query and a static “Apply To” list. Thus, one ACL rule can act as a dynamic rule and also as a traditional (static) rule.
- Because user attribute values can be dynamic, the administrator might want to limit the time that a rule can be cached in memory. This cache time, called “time to live,” can be set in number of minutes and the cache for that rule will be expired after the specified duration elapses. The Time to Live value can only be set for dynamic ACLs. If no value is entered in the Time to Live field, the cache will not expire until the overall cache refresh takes place for ACLCHECK. If the Time to Live value is set at zero, the rule is not cached.
- The administrator who creates a dynamic ACL must have write privileges on the ISO. (This is not required for non-dynamic ACLs.)

12.2.1 Setting Up Dynamic Access Control Rules

The dynamic ACL setup GUI allows you to create and test the search filter, then convert it to standard LDAP format to be stored in eDirectory. This query is later used to allow or disallow dynamic access control in iChain.

Dynamic ACLs can be defined when using the iChain Server Web Accelerator Wizard or the iChain Access Control snap-ins.

To create a dynamic ACL:

1 In ConsoleOne, open the ACL Rule object.

2 On the Access Control page click the Dynamic ACL button (represented as a magnifying glass) next to the Apply To dialog box.

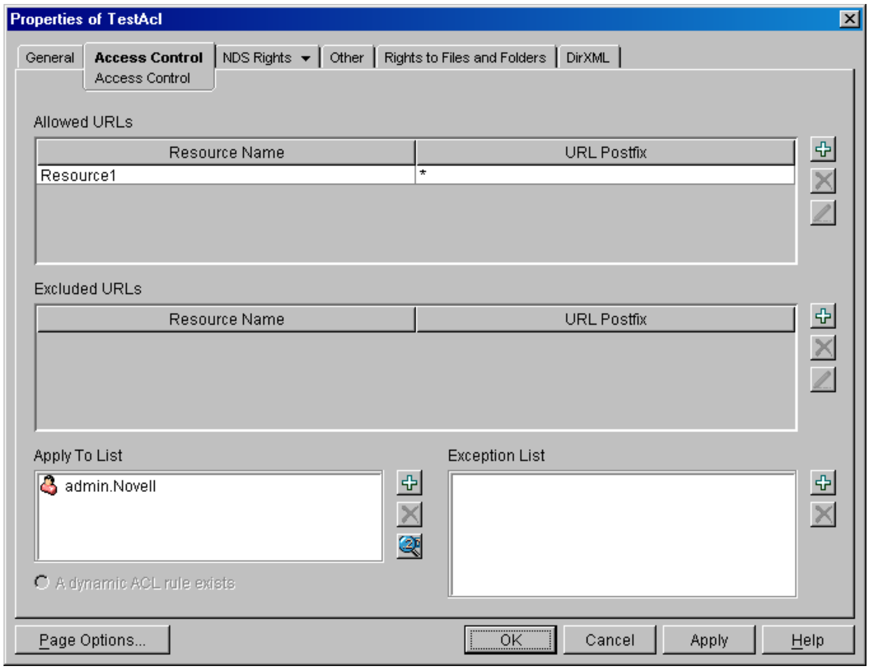

Allowed URLs

Resource Name	URL Postfix
Resource1	*

Excluded URLs

Resource Name	URL Postfix

Apply To List

admin.Novell

○ A dynamic ACL rule exists

Exception List

This opens the Dynamic ACL Query Setup dialog box.

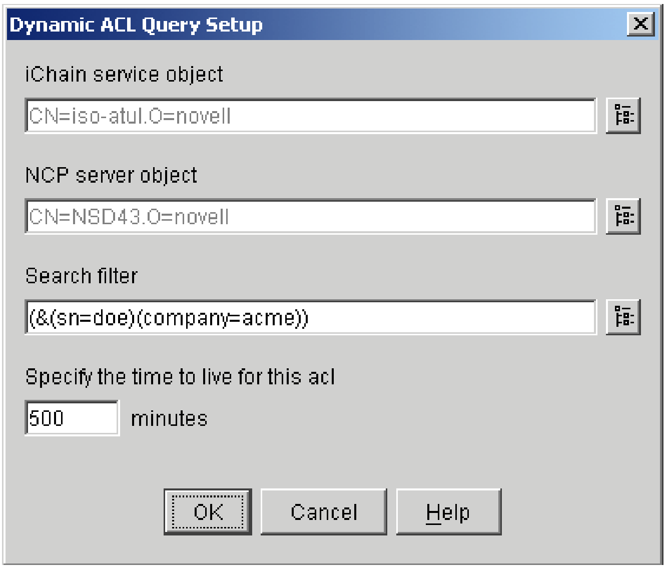
Dynamic ACL Query Setup

iChain service object: CN=iso-atul.O=novell

NCP server object: CN=NSD43.O=novell

Search filter: (&(sn=doe)(company=acme))

Specify the time to live for this acl: 500 minutes

3 Browse and select the active ISO and NCP Server object, if not pre-populated, which points to the appropriate LDAP group object having the updated eDirectory to LDAP attribute and class mappings.

4 Either update the existing LDAP Search filter field to form the query or click the Advanced Dynamic ACL query setup button located at the right of the Search filter field.

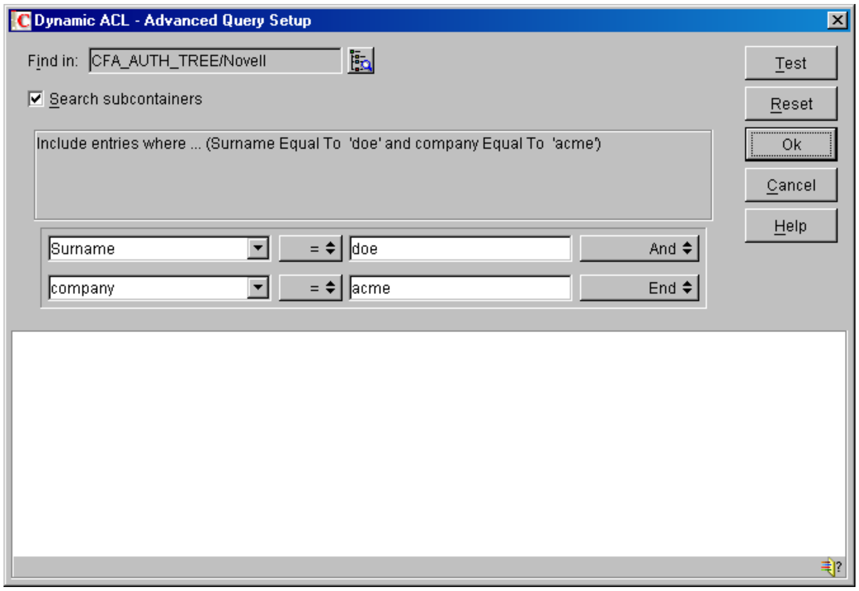
Dynamic ACL - Advanced Query Setup

Find in: CFA_AUTH_TREE/Novell

☑ Search subcontainers

Include entries where ... (Surname Equal To 'doe' and company Equal To 'acme')

Surname	=	doe	And
company	=	acme	End

5 Specify the time to live (in minutes) in cache for the dynamic ACL.

NOTE: This time to live is meant to keep track of changes in user's attributes rather than changes in the ACL rule value. If you delete an ACL rule, the user still has access to that resource. Dynamic ACL rules are read at startup time or at the cache refresh time, so this dynamic ACL rule is still cached until you perform an aclcheck refresh.

6 If you choose to build and test the query in the Advanced Dynamic ACL panel, the Find In field allows you to specify the container location to start searching for the objects.

7 The first query field allows you to select an object property to use as a search criterion, or select [Object Type] to search for the objects of a specific class.

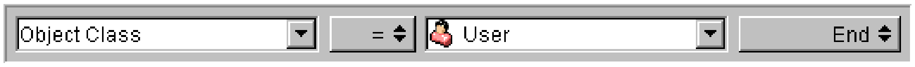

NOTE: The [Object Type] is required only to test this query in this panel. After the query is formed and tested, you should remove this search criteria before saving it because iChain currently allows dynamic access control on user objects only.

Click the comparison operator to select a logical operator to use when comparing the value of the specified attribute with the actual attribute value in the Authorization Server (eDirectory). For more information about the Authorization Server, see "Installing iChain Services Schema Extensions on the iChain Authorization Server" on page 27.

The second query field allows you to specify the attribute value to compare against the actual attribute value in the Authorization Server (eDirectory). The syntax is the type of data contained in the attribute, such as character string, or integer.

Click a statement connector keyword to select the keyword that specifies how this statement connects with the next statement or group of statements in the query. "And" specifies that both this and the next statement (or statement group) must be true for a match to occur. If there is no next statement, selecting this keyword adds one. "Or" specifies that either this or the next statement (or statement group) must be true for a match to occur. If there is no next statement, selecting this keyword adds one. "Insert Row" adds a new statement below this one. "Delete Row" deletes the statement from the query. "New Group" adds a new statement group below this one. "End" specifies that this statement ends the query.

NOTE: When creating a query, there are some limitations with the query tool. It does not allow a mixture of "and/or" conditions in the same group or between groups.

8 Click the Test button to test the query in the eDirectory namespace and display the results at the bottom of the dialog box.

You can right-click objects in the result list to perform actions as you do in the ConsoleOne right pane.

9 If you don't want to keep the query, click the Cancel button to discard the setup and exit the panel.

10 If you want to keep the query, click OK to convert the query (search criteria) to standard LDAP search filter format to be stored in the ACL object. Make sure your eDirectory to LDAP mappings are present in the right LDAP group object as pointed by the NCP Server object. This closes the Advanced window and allows you to edit the formed LDAP search filter before saving it along with its time-to-live-in-cache attribute.

IMPORTANT: It is important to understand the differences and associated limitations while testing your query. Query testing is done in the eDirectory namespace and stored in LDAP format. Thus, there may be situations where you need to specify the value in the second query field in eDirectory format while testing, but change it to LDAP format before saving it. A rule built with comma-delimited LDAP format (for example, commerceBudgetHolder = cn=CostCenter03,o=novell,c=us) fails to locate user objects matching the specified attribute when using the Test button, but works properly to allow or deny user access through

ACLCHECK. Changing the format to use eDirectory dot-delimited format (for example, commerceBudgetHolder = cn=CostCenter03.o=novell.c=us) allows the Test button to work as expected so that the rule can be verified. Remember to change the rule back to the comma-delimited LDAP format when saving so that ACLCHECK functions as expected.

12.3 Advanced Access Control Configuration

This section contains information about the following topics:

- "Enabling ACL Rule Checking for Community Objects" on page 149
- "Enabling Debugging Messages for Access Control" on page 149
- "Using ACLCHECK options" on page 150

12.3.1 Enabling ACL Rule Checking for Community Objects

iChain, by default, does not check community objects for ACL rules. Community objects existed in previous versions of iChain but are no longer provided; however, the functionality is provided to allow the use of pre-existing community objects.

To enable ACL rule checking for community objects:

1. Unlock the console.
2. Edit the appstart.ncf.
3. Change the load aclcheck entry to load aclcheck /m.
4. Restart the machine.

After specifying changes in the configuration, ACL rules are checked in the following sequence:

- OUs
- OUs' communities
- Groups
- Groups' communities
- User
- User's communities

If a specified option is not provided, checking for the italicized portions of the above list will not be performed for checking the ACL rules.

12.3.2 Enabling Debugging Messages for Access Control

The module that provides iChain's Access Control (aclcheck.nlm) can be configured to output debug information. The administrator can choose one of two levels of increasingly more detailed information. This information can be helpful to developers and consultants. By default, no debug information is output. To enable these debugging options, you can use either of two procedures:

To use the command line:

1. Use the command line option `ACLCHECK /D2` to temporarily enable the debug output (until the restart is performed or until the `/D0` command is issued to disable debug).

To use a configuration file:

1. Edit the appstart.ncf file on the iChain Proxy Server.
2. Find the line containing the `LOAD ACLCHECK` command and add a debug level switch at the end of that line, for example, `LOAD ACLCHECK /D2`.

   Enabling the /D2 option can impact performance and should only be used for troubleshooting aclcheck issues.
3. Shut down and restart the proxy server.

12.3.3 Using ACLCHECK options

You can configure the options of the ACLCHECK NLM from the System Console. These options are not case sensitive. When you change an ACLCHECK option, the update is stored in the appstart.ncf file. Use the following syntax to change an option:

```
aclcheck /<option>
```

To set multiple options at the same time, separate them with a space. For example:

```
aclcheck /<option> /<option>
```

Replace <option> with one of the options in the table below:

Option	Explanation	Example
/h or /?	Displays help information about the ACLCHECK utility.	`aclcheck /h`
/d<level>	Specifies the level of debug information. This information can be helpful to developers and consultants. Set the level at 1 or 2 for more detailed information. Default: 0	`aclcheck /d2`
/f<minutes>	Specifies, in minutes, how frequently the cache is refreshed. Keep this number higher if you are not likely to change DS information quickly. This can improve performance because ACLCHECK does not need to throw away the already built-up cache. Default: 180 minutes	`aclcheck /f300`
/g<number>	Determines whether dynamic group processing is enabled: 1 enables dynamic group processing 0 disables dynamic group processing You should disable this option if you do not have any dynamic groups or you are not using them with iChain. Default: 0	`aclcheck /g1`

Option	Explanation	Example
/k<number>	Determines whether iChain should try three times to read the membership of dynamic groups: 1 enables three retries 0 disables the retries Do not enable this option unless your users are experiencing membership-not-found errors, because this option slows down the system. For more information about this issue, see TID 10097124 (http://support.novell.com/cgi-bin/search/searchtid.cgi?10097124.htm). Default: 0	`aclcheck /k1`
/o<seconds>	Specifies, in seconds, when the cache for dynamic groups is refreshed. Certain values have specific meanings: 0 disables dynamic group caching -1 causes the cache to be refreshed when a user logs out or the idle count of the user times out Dynamic groups are cached when the system boots, and that cache is refreshed on every ACL check refresh. If you add or delete a dynamic group, this cache needs to be refreshed. If you change the membership of a dynamic group and want the change to be immediate, this cache needs to be refreshed. If you are not using dynamic groups, you should disable dynamic group caching. Default: 300 seconds (5 minutes)	`aclcheck /o360`
/p	Allows you to cancel the repeated display of IP address resolution error messages such as "Get IP addr failed for hostname: host.company.com," where host.company.com has a secured protected resource. 1 enables the display of these messages 0 disables the display of these messages Default: 0	`aclcheck /p1`
/q	By default, dynamic ACLs are checked after checking all traditional (static) ACLs. If this option is specified, ACLCHECK first checks for dynamic ACLs. This option should be used when you have mainly dynamic ACLs.	`aclcheck /q`

Option	Explanation	Example
/s<size>	Specifies, in kilobytes, the maximum size of a log file. Default: 1 MB **NOTE:** If you set this parameter to 7 KB or less, the logs files are not created.	`aclcheck /s2000`
/t<seconds>	Specifies the number of seconds for the semaphore timeout. Default: 10 seconds	`aclcheck /t9`
/v	Causes iChain to first verify that the DN in the ACL exists in the directory before checking rights.	`aclcheck /v`
/w<seconds>	Specifies, in seconds, how long an LDAP search can remain outstanding before timing out. The minimum value is 10 seconds. If set too high, this option can slow down the system. Default: 10 seconds	`aclcheck /w10`
/z	Determines whether information is sent to the LDAP pool screen: 1 enables the display of these messages 0 disables the display of these messages Default: 0	`aclcheck  /z1`

There are known timing issues with ACLCHECK. If the parameters are loaded via the appstart.ncf file, sometimes the parameters are not loaded. If you experience this problem, manually change the location of the load aclcheck string in the appstart.ncf file so that it loads before BRDSRV. For more information, see Cannot change ACLCHECK refresh interval with iChain 2.3 (http://support.novell.com/docs/Tids/Solutions/10100146.html)

Using Object-Level Access Control 13

The Novell® iChain® service enables you to integrate and allow access to Web-based applications. Sometimes these resources or objects need additional access control or application information about the user to be passed into the application. This additional information about the user can be stored in Novell® eDirectory® or some other database. Within iChain, these resources are called Protected Resources and access to them is set up through the Protected Resources page of the iChain Service object. Refer to Section 3.2.1, "Creating an iChain Service Object," on page 50 for basic setup information.

This section provides information on the following:

- Section 13.1, "Setting Up Object-Level Access Control," on page 153
- Section 13.2, "Using OLAC Custom Header Variables," on page 156
- Section 13.3, "Customizing the Authorization Header," on page 156
- Section 13.4, "Using OLAC Caching," on page 157
- Section 13.5, "Using OLAC for Shared Secrets," on page 157
- Section 13.6, "The Effects of Disabling OLAC," on page 159

13.1 Setting Up Object-Level Access Control

To access protected resources, a special iChain Object-Level Access Control (OLAC) plug-in (an LDAP plug-in) is available to access the database and retrieve the additional information. By default this plug-in allows you to define attributes in the LDAP datastore that are embedded and passed within the HTTP request header or as a query string. You can assign a name as the tag to the data.

When OLAC is configured to use a multivalued LDAP attribute, the values of that attribute are returned from the LDAP query as a comma-delimited list and forwarded by OLAC to the Web server in the same format. Some back-end Web applications might not be able to process this comma-delimited value. The LDAP cn attribute is actually a multivalued attribute. (In ConsoleOne®, the values under the user object's Other Name: field on the General > Identification page are actually stored as part of the multivalued cn.) To configure OLAC to send only the user's common name (for example, user1), even when the cn attribute has multiple values, specify the LDAP attribute uid instead of cn in the OLAC configuration.

iChain also supports additional plug-ins called CONSTANT, SECRETSTORE, and INTERNAL. The CONSTANT plug-in allows you to pass the same constant literal with every OLAC request. This is particularly valuable when an application requires a constant to be passed and the administrator does not want to include the constant in each user object (for easier setup and maintenance).

The following table lists the LDAP and CONSTANT plug-ins' corresponding entries for the Data Source and Value fields in ConsoleOne.

Plug-In	Description	Data Source	Value
LDAP	Adds user attributes from a directory with LDAP support.	ldap (case insensitive)	Any LDAP user attribute (for example, surname, givenName).
CONSTANT	Adds the constant literal for every OLAC request, where defined.	constant (case insensitive)	Constant Literal (for example, string123).

The INTERNAL OLAC data source obtains user information that is available in the proxy. This allows the login query string to be passed to the Web server. It displays content based on login information. The following table lists the OLAC values and corresponding entries for the INTERNAL data source.

OLAC Value	Sample OLAC Name	Description
AuthProfiles	allAuthProfiles	Build a tag-value pair for all of the authentication profiles used to authenticate the user. For example, if LDAP1 and RADIUS3 were both used, the OLAC string that is generated is allAuthProfiles=LDAP1,RADIUS3.
LDAPProfile	myLDAP	Build a tag-value pair where the value is the name of the LDAP authentication profile if an LDAP profile was used to authenticate the user. For example, myLDAP=LDAP1. If the user authenticated with just RADIUS, then the tag-value pair is myLDAP=.
RADIUSProfile		Lists the RADIUS authentication profile used to authenticate or no value.
MutualProfile		Lists the Mutual authentication profile used to authenticate, or no value.

The OLAC Parameters dialog box is shown below:

Figure 13-1 *OLAC Parameters*

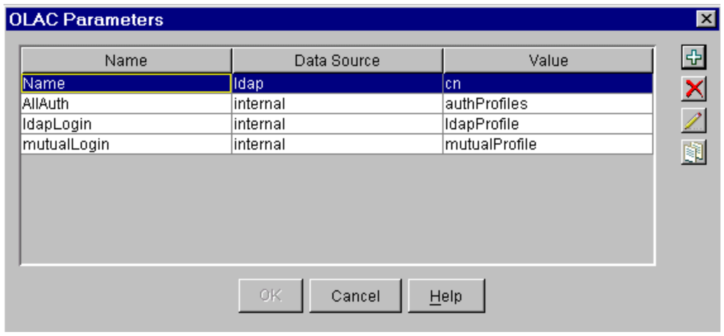

Name	Data Source	Value
Name	ldap	cn
AllAuth	internal	authProfiles
ldapLogin	internal	ldapProfile
mutualLogin	internal	mutualProfile

Because the LDAP plug-in is based on iChain APIs, you can customize iChain and create OLAC plug-ins to integrate your applications as needed. For more information about the APIs for customizing your iChain infrastructure, see the Novell appnote, *Developing a Custom OLAC Driver* (http://www.novell.com/coolsolutions/appnote/2544.html).

NOTE: Only administrators familiar with programming principles and Java programming syntax should attempt to customize OLAC plug-ins.

The settings for the OLAC Frameworks and its plug-ins are stored in the iChain Access Control profile and the oac.properties file, which is typically found in the sys:/ichain/oac directory on the iChain Proxy Server. The configuration file contains a section for the framework as well as one for the plug-in. The following table lists the valid OLAC options for each section:

Name	Description	Required?	Default Value
Object-Level Access ControlOptions [OAC] section			
Security Authentication	The method to use when authenticating to the LDAP server. Currently, only “simple” is supported.	No	simple
Server Port	The port on which the OLAC framework listens for lookup requests from the proxy server.	No	4444
Worker Count	The number of worker threads to create.	No	10
Refresh Time	The number of minutes after which the OLAC configuration is re-read from the ISO.	No	180
Value Delimiter	The delimiter used to separate multiple values assigned to the same name in the URL query string. For example, if Value Delimiter is specified as semi-colon (;), the resulting query string might look like COLORS=blue;green; yellow;orange&SHAPES=circle;square;triangle.	No	, (comma character)
LDAP Plug-In Options [LDAP Processor] Section			
Security Authentication	The method to use when authenticating to the LDAP server. Currently, only “simple” is supported.	No	simple
Class Name	The name of the class implementing the LDAP plug-in. Must be com.novell.ichain.oac. ldap.ParamListBuilder.	Yes	None

13.2 Using OLAC Custom Header Variables

You can pass the OLAC parameters defined for a protected resource as a part of the HTTP header itself. Thus, if needed, you can forward more of the OLAC parameters/values to the Web or application servers than would be possible by sending them as a part of the query string.You can specify if the OLAC parameters needed to be passed either as a part of the query string or header variables while defining the protected resource in ConsoleOne (as shown by radio buttons at the Add New Protected Resource dialog).OLAC parameters in the header variables must be taken care of while parsing and are specified in the following format:

```
x-<olac-param name>: <olac-param value>
```

NOTE: The letter x- is prefixed for all of the OLAC parameters (custom variables) in the header. This is done to negate or minimize name collisions between OLAC and non-OLAC parameters residing in the header.

The *<olac-param name>* cannot contain any extended characters.

When sending an attribute value from eDirectory through OLAC using LDAP, the total length of the query string should not exceed 255 characters. Exceeding this number of characters might cause the authentication header to fail to appear for some Web browsers/servers (or appears with the username and password instead of the attribute that was mapped by the ICHAIN_UID parameter). The HTTP headers might also fail to appear.

13.3 Customizing the Authorization Header

Using the iChain OLAC feature, you can customize the authorization header as described below.

By default, iChain puts the fully qualified distinguished name and user-entered password in the authorization header. However, administrators might want to change the content of this header by changing the value of the username or password. This customization might be required because some Web servers (Microsoft IIS) require the common name (CN) instead of DN as the username. Using OLAC, you can customize the authorization header.

For a particular protected resource, you can define special OLAC parameters, such as ICHAIN_UID and ICHAIN_PWD, to change the values of the authorization header. The values returned by OLAC will be placed in the authorization header as username and password, respectively.

For example, if you define the ICHAIN_UID=CN and ICHAIN_PWD=SSN OLAC parameters for a protected resource, OLAC returns the values of CN and SSN attributes of the logged-in user. iChain uses these values as the username and password to construct the authorization header and sends it to the Web server.

Both of these parameters are optional. If only one parameter is defined, such as ICHAIN_UID=CN, the other parameter value is filled with default behavior, such as a password, provided by the user via Forward Authorization.

IMPORTANT: If you have defined these special parameters and OLAC is not enabled or the value of the given attribute is NULL, iChain passes NULL in the authentication header.

If you define values for ICHAIN_UID or ICHAIN_PWD, the value you set for the Forward Authorization Information to Web Server option is ignored. A basic authentication header is always sent. For more information about this option, see "Add Authentication Profiles Dialog Box" on page 305.

13.4 Using OLAC Caching

OLAC has two levels of caching: The first level of OLAC caching occurs in OLAC server (OACJAVA), which caches the protected resource name and its associated OLAC parameters, data source, and value (the name of the attribute in case of LDAP) initially when the OLAC server is started and/or when it is issued an OACREFRESH command.

You can refresh OLAC from the Web GUI (on the Access Control page) or you can run OACREFRESH from the NetWare® prompt (not the iChain prompt) on the iChain console. OACREFRESH is a Java program that opens a socket to the OLAC server and sends the REFRESH command. After receiving the command, the OLAC server refreshes the cache. OLAC is also refreshed when the administrator opts to refresh the iChain Proxy configuration in the ISO snap-in in ConsoleOne after making changes. OLAC also auto-refreshes every *n* minutes, where *n* is specified in the oac.properties file. This value currently defaults to 180 minutes.

The second level of OLAC caching is for OLAC values cached in proxy for that user session, and it cannot be refreshed by any command. The life of this cache is only for that authenticated user session. OLAC refresh gets the changes in the protected resources and associated parameter definitions. It does not refresh any LDAP connections. The administrator must restart OLAC for those changes to take effect. Restarting OLAC can be done easily from the Access Control page in the Web GUI.

13.5 Using OLAC for Shared Secrets

More Web service applications are using single sign-on or are sharing some application information. Administrators might want to preconfigure user credentials so the user never sees what his/her credentials are for various applications. Novell products including Novell Portal Services, Novell Secure Login, and iChain OLAC might share login credentials (user name and password).

13.5.1 How Shared Secrets Works

Shared Secrets is a new feature that allows the sharing of a user's secret credentials across applications. Administrators can register users' credentials, and can only create and/or overwrite the credentials. The same as with a password, administrators cannot read and/or retrieve users' credentials.

The shared secret is identified by <type>:<name>, where <type> is the type of shared secret and <name> is the name of the shared secret.

NOTE: In order to use Shared Secrets, the administrator must install Novell SecretStore® version 3.*x* on the iChain Authorization Server.

13.5.2 Defining OLAC for Shared Secrets

The OLAC plug-in for Shared Secrets supports defining OLAC parameters at the level of the key name of the SS_CredSet shared secret. The datasource name is SECRETSTORE.

For example, this can be defined in the OLAC configuration screen in ConsoleOne as described in the following table:

OLAC Parameter Name	Data Source	Value
CreditCardNo	SECRETSTORE	SS_CredSet:ss1:CreditCardNo
ICHAIN_UID	SECRETSTORE	SS_CredSet:ss1:username
Wallet_Key	SECRETSTORE	SS_CredSet:https\\:// www.domain.com/sevlet/ webaccess:WalletKey

where SS_CredSet is the type of shared secret, and ss1 is the name of the shared secret, and CreditCardNo is the key name in this particular shared secret.

NOTE: Use a double backslash (\\) to escape. The double backslash is necessary because Novell Secure Login (and other products) creates the Share Secret containing an escape (backslash) character. To use this secret through OLAC, you must add an additional backslash.

13.5.3 Configuring OLAC for Shared Secrets

Shared Secrets allows only secure channels for applications that want to communicate with it. Starting with iChain 2.3 SP5, both the LDAP profile and the SecretStore processor use the LDAP servers which are configured for the iChain ACLCHECK profile. As a result, the algorithm that takes care of LDAP load balancing and failover now also provides fault tolerance for OLAC. (For more information about this LDAP feature, see Appendix B, “Using LDAP Server Load Balancing and Failover,” on page 387.)

To enable SecretStore, add the following lines to the oac.properties file:

```
[SECRETSTORE Processor]
Initial Context Factory = com.sun.jndi.ldap.LdapCtxFactory
Class Name = com.novell.ichain.oac.secretstore.ParamListBuilder
```

If you have configured the OLAC LDAP processor for a non-secure port, you need to add the following lines:

```
Security Authentication = ssl
Client Certificate File = svr243IP.der
```

The OLAC LDAP processor can run on the non-secure port, but the SecretStore processor must run over SSL. These two lines configure the SecretStore processor for SSL. Replace the value for the Client Certificate File option with the name of the trusted root certificate you have installed on iChain for accessing the LDAP server. This certificate is located in sys:\.

Having OLAC LAP use non-secure communication and SecretStore use secure communication is not a recommended configuration. The health check built into proxy cannot detect problems with SSL communication if the normal communication to the LDAP server is set up for non-secure communication. It checks the LDAP servers according to the authentication profiles defined in the system, which means it checks the non-secure communication process and does not check the secure communication process.

IMPORTANT: Enabling the Shared Secrets plug-in for OLAC might significantly affect the system's performance, because Shared Secrets requires SSL binds for each user to get the secrets.

13.5.4 Known Limitations With OLAC and Shared Secrets

The following are known limitations for using OLAC with Shared Secrets:

- For Mutual and RADIUS authentication, there is no password. This means you cannot use Shared Secrets for these types of authentication because they require a user name and password.
- In Shared Secrets, administrators can define duplicated names in name=value pairs. There is no way for OLAC and Form Fill to pick the right name=value because OLAC and Form Fill aren't interactive applications. This means if you have a duplicated case, you should not define it for OLAC or Form Fill.
- OLAC and Form Fill currently only support login credentials (for example, SS_CredSet).

13.6 The Effects of Disabling OLAC

If you initially enable OLAC with authentication profile parameters using a query string and then disable OLAC, the authentication profile parameters are still passed in the query string for existing users who are logged in.

Form Fill

14

iChain® Form Fill is a tool for automatically filling in data on a form for a user and posting it to the Web server. Form Fill allows you to select which fields are automatically filled and which fields require input from the user.

Form Fill uses two components: the HTML form and the Form Fill policy. The HTML form is created with HTML tags and consists of form elements such as fields, menus, check boxes, and buttons. The Form Fill policy is created with a set of Form Fill tags that you use to customize your form by specifying the following:

- Which information is entered automatically and not displayed to the user
- Which information is displayed so that the user, at least the first time, can enter the information
- What is done with the information (for example, is it saved so that the user doesn't need to enter it when accessing the form again)

The following sections explain how to create and use a Form Fill policy:

- Section 14.1, "Understanding an HTML Form," on page 161
- Section 14.2, "Designing a Basic Form Fill Policy," on page 165
- Section 14.3, "Creating a Form Fill Policy," on page 175
- Section 14.4, "Using Additional Form Fill Policy Options," on page 176
- Section 14.5, "Troubleshooting a Form Fill Policy," on page 183
- Section 14.6, "Setting Up SecretStore," on page 186

14.1 Understanding an HTML Form

The following figure is an example of a Web page containing an HTML form.

Figure 14-1 *Sample HTML Form*

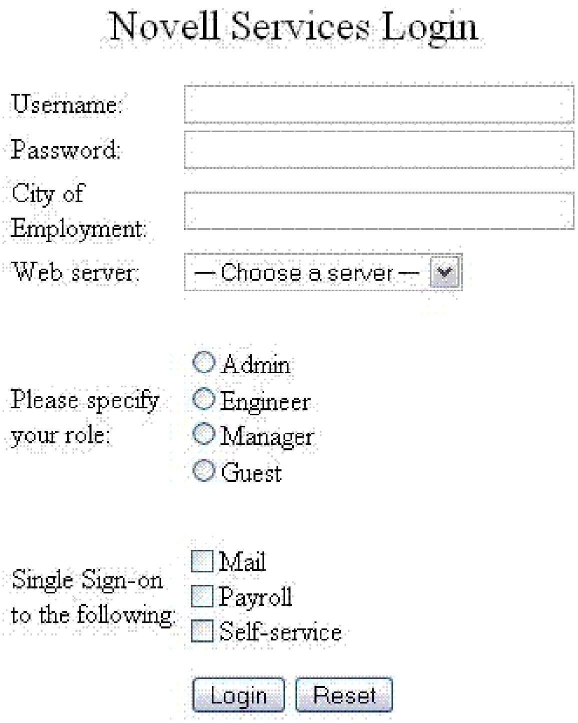

Novell Services Login

Username: □

Password: □

City of Employment: □

Web server: — Choose a server — ▾

Please specify your role:
○ Admin
○ Engineer
○ Manager
○ Guest

Single Sign-on to the following:
□ Mail
□ Payroll
□ Self-service

The information in this section uses this sample form to explain how to create a policy. This form deliberately contains a variety of field types:

- Input items for Username and Password
- Selection options for the Web server field
- Radio buttons for the role
- Check boxes for Single Sign-on

When analyzing a form, you need to decide if you want the policy to fill in all the fields or just some of them. You then need to look at the source HTML of the form to discover the names of the fields and their types.

14.1.1 Analyzing the HTML Form

An HTML form is created using a set of HTML tags. A form consists of elements (fields, menus, check boxes, radio buttons, push buttons, etc.) that control how the form is completed and submitted. For more detailed information about forms, see the Forms section at www.w3.org (http://www.w3.org/TR/html401/interact/forms.html).

The following HTML data corresponds to the sample form (see Figure 14-1). The lines that contain the information needed to create a Form Fill policy appear in bold type. Each line corresponds to a field in the form that requires information or allows the user to select information.

```
<!DOCTYPE html PUBLIC "-//W3C//DTD HTML 4.01 Transitional//EN"
   "http://www.w3.org/TR/html4/loose.dtd">
<html>
<head>
  <meta http-equiv="Content-Type" content="text/html; charset=UTF-8">
  <title>iChain Form Test Page</title>
</head>
<body>
  <form name="mylogin" action="validatepassword.php" method="post" id="mylogin">

    <table align="center" border="0" cellpadding="4" cellspacing="4">
      <tr align="center" valign="top">
        <td>
          <p align="center"><font size="5">Novell Services Login</font></p>

          <table align="center" border="0">

            <tr align="left">
              <td>Username:</td>
              <td><input type="text" name="username" size="30"></td>
            </tr>

            <tr align="left">
              <td>Password:</td>
              <td><input type="password" name="password" size="30"></td>
            </tr>

            <tr align="left">
              <td>City of<br>Employment:</td>
              <td><input type="text" name="city" size="30"></td>
            </tr>

            <tr align="left">
              <td>Web server:</td>
              <td>
                <select name="webserv" size="1">
                  <option value="default" selected>
                    --- Choose a server ---
                  </option>
                  <option value="Human Resources">
                    Human Resources
                  </option>
                  <option value="Development">
                    Development
                  </option>
                  <option value="Accounting">
                    Accounting
                  </option>
                  <option value="Sales">
                    Sales
                  </option>
                </select>
              </td>
            </tr>

            <tr>
              <td colspan="2" align="left" height="25" valign="top">
                <p></p>
              </td>
            </tr>

            <tr align="left">
              <td>Please specify<br>your role:</td>
              <td>
                <input name="role" value="admin" type="radio">Admin<br>
                <input name="role" value="engineer" type="radio">Engineer<br>
                <input name="role" value="manager" type="radio">Manager<br>
                <input name="role" value="guest" type="radio">Guest
              </td>
            </tr>

            <tr>
              <td colspan="2" align="left" height="25" valign="top" width="121">
                <p></p>
              </td>
            </tr>

            <tr align="left">
              <td>Single Sign-on<br>to the following:</td>
              <td>
                <input name="mail" type="checkbox">Mail<br>
                <input name="payroll" type="checkbox">Payroll<br>
                <input name="selfservice" type="checkbox">Self-service<br>
              </td>
            </tr>
          </table>
        </td>
      </tr>

      <tr>
        <td colspan="2" align="center">
          <input value="Login" type="submit">
          <input type="reset">
        </td>
      </tr>
    </table>
  </form>
</body>
</html>
```

Each bold line contains information about a field, its name, and type. You use this information in the policy to specify how the information in the field is filled. For example, if you want the Username field filled in automatically, your Form Fill policy would have a line similar to the following:

```
<input type=”TEXT” name=”username” value="~cn">
```

The Username field would automatically be filled in with the value of the cn attribute of the user. If you want the users to enter their usernames the first time the form is accessed, your Form Fill policy would have a line similar to the following:

```
<input type=”TEXT” name=”username” value="~">
```

The next section explains how to use the tilde operator and tags to fill in values.

14.2 Designing a Basic Form Fill Policy

The purpose of a Form Fill policy is to tell iChain what data to use when filling in the form. The following sections discuss the components you use to create a Form Fill policy:

- Section 14.2.1, "Basic Components of a Form Fill Policy," on page 165
- Section 14.2.2, "Sample Form Fill Policy," on page 174

For information about the advanced features and tags not described in this section, see Section 14.4, "Using Additional Form Fill Policy Options," on page 176.

14.2.1 Basic Components of a Form Fill Policy

Form Fill policies follow standard XML document structure. Unless otherwise specified, each tag requires that an ending tag also be present. The ending tag is simply </tag>, where "tag" is the name of the tag. For example, the ending tag for <urlPolicy> is </urlPolicy>. If no ending tag is required (because the tag requires no data), the tag name ends with a slash (/). For example, <post/> or <debugPost/>.

When defining a basic Form Fill policy, you need to decide the following:

- Determine how many <urlPolicy> elements are needed. See "<urlPolicy> Element" on page 165.
- Identify which form uses the policy. See "Header" on page 165.
- Determine which elements to use for filling in the fields. See "Input Actions" on page 168.
- Provide commands for posting the information. See "Post Actions" on page 171.
- Determine how the data should be saved and provide commands for saving it. See "Storing Form Fill Data" on page 172.
- Specify a redirection policy if values are saved. See "Handling Login Failures" on page 173.

<urlPolicy> Element

The <urlPolicy> is the main element in a Form Fill policy. You can have multiple <urlPolicy> sections in a Form Fill policy, each representing a separate set of criteria or actions to perform. If you are doing Form Fill for multiple URLs or HTML pages, you can have multiple <urlPolicy> sections in your Form Fill policy. You can also have multiple <urlPolicy> sections if you need to handle conditions such as login failures or error messages for the same URL.

Header

The header tags identify the policy and the form that uses it.

Tag	Purpose
<name>	Specifies the name of the policy. Although this tag is optional, you should give each of your policies a unique name. Having policy names is especially important when using <redirect> or <deleteRemembered> policies. **NOTE:** The tag value cannot contain spaces.

Tag	Purpose
<url>	(Required) Specifies the URL or form for the iChain Form Fill policy to match. The URL defined in this tag needs to match the name of the protected resource, not the name of the back end Web server. It is a common practice to include the scheme (http, https, ftp, gopher) with the URL in the <url> tag. For a Form Fill policy, do not include this part of the URL; specify only www.*address*.com **NOTE:** An asterisk at the end of the URL means the policy is not an exact match.

The following tags are used to distinguish among multiple forms that use the same URL. For best performance, use as few of these tags as possible to distinguish one form from another.

Tag	Purpose

<cgiCriteria>	Evaluates the query string of a URL (portion after the ?) to differentiate pages having the same URL.
	The <cgiCriteria> tag complements the <url> tag. If a <cgiCriteria> tag exists in a policy, the text in that section must appear on the query string for the specified URL in order for the policy to be deemed a match. Consider the following URL:
	`http://webaccess.novell.com/servlet/webacc?Action=User.login`
	For this URL, the following policy has been defined:
	`<url>webaccess.novell.com/servlet/webacc</url>`
	`<cgiCriteria>Action=User.login</cgiCriteria>`
	This policy limits the match to the login page and cannot be used to match either of the following URLs:
	`http://webaccess.novell.com/servlet/webacc`
	`http://webaccess.novell.com/servlet/webacc?Action=User.logout`

Tag	Purpose

<formCriteria>	Causes iChain to find a match if the text specified in the <formCriteria> tag is found on the page. The text contained between the beginning and ending tags must appear in the document returned by the origin server for the page to be a match. For example:
	`<formCriteria><TITLE>iChainTestPage</TITLE></formCriteria>`
	iChain searches the HTML page for this string. If the text is found in the document, the page is a match for the policy. If the text is not found, the policy does not match.
	If the <formCriteria> spans multiple lines, the text on each line must either be left-justified, or appear in the document exactly as specified in the policy. White space is not ignored. For example:
	`<formCriteria>`
	`          <TITLE>iChainTestPage</TITLE>`
	`loginForm`
	`</formCriteria>`
	In the example above, the text <TITLE>iChainTestPage</TITLE> must appear in the document preceded by ten spaces, but the text loginForm can appear anywhere in the document at any position in a line. Both lines must appear in the document for the page to be a match.
	NOTE: iChain searches all of the information in these tags. The more specific your information is, the faster iChain can match the form. Parsing <formCriteria> is a very intensive process for iChain to perform. If possible, using other criteria (such as <cgiCriteria>, <formName>, and <formNum>) is preferred and can increase performance.

<formName>	Causes the policy to match the page only when the name of the form is <name>.
	This tag has priority over the <formNum> tag. If both <formNum> tag and <formName> tags exist in the same policy, the <formName> tag is used, and you are notified (on the SSO Debug screen) that both tags are present. For example:
	`<urlPolicy>`
	`   <name>test</name>`
	`   <url>www.novell.com/signon_welcome/screen</url>`
	`   <formNum>2</formNum>`
	`   <formName>`*name*`</formName>`
	`</urlPolicy>`
	Using the <formName> tag is useful either as an additional way to determine whether or not to process the page with Form Fill (such as when the same URL is used for both login and content display), or when multiple forms exist on the same HTML page.

Tag	Purpose

<formNum>	Causes the policy to process form N, where N represents the form instance within the document. Forms instances are counted sequentially from the top of the document down. For example:
	`<formNum>1</formNum>`
	The first form on the page is number one; the second, number two; and so on. For example:
	`<urlPolicy>`
	`   <name>test</name>`
	`   <url>www.novell.com/signon_welcome/screen</url>`
	`   <formNum>2</formNum>`
	`</urlPolicy>`
	This example causes the second form on the page to be processed.
	If multiple forms exist in a page and neither the <formName> or <formNum> sections are specified in the policy, the first form in the page is processed by Form Fill and all other forms on the page are ignored.

Input Actions

The input tags identify which fields are to be filled and where the information can be obtained.

Tag or Attribute	Purpose
<injectStaticValue >	Specifies information that you want to inject into the form instead of having the user enter this information. The <injectStaticValue> tag overrides any information the user enters in the field. For example: `<injectStaticValue>city=Provo</injectStaticValue>` For more information about injecting static values, see Section 14.4.2, “Injecting Static Values,” on page 178.
<actions>	Specifies various actions for iChain to perform for this policy. The <actions> tag can contain <fill>, <post/>, <debugPost/>, <maskedPost/>, <redirect>, <errorRedirect> and <deleteRemembered> elements.

<fill>	Delimits a section of the form that requires input. It is a subordinate element of <actions> and must contain either an <input> element, a <select> element, or both. For example:
	`<actions>`
	`   <fill>`
	`      <input name="role" type="radio" value="~title">`
	`   </fill>`
	`   <post/>`
	`</actions>`
	This example fills in one field.

Tag or Attribute	Purpose

<input>	Specifies the form fields that need information and specifies how that information is to be obtained. The fill section of the policy can have multiple <input> fields. The <input> fields you place within the <fill> element must be specified in the order they appear on the form. The <input> element has three required attributes (name, type, and value) and one optional case modifier attribute (ff_lower_upper). For example:
	`<fill>`
	`      <input name="username" type="text" value="~cn">`
	`</fill>`

name	Specifies the name of the field. This must match the name specified in the form.
type	Specifies the type of field. This must match the type specified in the form and can be radio, checkbox, or text. The <select> element is used for list boxes.
value	Specifies how the information is to be obtained. The tilde operator allows you to determine whether the user must enter the information or whether an eDirectory™ attribute can be used to obtain the information. For more information on how to use this operator, see “The Tilde (~) Operator” on page 170.

ff_lower_upper	Converts the resulting value to either all uppercase or all lowercase. Valid values for this attribute are “lower” or “upper.” For example, to force lowercase conversion:
	`<fill>`
	`     <input name="username" type="text" value="~cn"`
	`ff_lower_upper="lower">`
	`</fill>`
	To force uppercase conversion:
	`<fill>`
	`     <input name="username" type="text" value="~cn"`
	`ff_lower_upper="upper">`
	`</fill>`
	If the attribute value is not specified as either “lower” or “upper”, no case conversion is performed on this field.

<select>	Sets up a list box field and specifies how that information is to be obtained. The fill section of the policy can have multiple <select> fields. The <select> fields you place within the <fill> element must be specified in the order they appear on the form. The <select> element has three required attributes (name, type, and value) and one optional case modifier attribute (ff_lower_upper). For example:
	`<fill>`
	`   <select name="webserv" type="listbox" value="~">`
	`</fill>`

name	Specifies the name of the select field. This must match the name specified in the form.
type	Specifies the type of field, and for a <select> element, it must be set to listbox.
value	Specifies how the information is to be obtained. The tilde operator allows you to determine whether the user must enter the information or whether an eDirectory attribute can be used to obtain the information. For more information on how to use this operator, see “The Tilde (~) Operator” on page 170.

Tag or Attribute	Purpose

ff_lower_upper	Converts the resulting value to either all uppercase or all lowercase. Valid values for this attribute are “lower” or “upper”. For example, to force lowercase conversion:
	`<select>`
	`  <input name="username" type="listbox" value="~cn"`
	`ff_lower_upper="lower">`
	`</select>`
	To force uppercase conversion:
	`<select>`
	` <input name="username" type="listbox" value="~cn"`
	`ff_lower_upper="upper">`
	`</select>`
	If the attribute value is not specified as either “lower” or “upper,” no case conversion is performed on this field.

The Tilde (~) Operator

The tilde (~) operator is used to designate how the value for the given form element is obtained. It specifies whether the user is prompted for the information, or if it is provided automatically by iChain. The tilde operator can be used in the following three ways:

- **~:** A tilde with no trailing text indicates a user-supplied value.

  The first time users access the form, they are presented with the original page and are allowed to enter values for each of the fields. For the users, the page behaves as though they are accessing the origin server directly or as if Form Fill is not active for the page.

  After the users submit the form, iChain stores (via LDAP) the values they entered for each of the fields in either the ichainFormFillCrib attribute or SecretStore, depending on how you have configured your system. When the users access this page again, iChain retrieves (via LDAP) the previously stored values for each of these elements.

  Suppose your form requests a user’s Social Security Number (<input type="text" name="SSN">, and no SSN attribute exists in eDirectory. You can use the ~ operator to have iChain store the user-entered value. For this example, the corresponding Form Fill tag would be:

  ```
  <input type="text" name="SSN" value="~">
  ```

  The ~ operator can be used with all supported form element types including text, password, checkbox, radio, and select. Using this operator is especially useful in situations where the username or password for an application differs from the eDirectory username and password, or when the required value for an element is not already stored in the directory, such as in the SSN example above.

- **~LDAP Attribute:** A tilde followed by the name of an LDAP attribute instructs iChain to obtain the value for this field from the named LDAP attribute on the user object.

  When the value for a particular form element is already stored in eDirectory, you can use this value type to obtain the information and use it for the value of that element. For example, if a login form requests the user’s email address (`<input type="text"`

name="EmailAddress">), and every user's email address is stored in the Internet Email Address attribute in eDirectory, then the administrator can use "~mail" as the value for that input field as follows:

```
<input type="text" name="EmailAddress" value="~mail">
```

This causes iChain to do an LDAP query for the mail attribute on the user object, then use the resulting value for the EmailAddress form element.

The tilde operator requires the LDAP name for the attribute. In eDirectory, the attribute must be mapped to the LDAP attribute. In this example, the mail attribute of LDAP is mapped to the Internet Email Address attribute of eDirectory. To view or change your current LDAP attribute mappings, see the LDAP Group object for your LDAP server.

- **~password:** A tilde followed by the word "password" instructs iChain to use the password the user supplied to authenticate to iChain as the value for this field.

  When iChain encounters this in a Form Fill policy, it uses the user's iChain password as the value for this field. The user's password is not pulled from eDirectory, and no LDAP traffic is generated. For example:

  ```
  <input type="password" name="password" value="~password">
  ```

  NOTE: If you are using mutual authentication (certificate based), iChain does not have a password for the user, so this method might not work as expected.

  If the user's password contains a double quote character ("), the form fill fails. Other special characters are allowed in the password. You'll need to instruct your users not to use the double quote character in their passwords.

The tilde operator is one method of adding values to fields. If you want to supply hard-coded values for fields in the form, you can use the <injectStaticValue> tag to do so. See Section 14.4.2, "Injecting Static Values," on page 178.

Post Actions

If you want the form to be submitted automatically, you need to use one or more of the following post tags. The <actions> tag only requires the post tags when filling a form. Other actions such as <redirect> and <deleteRemembered> do not require them.

Tag	Purpose
<post/>	Submits the form.
<maskedPost/>	When processing this tag, Form Fill replaces text input field values (username, password, etc.) with "nov-ss-ff-masked" instead of the value specified by the "value=" parameter. The appropriate corresponding values are replaced by iChain when the form is submitted to the origin server. The user's browser never sees the actual values for these fields.
<debugPost/>	Requires a <post/> or <maskedPost/> element. This is a troubleshooting tag that lets you verify that your information is correct before you submit the form. We recommend that you use this tag when creating a new policy and that you remove it when you have determined that the policy is behaving as expected.

If you are not using SSL, you can increase your data security during a Form Fill operation by using the <maskedPost/> tag in the policy. When the actual posting of the data is done by the browser, a regular <post/> tag transmits the values for these fields across the wire from iChain to the browser. The data is then transmitted from the browser back to iChain as part of the POST request. This can be a security risk, especially if you are not using SSL for your site.

To use this feature, simply replace the <post/> tag with <maskedPost/> in the policy. A Form Fill policy should not have both of these tags in the same <urlPolicy> section.

NOTE: If your users can see the form on the browser (for example, during the first usage, or if there is an error), iChain displays the actual data without masked values.

For more information about the methods of submitting forms, see Section 14.4.3, "Submitting Form Data to the Server," on page 178.

Storing Form Fill Data

You can choose one of three ways to store Form Fill data:

- "Using the ichainFormFillCrib Attribute" on page 172
- "Using SecretStore Shared Secrets" on page 172
- Section 14.6.3, "Using the SecretStore Proxy Service," on page 191

In iChain, the ichainFormFillCrib is the default attribute that stores the user-entered data. For more secure storage of credentials, use Novell® SecretStore®. If you use SecretStore, you need to first install it. For instructions on installing SecretStore, see Section 14.6.1, "Installing SecretStore," on page 187.

Using the ichainFormFillCrib Attribute

The ichainFormFillCrib attribute is a user attribute in eDirectory that stores Form Fill credential information for the user in an encrypted format. It is created only when the Form Fill policy is configured to save the values associated with one of the login page's input fields. For example, if the Form Fill policy is configured with a <input type="text" name="userID" value="~"> statement, it saves the value of the userID field in the ichainFormFillCrib attribute for that user.

This attribute is created only when the following occurs:

- iChain is not using Novell Shared Secret (NSSO).
- The value type is defined as ~ with no attributes.

If you use this storage method, no other applications can access and use the data.

The default setting is to use the ichainFormFillCrib attribute. It works with both LDAP and Secure LDAP.

Using SecretStore Shared Secrets

To register user credentials for Form Fill and other applications, you can use the Form Fill Shared Secrets feature as long as the other applications also use Shared Secrets.

Some Web service applications use single sign-on or share application information. For some applications, you might want to preconfigure user credentials so your users never see them. Some Novell products, including Novell exteNd™ Director™, Novell SecureLogin, and iChain Form Fill need to share login credentials (usename and password).

You need to use the SecretStore interface to set the naming convention for your secret name for iChain. To function properly, the iChain naming convention must match the naming convention of the other applications. For example, the naming convention for Novell exteNd Director is //novell.com/nps/servlet. When you change this information in one application, such as Novell Secure Login, the change appears in all the other applications.

Novell SecretStore lets you overwrite credentials, but not read or retrieve the users' credentials.

IMPORTANT: Before you use Shared Secrets, you must install SecretStore version 3.x on the iChain Authorization Server. For installation instructions, see "Installing SecretStore On NetWare" on page 187 or "Installing SecretStore on Windows" on page 187.

If you are using SecretStore, use the following tags to specify how the data is stored.

Tag	Purpose
<secretName>	Stores users' login credentials.
<sharedSecret>	Shares users' credentials with other applications. The sharedSecret tag contains a <migration/> sub-tag. The <migration/> tag means the original credentials are migrated to Shared Secrets. The original credentials are stored in eDirectory or in SecretStore, as specified in the Use Novell SecretStore for Form Fill check box in the Form Fill Policy tag (under the ISO object).
<migration/>	Migrates the original credentials to Shared Secrets. The original credentials are stored in eDirectory or in Novell SecretStore.

Handling Login Failures

When doing form fill, you need to account for login failure behaviors. For example, you might be storing credential data for a user whose password was changed on the back-end server. When such a user authenticates through iChain, their old credential data is still used. Or you might be passing the user's iChain password to a back-end server, but the user's eDirectory password was changed recently. In this case, the user's password on the back-end server might not match the iChain password any longer. When login fails, the user is presented with an error page, redirected back to the login page, or some other behavior. If this is not accounted for with a login failure policy, the login process appears to hang or loop, never taking the user to the desired application or giving any other indication of a problem with the login.

To prevent this from happening, you should create a login failure policy. The Login Failure policy can be designed to recognize the URL of an error page or can utilize <formCriteria> or <cgiCriteria> to differentiate when the error page is displayed on the same URL as the login page. You can use any of the tags in "Header" on page 165 to identity the page the user receives on login failure.

The following Form Fill tags, which are subcomponents of the <actions> element, are useful in designing a Login Failure policy:

Tag	Purpose
<deleteRemembered>	Deletes the user’s stored data for the named policy from either the ichainFormFillCrib attribute or SecretStore, depending on your system configuration. For example: `<deleteRemembered>GroupWise</deleteRemembered>` This example deletes the user’s credentials stored for the GroupWise® Form Fill policy.

<redirect>	Redirects the user to the specified URI. For example:
	`<redirect>http://my.server.com/loginfailed.html</redirect>`
	This example redirects the user to the http://my.server.com/loginfailed.html page. The format for the URI is
	`scheme://host/path`
	The scheme component is optional. It can be any valid scheme that the browser understands, such as http:// or https://. If not specified, iChain defaults to http://.
	The host component is the target server and can be specified by either IP address or DNS name.
	The path component is optional. It can be any valid path for the designated host and the resulting URI.
	Example URIs for the <redirect> tag and the resulting URI:
	`https://www.company.com/  ->  https://www.company.com/`
	`www.company.com/loginfail.html  ->  http://www.company.com/loginfail.html`
	`/loginfailure.html  ->  http:///loginfailure.html (this is an invalid URL)`
	`GroupWise  ->  http:///GroupWise (this is an invalid URL)`
	We recommend that you specify the complete URI for the new location.

If a Login Failure policy references the same URL as a Login policy, the Login Failure policy should appear first in the Form Fill policy. iChain Form Fill processing is from the top of the policy down, and the first matching policy is used.

For a sample Login Failure policy, see the iChainTestFailure policy in Section 14.2.2, “Sample Form Fill Policy,” on page 174.

14.2.2 Sample Form Fill Policy

This sample policy can be used for the form displayed in Figure 14-1 on page 162. The actions section follows the order of the fields in the form and uses the names and types defined in the HTML source.

This policy is designed so the user must enter the data once, but after the data has been entered, iChain stores the data and automatically logs the user in. If the data on the back-end server has changed and the stored data needs to be updated, the second <urlPolicy> element defines a Login Failure policy that deletes the stored data and allows the user to fill in the form again.

```
<urlPolicy>
   <name>iChainTest</name>
   <url>iChainTestserver.com/FullTest.html</url>
   <injectStaticValue>
      city=Provo
   </injectStaticValue>
   <sharedSecret>
   </sharedSecret>
   <formCriteria>iChain Form Test Page</formCriteria>
   <actions>
      <fill>
         <input name="username" value="~" ff_lower_upper="upper">
         <input name="password" value="~">
         <select name="webserv" type="listbox" value="~">
         <input name="role" type="radio" value="~">
         <input name="mail" type="checkbox" value="~">
         <input name="payroll" type="checkbox" value="~">
         <input name="selfservice" type="checkbox" value="~">
      </fill>
      <post/>
   </actions>
</urlPolicy>

   <name>iChainTestFailure</name>
   <url>iChainTestserver.com/FullTestFailure.html</url>
   <actions>
      <deleteRemembered>iChainTest</deleteRemembered>
      <redirect>iChainTestserver.com/FullTest.html</redirect>
   </actions>
</urlPolicy>
```

14.3 Creating a Form Fill Policy

After you have designed your policy, you are ready to create it.

1 In ConsoleOne®, select the ISO object and click Form Fill Policy.

  This opens an XML editor.

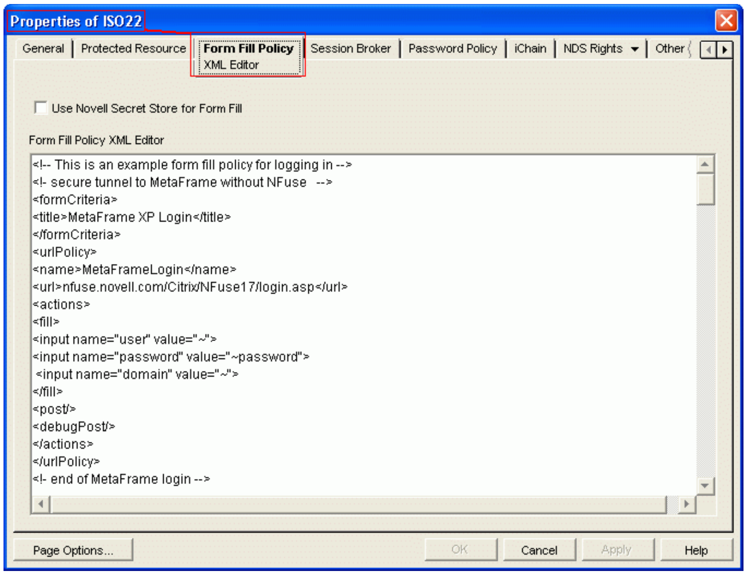

□ Use Novell Secret Store for Form Fill

Form Fill Policy XML Editor

```
<!-- This is an example form fill policy for logging in -->
<!- secure tunnel to MetaFrame without NFuse   -->
<formCriteria>
<title>MetaFrame XP Login</title>
</formCriteria>
<urlPolicy>
<name>MetaFrameLogin</name>
<url>nfuse.novell.com/Citrix/NFuse17/login.asp</url>
<actions>
<fill>
<input name="user" value="~">
<input name="password" value="~password">
 <input name="domain" value="~">
</fill>
<post/>
<debugPost/>
</actions>
</urlPolicy>
<!- end of MetaFrame login -->
```

2 Enter the tags you have selected for your policy. For information on designing a policy, see Section 14.2, "Designing a Basic Form Fill Policy," on page 165.

For information on fixing a policy that isn't working correctly, see Section 14.5, "Troubleshooting a Form Fill Policy," on page 183.

14.4 Using Additional Form Fill Policy Options

This section discusses some of the different items you can include in a Form Fill policy. This section explains the following Form Fill options:

- Section 14.4.1, "Using Other Form Fill Tags," on page 176
- Section 14.4.2, "Injecting Static Values," on page 178
- Section 14.4.3, "Submitting Form Data to the Server," on page 178
- Section 14.4.4, "Using JavaScript in Forms," on page 179
- Section 14.4.5, "Using Login Pages in Multiple Language Character Sets," on page 182
- Section 14.4.6, "Specifying Form Fill Switches," on page 182

14.4.1 Using Other Form Fill Tags

The following is a list of the tags that were not described in Section 14.2.1, "Basic Components of a Form Fill Policy," on page 165.

Tag	Purpose
<errorRedirect>	Must be inside the <action></action> tags. In the event of an LDAP or NSSS error, the user is redirected to the specified URL with the following query string parameters: ◆ Stage. The value for Stage is Fill | Post. Fill applies to building the page. Post applies to the data that is processed and returned. ◆ URL. The URL of the original page. ◆ Policy. The policy name. Unknown if no policy is found. ◆ Error. The decimal error number, like -811. ◆ Class. The subsystem that had the error, such as NSSS for SecretStore or LDAP for LDAP problems. See the sys:system\nls\err*.cfg for a list of possible errors. ◆ Define. The corresponding #define in the source code. ◆ Description. A human readable form of the detected error. The CFG files contain more information. For example if the specified URL is http://server/sample_URL/error.html, the user is redirected to http://server/sample_URLerror.html?Stage=%s&Url=%s&Policy=%s&Error=%d&Class=%d &Define=%s&Description=%s If this tag is not used and an LDAP or NSSS error is encountered, the user is presented with the standard iChain error page with corresponding information.
<javaScript>	Retains JavaScript from the original page. Individual functions can be listed for retention, or the entry can be left blank to retain all JavaScript in the page. For an example of the <javascript> tag, see Section 14.4.4, “Using JavaScript in Forms,” on page 179.
<scriptForPost>	Works in conjunction with the JavaScript function and specifies additional functions to be executed prior to the posting of the form. For an example of the <scriptForPost> tag, see Section 14.4.4, “Using JavaScript in Forms,” on page 179.
<LocalPolicy>	Because the buffer read limit in Form Fill is 50 KB, iChain lets you store additional Form Fill Policies on the local file system of the iChain machine. Syntax: Add the following to the Form Fill Policy on the ISO object: `<LocalPolicy>{Filename}</LocalPolicy>` **NOTE:** If {FileName} does not contain a slash (\ /) or a colon (:), it is expected to reference a file in the sys:etc\proxy\appliance\config\user\formfill directory; otherwise, iChain reads the information as an absolute path. Multiple <LocalPolicy> tags can be included, but the maximum size of the resulting Form Fill policy is limited to 1 MB. This tag can appear anywhere in the policy, and is not limited to a particular section.

14.4.2 Injecting Static Values

You might want to inject static values to the data of a POST request or to the URL of a GET request. This can be done using the <injectStaticValue> tag in your Form Fill policy. Multiple values can be separated with the ampersand (&) character. The <injectStaticValue> tag is a subelement of the <urlPolicy> tag.

IMPORTANT: You cannot specify a static value for the value attribute. The following <input> example does not work.

```
<input type="text" name="city" value="Provo">
```

The following line illustrates the valid way to inject a static value in a Form Fill policy:

```
<injectStaticValue>city=Provo</injectStaticValue>
```

The following example illustrates how to supply multiple values:

```
<formCriteria>string</formCriteria>
<injectStaticValue>
    A=b&C=def&city=Provo
</injectStaticValue>
```

In this example, three separate values are injected:

- The value of “A” is set to “b”.
- The value of “C” is set to “def”.
- The value of “city” is set to “Provo”.

Injected data overrides any values entered by the user, or any data stored in the ichainFormFillCrib, the SecretStore, or retrieved from the user object via LDAP.

NOTE: Data supplied with the <injectStaticValue> section is not visible in either the original HTML page (when storing data upon first access to a page), or in the source of the HTML page generated by a <debugPost/>. This data is only injected when the request is sent from iChain to the origin server.

14.4.3 Submitting Form Data to the Server

Form data is sent to the origin server using one of two HTTP methods: POST or GET. The method used is specified in the <FORM> tag of the original HTML page, and must match the method required by the back-end application that receives the data (the URI specified in the ACTION attribute of the <FORM> tag). These methods have the following effects:

- **POST:** With the POST method, the form data set is included in the body of the form and sent to the URI listed in the ACTION attribute. The POST method is used for applications such as database modification or subscription to a service. This is the most common method of submitting data.
- **GET:** With the GET method, the form data set is appended to the URI specified by the ACTION attribute with a question mark (?) as separator, and this new URI is sent to the server. The GET method is not as secure as POST. It is best used with single responses, such as a single text box.

For complete information regarding HTML form directives, see the HTML specification for Forms (http://www.w3.org/TR/html4/interact/forms.html).

NOTE: iChain Form Fill only supports the application/x-www-form-urlencoded encoding type (enctype). The multipart/form-data enctype is not supported. For more information, see TID 10090327 (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10090327.htm).

14.4.4 Using JavaScript in Forms

Many Web pages include JavaScript to validate data and set values in the form. Form Fill provides the <javaScript> tag for including JavaScript for form processing.

If JavaScript is validating or completing form fields before submitting, you do not need JavaScript in the URL policy. The first time a user goes to a page that needs to be form filled, the user completes the form and clicks Submit. This completes and executes the field values. Form Fill captures and stores this data to use the next time the user accesses the page. When the user opens the page, Form Fill removes all of the JavaScript and creates a new form with all of the data stored as hidden values. It then uses its own JavaScript function to submit the form. Because all of the data was validated and completed the first time the user completed the form, JavaScript is not necessary for Form Fill to submit the form.

Sometimes, the Web page uses JavaScript to set a cookie required by the Web server for the form to be submitted. In this case, you need to include the JavaScript to set the cookie in the URL policy so that Form Fill can set the cookie each time it submits the form.

JavaScript might also be used to set a time stamp that the Web server depends on when the form is submitted.

The JavaScript Examples section provides examples of pages that use the JavaScript tags. The first example sets a cookie before submitting the form. The subsequent examples show how to use the Form Fill JavaScript tags in the URL policy to ensure that Form Fill sets the cookie each time it submits the form.

The following example has two JavaScript functions: Validate and setCookie. You do not need to include the validate function in the URL policy because the data that Form Fill is storing has been validated. It is necessary for setCookie, because the cookie needs to be set each time this form is submitted.

JavaScript Examples

```
<html>
<head>
   <title> Login Page </title>
</head>

<body>
   <h1 align=”center”> Login Page </H1>
   <script language=”JavaScript”>
      function setCookie(){
        document.cookie=”myCookieName=myCookieValue”;
      }
      function validate(){
         if(document.mylogin.ldap.value.length == 0){
             alert(“You must provide the IP address of the LDAP server you
                    wish to login to!”);
             return false;
             }
         return true;
      }
   </script>

   <form name=”mylogin” action=”form3.php” method=”POST”
               onsubmit=”setCookie()” >
   <center>
   <table border=”1” cellpadding=”4” cellspacing=”4”>
   <tr>
      <td>Username:</td>
      <td><input type=”TEXT” name=”username” size=”30”/></td>
   </tr>

   <tr>
      <td>Password:</td>
      <td><input type=”PASSWORD” name=”password” size=”30”/></td>
   </tr>

   <tr>
      <td>LDAP Server IP:</td>
      <td><input type=”TEXT” name=”ldap” size=”30”/></td>
   </tr>

   <tr>
      <td colspan=2 align=center>
         <input type=”submit” value=”Login” onclick=”return validate();”>
      </td>
   </tr>
   </table>
   </center>
   </form>
</body>
</html>
```

To include JavaScript code that can be used by Form Fill, use the XML <javaScript> tags.

One way to include JavaScript is to include nothing between the <javaScript></javaScript> tags, as in the following example:

```
<urlPolicy>
   <name>forms3</name>
   <url>form.formfill.com/dtest/login.php</url>
   <javaScript>
   </javaScript>
   <scriptForPost>
   </scriptForPost>
   <actions>
      <fill>
         <input name=”username” value=”~cn”>
         <input name=”password” value=”~”>
         <select name=”ldap” value=”~”>
      </fill>
      <post/>
   </actions>
</urlPolicy>
```

This URL policy results in all JavaScript code contained in the original form to be included with the form submitted by Form Fill as shown in the following example:

```
<script language=”JavaScript”>
function setCookie(){
   document.cookie=”myCookieName=myCookieValue”;
   }
function validate(){
   if(document.mylogin.ldap.value.length == 0){
      alert(“You must provide the IP address of the LDAP server you wish
             to login to!”);
      return false;
      }
      return true;
   }
</script>
```

Another way to include JavaScript code in a form is to use the setCookie function only, as in the following example:

```
<javaScript>
   function setCookie()
</javaScript>
```

This example shows how to include only the setCookie function by putting the setCookie function name in between the JavaScript tags. Remember to include the *function* keyword in the tags. The page displays the following result:

```
<script language=”JavaScript”>
   function setCookie(){
      document.cookie=”myCookieName=myCookieValue”;
   }
</script>
```

The <javaScript> tag lets you include only the JavaScript from the original form with the form that Form Fill submits. To actually execute the JavaScript code, you must use the <scriptForPost> </scriptForPost> tags.

In the example URL policy, the <scriptForPost></scriptForPost> tags appear just below the <javaScript></javaScript> tags. When you include the <scriptForPost></scriptForPost> tags without any code in the middle, the JavaScript function specified in the onsubmit field of the original form tag executes before the form is posted. In this case, the setCookie function executes just before Form Fill submits the form, which is the desired behavior.

This ensures that the setCookie() function is called regardless of when it is called in the original form. Additionally, you can supply custom JavaScript code that does not appear in the original document inside the <scriptForPost> tags. For example, you can omit the <javaScript> tags from the URL policy and put only the code found inside the setCookie function directly in the <scriptForPost> tags:

```
<scriptForPost>
   document.cookie=”myCookieName=myCookieValue”;
</scriptForPost>
```

14.4.5 Using Login Pages in Multiple Language Character Sets

You can use many character sets in an HTML form. However, Form Fill supports only ISO-8859-1 and UTF-8. The ISO-8859-1 character set is mainly for single-byte characters. The UTF-8 character set is for double-byte characters.

NOTE: The UTF-8 ASCII character sets are the same as ISO-8859-1.

The following examples show how the HEAD element should appear in a form:

Example One:

```
<head>
   <title>Untitled for ASCII</title>
   <meta http-equiv=”Content_Type” content=”text/html;charset=iso-8859-1”>
</head>
```

Example Two:

```
<head>
   <meta http-equiv=”content-type” CONTENT=”utf-8”>
   <title>Untitled for UTF-8</title>
</head>
```

When you use a character set that is different from these examples, such as shift-jis, or ISO-2022-jp, you need to do the following:

1 Verify that the back-end Web application, which handles the login credential, supports the UTF-8 format.

2 Change the charset to UTF-8 and save the form in the UTF-8 format.

   NOTE: If the content of the title is in Japanese, it should be added after the META declaration as in Example Two above.

For more multiple language character set options and information about changing the backend Apache server to use UTF-8, see the Technical Information Document (TID) 10090464 (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10090464.htm).

14.4.6 Specifying Form Fill Switches

When Form Fill is enabled on the iChain box, the sso.nlm file loads automatically. No command-line options are supplied, and all default option values are used. If you want to load the sso.nlm file with options every time the system starts, add a “load sso” statement to the end of the sys:\system\ap_start.ncf file with the appropriate options.

The sso.nlm file can be loaded re-entrantly, so the user can also change the current options from the command line (ICS_SERVER: prompt).

The following table lists and defines the sso.nlm switches that you can apply when loading iChain:

Switch	Purpose
/E<number>	Enables or disables the enhanced protection on Novell SecretStore data. An administrator typically has read and write access to users' data. When Enhanced Protection is enabled, the administrator can write data to SecretStore, but he or she cannot read the data. 0 disables Enhanced Protection on SecretStore data 1 enables Enhanced Protection on SecretStore data Default: 1
/D<number>	Enables Form Fill debugging as the specified level: 0 enables standard SSO output 1-5 enables debugging output, with each higher level providing more information. Default: 0
/L<number>	Enables or disables the logging of Form Fill debugging information. 0 disables logging of Form Fill debug information 1 enables logging to the Logger screen 2 enables logging to the Extended log file 3 enables logging to both the Logger screen and the Extended log file Default: 0
/X	Scans the forms for select and input tags for attribute values that iChain saves. iChain then scans the page tags to ensure that it writes out the correct attributes. This is usually required when you have multiple <form> sections with <select> tags on a page.

14.5 Troubleshooting a Form Fill Policy

This section provides Form Fill troubleshooting information for the following items:

- Section 14.5.1, "Tags," on page 183
- Section 14.5.2, "Tools and Documentation," on page 185

NOTE: For assistance in building Form Fill scripts, see the automatic Form Fill script generator (http://www.novell.com/coolsolutions/tools/1788.html).

14.5.1 Tags

- Make sure that the URL defined in the <url></url> tag matches the name of the protected resource, not the name of the back-end server.
- If you use wildcards for the <url> tag, the content-length field of a response might get modified for data that is not specific to an application login form. Always try to make the <url></url> entry in the Form Fill policy as specific as possible. If this is not possible, use the <formCriteria> tag and/or the <cgiCriteria> tag to narrow the list.
- Do not include the http:// scheme in front of the URL in the <url> tag. This prevents the link from finding a matching profile. It also causes the SSO to fail.

- Always copy and paste the URL into the Form Fill policy URL tag. This reduces the chances of incorrectly typing a URL.
- Form Fill does not preserve the input for type="image". Check the <form> tag in the application login form to make sure that this type does not exist. If it does, try modifying it. If it does exist, you might be able to use the custom rewriter to substitute the required data. However, we do not recommend using this method.

  You might be able to supply the required data with the <injectStaticValue> tag.
- Remove the <post/> tag and leave only the <fill> tag in the URL policy.

  This lets you confirm that the credentials were available via LDAP and that LDAP was set up correctly.
- The <debugPost/> Form Fill tag allows single sign-on usage to modify the HTML page with changes needed for auto post.

  When the <debugPost/> tag is active and the URL matches the <url></url> in the Form Fill policy, iChain displays a screen with the message "Please look at HTML source for Form Fill modifications." In this screen, you can select the View Source option in the browser. This shows the variable names and values that iChain injects into the login form. It also lets you check the JavaScript methods sent back to the browser.

  The following example shows the values that iChain injects for a form that requests the user's name, password, and domain field:

```
<html>
<body>
   <b>Please look at HTML Source for Formfill modifications</b>
   <form name=”NFuseForm” action=”login.asp” method=”POST”>
      <input value=”Explicit” name=”LoginType” TYPE=”HIDDEN”>
      <input value=”administrator” maxlength=”256”
                    onFocus=”focus_UPD(this.form);” class=”loginEntries”
                    name=”user” type=”hidden”>
      <input value=”novell” maxlength=”254”
                    onFocus=”focus_UPD(this.form);”
                    class=”loginEntries” name=”password” type=”hidden”>
      <input value=”ICHAINFARM” maxlength=”256”
                    onFocus=”focus_UPD(this.form);” class=”loginEntries”
                    name=”domain” type=”hidden”>
   </form>

<script language=”JavaScript”>
   <!--
   function iChainPostForm()
     {
         document.forms[0].submit();
     }
   //-->
</script>
<a href=”JavaScript:iChainPostForm()”>Click to submit</a>
</body>
</html>
```

- Make sure that the form declaration contains an action attribute. The action attribute specifies the URL that will process the form submission. This is a required attribute and Form Fill will not process the form if no action exists for the form.

  The attribute doesn't actually need to have a value, but the attribute needs to exist. For example:

```
<form action method="post">
```

The above example form declaration is sufficient for FormFill to process the form. If action has no value (as in the example above), the form is submitted back to the same URL as the page containing the form.

- Make sure that you do not have a form element (input, select, button, etc.) named “submit.” This name causes problems with the auto-post feature of iChain. You can use a name of “Submit” (with a capital S), or any other value for the name of the form element. This needs to be changed in the original HTML source page.

14.5.2 Tools and Documentation

- The SecretStore Software Developer Kit (SDK) contains helpful debugging tools, such as the SSManager. To access the SecretStore SDK, see the SecretStore Developer Kit for C (http://developer.novell.com/ndk/ssocomp.htm) or the SecretStore Developer Kit for Java (http://developer.novell.com/ndk/nssoj.htm).
- For more SecretStore information, see the *SecretStore Administration Guide* (http://www.novell.com/documentation/secretstore33/nssadm/data/admu1ef.html).
- The pktscan.nlm, which ships with iChain, gathers traces.

  This tool prevents you from needing to replicate a port on a switch to gather a LAN trace of traffic in and out of the iChain box. You might be able to use Sniffer or Ethereal to save and view this trace. To do this, you need to temporarily set up the accelerator for HTTP. HTTPS is not decodable with Sniffer or Ethereal.
- The sso.nlm is a Form Fill module that includes the /Dx /Lx option (where x ranges from debug level 1 through 5, where level 5 contains the most detail.) For more information about loading the sso.nlm, see Section 14.4.6, “Specifying Form Fill Switches,” on page 182.

  The Form Fill operations are logged to a file through the iChain Web GUI Cache Logs tab.

Figure 14-2 iChain GUI Cache Logs

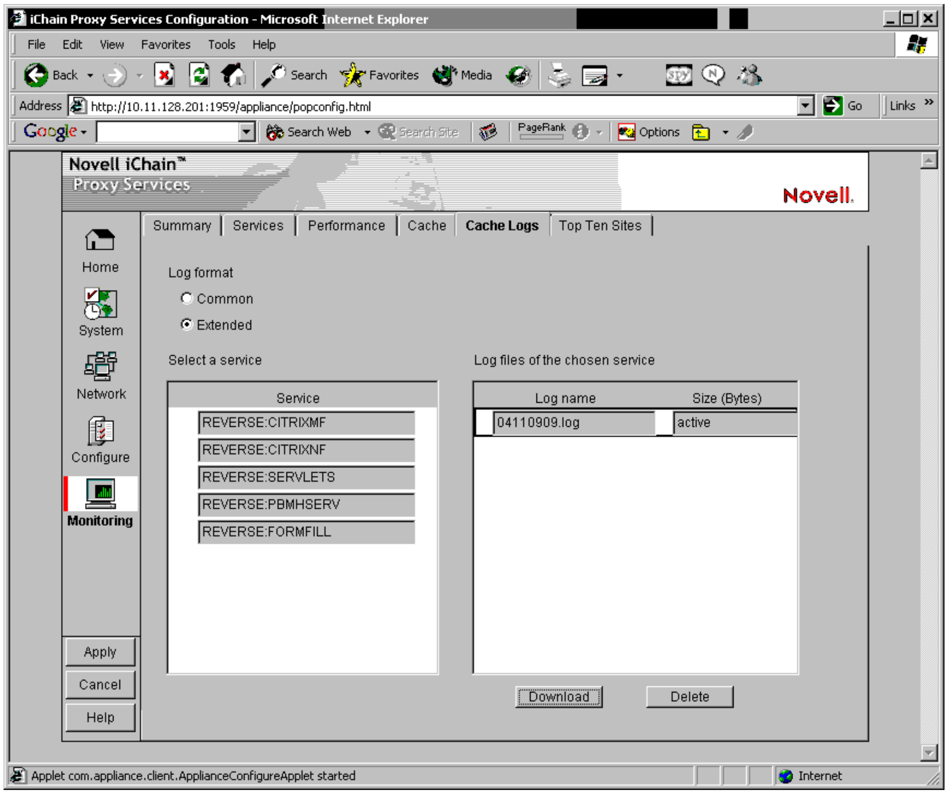
iChain Proxy Services Configuration - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://10.11.128.201:1959/appliance/popconfig.html
Go Links
Google Search Web Search Site PageRank Options
Novell iChain™
Proxy Services
Novell.
Summary Services Performance Cache Cache Logs Top Ten Sites
Home
System
Network
Configure
Monitoring
Apply
Cancel
Help
Download
Delete
Applet com.appliance.client.ApplianceConfigureApplet started
Internet

Log format

○ Common

● Extended

Select a service

Service
REVERSE:CITRIXMF
REVERSE:CITRIXNF
REVERSE:SERVLETS
REVERSE:PBMHSERV
REVERSE:FORMFILL

Log files of the chosen service

Log name	Size (Bytes)
04110909.log	active

You can use this file to see if a policy has been matched and if any errors occurred during the processing of that policy.

The following example shows that a request came in for http://nfuse.novell.com/Citrix/NFuse17 /login.asp, and that a matching policy named metaFrameLogin was located. Another request came in for http://nfuse.novell.com/Citrix/NFuse17framseset.asp, but no matching policy was found:

```
[09/Nov/2004:11:51:39 +0100] SSO_1: No policy: 'nfuse.novell.com/Citrix/
NFuse17/login.asp' [09/Nov/2004:11:51:39 +0100] SSO_4: nfuse.novell.com/
Citrix/NFuse17/login.asp[09/Nov/2004:11:51:39 +0100] SSO_4: Policy
'MetaFrameLogin': nfuse.novell.com/Citrix/NFuse17/login.asp [09/Nov/
2004:11:51:39 +0100] fillHtml: Start [09/Nov/2004:11:51:39 +0100] SSO_4:
formfill no cache flag is turned on [09/Nov/2004:11:51:39 +0100] SSO_4: New
page data [09/Nov/2004:11:51:39 +0100] SSO_1: No policy: 'nfuse.novell.com/
Citrix/NFuse17/frameset.asp'
```

For the most current information about configuring iChain with Citrix, see the *How to Configure and Troubleshoot iChain 2.3 Issues Accelerating a Citrix Metaframe server* Appnote (http://www.novell.com/coolsolutions/appnote/2562.html).

14.6 Setting Up SecretStore

The following sections explain how to install and configure iChain so that you can use SecretStore to save Form Fill data:

- Section 14.6.1, "Installing SecretStore," on page 187

- Section 14.6.2, “Configuring SecretStore to Use Shared Secrets,” on page 189
- Section 14.6.3, “Using the SecretStore Proxy Service,” on page 191

14.6.1 Installing SecretStore

Before you use SecretStore, you need to upgrade to Novell International Cryptographic Infrastructure (NICI) 2.6 or later, install SecretStore, then enable SecretStore. These procedures are explained in the following sections:

- “Installing SecretStore On NetWare” on page 187
- “Installing SecretStore on Windows” on page 187
- “Enabling SecretStore With eDirectory” on page 188

For information about secure login and single sign-on, see the *Nsure SecureLogin Administration Guide* (http://www.novell.com/documentation/securelogin3511/nsladm/data/a602kny.html).

Installing SecretStore On NetWare

Before you install SecretStore, you need to upgrade your version of NICI. The SecretStore procedures include instructions for this upgrade.

1 Insert the *iChain Authorization Server* CD into your CD drive.

2 Open a remote console prompt and enter `CDROM` to mount the *iChain Authorization Server* CD.

3 To upgrade to NICI 2.6, enter `nwconfig`.

   3a Click Product Options, then click Install a Product Not Listed.

   3b Press Esc, press F3, then type the path to the NICI upgrade: `ichain_auth_sr:/nici/nwserver`. Press Enter.

4 To install SecretStore, enter `nwconfig` at a console prompt, click Product Options, then click Install a Product Not Listed.

5 Press Esc, press F3, then specify the path to the SecretStore installation: `ichain_auth_sr:/sso/netware`

6 Follow the installation prompts.

IMPORTANT: You must have administrator or equivalent rights to change the directory schema.

7 Press Enter to complete the SecretStore installation.

8 From the *Novell iChain Authorization Server CD*, using the NetWare server as the default server, run /sso/utils/ssinit.exe.

9 Continue with “Enabling SecretStore With eDirectory” on page 188.

Installing SecretStore on Windows

1 Insert the *iChain Authorization Server CD* into your CD drive.

2 Run nici/wcniciu0.exe.

3 Follow the installation prompts to complete the installation.

4 Run sso/nt.setup.exe.

During the SecretStore installation, you need to log into your tree as an Administrator in order for the schema to be modified.

5 Start the eDirectory Services Console.

6 The following services start automatically:

- ssldp.dlm
- ssncp.dlm
- sss.dlm

If they do not start automatically, you need to manually start them.

7 Continue with "Enabling SecretStore With eDirectory" on page 188

Enabling SecretStore With eDirectory

1 Log onto the LDAP server.

2 Export your trusted root for the LDAP server.

2a From ConsoleOne, view the properties of the key material object (usually named SSLCertificate*).

2b Select the Certificates tab, select Trusted Root, then click Export.

2c Save the trusted root to the local drive in Base64 format.

3 Click the ISO object.

4 From ConsoleOne, click Form Fill Policy, then paste your Form Fill policy into the Form Fill policy text editor.

Make sure your policy contains the <sharedSecret> tags so that the SecretStore functionality is enabled.

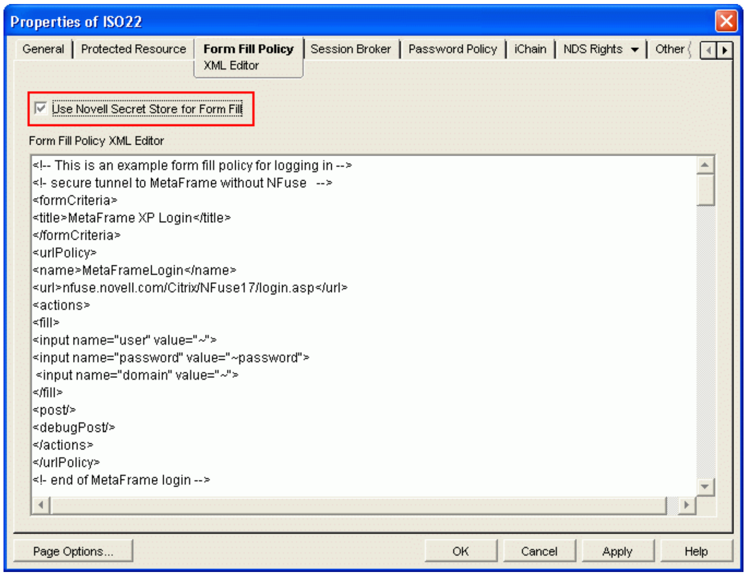

☑ Use Novell Secret Store for Form Fill

Form Fill Policy XML Editor

```
<!-- This is an example form fill policy for logging in -->
<!- secure tunnel to MetaFrame without NFuse   -->
<formCriteria>
<title>MetaFrame XP Login</title>
</formCriteria>
<urlPolicy>
<name>MetaFrameLogin</name>
<url>nfuse.novell.com/Citrix/NFuse17/login.asp</url>
<actions>
<fill>
<input name="user" value="~">
<input name="password" value="~password">
 <input name="domain" value="~">
</fill>
<post/>
<debugPost/>
</actions>
</urlPolicy>
<!- end of MetaFrame login -->
```

5 Click OK.

6 From the iChain Proxy Services Administration GUI, import the trusted root:

6a Select Configure (this icon is in the bottom left corner), select the Access Control tab, then select Enable Secure Access to LDAP Server.

6b Select Import Trusted Root and specify a name for the imported file.

The filename must adhere to the 8.3 naming conventions.

6c Using a text editor, open the file that was exported from the LDAP server to a local drive, copy all the contents of the file to the text field in the Trusted Root dialog box, then click OK.

6d Select Enable Form Fill Authentication, then click Apply.

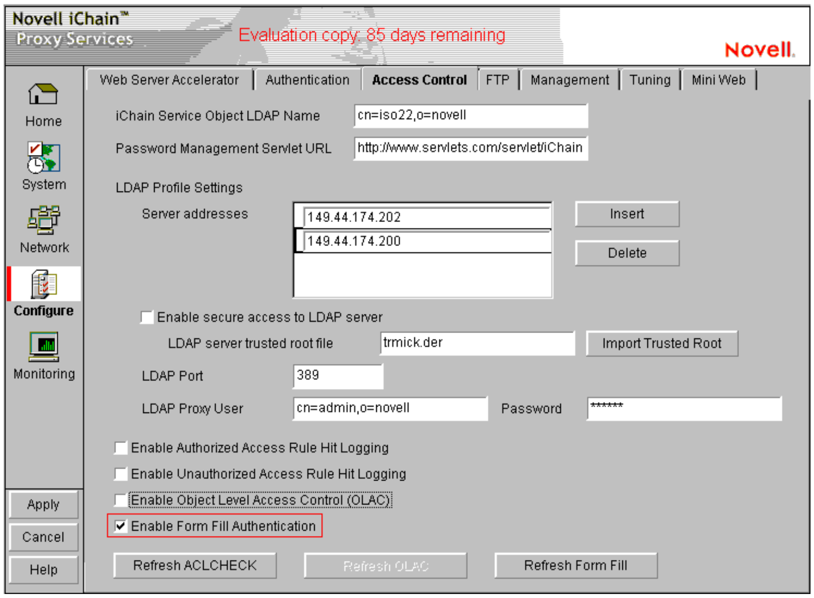

iChain Service Object LDAP Name: cn=iso22,o=novell

Password Management Servlet URL: http://www.servlets.com/servlet/iChain

LDAP Profile Settings

Server addresses: 149.44.174.202
149.44.174.200

☐ Enable secure access to LDAP server

LDAP server trusted root file: trmick.der

LDAP Port: 389

LDAP Proxy User: cn=admin,o=novell Password: ******

☐ Enable Authorized Access Rule Hit Logging

☐ Enable Unauthorized Access Rule Hit Logging

☐ Enable Object Level Access Control (OLAC)

☑ Enable Form Fill Authentication

If Form Fill was already enabled, refresh Form Fill from the Access Control tab.

7 If you use the SecretStore proxy service, or if you plan to use Shared Secrets, continue with Section 14.6.2, "Configuring SecretStore to Use Shared Secrets," on page 189

14.6.2 Configuring SecretStore to Use Shared Secrets

The following information explains the Shared Secrets feature:

- **Shared Secrets uses a secret ID as a key/index to retrieve information.** The format is <type>:<name>. There are only two types of secret IDs: SS_App (for application) or SS_CredSet (for login). Currently, Form Fill and OLAC support only the SS_CredSet type. The <name> is the name of the Form Fill policy. Shared Secret data consists of <name><delimiter><value> pairs, where the name and value are the same as in the <fill> section of the Form Fill policy.
- **Form Fill tags for Shared Secrets.** Form Fill adds the <sharedSecret> and </sharedSecret> tags for sharing users' credentials with other applications. These tags contain a sub-tag, <migration/>, which means the original credentials are migrated to Shared Secrets. The original credentials are stored in eDirectory or in SecretStore, as specified in the Use Novell SecretStore for Form Fill check box in the Form Fill Policy tag (under the ISO object).

  The following are two examples of how the <sharedSecret></sharedSecret> tags are used:

  Example one:

```
<url>www.novell.com/formfill/test/*<url>
 ............
 <sharedSecret>
 </sharedSecret>
 <actions>
 ..............
```

Example two:

```
 <sharedSecret>
 <migration/>
 </sharedSecret>
```

NOTE: After migration, the original credential is transferred to Shared Secrets and is removed from the old storage (eDirectory or SecretStore).

- **Form Fill supports two scenarios for Shared Secrets.** There is a reserved field name, ff_shared_name, in the policy for Shared Secrets. Some Shared Secrets applications might use a common name to share the value. This field name is applied to Input and Select tags in the <fill> section of the a policy. This option works only when secure LDAP is enabled.

  The policy name, for example, could be:

  <input name="user.id" value="~">,

  but in Shared Secrets, the field name is not User.id, it is Username.

  In this case, the field is

  <input name="user.id" value="~" ff_shared_name= "Username">.

  NOTE: The value of the ff_shared_name applies to reading and writing to or from Shared Secrets. It has nothing to do with the HTML form. For example, the name in the name=value pair of the POST packet is still User.id, not Username, but the name in the name=value pair in Shared Secrets is Username, not User,id. If ff_shared_name is missing, Form Fill uses the value of name (User.id) as a default value.

- **You set up the error redirect for a policy.** As the administrator, you must use a login failure policy to redirect a login failure by using the <redirect> tag. Do not redirect to the original policy or the user can get caught in an infinite loop. The policy should redirect to an error page with an error message similar to the following: "Cannot access Web application. Please report this error to the administrator."
- **You register all user credentials.** For information on how you can register all users' credentials, contact a Novell consultant. You can learn more about contacting a consultant in the Consulting section of the Novell Web site (http://www.novell.com/consulting/)
- **Use <maskedPost/>.** This is necessary to protect the secret.
- **Do not use the <deleteRemembered> tag.** If there is a login failure, users can see the login form and might attempt to log in themselves.
- **Users might need to know their credentials.** If users need to know their credentials, use the regular Form Fill feature. This lets users register their own credentials.

- **Form Fill shares credentials with OLAC and other Novell products.** This includes Novell exteNd Director and Novell SecureLogin. For information on setting up the OLAC sharing configuration, see Chapter 13, “Using Object-Level Access Control,” on page 153.
- **You can use the <secretName> tag for Shared Secrets.** Using Form Fill, you can add a <secretName> tag inside the <sharedSecret></sharedSecret> tag to make the <name> </name> field more logical and transparent to the secretID. If there is no <secretName> </secretName> tag, the default secret name is taken from the <name> field.

  Some applications use the secret name, which is a portion of the secretID (shared among Web applications), and which contains a question mark (?) character. This character causes Form Fill to act abnormally if it is found in the <name></name> field. The following are two examples of using the <secretName></secretName> Form Fill tags:

  Example One:

  ```
  <name>FormFillSharedWithNSL</name>
  <url>www.novell.com./formfill/test/*</url>
   .........
  <sharedSecret>
  <secretName>https\://novell.com/nps/servlet/
                      webacc?task\=abc&merg\=def</secretName>
  </sharedSecret>
   <actions>
  ```

  NOTE: In the <secretName> field, if there is a colon(:), equal sign (=), and/or a backslash(\) character (these are reserved for the secret ID), they must be preceded with a backslash(\) character.

  Example Two:

  ```
  <name>FormFillOwnSharedRule</name>
  <url>www.novell.com/formfill/test.*</url>
  ..........
  <sharedSecret>
  </sharedSecret>
  <actions>
  ...........
  ```

  Example Two does not have a <secretName>, so the secret name is the same as the <name> field. It is equivalent to having the following line:

  ```
  <secretname>FormFillOwnSharedRule</secretName>
  ```

14.6.3 Using the SecretStore Proxy Service

SecretStore is a service of eDirectory that allows applications to securely store sensitive data. This service guarantees protection of application data that is both in storage and being transmitted between the client and SecretStore. SecretStore is useful for storing information that is shared with other applications, such as your users' credentials.

IMPORTANT: You should not select the SecretStore check box for any configuration other than the legacy SecretStore installation that shipped with iChain 2.1. Do not select this box when using the ichainFormFillCrib attribute or Shared Secrets.

The Use Novell SecretStore for Form Fill check box refers to the SecretStore Proxy Store. This proxy store functionality is outdated. It is included only for backward compatibility. We recommend that you migrate to the new Shared Secrets format. For instructions on migrating to the new Shared Secrets format, see Sharing Secrets (http://www.novell.com/documentation/secretstore33/nssadm/data/bsqdjtt.html) in the *Novell SecretStore Administration Guide.* Also, see "Enabling SecretStore With eDirectory" on page 188.

Using iChain to Manage Certificates

15

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

- An appliance-specific certificate authority (CA) that 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.
- 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.

This section discusses how to create and modify certificates using Novell® iChain®. The following topics are included:

- Section 15.1, “Naming Certificates,” on page 193
- Section 15.2, “Creating Certificates Using the Appliance CA,” on page 194
- Section 15.3, “Obtaining a Certificate from an External CA,” on page 195
- Section 15.4, “Viewing (Exporting) a Certificate’s CA,” on page 198
- Section 15.5, “Modifying a Certificate,” on page 199
- Section 15.6, “Deleting a Certificate,” on page 199
- Section 15.7, “Backing Up a Certificate,” on page 199
- Section 15.8, “Restoring a Certificate,” on page 200
- Section 15.9, “Renewing a Third-Party Certificate,” on page 201

15.1 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. Choose the subject name that the browser expects to find in the certificate.

   For accelerator services, the Subject Name field must contain the DNS name, with the fields separated by periods (.).

   For example, the www.foo.gov Web server accelerator certificate must have a Subject Name of www.foo.gov.

15.2 Creating Certificates Using the Appliance CA

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

If this is not done, users 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 Section 15.1, "Naming Certificates," on page 193.
3. Type an appropriate subject name as explained in Section 15.1, "Naming Certificates," on page 193.
4. Click the Signature Algorithm drop-down list, then select the algorithm you want to use (SHA-1 or MD-5).
5. Click the RSA Key Size drop-down list, then 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, then 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, then click OK.
13. Click Apply and repeat the modification process until the Status field displays the word Active.

15.3 Obtaining a Certificate from an External CA

- Section 15.3.1, “Requesting the CSR,” on page 195
- Section 15.3.2, “Sending the CSR,” on page 196
- Section 15.3.3, “Storing the Certificate,” on page 197
- Section 15.3.4, “Importing a CSR Signed by Intermediates,” on page 197

15.3.1 Requesting the CSR

1 In the browser-based management tool, click Home, click Certificate Maintenance, then click Create.

The Create Certificate screen displays:

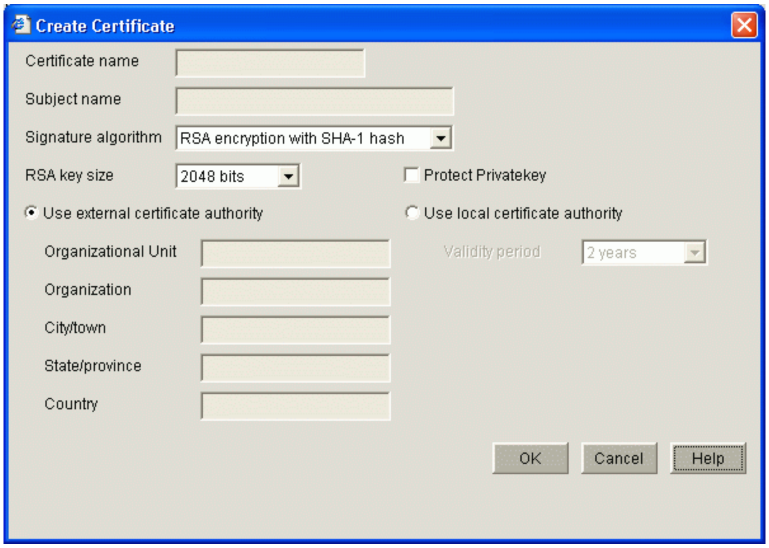

Certificate name

Subject name

Signature algorithm: RSA encryption with SHA-1 hash

RSA key size: 2048 bits

□ Protect Privatekey

● Use external certificate authority

○ Use local certificate authority

Organizational Unit

Validity period: 2 years

Organization

City/town

State/province

Country

2 Specify an appropriate name for the certificate as explained in Section 15.1, “Naming Certificates,” on page 193.

3 Type an appropriate subject name as explained in Section 15.1, “Naming Certificates,” on page 193.

4 Click the Signature Algorithm drop-down list, then select the algorithm you want to use (SHA-1 or MD-5).

5 Click the RSA Key Size drop-down list, then 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 (Optional) Select PrivateKey.

This flag is disabled by default. For added security, when it is enabled, you cannot backup the certificate.

7 Click Use External Certificate Authority.

The external certificate authority sets the validity period. You cannot set it using the Validity period option.

8 Type a name for the Organizational Unit.

This is used to describe departments or divisions.

9 Type a name for the Organization.

This is used to differentiate between organizational divisions.

10 Type the city or town where your organization does business.

This is commonly referred to as the Locality.

11 Type the non-abbreviated name of the state or province where the organization does business.

This is commonly referred to as the State.

12 Type the International Standards Organization (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.

13 Click OK.

14 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.

15 Click Apply.

If any errors occur during the certificate request process, they are displayed in the Error field on a red background.

16 If an error occurs, click Modify.

17 In the Modify Certificate dialog box, make the changes necessary to resolve the errors, click OK.

18 Click Apply and repeat the modification process until the Status field displays the words CSR in Progress on a yellow background.

NOTE: As an added precaution, Update Clone can be used to help safeguard the private key of the certificate until the certificate is returned and stored. After the certificate is returned and stored, it can then be backed up. Update Clone is found in the iChain Proxy Server browser-based administration tool under System > Actions.

15.3.2 Sending the CSR

1 Click View CSR to open a new browser window that displays the CSR contents.

2 Select and copy the complete CSR text into your computer’s Clipboard. 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 varies 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.

15.3.3 Storing the Certificate

After the external CA responds with the certificate:

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 into the CA Certificate Contents box.

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, then click OK.

9 Click Apply.

15.3.4 Importing a CSR Signed by Intermediates

1 Convert the response file into PKCS #7 format using a current, patched version of Internet Explorer*.

1a Save the response file (certificate) sent from Verisign* using Wordpad as a .cer file.

1b Open Tools, select Internet Options, select Content, then select Certificates.

1c Click Import, click Next, then browse to the response file sent from Verisign.

1d Choose “Automatically select the certificate based on the type of certificate”.

1e Click Next, then click Finish.

1f Open Tools, select Internet Options, click Content, click Certificates, then click Other People.

1g Highlight the imported certificate, click Export, then click Next.

1h Select Cryptographic message syntax standard - PKSC #7 (.P7b) and select Include all certificates in the certification path if possible.

1i Select Next, then save and export the file.

2 Enable NCP on the iChain box and login to the ICS_tree.

2a Either load c:\nwserver\ncpip.old or rename to c:\nwserver\ncpip.nlm and load ncpip.nlm.

2b Follow the editing instructions in the Tune.ncf file or specify the following set parameters at the server console:

```
SET NCP include IP addresses = All
SET NCP exclude IP addresses = None
SET NCP over UDP = On
```

2c Log in to the ICS_Tree using the IP address of the iChain box.

NOTE: For security reasons, should be disabled when you finish.

3 Import the certificate.

3a Launch ConsoleOne and right-click on the Key Material Object for the respective certificate in the ICS_tree.

The Key Material Object has the same name that was given when you created the CSR in the iChain browser-based administrative tool.

3b Select Properties, select Public Key Certificate on the drop-down of the Certificate tab.

3c Click Import, check the mark "No Trusted Root Certificate Available" check box, then click Next.

NOTE: The trusted root was included in the PKCS #7 file that was created in Step 1.

3d At the "Paste your Server Certificate here or read it from a file" screen, select Read from file, browse to the .pb7 file and finish importing the file.

NOTE: An informational message might appear when finishing the import stating that the issuer or subject do not match. Click Continue.

4 Restart the server and back up the certificate using the iChain browser-based administration tool.

When the certificate is successfully imported, restart the server, open the iChain browser-based wizard, and back up the certificate

4a Open the browser-based administration utility, click Home, then click Certificate Maintenance.

4b Highlight the certificate.

4c Click Backup.

15.4 Viewing (Exporting) a Certificate's CA

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 are displayed in a new browser window.

15.5 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. In the Modify Certificate dialog box, make the desired changes.
3. After making the necessary changes, click OK to accept the changed values.
4. If the Action field displays the word Request or Create on a red background, you must click Apply to make the changes.

15.6 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.
4. The certificate is removed from the certificates list.

   If you have deleted the certificate in error, click Cancel.
5. Click Apply to remove the certificate from the appliance.

   After clicking Apply, the certificate cannot be restored unless you have created a backup copy.

15.7 Backing Up a Certificate

Only active certificates can be backed up.

1. In the browser-based management 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. Select either Disk (hard drive) or Floppy to indicate where the backup file should be placed.

   If you select Floppy, the backup file is saved to the floppy drive of the iChain server.
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.

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. (If you save to a floppy, the name is limited to 8 characters.)

IMPORTANT: If the certificate was backed up to the appliance hard disk, you should transfer the file from the appliance to another secure location, or the backup copy will be lost if the appliance fails and must be reimaged.

Certificate backup files are stored in sys:\etc/proxy/appliance/config/user/cert/backup. See "Using FTP" on page 348 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.

15.8 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 sys:\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. Click OK.
5. Click Disk or Floppy to indicate where the backup file is.
6. 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 by the Action field.
7. Click Apply.

If any errors occur during the restore process, they are displayed on the Error line and 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.

15.9 Renewing a Third-Party Certificate

Every certificate has a validity period. When that validity period expires, the certificate is no longer considered an acceptable or usable credential. You can renew the certificate with either the same key set you used before or with a new key set.

Before you renew a certificate, you need to know the following information:

- The issuing certification authority.
- (Optional) If you want a new public key and private key pair for the certificate, the cryptographic service provider (CSP) that should be used to generate the key pair.

You might want to return a certificate to its original state. To do this, you remove the following attributes from the Key Material Object (KMO):

- ? NDSPKI:Certificate Chain
- ? NDSPKI:Key File
- ?KDSPKI:Public Key Certificate

1. Open the iChain GUI administration utility.
2. Back up the certificates you want to renew.

   Assign a different name to the certificates in case you need to restore them.
3. Restore the certificate you backed up.
4. Rename the pki.jar file to pki.org in the 1.3.X ConsoleOne® snapin directory.
5. Log in to the ICS_Tree on the iChain server using user ichainadmin.ics, with password novell.
6. Open ConsoleOne and find the corresponding KMO in the .ICS container in the ICS_Tree.
7. Delete the following attributes under the Other tab on the KMO:
   - NDSPKI:Certificate Chain
   - NDSPKI:Key File
   - NDSPKI:Public Key Certificate
8. Close ConsoleOne and rename the pki.org file back to pki.jar.

   The certificate should be ready to store the renewal certificate using ConsoleOne. Before you import the certificate, convert the renewal certificate to .p7b format by selecting the pb7 option and importing the certificate. For more instructions on converting the renewal certificate, see TID 10073709 (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10073709.htm).

   NOTE: Do not attempt to use the iChain Administration GUI to store the updated response file because the GUI process might fail.

If any errors occur during the storage process, check the following:

- Ensure that the certificate has an intermediate CA, convert the response file to .p7b format before you store the renewed certificate with ConsoleOne. For instructions on converting the response file, see TID 10073709 (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10073709.htm).
- Ensure that you have the latest certificate server snapins for ConsoleOne.
- Install the .p7b certificate into Internet Explorer, then view it to make sure there are no problems in the certificate chain.
- Ensure that the time on the iChain server is accurate and that there are no timesync errors in the iChain/NetWare debug consoles. Make sure that your timezone is set correctly. Compare the validity period of your new certificate to the time and date on iChain.
- Remove the ndspkiAdditionalRoots attributes and their associated values on the KMO.

For more information about certificates and renewing certificates that have been issued by an external certificate authority, see How to Renew a Novell Certificate that Was Issued by an External CA (http://www.novell.com/coolsolutions/appnote/2571.html).

Using Advanced Accelerator Features 16

This section describes advanced Novell® iChain® accelerator features. The following topics are discussed:

- Section 16.1, “Using Secure Thin Client Services,” on page 203
- Section 16.2, “Using iChain With Novell Nsure Audit,” on page 216
- Section 16.3, “Logging,” on page 219
- Section 16.4, “Object Pinning,” on page 234
- Section 16.6, “Concurrent Login Restriction,” on page 240

16.1 Using Secure Thin Client Services

The secure thin client provides secure access to thin-client server applications through the iChain Proxy server. iChain provides a secure gatekeeper functionality that protects HTTP applications. By enabling you to move all Web resources out of a demilitarized zone (DMZ) from your secure network to the Internet, iChain effectively removes users' direct access to Web applications and directory information.

As you deploy Web-based solutions to your users, one of the common services is thin client application delivery from vendors such as Citrix (terminal server based on ICA protocol). Using iChain to extend security to these thin client applications requires an iChain authentication before access is granted to the thin client server.

The advantages of using iChain to secure your Citrix thin-client applications include the following:

- Authentication is performed by iChain before any access is allowed to the MetaFrame server or its published applications.
- You don't need to open port 1494 (the MetaFrame server's well-known port) on the external firewall.
- iChain user credentials are not exposed and do not get passed on the wire.
- A one-time iChain token is encrypted and is time-sensitive.
- Single sign-on to MetaFrame servers can be done either with or without Nfuse*.
- You can handle multiple back-end MetaFrame servers through a single secure iChain tunnel.
- Private IP addresses or domain names of back-end MetaFrame servers are not exposed on the wire or in the ICA session file.

The following sections are included:

- Secure Access to Citrix Thin Clients
- Configuring iChain to Accelerate Citrix MetaFrame Servers with Nfuse
- Configuring iChain to Accelerate Citrix MetaFrame Servers without Nfuse
- Using the Citrix Java Client

16.1.1 Secure Access to Citrix Thin Clients

iChain provides secure access to Citrix-based thin-client services that use the Citrix ICA protocol. This includes single sign-on to Citrix Nfuse and iChain-authenticated access to Citrix MetaFrame servers.

Before a session can be established through iChain to a protected Citrix MetaFrame server, the user must have an iChain authentication token. The only way for the user to receive this token is to be authenticated to iChain at the time the Citrix client attempts to connect to the MetaFrame server.

iChain's secure access for the Citrix thin client also provides the additional benefit of supporting Citrix client connections over standard port:80 (client to iChain proxy).

NOTE: If you are running the Citrix Secure Gateway (CSG) solution, you do not need to enable all of the Citrix-related iChain settings. You only need to configure iChain to accelerate the Citrix CSG server in tunnel mode. Tunnel mode is required since unlike HTTP, Secure Gateway ICA connections cannot have their SSL layer terminated and recreated by a proxy server inserted between the ICA Client and the secure gateway server.

16.1.2 Configuring iChain to Accelerate Citrix MetaFrame Servers with Nfuse

To enable a secure iChain tunnel:

1. Create an accelerator (for example, www.metaframe.com).

   This will be the accelerator used for the ICA tunnel.

   If there are multiple MetaFrame servers, you must include all of the IP addresses for all of the MetaFrame servers in the Web server addresses field. Do not enable authentication on this accelerator, because this will be a special MetaFrame accelerator. Do not select Act as a Tunnel. Because this is a special MetaFrame accelerator, refer to Step 3 on how to enable the ICA tunnel.
2. The ICA tunnel accelerator (for example, www.metaframe.com) is a special type of Citrix tunnel that requires a forward proxy authentication between the client and iChain. It acts as a tunnel from iChain to the backend MetaFrame server. You can configure it at the iChain Proxy console. Enter the following command to enable the ICA thin client tunnel:

   ```
   set accelerator accelerator name tunnelauthforica = yes
   ```

   This command is only used on the ICA tunnel accelerator (www.metaframe.com) and not on the Nfuse accelerator (www.nfuse.com).
3. Create another accelerator (for example, www.nfuse.com) that secures the Nfuse Web server. Enable authentication or any of iChain's other advanced features.
4. Configure Form Fill rules for single sign-on to the Nfuse Web server so that Form Fill can fill the redirect information and one-time token into an ICA file. A Citrix client will use this rendered ICA file to launch a session to the iChain accelerator you created in Step 1 (for example, www.metaframe.com).

The following section is an example of how you can configure the Form Fill rules for single sign-on to the Nfuse Web server (www.nfuse.com) and rewrite the ICA file so that the client connects to the ICA tunnel accelerator (www.metaframe.com).

IMPORTANT: Form Fill tags are case sensitive. If you disregard the case with your tags (for example, if you use <icafill> instead of <icaFill>), single sign-on to your back-end applications might fail.

Sample Form Fill Configuration File for Citrix MetaFrame Server with Nfuse

```
<!-- Nfuse Single Sign-On login -->
<urlPolicy>
    <name>MyNfuseLogin</name>
   <url>www.Nfuse.com/Citrix/Metaframe/login.asp</url>
   <actions>
       <fill>
            <input name=”user” value=”~”>
            <input name=”password” value=”~”>
            <input name=”domain” value=”~”>
       </fill>
       <post/>
   </actions>
</urlPolicy>
<!- end of Nfuse login -->

<!-- Nfuse Single Sign-On logout -->
<urlPolicy>
    <name>MyNfuseLogout</name>
    <url>www.Nfuse.com/Citrix/Metaframe/logout.asp</url>
    <actions>
        <redirect>/cmd/BM-Logout.html</redirect>
    </actions>
</urlPolicy>
<!- end of Nfuse logout -->

<!- secure tunnel to MetaFrame with Nfuse -->
<urlPolicy>
<name>NfuseTest</name>
<url>www.Nfuse.com/Citrix/MetaFrame/launch.asp*</url>
<actions>
 <icaFill>
     <icaOriginal>
       [WFClient]
     </icaOriginal>
     <icaReplace>
       [WFClient]
       ProxyHost=www.metaframe.com:80
       ICHAIN-TOKEN
       </icaReplace>
       <icaMetaPrivateAddress>
           Address=1.1.1.1
       </icaMetaPrivateAddress>
       <icaMetaPublicAddress>
           Address=www.metaframe.com
       </icaMetaPublicAddress>
     </icaFill>
</actions>
</urlPolicy>
<!- end of secure tunnel to MetaFrame through Nfuse -->
```

WARNING: This Form Fill example does not work for the newer 9.x ICA clients. See TID 3663363 (https://secure-support.novell.com/KanisaPlatform/Publishing/160/3663363_f.SAL_Public.html)

Understanding the ICA Form Fill Policy

The above-listed NfuseTest Form Fill policy rewrites the ICA file that is generated and sent to the client. When the client connects to the url within the <url> tags, Form Fill will perform the actions listed under the <actions> tag:

```
<url>www.Nfuse.com/Citrix/MetaFrame/launch.asp*</url>
```

The ICA Form Fill begins with the following tag value:

```
<icaFill>
```

The lines following the <icaFill> tag tell Form Fill to find and replace part of the ICA file. This inserts an encrypted one-time, time-sensitive (for example, 60 seconds) iChain token and redirects Citrix clients to www.metaframe.com port 80:

```
   <icaOriginal>
      [WFClient]
   </icaOriginal>
   <icaReplace>
      [WFClient]
      ProxyHost=www.metaframe.com:80
      ICHAIN-TOKEN=60
   </icaReplace>
```

Form Fill parses the ICA file, looking for the [WFClient] keyword, and replaces it with everything in the <icaReplace> section.

The following line is used to redirect a Citrix client to the ICA tunnel accelerator (www.metaframe.com) on port 80:

```
          ProxyHost=www.metaframe.com:80
```

The following ICHAIN-TOKEN keyword instructs Form Fill to inject an encrypted one-time iChain token in seconds, where the default is 180 seconds (the minimum is 60 seconds, and the maximum is 600 seconds). This amount of time is configured by including an =<value> after ICHAIN-TOKEN. This ICHAIN-TOKEN is what the ICA client uses to authenticate to the ICA tunnel accelerator on www.metaframe.com. For example:

```
          ICHAIN-TOKEN=60
```

The next part of the ICA Form Fill policy tells Form Fill to hide the private metaframe server IP address or the domain name with an iChain accelerator name. All of the <icaMetaPrivateAddress> Address= lines are replaced with the Address= line in the <icaMetaPublicAddress>.

```
      <icaMetaPrivateAddress>
         Address=1.1.1.1
      </icaMetaPrivateAddress>
      <icaMetaPublicAddress>
         Address=www.metaframe.com
      </icaMetaPublicAddress>
```

For *1.1.1.1* above, specify the domain name or IP address of the MetaFrame server from the ICA file.

```
      <icaMetaPrivateAddress>
         Address=1.1.1.1
      </icaMetaPrivateAddress>
```

If there are multiple MetaFrame servers, you can add multiple entries within the <icaMetaPrivateAddress> field:

```
      <icaMetaPrivateAddress>
         Address=1.1.1.1
         Address=2.2.2.2
      </icaMetaPrivateAddress>
```

The <icaMetaPublicAddress> field is the ICA tunnel accelerator. Specify the domain name.

For example, in the *www.metaframe.com* above, enter the accelerator name that secures the ICA tunnel.

```
<icaMetaPublicAddress>
   Address=www.metaframe.com
</icaMetaPublicAddress>
```

For troubleshooting information on configuring iChain to accelerate Citrix Metaframe servers, see the Appnote: setting up and troubleshooting the iChain/Citrix integration (http://www.novell.com/coolsolutions/appnote/2562.html).

16.1.3 Configuring iChain to Accelerate Citrix MetaFrame Servers without Nfuse

If you configure iChain to accelerate Citrix Metaframe servers without Nfuse by making use of a static ICA file, be aware that load balancing/failover does not work in this scenario. The controlling entity for load balancing/failover becomes Citrix, which as part of this process, dynamically adds the IP address to connect to in the ICA file. On the Citrix side there is a check on what is the least busy and active server and based on the outcome, the IP address of that server is passed over to be injected in the ICA file. In case of a static ICA file, there are no dynamic updates to the file, so there is no load balancing/failover.

When Nfuse is not deployed, the following steps describe how you can configure iChain in order to accelerate sessions to your MetaFrame servers or applications on your MetaFrame servers.

On the MetaFrame server side:

1. Go to the MetaFrame servers where you want your applications to be secured by iChain.
2. Launch the Citrix Management Console on those servers.
3. Click Applications to display applications currently being published on the server.
4. Right-click the application or the desktop icons that you want iChain to secure and select Create ICA File.

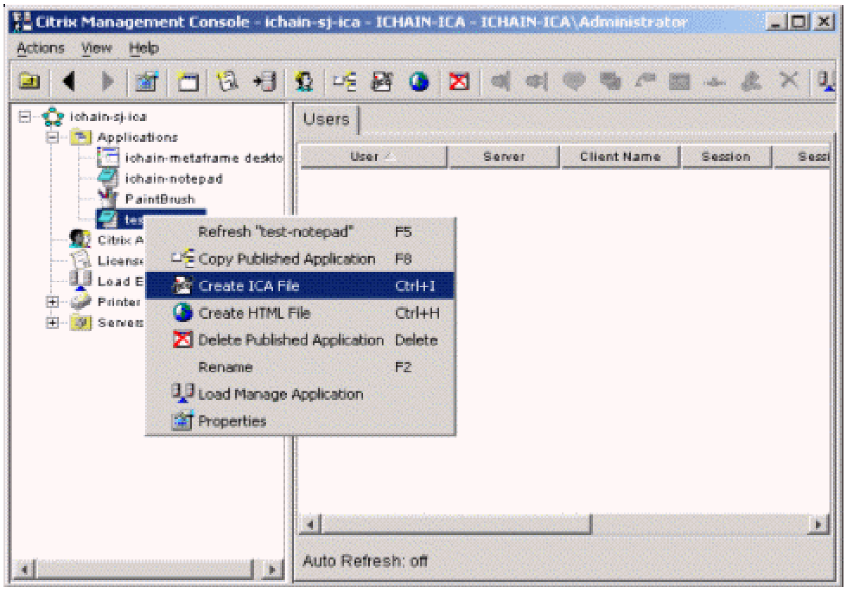

The ICA File Wizard launches.

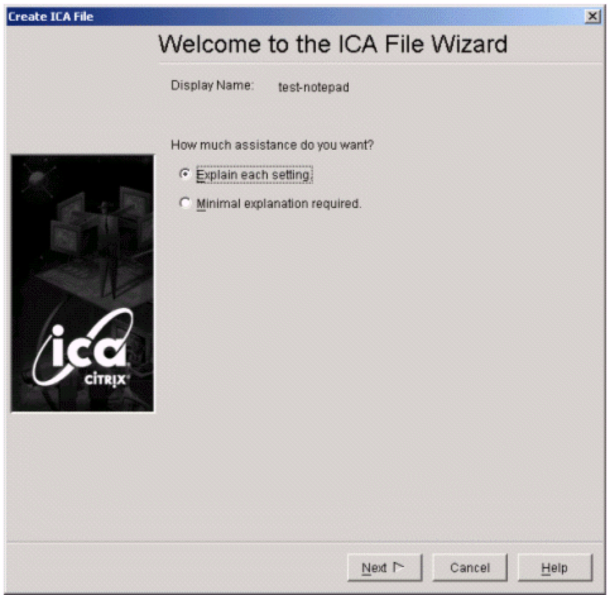

Display Name: test-notepad

How much assistance do you want?

● Explain each setting.

○ Minimal explanation required.

5 Complete the following options in the ICA File Wizard:

5a Specify the session settings.

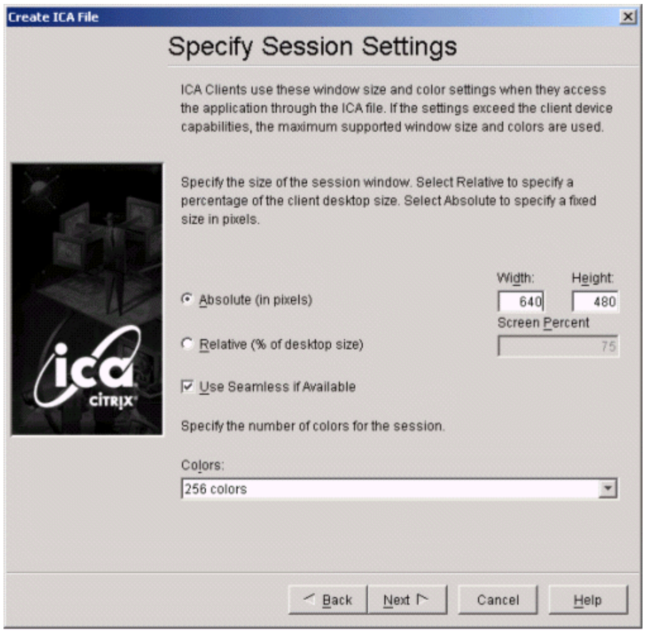

Create ICA File

Specify Session Settings

ICA Clients use these window size and color settings when they access the application through the ICA file. If the settings exceed the client device capabilities, the maximum supported window size and colors are used.

Specify the size of the session window. Select Relative to specify a percentage of the client desktop size. Select Absolute to specify a fixed size in pixels.

● Absolute (in pixels) Width: 640 Height: 480

○ Relative (% of desktop size) Screen Percent 75

☑ Use Seamless if Available

Specify the number of colors for the session.

Colors: 256 colors

5b At the Encryption page, select RC5(128-bit).

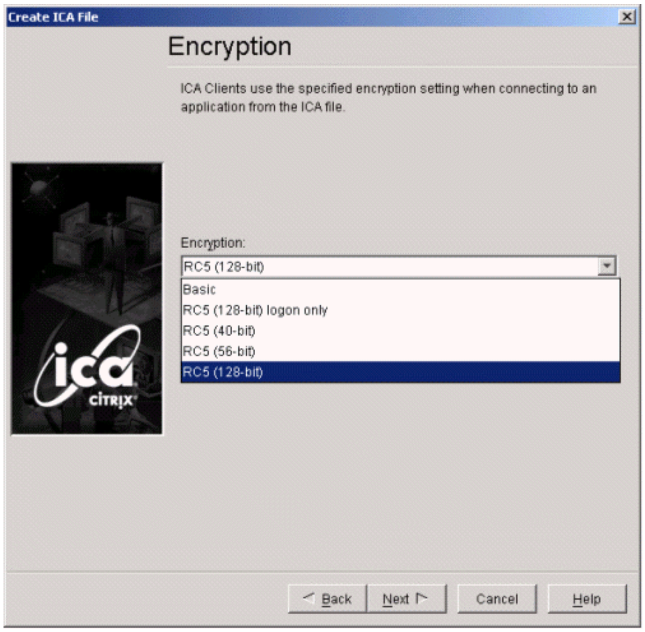

Create ICA File

Encryption

ICA Clients use the specified encryption setting when connecting to an application from the ICA file.

Encryption: RC5 (128-bit)

- Basic
- RC5 (128-bit) logon only
- RC5 (40-bit)
- RC5 (56-bit)
- RC5 (128-bit)

5c Select Compress ICA data stream.

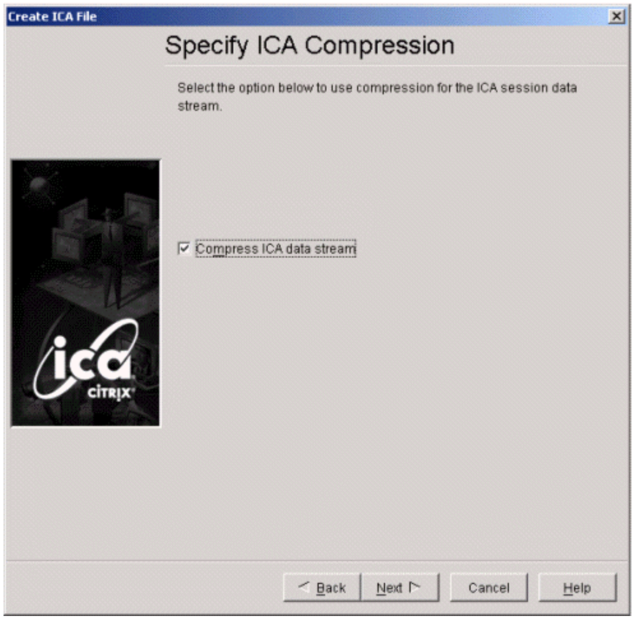

Specify ICA Compression

Select the option below to use compression for the ICA session data stream.

☑ Compress ICA data stream

5d Select the audio setting.

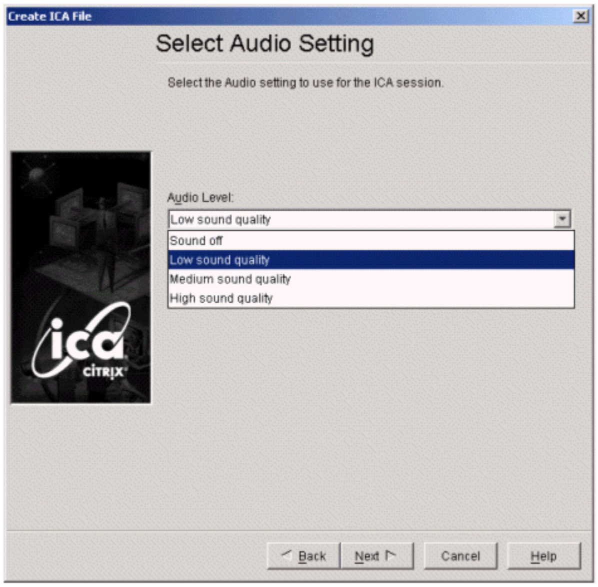

Select Audio Setting

Select the Audio setting to use for the ICA session.

Audio Level: Low sound quality

Sound off
Low sound quality
Medium sound quality
High sound quality

5e On the TCP/IP+HTTP Server page, do not check Use TCP/IP+HTTP Browsing.

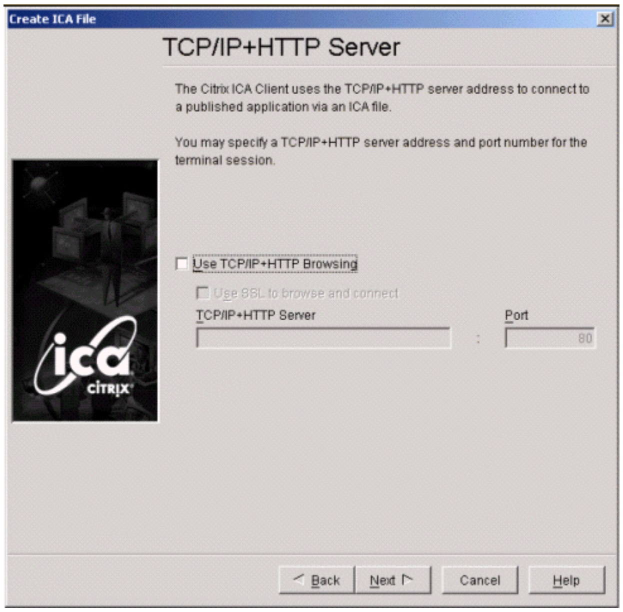

TCP/IP+HTTP Server

The Citrix ICA Client uses the TCP/IP+HTTP server address to connect to a published application via an ICA file.

You may specify a TCP/IP+HTTP server address and port number for the terminal session.

□ Use TCP/IP+HTTP Browsing

□ Use SSL to browse and connect

TCP/IP+HTTP Server : Port 80

5f On the Specify ICA File Name page, specify the path where you want this file to be saved. Make note of the path that you enter.

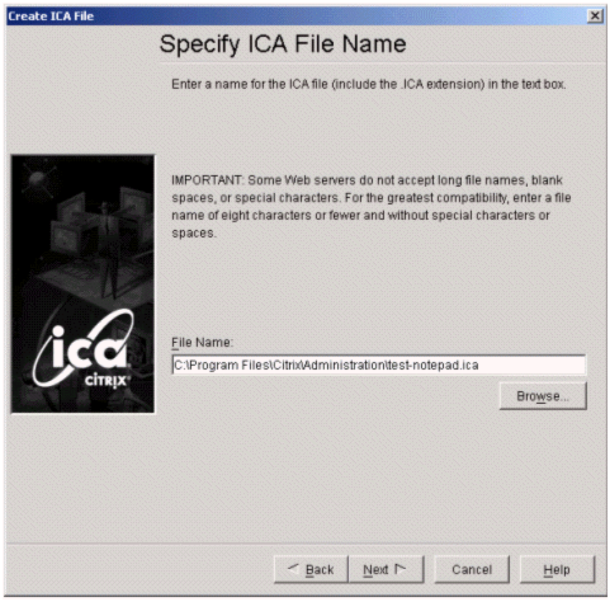

Specify ICA File Name

Enter a name for the ICA file (include the .ICA extension) in the text box.

IMPORTANT: Some Web servers do not accept long file names, blank spaces, or special characters. For the greatest compatibility, enter a file name of eight characters or fewer and without special characters or spaces.

File Name:
C:\Program Files\Citrix\Administration\test-notepad.ica

5g On the Create HTML File page, select No.

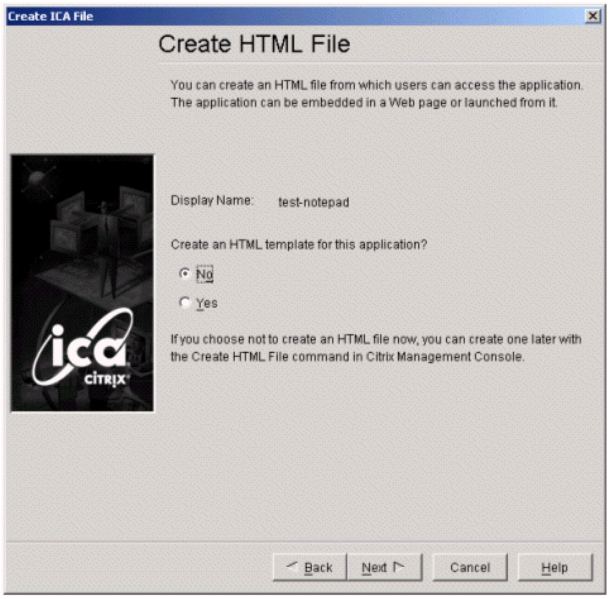

Create HTML File

You can create an HTML file from which users can access the application. The application can be embedded in a Web page or launched from it.

Display Name: test-notepad

Create an HTML template for this application?

● No

○ Yes

If you choose not to create an HTML file now, you can create one later with the Create HTML File command in Citrix Management Console.

5h On the ICA File Summary page, click Finish.

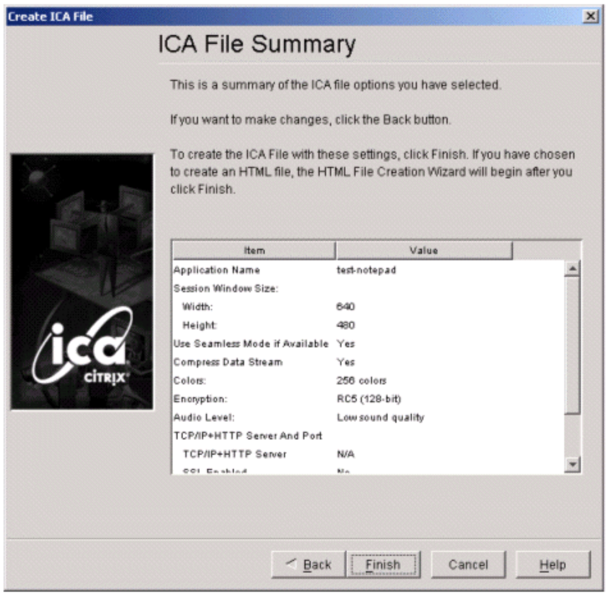

ICA File Summary

This is a summary of the ICA file options you have selected.

If you want to make changes, click the Back button.

To create the ICA File with these settings, click Finish. If you have chosen to create an HTML file, the HTML File Creation Wizard will begin after you click Finish.

Item	Value
Application Name	test-notepad
Session Window Size:	
Width:	640
Height:	480
Use Seamless Mode if Available	Yes
Compress Data Stream	Yes
Colors:	256 colors
Encryption:	RC5 (128-bit)
Audio Level:	Low sound quality
TCP/IP+HTTP Server And Port	
TCP/IP+HTTP Server	N/A
SSL Enabled	No

6 Edit the file you created in Step 5f:

In section [test-notepad], replace the Address=test-notepad with the IP address or domain name from section [WFClient] TCPBrowserAddress=1.2.3.4. Delete the TCPBrowserAddress=1.2.3.4 after forwards.

For example, a test-notepad.ica file is created with the following content:

```
[WFClient]
  Version=2
  TCPBrowserAddress=1.2.3.4
  [ApplicationServers]
  test-notepad=
  [test-notepad]
  Address=test-notepad
```

With the above sample ICA file, you would change it to:

```
[WFClient]
Version=2
  [ApplicationServers]
  test-notepad=
  [test-notepad]
  Address=1.2.3.4
```

7 Edit the file again by adding the following section of code before [WFClient] section. The added code is to make browser launch the Citrix's plug-in accordingly.

```
<%@Language=VBScript%>
<%
Response.ContentType=”application/x-ica”
%>
  [WFClient]
```

8 Save the modified file with an .asp extension and copy it to the Web server that you want iChain to secure.

The .asp file extension is necessary so that browser does not cache the page. This ensures that the browser asks for a new ICA file from iChain if a user clicks the same application again.

On the Web server:

1 Create your own custom HTML page by putting the above .asp file as a reference link.

2 Set up your ACL rules for the custom HTML page you created (including the ACL rules for your .asp file).

3 Configure the Form Fill rules when the .asp file in the HTML page is accessed.

In ConsoleOne®:

```
<!-- MetaFrame login -->
<urlPolicy>
    <name>loginWithoutNfuse</name>
    <url>www.native.com/ica/icalogin.html</url>
   <formCriteria>
         <title>CitrixNativeLogin</title>
    </formCriteria>
    <actions>
         <fill>
            <input name=”Username” value=”~cn”>
            <input name=”ClearPassword” value=”~password”>
        </fill>
    </actions>
</urlPolicy>
<!- end of MetaFrame login -->
<!- secure tunnel to MetaFrame without Nfuse -->
<urlPolicy>
<name>NativeMFTest</name>
```

```
<url>www.native.com/ica/c-*</url>
<actions>
 <icaFill>
      <icaOriginal>
        [WFClient]
      </icaOriginal>
      <icaReplace>
        [WFClient]
        ProxyHost=www.metaframe.com:80
        ICHAIN-TOKEN
      </icaReplace>
      <icaMetaPrivateAddress>
          Address=10.10.0.5
      </icaMetaPrivateAddress>
      <icaMetaPublicAddress>
          Address=www.metaframe.com
      </icaMetaPublicAddress>
       <icaOriginal>
          TransportDriver=TCP/IP
       </icaOriginal>
       <icaReplace>
          TransportDriver=TCP/IP
          ICHAIN-ICA-SSO-POLICY=loginWithoutNfuse
       </icaReplace>
  </icaFill>
</actions>
</urlPolicy>
<!- end of secure tunnel to MetaFrame without Nfuse -->
```

WARNING: This Form Fill example does not work for the newer 9.x ICA clients. See TID 3663363 (https://secure-support.novell.com/KanisaPlatform/Publishing/160/3663363_f.SAL_Public.html)

There needs to be a Form Fill policy that is used for holding the username and password that will be added into the ICA file. In the above example, this is the <loginWithoutNFuse> Form Fill policy. See Chapter 14, “Form Fill,” on page 161 for advanced Form Fill functionality and configuration. The second <name>NativeMFTest</name> Form Fill policy looks similar to the above Form Fill policy used with Nfuse. See the above policy for the examination of the ICA Form Fill Policy. The difference with the Nfuse Form Fill policy and the Form Fill policy without Nfuse is another replacement within the ICA file listed below.

IMPORTANT: Login credentials to the MetaFrame server are inserted as clear text to an ICA file. The Citrix client saves this ICA file locally to launch a session.

```
<icaOriginal>
  TransportDriver=TCP/IP
</icaOriginal>
<icaReplace>
 TransportDriver=TCP/IP
 ICHAIN-ICA-SSO-POLICY=loginWithoutNfuse
</icaReplace>
```

Form Fill looks for the following entries within the ICA file generated in the *.asp file created previously.

```
   <icaOriginal>
     TransportDriver=TCP/IP
     </icaOriginal>
```

When Form Fill finds this entry, it replaces the it with the following:

```
     <icaReplace>
      TransportDriver=TCP/IP
      ICHAIN-ICA-SSO-POLICY=loginWithoutNfuse
     </icaReplace>
```

This instructs Form Fill to add the credentials within the <name>loginWithoutNFuse</name> Form Fill policy in the ICA file. These credentials are inserted into the ICA file as clear text. The user must be authenticated to iChain and the Form Fill Crib attributes must be available in order for Form Fill to insert the credentials listed in the ICHAIN-ICA-SSO-POLICY tag. When the credentials are added to the ICA file, iChain returns the ICA file to the client. The client then uses iChain's rewritten ICA file to connect to the MetaFrame accelerator on iChain with the credentials passed by the ICHAIN-ICA-SSO-POLICY.

For troubleshooting information on configuring iChain to accelerate Citrix Metaframe servers, see the Appnote: setting up and troubleshooting the iChain/Citrix integration (http://www.novell.com/coolsolutions/appnote/2562.html).

16.1.4 Using the Citrix Java Client

To use the Citrix Java Client, your system must meet the following requirements:

- ❑ The Citrix Presentation Server Client for JAVA 9.4 or later. Earlier versions incorrectly handle the encoding of the basic authentication header.
- ❑ iChain 2.3 Service Pack 4 Interim Release IR1a or later. Earlier versions cannot retrieve the IP address which is enclosed in quotes.

You also need to create a form fill policy similar to the following sample policy:

```
<!-start of secure tunnel to MetaFrame with Java ICA clients-->
<urlPolicy>
    <name>JavaCitrix</name>
    <url>nw65.ichainsite.com/citrix/JavaClient/examples/desktop.html</url>
    <actions>
    <icaFill>
        <icaOriginal>
            <param name="Start" value="auto">
        </icaOriginal>
        <icaReplace>
            <param name="Start" value="auto">
            <param name="ProxyHost" value="citrix.ichainsite.com:80">
            ICHAIN-JAVASCRIPT-TOKEN
        </icaReplace>
            <icaMetaPrivateAddress>
                value="10.1.1.195"
            </icaMetaPrivateAddress>
            <icaMetaPublicAddress>
```

```
                value="citrix.ichainsite.com"
            </icaMetaPublicAddress>
    </icaFill>
</actions>
</urlPolicy>
```

WARNING: This Form Fill example does not work for the newer 9.x ICA clients. See TID 3663363 (https://secure-support.novell.com/KanisaPlatform/Publishing/160/3663363_f.SAL_Public.html)

16.2 Using iChain With Novell Nsure Audit

This section provides an overview of the Novell Nsure™ Audit Report auditing system and reviews auditing fundamentals.

The following topics are included:

- Section 16.2.1, "Nsure Audit Overview," on page 216
- Section 16.2.2, "Auditing Background and Fundamentals," on page 216
- Section 16.2.3, "Accessing Novell Nsure Audit Documentation," on page 217
- Section 16.2.4, "Using iChain With Nsure Audit," on page 217

16.2.1 Nsure Audit Overview

Novell Nsure Audit is a centralized, cross-platform auditing service. It collects event data from multiple applications across multiple platforms and writes the data to a single, non-repudiable data store. Nsure Audit is also capable of creating filtered data stores. Based on criteria you define, Nsure Audit captures specific types of events and writes those events to secondary data stores.

Using the query and report generating tools included with Nsure Audit Report, you can then evaluate the information in your data stores to determine resource access, usage patterns, and overall compliance with organizational policies and regulations.Although queries and reports are invaluable in reviewing system activity, sometimes you need to know what is happening on your system as it happens. Therefore, Nsure Audit provides real time notifications and real-time monitoring so you can assess and act on events as they occur.

To some extent, Nsure Audit can even automate the process of responding to events in real-time. The Critical Value Reset (CVR) channel allows you to flag Directory attributes with reset policies. If the value of a given attribute is changed, the CVR channel resets the value as per the policy defined in the CVR Channel object. For example, if your organization has a policy prohibiting security equivalence, you can create a CVR Channel object that automatically resets the Security Equals attribute to a null value if it is ever reset by an administrator.

16.2.2 Auditing Background and Fundamentals

Novell Nsure Audit provides the tools you need to audit your organization's compliance with internal and external policies and regulations; however, the use of secure logging technology such as Novell Nsure Audit does not, in itself, provide a complete auditing solution. Auditing is actually a human-driven process and Novell Nsure Audit is simply a tool that facilitates that process.

Therefore, a complete auditing strategy requires several actions:

1. Define your organization's security and usage policies. That is, determine what resources your users are allowed to access, what rights they have to those resources, and so forth.
2. Log the events relevant to those policies. Configure notification filters to notify you in real time when a policy violation occurs. You can also use notification filters to route the events to the Critical Value Reset (CVR) channel to trigger an automated response to the violation. Perform regular compliance audits. This entails querying the data store for events relevant to your policies and then manually reviewing those events to determine if there are any violations of your corporate policies, when the violations occurred, and who was responsible.

After you have implemented your auditing strategy, Novell Nsure Audit provides the information you need to assess overall compliance with organizational policies and to respond to policy violations in a timely manner.

For example, in a secure environment, you might have a policy that prohibits assigning user rights using the Security Equals attribute because it makes it difficult to track and manage user rights. To audit this policy, you first configure Novell Nsure Audit to log the Change Security Equals event.

To facilitate a timely response to policy violations, you configure a notification filter to send a message to your mailbox any time the Change Security Equals event occurs. You also have the notification filter route the event to the CVR channel, which is configured to automatically reset the Security Equals attribute on User objects to a null value.

You can monitor your organization’s compliance with this policy by using iManager or Nsure Audit Report to query the data store for Change Security Equals events. You then review the query results to determine when violations occurred and who the perpetrators were.

16.2.3 Accessing Novell Nsure Audit Documentation

For the latest Nsure Audit documentation, including information on Nsure Audit setup and administration, go to the Nsure Audit documentation page (http://www.novell.com/documentation/lg/nsureaudit/index.html).

16.2.4 Using iChain With Nsure Audit

iChain 2.3 includes NSure Audit functionality. This section describes how to enable the logging feature within iChain, as well as a description of the events that are available to be logged.

The Nsure Audit configuration functionality is managed through the iChain Command Line Interface (CLI). The configuration can be set and viewed using get log and set log commands. The following two tables list the commands and events.

Command	Description
help get log	Lists a description of the get log command
get log	Lists the available events along with whether they are enabled.
help set log	Lists a description of the set log command.

Command	Description
set log *event* = *yes/no*	Activates or deactivates a given event. For example, set log AuthSuccess = yes turns on the event that notifies when a successful authentication has occurred.
set log all = *yes/no*	Activates or deactivates all events.
set log server address = *ip address*	Configures the IP address of the Nsure Audit server. For example, set log server = 151.155.115.155.
set log server port = *port*	Configures the port number of the Nsure Audit server. By default, the port number is 289.
set log server port - default	Configures iChain to use the default port number of the Nsure Audit server (289).

NSure Audit provides tools to view the events generated by iChain. NSure Audit requires an LSC file that describes the schema associated with the events generated by each product that is instrumented for NSure Audit. The LSC file for iChain is included in the installation of NSure Audit, and is installed as part of that system.

Event	Description
AuthSuccess	A user has successfully authenticated to iChain.
AuthFailed	A user has failed to authenticate to iChain.
IntruderLockout	A user has tripped the intruder lockout by failing to authenticate multiple times (as defined in eDirectory™).
AccessAllowed	Access control has allowed access to a given URL.
AccessDenied	Access control has denied access to a given URL.
CertificateRevoked	The certificate used for mutual authentication has been revoked.
NoCRLAccess	iChain does not have access to the CRL distribution point.
URLNotFound	The user tried to access a non-existent URL.
SystemStarted	iChain has been started.
SystemShutDown	iChain has been shut down.
TimeRestricted	The user does not have access because of a time restriction.
OLACParameters	An OLAC parameter was accessed.
OLACFailed	OLAC failed to produce a given parameter.
FormFillSuccess	A Form Fill form was successfully filled.
FormFillFailed	A Form Fill form was not filled correctly.
PasswordExpired	The user's password has expired.
CertificateExpired	The certificate used for mutual authentication has expired.
URLAccessed	The given URL was accessed.

Event	Description
IPAccessAttempted	The user attempted to access a URL that was specified by an IP address instead of the host name configured in iChain.

Events that correspond with mutual authentication using revoked certificates (CertificateExpired and CertificateRevoked) might not be logged. This occurs because nothing is logged when certificate error pages are enabled. When certificate error pages are disabled, a log entry is created, but it uses the information from a previous successful login and not the current data.

16.3 Logging

Logging of appliance caching activity can be useful for a number of reasons, such as billing for services rendered. The iChain Proxy Server lets you specify how often a new log file is started (rolled over), how long old log files are retained, and the format of the log files.

iChain offers the following logging services:

- You can turn on logging for 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 deletion 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 exists for manipulating and processing these files.

16.3.1 Using Appliance Logging

iChain Proxy Services provides a high performance proxy cache system capable of handling thousands of transactions per second. Although the iChain Proxy Server can log extensive details for each transaction, and although the disk space reserved for log files is quite generous on most appliances, if transaction volume is high and log entries consume a few hundred bytes each, iChain Proxy Services can fill up the available disk space in a matter of minutes.

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
Web Server Accelerator	Yes	Yes
Dynamic Bypass	No*	Yes

* The common log formats used here differ from the industry standard proxy cache common log format.

The Costs of Logging

- "Performance" on page 220
- "Disk Space" on page 220

Performance

Turning on logging for a given service increases system overhead and causes some degradation of performance. 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 222 for formulas you can use to estimate how quickly your logging disk will fill.

System Constraints for Logging

To plan a logging strategy you must know the capacity and limitations of your appliance.

Preset Disk Space

Logging disk space is not user-configurable; it is preset to 1 gigabyte on the iChain Proxy Server. Plan to download and delete log files before the disk space is filled.

Rolling Over Log Files Before Deletion

iChain Proxy Services does 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 that there are 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 for log files to roll over on a schedule that coincides with your download and deletion schedule. This is to ensure that there is 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

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. 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.

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 ensures 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	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. 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 of less than 6 hours, so at least one log file per service is closed and ready to be downloaded and deleted.
Roll Over by Size	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. 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.

16.3.2 Planning Your Logging Strategy

As explained in "The Costs of Logging" on page 220 and "System Constraints for Logging" on page 220, 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 ensures 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 behind the need for logging.

1. Identify the business and other reasons for tracking service transactions.

   Possible examples might include customer billing requirements, statistical analysis, growth planning, etc.
2. Determine which services you need to track.
3. Record this information for further 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. 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.

   For example, if you select URI, selecting URI-stem and URI-query would be redundant because URI = URI-stem + URI-query. Also, logging cookie information can consume a lot of space and might not provide critical information.

   A few bytes can add up quickly when the appliance is tracking thousands of hits every second.

Planning Step 3: Calculating Log Rollover Requirements

As explained in "Log Rollover Options" on page 221, 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 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.

  The default size is 1 gigabyte.

- **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 a service and get more accurate data by accessing the browser-based management tool's Monitoring pages.

- **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 log files per service before the disk is filled. See "Rolling Over Log Files Before Deletion" on page 220.

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* = 1 KB
- *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 or limit the requests being logged.

To calculate the *diskfull_time* for your appliance:

1. Determine the values of the four variables listed in the bullet list above.

   For more information, refer to "Variable Definitions" on page 223.

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.

```
max_roll_time = diskfull_time / 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 over your logs 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:

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 223, 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.

```
max_log_roll_size = logvol_size / (num_services * 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 over your logs when they reach a specific size, the file size must be no more than 100 MB. Otherwise, the system runs out of disk space before you have three complete log files, and scheduling the download and deletion of log files is much more complex.

To calculate the *max_log_roll_size* for your appliance:

1. Determine the values of the three variables in the bulleted list above.
2. Using the *max_log_roll_size* formula, calculate the maximum size a log file should reach before the appliance rolls it over.
3. Record the *max_roll_time* result on your planning sheet.

16.3.3 Configuring Logging Options

Based on the planning you have completed in “Planning Your Logging Strategy” on page 221, 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 222.

   The path for the Web Server Accelerator service is Configure > Web Server Accelerator > Insert > Enable Logging for This Accelerator > Log Options.

The following sections discuss each of the areas within the Log Options dialog box.

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 222.

   Remember that each bit of information you log increases the size of each log entry, and affects 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 222.

Configuration Step 4: Specifying Handling of Older Files

You must schedule 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 (see “About the FTP Log Push Feature” on page 226). 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 228.

The appliance also provides three options for dealing with old files as a failover precaution.

- **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 preserves 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:

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 227, then return to this procedure.

3 In the Log Options dialog box (see "Configuring Logging Options" on page 225), 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 225; otherwise, continue with the next section.

Ideally, iChain Proxy Services will never actually use the old file option you select because you scheduled 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.

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.

- 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.
- Ensure that 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.
- Ensure that your logging strategy is leaving a buffer of free log disk space adequate for possible surges in appliance traffic.
- Ensure that the external storage capacity (FTP server or other storage) is adequate.
- Ensure that all aspects of your logging strategy are keeping pace with increases in traffic through the appliance.

16.3.4 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 and/or on specified days of the week. However, log files cannot be pushed more often than once a day.

The feature operates within the following parameters:

- iChain Proxy Services tries as many times as necessary to establish one connection with the FTP server during the hour of the scheduled push. When the hour changes, iChain Proxy Services stops trying until the next interval you have specified.
- When a connection with the FTP server is established, iChain Proxy Services assumes that the pushing of log files is successful. Any errors that prevent the successful pushing of log files are not detected by iChain Proxy Services.

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, iChain Proxy Services begins trying to establish a connection with the FTP server.

If a connection cannot be established before the hour changes to 1 a.m., iChain Proxy Services 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 iChain Proxy Services doesn't try to connect again until 12 midnight the next day.

16.3.5 Using FTP Push to Automatically Download and Delete Log Files

To configure your appliance to use the FTP Log Push feature:

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 225.

   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 select 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 of the log types to be managed through FTP push must be selected).
   - 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 protects 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 after they have been pushed.

     We recommend deleting log files after they have been pushed unless there is a compelling reason for manually deleting them. Automatic deletion also protects your appliance from sudden surges in traffic.

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.

16.3.6 Manually Downloading and Deleting Log Files

Whenever possible, we recommend that you use the FTP Log Push feature to automatically download and delete log files. See "Using FTP Push to Automatically Download and Delete Log Files" on page 227.

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 information:

- 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 should develop specific procedures for your situation. The following sections contain general ideas for accomplishing these tasks.

- "When to Download and Delete Log Files" on page 228
- "Getting Log Filenames" on page 228
- "Downloading Log Files" on page 230
- "Deleting Downloaded Log Files" on page 230

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
- 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 is not closed until the start of the next day.

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/

File	Location
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/*name* The variable *name* is the name of the Web server accelerator.
Web server accelerator log files in extended format	log:etc/proxy/data/logs/reverse/extended/*name* The variable *name* is the name of the Web server accelerator.
Dynamic Bypass log files in extended format	log:etc/proxy/data/logs/dbypass/extended/

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 supports the CWD command for changing to the target log directories. You can also 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 228.

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	`get comlog forward`
See a list of available Web server accelerator log files in common format	`get comlog reverse:`*`name`* (The variable *name* is the name of the Web server accelerator.)
See a list of available filtering log files in proxy server filtering common format	`get comlog filter`

If You Want To	Then Enter
See a list of available forward/transparent log files in extended format	`get extlog forward`
See a list of available Web server accelerator log files in extended format	`get extlog reverse:`*`name`* (The variable *name* is the name of the Web server accelerator.)

Downloading Log Files

- "Using the Browser-Based Management Tool to Download Log Files" on page 230
- "Using FTP to Download Log Files" on page 230

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 with the get command. You must first obtain each filename by using one of the options explained in "Getting Log Filenames" on page 228.

After you have the log filename, you can transfer the file 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 this command also downloads active log files that are not complete.

The appliance doesn't currently support the FTP server put command.

The following is a list of supported FTP commands:

```
TYPEUSERDELEQUITPWDSYST
```

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

16.3.7 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 can be one of two types: string or 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 (http://www.w3.org/TR/WD-logfile) Web site. 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	The URL stem is everything up to the first question mark (?). If the URL has no question mark, the cs-uri-stem is the same as the cs-uri. This field is redundant if cs-uri is selected.

Name	Description	Type	Selectable	Comments
cs-uri-query	The query portion of the HTTP URL the browser sent to the appliance	non-string	Yes	The query portion is the first question mark through the end of the URL. If the URL has no question mark, cs-uri-query has no value. This field is redundant if cs-uri is selected.
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.
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 brought the request from that proxy
x-origin-ip	The IP address of the origin server	non-string	Yes	Assumes the appliance brought the request directly from the origin server

16.3.8 Enabling and Viewing the ACL Rule Log File

The logs for authorized access attempts and unauthorized access attempts can be turned on or off globally. The rules for logging authorized access attempts for an individual access control list (ACL) can also be turned on or off. However, because unauthorized access attempts are usually the result of a user not being defined in any ACL rule, logging of unauthorized access attempts cannot be turned on or off for individual ACL rules.

To enable or disable ACL rule logging on a global level:

1. Access the URL of the iChain Proxy Server where you installed the iChain Proxy Services software to launch the proxy server browser-based administration tool.

   For example, http://*xx.xx.xx.xx*:1959/appliance/config.html where *xx.xx.xx.xx* is the IP address.
2. Click Configure > Web Server Accelerator > Modify.
3. Select the Enable Logging check box.

To enable or disable ACL rule logging for an individual ACL rule:

1. In ConsoleOne, right-click an ACL Rule object.
2. Select Properties.
3. Check the Authorized Logging check box.

The ACL log files for each 24-hour period are saved to log:\etc\proxy\data\logs\reverse\extended\aclcheck\*yymmdd-a*.log, where *yymmdd* is today's date represented by two digits for the year, month and date. The default maximum size of the file is 1 MB. The default size can be changed; see "Using ACLCHECK options" on page 150 for more information. If a log exceeds the maximum size, a new file named *yymmdd-b*.log is created.

Each file contains the following fields:

- Date
- Time
- Source IP address
- Destination IP address
- Protocol
- Source port
- Destination port
- TCP flag
- Access (Allow=1, Deny=0)
- IP headers
- IP payload
- Username
- Destination host name
- URL (the user requested)
- Rule Object name (if access was allowed; if denied, the field displays a —)

16.4 Object Pinning

This section contains the following topics:

- "The Pin List" on page 234
- "URL Mask" on page 234
- "Pin Type" on page 234
- "Pin Links" on page 235
- "Pin Images" on page 235
- "Refresh Frequency/Time" on page 235
- "Processing URL Masks" on page 236
- "Wildcards in Pin Lists" on page 238
- "Pin List Examples" on page 238

16.4.1 The Pin List

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.

Pinned objects remain in the cache indefinitely unless it fills up. This ensures that the lists are available from cache and are not bumped out by more recently requested objects.

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 238.

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 "Processing URL Masks" on page 236.

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 specifies whether and how the appliance caches objects that match the URL mask.

- **Normal:** iChain Proxy Services 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/* with a pin type of normal. This causes all files, including .jpg files, in the graphics directory on the foo.gov Web site to be cached as requested.

They are not, however, pinned in cache because of the normal pin type. Assuming there are no other URL masks in the pin list, all other JPG graphics are not cached because of the /*.jpg mask.

- **Cache:** iChain Proxy Services keeps the pinned objects in cache as long as possible, although they might be written to the appliance's hard disk.
- **Memory:** iChain Proxy Services 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:** iChain Proxy Services does not cache the objects. In other words, you can use this option to prevent objects from being cached.

  For more information about the cache bypass list, see TID 10097536 (http://support.novell.com/docs/Tids/Solutions/10097536.html).

Pin Links

This specifies how many link levels iChain Proxy Services will follow for 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 is downloaded and cached when the pin list is applied along with all the items linked from the page. These cached objects are also refreshed at the frequency and time specified.

To use levels 1 or 2 you must specify an absolute address, including the scheme, host, and path for the URL mask, for example, http://www.foo.gov/documents/. The tool lets you insert masks that do not meet this requirement, but the entries are removed when you click Apply.

Attempting to include an asterisk wildcard immediately hides this option.

Pin Images

This 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 iChain Proxy Services will follow when downloading a page.

For example, if the requested HTML page uses images that reside on another host and you have selected this option, the HTML page is 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 are also pinned.

On the other hand, if the Pin Images option is not checked, iChain Proxy Services only pins the images that reside on the same host as the requested page.

Refresh Frequency/Time

This lets you specify a refresh frequency and time for the URL that is different from the default values shown above the pin list.

Processing URL Masks

There are four basic types of URL masks you can enter 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 iChain Proxy Services.

Type	URL Mask Examples by Specificity	Notes
Hostname	http://www.foo.gov/documents/picture.gif http://www.foo.gov/documents/ http://www.foo.gov foo.gov/documents/ foo.gov/ *.foo.gov/	Although these entries can include the protocol or scheme, the DNS name, the path, and the filename, only the DNS or hostname must be present in the mask. All DNS label portions must be indicated, if only by an asterisk wildcard. iChain Proxy Services processes hostname entries before it processes other mask types. It also processes the most specific URL mask entries first. When an object match occurs, iChain Proxy Services applies the pin type rule, and processing of the object is finished. For example, if the first URL mask in the examples column has a pin type rule of bypass, picture.gif is not cached regardless of the pin type rules for the other URL masks. Hostname entries can have a dramatic impact on object pinning and cache bypassing. 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. 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. Finally, Configure 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.

Type	URL Mask Examples by Specificity	Notes
Path	/documents/picture.gif /documents/picture.gif/ /documents/	iChain Proxy Services processes path entries after all hostname entries have been considered. It assumes that the first forward slash immediately follows a hostname. A leading forward slash must always be used when specifying a directory. The leading slash always references the root directory of the Web server. For example, the first entry applies only to a graphics file named picture.gif that is located in a documents directory at the root of the host. The forward slash in the second entry causes iChain Proxy Services 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. The third entry applies to any matched objects that contain a documents directory at the Web server's specified root directory.
Filename	/picture.gif /widget.js /default.htm	After the path entries have all been processed, iChain Proxy Services looks for specific filenames. A leading forward slash must be used and, as opposed to a path-based mask, does not reference the root directory of the Web server. For example, if requested files named picture.gif, widget.js, and default.htm have not been covered by one of the hostname or path entries above, the files have the pin type rule for their respective filename mask applied to them. 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 are not cached.
File Extension	/*.gif /*.js /*.htm	File extension entries are processed last. These are simply filename entries with the root of the filename replaced by an asterisk, which makes them less specific that complete filenames. A leading forward slash must be used and, as opposed to a path-based mask, does not reference the root directory of the Web server. 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.

Wildcards in Pin Lists

Only the asterisk (*) wildcard is allowed in pin list entries.

iChain Proxy Services 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	As a general rule, you should always include fully qualified DNS or hostnames in the pin list. iChain Proxy Services resolves these more quickly than other masks, and you will be able to track the effects on pinning more easily. For this URL mask, iChain Proxy Services downloads, caches, and pins all objects whose URL starts with the mask. In other words, all objects below the documents directory are downloaded, cached, and pinned. Also, all objects that are linked from one of the pinned objects are downloaded, cached, and pinned. And finally, images that reside on other hosts are downloaded, cached, and pinned. Objects are refreshed according to the refresh settings (default or specific) as specified in the pin list entry.
www.foo.gov/groups.html	cache	1	No	iChain Proxy Services 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.
www.foo.gov/groups.html/	normal	1	Yes	iChain Proxy Services downloads and caches objects in the subdirectory named groups.html and in pages linked from any of those objects. The forward slash at the end of the path tells iChain Proxy Services that this is a directory rather than a file. Objects are cached but not pinned in cache, meaning they might be bumped by more frequently accessed objects or objects that are pinned. Images linked from other hosts are downloaded and cached.

URL Mask	Pin Type	Pin Links	Pin Images	Effect on Cache
www.foo.*	bypass	n/a	n/a	iChain Proxy Services doesn't cache objects from any URLs whose DNS names begin with www.foo. All domain extensions (.com, .net, .org, etc.) are covered by the asterisk wildcard. Link and image pinning is not available for bypass pin types. If this entry appeared in a pin list with either of the previous two entries, it does not prevent caching of objects covered by them because it is less specific than they are.
w*.f*.com	bypass	n/a	n/a	iChain Proxy Services 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. This mask doesn't prevent caching of objects on other domains such as .net, .gov, etc.
w*.f*.*	bypass	n/a	n/a	This mask functions like the previous entry, but the domain is not limited to .com.
.foo.	cache	n/a	n/a	This causes all objects on any Web server whose second domain label is foo to be pinned in cache. Link and image pinning are not available because the mask contains asterisks. This mask does 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 happens to resolve in DNS to the same IP address as www.foo.gov, the iChain Proxy Server applies the pinning rules specified for www.foo.gov to foo.gov. To understand more about IP addresses and URL masks, see Section 16.5, “Using the Proxy Server to Record IP Addresses When Resolving URL Masks,” on page 239.

16.5 Using the Proxy Server to Record IP Addresses When Resolving URL Masks

As stated earlier, you should include fully qualified DNS or hostnames in URL masks whenever possible.

The iChain Proxy Server 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.

Because both URLs resolve to the same IP address, iChain Proxy Services treats objects for both DNS names the same.

On the other hand, if www.foo.gov and foo.gov resolve to different IP addresses, separate pin list entries are required to cover both sites.

16.5.1 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.

16.5.2 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 the following 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 Configure > 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.

16.6 Concurrent Login Restriction

You might want to limit the number of times a specific user can log in. For example, if you set a user's login limit to three, iChain does not let that person log in again after the login limit is reached. If the user opens multiple sessions, iChain closes the first login instance when the limit is reached.

To set up the concurrent login option, enter the login restriction settings in the iChain console.

NOTE: Concurrent Login Restrictions should not be used in a Session Broker setting.

1. Open the iChain console screen.
2. Type the following command:

```
set authentication limitconcurrentlogins = (yes/no)
```

This turns on the concurrent login restrictionfeature. For example

set authentication limitconcurrentlylogin=yes

When it is set to No, the concurrent login setting is disabled. When it is set to Yes, use the commands in Step 3 to control the functioning of the feature.

If you are using SSL as an authentication method for your accelerators, make sure that the Send and Error Page When a Mutual SSL Error option is enabled. Otherwise, a blank page appears when users reach their authentication limit.

3 Set up the concurrent login option using the following commands:

3a Type `set authentication maxlogins =` *`(nonzero positive integer)`*`,` then click Apply.

This sets the number of concurrent logins that are allowed. For example

set authentication maxlogins=4 (or a number you choose).

After the maximum number of logins is reached, the user is denied access or an older instance is logged out.

3b Type `set authentication logoutoldest= (yes/no).`

This determines what action to take when the maximum number of logins is reached. For example, `set authentication logoutoldest=yes.`

When it is set to Yes, the least recently accessed connection of the user is logged out and a new login occurs. When it is set to No, the new login is rejected with a message indicating that the maximum number of logins has been exceeded. The default is No.

4 After you set up the concurrent login option, reboot the iChain Proxy Server.

Rewriter Support 17

Novell® iChain® provides an internal rewriter and a custom rewriter. This section provides information on the purpose and use of both. The following topics are discussed:

- Section 17.1, "The Internal Rewriter," on page 243
- Section 17.2, "The Custom Rewriter," on page 252

17.1 The Internal Rewriter

The iChain internal rewriter is used to accomplish the following:

- To rewrite URL references with the proper scheme (HTTP or HTTPS).

  For example, an HTML file being accessed through an iChain accelerator for the Web site mynovell.com might contain a URL reference to http://mynovell.com/file1.html. If the accelerator for mynovell.com is using SSL sessions between the browser and iChain, the URL reference http://mynovell.com/file1.html must be rewritten to https://mynovell.com/file1.html. Otherwise, when the user clicks this link, the browser bounces between HTTP and HTTPS to establish a new SSL session.

- To rewrite URL references that contain private IP addresses or private DNS names with the public DNS name of the iChain accelerator.

  For example, suppose that a company has an internal Web site, internal.web.site.com, and wants to expose this site to Internet users through an iChain accelerator using a public DNS name of mynovell.com. Many of the HTML pages on this Web site have URL references that contain the private DNS name, such as http://internal.web.site.com/docs/file1.html. Because Internet users are unable to resolve internal.web.site.com, links using this URL reference would return DNS errors in the browser.

  The internal rewriter can resolve this issue. The DNS name field in the accelerator configuration is set to mynovell.com, which users can resolve through a public DNS server to the accelerator's public IP address. The rewriter parses Web content retrieved through the accelerator, and any URL references matching the private DNS name or private IP address listed in the Web server address field of the accelerator are changed (rewritten) with the public DNS name mynovell.com and port number of the accelerator.

  Rewriting URL references addresses two issues: 1) URL references that are unreachable because of the use of private DNS names or IP addresses are now made accessible and 2) Rewriting prevents the exposure of private IP addresses and DNS names that might be viewed as sensitive information.

- To rewrite the Host header in incoming HTTP packets to the name expected by the internal Web server.

  Using the example above, suppose that the internal Web server expects all HTTP or HTTPS requests to have the Host field set to internal.web.site.com. When users send requests using the public DNS name mynovell.com, the Host field of the packets in those requests received by iChain is set to mynovell.com. iChain can be configured to rewrite this public name to the private name expected by the Web server by enabling the Alternate Host Name option, then entering the value internal.web.site.com in the adjacent field. Before iChain forwards packets to the Web server, the Host field is changed (rewritten) from mynovell.com to internal.web.site.com.

The following sections describe how the internal rewriter works and how to configure it:

- Section 17.1.1, "Which URL References Are Rewritten?," on page 244
- Section 17.1.2, "What Other Criteria Are Considered?," on page 244
- Section 17.1.3, "Configuring the Internal Rewriter," on page 246
- Section 17.1.4, "Sample Rewriter Scenario," on page 249
- Section 17.1.5, "Disabling the Internal Rewriter," on page 251
- Section 17.1.6, "Internal Rewriter Summary," on page 252

17.1.1 Which URL References Are Rewritten?

The internal rewriter searches and parses Web content that passes through the accelerator and that meets certain criteria (see "What Other Criteria Are Considered?" on page 244) for URL references qualified to be rewritten. URL references are rewritten only under the following conditions:

- URL references containing DNS names or IP addresses matching those in the accelerator's Web server address list are rewritten with the accelerator's DNS name.
- URL references matching the accelerator's Alternate Host Name field are rewritten with the accelerator's DNS name.
- URL references matching entries in the [Alias Host Names] section of the rewriter.cfg configuration file are rewritten with the accelerator's DNS name. Details on the use of this file can be found in "Configuring the Internal Rewriter" on page 246.

17.1.2 What Other Criteria Are Considered?

The following criteria are considered when determining whether URL references should be rewritten:

- Query Strings
- HTTP Headers
- JavaScript
- HTML Tags
- Mime Types
- Absolute and Relative References
- Path-Based Multi-Homing

Query Strings

The internal rewriter does not rewrite URL references contained within query strings. Only the hostname portion of the reference is evaluated for rewriting.

HTTP Headers

The internal rewriter rewrites qualified URL references occurring within certain types of HTTP response headers such as Location and Content-Location. The Location header is used to redirect the browser to where the resource can be found. The Content-Location header is used to provide an alternate location where the resource can be found.

JavaScript

Within JavaScript*, only absolute references are evaluated for rewriting. Relative references and absolute paths are not attempted. Absolute paths (/path/file.html) are evaluated if the file is read from a path-based multi-homing accelerator's origin Web server and the reference follows an HTML tag. For example, the string href=‘/path/file.html' is rewritten to href=‘/accelPath/path/file.html'.

HTML Tags

URL references occurring within the following HTML tags are evaluated for rewriting:

- action
- archive
- background
- base
- cite
- code
- codebase
- data
- dynsrc
- href
- longdesc
- lowsrc
- onclick
- pluginspage
- src
- usemap
- usemapborderimage

NOTE: Value is not a default tag.

Mime Types

The rewriter parses pages with certain Mime Content-Types regardless of the file extension. By default, the internal rewriter parses pages with the following Mime Content-Types:

- text/html
- text/css
- application/x-javascript

If an HTTP or HTTPS response has a Mime Content-Type set to any of the above types, or if the file extension is html, htm, shtml, jhtml, asp, jsp, or NO EXTENSION, the page is parsed for possible rewriting.

Absolute and Relative References

An absolute reference is a reference that has all the information needed to locate a resource, including the hostname, such as http://internal.web.site.com/index.html. The internal rewriter always attempts to rewrite absolute references.

A relative reference is a reference that assumes the host/path and provides only the resource of the URI, such as index.htm. The internal rewriter does not attempt to rewrite a relative reference.

An absolute path is a reference that assumes the host. It provides the complete path, including the resource. The internal rewriter attempts to rewrite an absolute path only when it is defined in a path-based multi-homed accelerator.

Path-Based Multi-Homing

With an accelerator configured for path-based multi-homing, absolute references and absolute paths are evaluated for rewriting. Relative references are not attempted.

17.1.3 Configuring the Internal Rewriter

The behavior of the internal rewriter can be controlled through use of the sys:/etc/proxy/rewriter.cfg configuration file. This section provides information on the parameters that can be used within this file.

Several configuration sections can be added to rewriter.cfg, including [Mime Content-Type], [Extension], [Exclude], [Javascript Variables], [Javascript Calls], and [Alias Host Names]. Each section is detailed below.

Remember the following conditions when configuring the internal rewriter:

- Sections within the file must be separated with two returns or empty lines.
- If the first part of the line contains a pound sign (#) or semi-colon (;), the line is considered a comment line.
- For the changes made to become effective, you must apply the changes. To have the changes apply to previously cached pages, you need to purge the cache.

[Mime Content-Type]

In addition to files with extensions listed in the [Extension] section below, the rewriter parses pages with certain Mime Content-Types regardless of the file extension. By default, the internal rewriter parses pages with the following Mime Content-Types:

```
[Mime Content-Type]           text/html
text/css
text/xml
text/javascript
application/javascript
application/x-javascript
```

Text/plain is not a default content type.

[Extension]

It is unusual for data to come back from a Web server without a content type. If it does happen, the file extension is compared. If a file doesn't have an extension, it always matches the null case, and the rewriter parses the page. The default extensions are as follows:

```
[Extension]
html, htm
shtml, jhtml
asp, jsp
js
css
```

Additional file extensions can be added by using the [Extension] section of rewriter.cfg, as shown below:

```
[Extension]
home, myNewExtension
anotherLongExtension
a,b,c,d,e,f,g
```

As shown in this example, additional extensions can be specified on individual lines or by using commas to separate multiple extensions specified on a single line. All of the extensions listed are appended to the default list shown above. You cannot remove the default extensions.

[Exclude]

Exclude keeps the entity data from being rewritten if the requesting URL has a match in the exclude list. It checks the requested URL and doesn't check the URL returned from the Web server. For example:

```
[Exclude]
http://www.a.com/dont_rewrite            ; Match without ending slash
http://www.a.com/dont_rewrite/           ; Match with ending slash
http://www.a.com/dont_rewrite/index.html  ; Match specific file name
http://www.a.com/dont_rewrite/*          ; Includes all files and
                                          ; subfiles and the three
                                          ; examples above
http://www.a.com/*                        ; Turn off rewriting for
                                          ; this accelerator and all
                                          ; path-based children.
```

[Javascript Variables]

You can add JavaScript variables to look for URL references. This adds to the HTML variable list (href=, src=, onclick=). The word parser edits the variable to the form [w]variable[ow]=[ow], where [w] represents a space, [ow] represents an optional space, and the URL reference follows.

For example, suppose you use the following JavaScript variables:

```
[Javascript Variables]
headingGif
plusGifUrl
minusGifUrl
stylesheetUrl
```

A JavaScript example would be:

```
var headingGif = "/path/path/file.gif";
```

[Javascript Calls]

A JavaScript call can be made within JavaScript code or in HTML code, such as:

```
onclick='open("/path/file.html")'
```

By default, all parameters of a JavaScript call are looked at for rewriting. The first JavaScript call within an HTML tag (href, src, onclick, etc.) has all of its parameters parsed for rewriting.

This adds to the word parser the call to parse the parameters within a JavaScript call that appears in JavaScript code, not within the HTML URL tags.

The word parser edits the variable to the form "[w]call([ow][url]...)" where the URL references are parameters within the calls.

For example:

```
[Javascript Calls]
openWindow
```

[Alias Host Names]

Sometimes a URL reference specifies a hostname that does not meet default criteria for being rewritten; that is, it does not match the accelerator's alternate hostname or any value in the Web server address list. For example, assume that a URL reference contains the hostname of home (http://home/index.html), and home is not included in the Web server address list because it is not resolvable, nor is it the value of the accelerator's Alternate Host Name field. By default, rewriting of the URL reference http://home/index.html would not occur. The [Alias Host Names] section of rewriter.cfg can be used to specify additional hostnames to be rewritten. Using alias hostnames only applies for absolute URLs. You cannot use alias hostnames for relative URLs.

The following is an example of how to use the [Alias Host Names] section. It has the following syntax:

```
[Alias Host Names]
AcceleratorName=aliasName
```

where AcceleratorName is the value specified in the accelerator's Name field, and aliasName is the string that is rewritten with the value specified in the accelerator's DNS name field.

For the home example used above, if the accelerator name is accel2 and the URL reference to be rewritten is the hostname home, the correct syntax is:

```
[Alias Host Names]
accel2=home
```

NOTE: The alias names are not case sensitive because hostnames should not be case sensitive.

You can use [Alias Host Names] to add and remove the set of hostnames, schemes, ports, and paths that are used to represent the identity of an accelerator's Web server.

The following example illustrates the syntax to use to rewrite these items and then provides examples:

```
#add an alias hostname
acceleratorName=aliasHostName

#add a full host reference
# http://alias
acceleratorName=scheme://aliasHostName
# http://alias:80
accleratorName=scheme://aliasHostName

#add an addition subpath for a path-based multi-homed accelerator
# that does NOT have the remove subpath option checked
acceleratorName=/additionalPath

#Remove a hostname from being used to rewriting to this accelerator.
# Use this option when there are multiple accelerators that contain
# the same alternate hostname and Web server port.
acceleratorName!=alternateHostName

[Alias Host Names]
novell=HOME
novell=https://www.backend.com:444
download=/fileDownloads
novell!=testserver
```

17.1.4 Sample Rewriter Scenario

Assume that an accelerator is set up with the following information:

```
Accelerator Name:       novell
DNS Host Name:          www.novell.com
Alternate Host Name:    www.backend.com
Web Server Address:     151.155.1.1
Web Server Port:        80
```

- The Web server delivers content through schemes HTTP and HTTPS on ports 80 and 443, respectively. The accelerator is set up to talk to Web server port 80, and to rewrite the https references.

  ```
  [Alias Host Names]
  novell=https://www.backend.com
  ```

  This is a common issue. If you use two accelerators, one on port 80 and the other on port 443, you are forced to have two listeners on the public side also listening to ports 80 and 443. This makes authentication to this site very difficult. The public side listening on port 80 must authenticate on a port other than 443 because two accelerators cannot share the same listening port.

  If you want only an HTTPS (secure) connection to the browser, using an alias works as long as the Web server can serve up the content on port 80.

  The setting of

  ```
  [Alias Host Names]
  novell=https://www.backend.com:443
  ```

  gives the same results because 443 is the default port for HTTPS.

- On an Oracle Web server, the hostname of HOME appears in some URL references (for example, http://HOME/path/file). This hostname does not appear in any DNS table and should not be used as an alternate hostname because it is a real hostname for the Web server in addition to HOME.

  For example:

  ```
  [Alias Host Names]
  novell=HOME
  ```

- Suppose that an IS department is phasing out a DNS hostname of www.novell.com and prefers that users use the new DNS hostname of thenew.novell.com. This means that there are two accelerators with two different sets of DNS hostnames that are both reading from the same back-end Web server name of www.backend.com. All references for http://www.backend.com are rewritten to the original requesting hostname. This should not be an issue.

  The issue for the rewriter is when a third Web server has a reference such as http://www.backend.com. This could be rewritten as http://www.novell.com or http://thenew.novell.com. The following setting removes the back-end name of www.backend.com from the rewriter list for the accelerator that has the public name of www.novell.com. In this example, oldname! is the name of the accelerator.

  ```
  [Alias Host Names]
  oldname!=www.backend.com
  ```

This has a side effect of always rewriting http://www.backend.com to http://thenew.novell.com, even if you are reading from the www.novell.com accelerator.

- Consider that an accelerator is a path-based multihomed child that does not remove the child sub-path from the URL. To add additional paths to re-route requests to this multihomed child, you would do the following:

```
{Alias Host Names]
downloads=/downloads1
downloads=downloads2
```

17.1.5 Disabling the Internal Rewriter

There are three methods you can use to disable the internal rewriter:

- Disabling Per Accelerator
- Disabling Per URL
- Disabling In a Page

Disabling Per Accelerator

By default, the internal rewriter is enabled for all accelerators. The internal rewriter can slow performance because of the parsing overhead. In some cases, a Web site might not have content with URL references that need to be rewritten. The internal rewriter can be disabled on a per-accelerator basis using the SET command on the command line of the iChain machine. The following is an example of how you would use this command:

```
SET ACCELERATOR <name> DisableRewriter=Yes
```

where <*name*> is the name of the accelerator for which you want to disable rewriting. This action is permanent upon reboot and is exported to the .nas file.

Disabling Per URL

The rewriter.cfg file also allows you to specify a list of URLs which are to be excluded by the rewriter. For example:

```
[exclude]
http://www.abc.com/xyz/*
http://www.abc.com/donotrewrite.html
```

As shown in this example, the exclusion causes all pages in the xyz subdirectory and the donotrewrite.html file to be left untouched. The syntax of the URLs requires them to be prefixed by http, and the domain name of the accelerator also must be defined.

Disabling In a Page

In some circumstances, you might find that you need more granularity. There are cases when only part of a page cannot or should not be rewritten. Although this deviates from the premise of iChain that you shouldn’t have to modify the origin server, you might encounter circumstances where it cannot be avoided.

In these cases, you can use the following tags in your origin pages.

For example:

```
<!--NOVELL_REWRITER_OFF-->
.
.
HTML data not to be rewritten
.
.
<!--NOVELL_REWRITER_ON-->
```

These tags are seen by browsers as a comment mark, and do show up on the screen (except possibly on older browser versions). Also, the last tag is optional, and if omitted, it prevents the rest of the page from being rewritten after the initial tag is encountered.

NOTE: If the page has been cached before you add the comment to turn off the rewriter, the page must be purged from cache.

17.1.6 Internal Rewriter Summary

- The internal rewriter is on by default for all accelerators.
- It reads sys://etc/proxy/rewriter.cfg for configuration settings.
- The accelerator's DNS Name and appropriate port number are always rewritten by the rewriter.
- The rewriter compares URL references to values in the accelerator's Alternate Host Name, Web server address fields, and to the [Alias Host Names] section of rewriter.cfg to determine whether a rewrite should occur.
- It looks for URL references within files that pass the MIME type check.
- If the MIME type is not found in the packet, the file extension is examined. By default, rewriting will be attempted for files with no extension and for files with the following extensions: html, htm, shtml, jhtml, asp, jsp.
- Within JavaScript, only absolute references are rewritten.
- It looks for URL references in HTTP header types content-location and location.
- It rewrites absolute references and absolute paths from a path-based multi-homing accelerator.
- It does not support rewriting non-UTF-encoded nested URLs, such as the following:

  ```
  <a href="javascript:document.forms.queryPropertyForm.action="url""></a>
  ```

  The second double quote character triggers iChain to cut off the URL completely, for example:

  ```
  "javascript:document.forms.queryPropertyForm.action="
  ```

  To avoid this problem, use one of the following formats for the nested URL:

  ```
  <a href="javascript:document.forms.queryPropertyForm.action='url'">
  <a href='javascript:document.forms.queryPropertyForm.action="url"'>
  ```

17.2 The Custom Rewriter

The Novell iChain custom rewriter (rewrite.nlm) was developed because of the need to replace URL text pieces. For example, a document could have the following JavaScript variables:

```
var scheme = “http://”
var host = “internal.web.site.com”
var path = “/path/”
var file = “file1.html”
var URL = scheme + host + path + file;
```

These variables are used to build URL strings for the user. The internal rewriter can only rewrite true URL references such as http://internal.web.site.com/path/file1.html to the schema and hostname of the accelerator that accesses the internal Web site (https://mynovell.com/path/file1.html).

The custom rewriter also enables administrators to search and replace user-specified strings with new strings. For example, a user might want to change all occurrences of “President” to “Secretary” in the data portion of the page.

The custom rewriter has an overlapping region problem when used in conjunction with the internal rewriter. If both rewriters want to change the same characters of data within a file, it is undefined which rewriter takes precedence. In some cases, a few of characters around the rewritten area are lost or added. This problem happens when the custom rewriter is used to rewrite the same URLs that the internal rewriter also rewrites. The custom rewriter was not developed to search and replace full URL references; only URL pieces (as shown in the above example)

The following sections explain how the custom rewriter works and how to configure it:

- Section 17.2.1, “Rewrite Filter Configuration File,” on page 253
- Section 17.2.2, “Enabling the Custom Rewriter,” on page 255
- Section 17.2.3, “Replacement Rules,” on page 257

17.2.1 Rewrite Filter Configuration File

You configure the custom rewriter by creating a rewrite filter in a text file. The custom rewriter reads the configuration file and replaces all occurrences of a user-specified string with the specified replacement string, irrespective of the location of the original string in the data. An example of a configuration file is given below:

```
[Name=oraclefilter]

[Extension]
html, htm

[Mime Content-Type]
Text/html, text/css

[URL]
sjf-siva.sjf.novell.com/oracle/*
sjf-siva.sjf.novell.com/database/OracleQuery.html
sjf-siva.sjf.novell.com/oraclereports/

[Replace]
var xsport     =”9003”<<>> var xsport    =”443”
var xsname     =”xyz01”<<>> var xsname   =”www”
var xconnectMode = “http”<<>> var xconnectMode = “https”
<PARAM name=serverHost   value=”xyz01.novell.com”><<>><PARAM name=serverHost
value=”www.novell.com”>
<PARAM name=serverPort value=”9003”><<>><PARAM name=serverPort value=”443”>
<PARAM name=connectMode value=”https”><<>><PARAM name=connectMode  value=”http”
```

This file has the following elements:

[Name=<filtername>]: The name of the rewriter filter. It is used for filter identification in rwfilter commands.

[Extension]: The user can specify file extensions (in a single line) that need to be parsed. This section is optional, but if this section is not specified, all extension types are parsed. Specifying this section results in parsing only specified extensions.

NOTE: If possible, you should use an [extension] section. Otherwise, the custom rewriter processes all kinds of files (including .zip files and .iso files).

All extensions should be specified in a single line (just after the Extensions section header). Individual extension elements should be separated by a comma.

For example:

```
[Extension]
html
```

This entry enables the following behavior:

- Parses files with html extensions.
- Parses http://www.novell.com/index.html.
- Does not parse http://www.novell.com/logo.gif.
- Does not parse http://www.novell.com.

For example:

```
[Extension]
html, htm, txt
```

This entry enables the following behavior:

- Parses files with extensions html, htm, txt, and any default files implied by a trailing slash (/).
- Parses http://www.novell.com/license.txt.
- Parses http://www.novell.com/.
- Does not parse http://www.novell.com/logo.gif.

[Mime Content-Type]: The user can specify mime types to be rewritten by the custom rewriter. If this section is present, it takes precedence over the [Extension] section. Any files that match the mime type are rewritten, regardless of their extension.

All mime types should be specified in a single line (just after the mime content-type section header). Individual mime-type elements should be separated by a comma.

For example:

```
[Mime Content-Type]
text/html
```

This entry enables the following behavior:

- Parses files with extensions as a mime type text/html.

- Parses http://www.novell.com/index.html.
- Does not parse http://www.novell.com/logo.gif.
- Does not parse http://www.novell.com.

[URL]: This section lists all URLs to be parsed. The rules of URL specification are similar to the ACLCheck module as follows:

```
www.novell.com              : Exact match
www.novell.com/contact/     : All files in contact directory, but
                              not in subdirectories.
www.novell.com/contact/*    : All URLs starting with www.novell.com/
                              contact, including subdirectories of
                             contact.
```

[Replace]: The user can specify replacement pairs in this section. The string <<>> is used as a separator between the original string an the new string. The format is as follows:

```
[Replace]
<search string><<>><replacement string>
```

For example, to replace www.novell.com with support.novell.com, you would use the following format:

```
www.novell.com<<>>support.novell.com
```

All the lines after this Replacement section header are treated as replacement string pairs, so this should be the last section in the file.

Each replacement pair (including the last one) must be terminated with the end of the line.

17.2.2 Enabling the Custom Rewriter

You can enable the custom rewriter at the system console. The name of the module is rewrite.nlm. The configuration file or files can be supplied as command line parameters to the load command.

Filename restrictions: The rewrite.nlm module does not place any restrictions on the configuration filename format, but files that do not conform to DOS filename format (8.3) could cause errors while loading the file.

Load Time Command Line Parameters: The load command uses the following parameters.

```
rewrite [-s] -f filename [ -f filename]*
```

For example:

```
rewrite [-s] -f system/accel1.rw -f etc/accel2.rw
```

The rewrite command supports the following parameters:

Parameter	Description
-s	Creates a separate screen for the rewrite filter. This parameter is optional

Parameter	Description
`-f` *`filename`*	Specifies the configuration file. This is the file from which configuration information is loaded. The user can specify as many configuration files as required. In the example above, accell.rw and accel2.rw are two configuration files in the specified format.

Runtime Console Commands

The following runtime console commands are also available to activate filters, stop filters, and to view current configuration information.

Command	Description
`rwfilter unload` *`filtername`*	Stops a filter (the filter name is as given in the configuration file).
`rwfilter load` *`configFile`*	Loads a filter configuration file.
`rwfilter list`	Lists all loaded rewriter filters by name.
`rwfilter print` *`filtername`*	Displays rewrite configuration information (the filter name is as given in the configuration). Printed information should be similar to the configuration file.

NOTE: The rwfilter list and rwfilter print *filtername* work only if you loaded rewrite.nlm with the -s option.

Configuring the Custom Rewriter to Start on Reboot

iChain 2.3 SP4 IR3 introduces a new way to automatically load custom rewriter configurations on reboot. Previously, custom rewriter commands were added to the end of the appstart.ncf file. As this file is replaced during the upgrade process, the commands to load the custom rewriter would have to be manually added after each update.

The appstart.ncf file now includes the following line:

```
sys:\etc\custom\customrw.ncf
```

If this file exists, the commands in it are executed automatically on reboot.

Additionally, all files in the sys:\etc\custom\ directory are added to the end of the current.nas file for backup and restore purposes. It is recommended that all custom rewriter configuration files be placed in this directory.

To have the custom rewriter automatically start when you reboot iChain, put the following line in the sys:\etc\custom\customrw.ncf file:

```
load rewrite -f filename
```

Replace filename with the path and name of your configuration file.

Multiple configuration files can be specified as follows.

```
load rewrite -f file1 -f file2 -f file3
```

The optional -s parameter can also be specified to create a separate rewrite filter screen.

If you are using a build of iChain prior to 2.3.329c (iChain 2.3 SP4 IR3), put the commands to load the custom rewriter files at the end of the sys:\system\appstart.ncf file. Because this file is replaced during the upgrade process, you need to manually add these modifications after each update.

17.2.3 Replacement Rules

The following are simple rules that the iChain rewrite filter uses:

- String replacement is done as a single pass.
- String replacement is not performed recursively. For example:

  ```
  [Replace]
  DOG<<>>CAT
  A<<>>O
  ```

  If the original string is DOG, the rewritten string is CAT. That is, only one replacement will occur (CAT will not be further replaced with COT).
- Because string replacement is done in one pass, the string that matches first takes precedence. For example:

  ```
  [Replace]
  ABC<<>>XYZ
  BCDEF<<>>PQRSTUVWXYZ
  ```

  If the original string is ABCDEFGH, the replaced string is XYZDEFGH.
- If two user-specified strings match the data portion, the original string of longer length is used for replacement except in cases detailed above. For example:

  ```
  [Replace]
  ABC<<>>XYZ
  ABCDEF<<>>PQRSTUVWXYZ
  ```

  If the original string is ABCDEFGH, the replaced string is PQRSTUVWXYZGH.
- The rewrite filter module, rewrite.nlm, applies only one matching filter (each configuration filter file constitutes a filter) for a request. If multiple filters match a request URL, the matching filter that was configured most recently is used for data parsing.

Using Start Up Options and the Command Line Interface 18

This section provides information about the Command Line Interface (CLI).

- Section 18.1, “Customizing the Command Line Prompt,” on page 259
- Section 18.2, “Using the Enhanced Configuration Export,” on page 260
- Section 18.3, “Managing Appliance Security Features,” on page 263

18.1 Customizing the Command Line Prompt

You can change the prompt displayed on the command line interface. This allows you to customize the prompt so you can more easily recognize which machine you are connected to when you are using a switchbox. This allows a limited number of variables, such as time and version information. To view information about customizing the prompt, enter the command:

```
set prompt?
```

The syntax of the command is as follows:

```
Set the appliance Prompt
 syntax:  set prompt = prompt
          set the prompt using $ parameters
          Default prompt is set to $G
 parameters:
   <prompt> can contain the following special characters
  $B      : Build Number
  $D      : Current Date
  $G      : '>'
  $L      : '<'
  $N      : Appliance DNS Alias
  $T      : Current time (24H format)
  $S      : ' '
  $V      : Version
  $_      : Carriage Return
  $$      : '$'
  Any other $ combination will delete the $
  example: set prompt = $T$S$V$S($B)$_$G
          gives the following prompt
          '16:55:03 iChain 2.3 (2.2.217d)'
          '>'
```

This prompt can contain any other character. For example, if you set the prompt as:

```
Set prompt=iChain$G
```

it displays as iChain>

This prompt is only shown on an unlocked machine. On a locked machine, > is still the default. This is to protect information on the locked machine.

This setting is not available from the Proxy Administration Tool. It is stored in the .nas file.

18.2 Using the Enhanced Configuration Export

You can export all of your certificates, trusted root files, rewriter.cfg, and all files in sys:\etc\custom to the .nas file. Sys:\etc\custom is meant for custom rewriter configuration use. The behavior of the export command is controlled by the following set export ? commands.

```
syntax1: set export certificate = <no|yes|auto>
           - enables / disables the export of certificates into
             the NAS file.
                 - no  = disable
                 - yes = enable, take the certificates from the
                         backup dir
                 - auto= enable, export the certs to the backup,
                         then include
syntax2: set export trustedroot = <no|yes>
           - enables / disables the export of trusted roots
             the NAS file.
                 - no  = disable
                 - yes = enable, take the trustedroots from SYS:
syntax3: set export password = <password>
           - sets the password for the certificates export in case
             auto is enabled.
```

This password is not exported to the .nas file.

Certificates and trusted roots were not included as defaults in previous versions; however, the sys:etc\proxy\rewriter.cfg and all files in sys:etc\custom are now included. This allows you to export your entire configuration in a single, convenient location inside one file.

This setting is not available from the Proxy Administration Tool. It is stored in the .nas file.

IMPORTANT: If you export a .nas file from a machine that has a trusted root configured on the Access Control tab, the .nas environment does not work when you restore the file. To prevent this issue, do not export the current.nas file unless you set the export trustedroot=yes and make it the default setting.

Importing an iChain 2.2 .nas File

If you import an iChain 2.2 .nas file that contains password management servlet information to an iChain 2.3 server, the password information might be lost because the SNMP files change. If this occurs, you need to reboot.

You can import iChain 2.3 .nas files with the same information , and the password information is not lost.

18.2.1 File Format

The .nas file format is shown as follows:

```
# iChain(r) Configuration file
#
# Build: iChain 2.3 (2.3.222d)
# Date:  Wed Dec 10 12:10:29 MST 2003
# File:  abc.nas
```

```
#
# ====================== WARNING ======================
#  This file contains security sensitive information!
#  - LDAP usernames and passwords
#  - Certificates
#  - Trusted roots
#  - IP addresses and infrastructure information
#
#      PLEASE TREAT THIS FILE WITH EXTREME CAUTION
#
# ====================== WARNING ======================
#
clear accelerator
.
.
.        < All kinds of settings - as before >
.
.
set export certificate=auto
set export trustedroot=on
.
.
.        < All kinds of settings - as before >
.
set prompt=$G
.
.
.        < All kinds of settings - as before >
.
restore initialize
restore begin \etc\proxy\rewriter.cfg
restore end
#
# This section contains the Trusted roots of the iChain server
# They will be restored and automatically active.
#
restore trustedroot begin 175TR 1285 FF1A9CC1
*8G0CG205XCG273fe030YX0YY0vY0K5UY5ma9EGKDQpGiN7ODGgH1C(NLxeKY0XA8Gj069CgcK8cA4l
*8GXutj0XX05510G9Dn1wG0O623L0ah0pHJlIRdXRk9RwXRK9RFkRXCAWZH1n1FG0j623L0aA0p6lKG
*8GYIgiIbtIfJ8GU0Nj9Go9Gn9oL9nL9qn9GGGQN1jn9oG9no9Ln9tq9nGAGQ9GD8nw8GO063GLaxLY
*8G30hp2HlRIdRXkR9wRXKR9FRkXPCWIZ18nF8Gj063GLa0Ap26gIibItfHJGW2X92G0j619gYc8cgO
*8Gadct0jX0XX0504320XF10GW2X0AYW2X0X0hjO2pnpR5pjaSTW5a)tmErI)4pgnj4NoMZE5W2mhN)
*8G5ShtZS3eOy2wDaWDf)o0TOiTqkpj09hUOsrXRejhTYXJE5p5hOOXMEesXO1Hl9Giy2qZL0Qa(Z7T
*8G6QKzuBYqyuY44jY3GZrcbC(2TBDT8(eEE)rjhlRgMZHP0xlyJqhJul)UedWqemUphhd4KloHffmL
*8GdxoX03e)243AYeXIAJuc(spCeKZ7Wl6Y1QN4gohgPOQ4YQlR1BxaBMyU)rGfRn9Jjs5WD)h9jVaL
*8Ge0r9k9eIlngyMCBVP6JTNHycDlplnUZpCHEVXns2cVsZ8Ae6kc7dIMgPHD(iY2UJELhKb(ALVwtP
*8G9OO3aqPLQTc8OZVCJJCNBVIGBrI3Bza(nmY7cl9WeHwbgY03jfNndgWISB0Y30X05X3W2Y1IG16P
*8GAW2Y1kG0A623L0Tk0a30aX1XG0C623L1TZ0a58G3WWX1XG0F623L0Tp0XX00a15G03X7XV8GkC2n
*8Gh063GLT0FX0X00aa03Y0X6CG26XF06hS0cG8Xdcu8tX09a0XX0X04a25XRCG25XN0aY0X00XX0Ie
*8GCuVp2TERFsR5CPCWJJ5R3rRI9RKvAW1RKKRI9RYrRK5B8KPjf2sZReKRKm9QlBlaR5sR5CRFmkv3
*8GjR5IBEkRFsR5CPCER3FPjlRI5RmFRp9RKFRIvBlXRKKRI9RYrRK5PplR35RIKRXKRKIQpVPsnN88
*8Gk9GEReKPjGW2XL800wX0X08Ge8G60YX0XY2Xc8Ge8G60YX0XY0XA0YXT9X0wX0X08Ge8G60YXjIy
*8GF0XY2Xc8Ge8G60YX0XY0XA0YXT9Y06Y0XO0XXzV3W2X5a0GuY0XY0YY70V0YX0030j0WW0000LDX
*8Gm0000000000003940W00000000010G1OG0mY0X00YeV)V)VV)VV)VV0XX00Y0a6(G)H8G1OGMGv
*8GH0mY0X00YeV)V)VV)VV)VV0XX00Y0a6(G)L8XGuY0XY0YY70V0YX0030j0GW0000000000000G1H
*8GI00003920W00000000010G1OG0mY0X00YeV)V)VV)VV)VV0XX00Y0a5UY5naG1OG0mY0X00Yezuh
*8GpV)V)VV)VV)VV0XX00Y0a5UY5raYHEGGiY0XY0YY70V0YX0030j0dWV)VV)VV)VV)VV)VV039IuD
*8GK40W)VV)VV)VVvVG1IG0mY0X00YeV)V)VV)VV)VV0XXvVG1IG0mY0X00YeV)V)VV)VV)VV0XXw8W
*8GrvVG0j619gYc8dct0jX0X50504320XX00HLbvm6OvcVUmXAc1O5KnAqTrHIcVHB1iNibnh19i9PS
```

```
*8GseY5l6bU0kz(moySMuItkin6BdPofyqxZePj1EVUk0kVUZxPmz23XEDqInRmaZlKE7jyV7At89Dn
*8GNUAvTF9QOBjvkUdapNaXVZIWgI(fzzwXDLmkjT)NF4iwBltC8iroTkO0HjcV9CqgDVNja6lTBlNx
*8GOZca4AGM29sN6dQ1WzCS5xDW0sX)fA71YV4AEHuEz8KAJ8xQmh)Y(W65X)VabCzi3)K10VanPxEm
*8GvKbE8gFdJtuBXikGS)xty2FloG(6U606UO0X9zoObodaWB)bioyL4sQ8(FnTgepr2X77jSOYFiez
*8bwneQ)87kOhe0TZ1huqdYiZdJ492lHWB0pOmOwTHAXbi8P6lNGEbhD4m0w0ORh
restore trustedroot end
apply
```

18.2.2 New Commands

The .nas file has the following new commands:

```
restore trustedroot begin name size crc
data
.
.
.
restore trusted root end
```

and

```
restore trustedroot begin name size crc data
.
.
.
restore trusted root end
```

and

```
restore certificate begin name size crcdata
.
.
.
data
restore certificate end
```

Also updated are:

```
restore begin filenamedata...**BLANK_LINE**
.
data
restore end
```

where **BLANK_LINE** is inserted when a blank line is found. This is a limitation of the parser.

The certificate and trusted root data are protected by a double CRC (cyclic redundancy code). A CRC is calculated on a line-by-line basis, as well as on the entire file. If either of the CRCs do not match, or the resulting file size does not match, the file is not restored to its original location. In such a case, the old file is left in place.

The trusted root is restored to sys:\ and is available for immediate use. The certificate date is restored to the backup directory on the sys: volume and is not immediately active so that it does not overwrite active certificates. It can be restored either by using the Proxy Administration Tool in the Certificate Maintenance tab, or by using the following instructions as outlined in the .nas file if the server contains a certificate:

```
#
# This section contains the Certificates of the iChain server.
# *.pfx files will be restored to the backup directory.
# To activate them you need to restore the certificates using
# the certificate restore option on the certificate menu.
#
# If you want to automate this process, please include the
# following commands after the last "Restore Certificate End"
# command, before the apply. (without the "#")
#
# add certificate name <name>
# set certificate name <name> target=disk, action=restore,
#     action=<password>
#
# where <name> is the name of the certificate
# and <password> is the certificate password.
# Repeat this for every certificate in the list.
#
# WARNING: In that case, anyone can dupe your server.
#          This is a security risk!!!
#
```

WARNING: This allows automated staging setups or lab setups. If you use this feature in production environments, you must guard this file since it contains all information needed to clone your iChain server.

The Get EXPORT PASSWORD command is not exported to the .nas file for security reasons.

18.2.3 Using a Comment Specifier in NAS Files

You can use the pound sign (#) and semi-colon (;) to start a comment line in the file. Anything containing these characters is ignored, except if they are specified between a `restore` *`certificate|trustedroot`* `begin` and `restore |`*`certificate|trustedroot`* `end` specifier.

18.3 Managing Appliance Security Features

This section contains the following topics:

- "Using the Console Lock Feature" on page 263
- "Accessing Proxy Internals" on page 264

18.3.1 Using the Console Lock Feature

The iChain Proxy Services 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 is unlocked, it remains unlocked until you lock it using the lock command.

18.3.2 Accessing Proxy Internals

A few iChain features require the administrator to access the internals of the iChain Proxy Server.

WARNING: Changes to the internal proxy server configuration should be limited to those items specified in authorized iChain documentation or as directed by Novell support personnel. Undocumented changes might result in proxy server malfunction and might require Novell support personnel to request a software re-image for that server.

To access the internals of the proxy server, enter the following command at the proxy server console:

```
debug
```

```
proxy_debug_password
```

The proxy_debug_password is "proxydebug".

Editing the Tune.ncf File

Edit the tune.ncf as instructed in the tune.ncf, or enter the following SET parameters at the server console:

```
SET NCP INCLUDE IP ADDRESSES = ALL
SET NCP EXCLUDE IP ADDRESSES = NONE
SET NCP OVER UDP = ON
```

WARNING: Remember that editing this file can create a security hole. One way to reduce this risk would be to enable a single interface for NCP™ access using the set parameter SET NCP INCLUDE IP ADDRESSES = *IP_address_of_private_interface*. This provides access only on the specified interface. You should disable login when you are finished editing the file.

To edit the sys:\system\tune.ncf file, complete the following steps:

1. At the NetWare® System Console, load EDIT.

   To get to the NetWare System Console, you must first unlock the ICS console by entering `unlock`, followed by the password when prompted. (The password is the config user's password.)

2. Enter `debug`. The debug password is proxydebug.
3. After NCP is enabled on the server, use the following credentials:

   ```
   User: ichainadmin
   Password: novell
   ```

For information on how to enable NCP on the iChain server in order to edit the tune.ncf file, see the Novell Technical Information Document (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10065889.htm).

18.3.3 Setting Up Enhanced Security Within the Authentication Cookie

iChain can require at least 48 bits of random data to be matched for authentication instead of only 16 bits. You can set up this secure cookie requirement by entering a command line startup when loading the proxy. The following is the switch you need to enter:

```
Proxy -cv 2
```

Using the Proxy Services Interface 19

The Novell® iChain® Proxy Server is a key component of the iChain infrastructure. The iChain Proxy Server is optimized to perform the functions needed for your iChain infrastructure. This sections details the configuration and management of the iChain Proxy Server.

The following topics are discussed:

- Section 19.1, "Using the Proxy Administration Tool," on page 267
- Section 19.2, "The Home Panel Options," on page 269
- Section 19.3, "The System Panel Options," on page 273
- Section 19.4, "The Network Panel Options," on page 285
- Section 19.5, "The Configure Panel Options," on page 295
- Section 19.6, "The Monitoring Panel Options," on page 328
- Section 19.7, "Using Appliance Commands," on page 333
- Section 19.8, "Performing Patch Upgrades," on page 341

19.1 Using the Proxy Administration Tool

The iChain Proxy Server supports a Proxy Administration Tool, allowing you to manage and administer your iChain Proxy Server from a browser. To launch the tool, you simply access a special management URL on the iChain Proxy Server. The URL must contain either the 172-net management address or the IP address you have already configured for the server, followed by:1959/appliance/config.html. For example,

http://172.16.0.1:1959/appliance/config.html

NOTE: If the iChain Proxy Server is located behind a firewall and you are accessing the Proxy Administration Tool from a browser outside that firewall, you must open ports 1959, 2222, and 51100 on the firewall to administer the proxy server.

19.1.1 Prerequisites for Running the Administration Tool

You need the following:

- ❑ A proxy server that has been initialized and is currently running
- ❑ A Java-enabled browser, such as Internet Explorer 5.5 (or higher) running on your workstation
- ❑ Sun JRE 1.5 SP4 on your workstation. If you use an older version, you might experience the following problems:
  - A delay when selecting Configure > Authentication. For more details, see "Windows XP" on page 268.
  - If you use a large PIN list, the entire list might not display. Earlier versions can handle approximately 762 entries.
- ❑ SSL 2.0 and SSL 3.0 (where available) enabled on the browser

- ❑ A network or cross-over cable connection to the proxy server
- ❑ The IP address of the proxy server

After the appliance has been configured with an IP address and mask, a gateway server, and a DNS server, you can administer it over the network via any client that can communicate with it over IP.

Until you have completed that configuration, you must use a cross-over cable and a client with the following constraints:

- Client IP address set to 172.16.0.2 (or another available 172-net IP address) with a mask of 255.255.255.255
- Client gateway address set to 172.16.0.1 (the management address of the appliance)
- Client DNS server address set to 172.16.0.1

Windows XP

If you are running the Proxy Administration Tool on a Windows XP machine with Sun JRE, you might experience a delay when selecting Configure > Authentication. While the page is loading, you cannot access the Proxy Administration Tool until all of the authentication profiles are added.

A possible workaround is to go to the control panel and double-click the Java Plug-in, then select the Browser tab and deselect Microsoft Internet Explorer so that the Sun JRE is not used.

If you need to use the Sun JRE, we recommend that you use the Sun JRE 1.5 SP4 version to avoid the slowness issue on Windows platforms. You can use the Microsoft* JVM, but you need to already have it downloaded because it is no longer available for download.

19.1.2 Starting the Administration 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 172-net management address or an IP address you have already configured on the appliance, followed by :1959/appliance/config.html, for example:

   http://172.16.0.1.: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 displays an error indicating that the page cannot be displayed.
4. Enter a password if you have previously specified one for the appliance.

19.1.3 Applying and Cancelling Changes

As you make changes to appliance parameters in the Proxy Administration 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 immediately effective. If you change the NTP server, the appliance time changes with the next synchronization cycle (normally about 15 minutes).

Except in the cases just mentioned, clicking Apply commits all changes made in any page 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.

When you click Apply or Cancel, the action cannot be undone.

19.1.4 The Help Button

Click the Help button in the left frame to display the online documentation with a table of contents in the left frame. To navigate through the documentation, click the titles in the table of contents.

19.1.5 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.

19.2 The Home Panel Options

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.

- Section 19.2.1, “Introduction Page,” on page 269
- Section 19.2.2, “Health Status Page,” on page 270
- Section 19.2.3, “Certificate Maintenance Page,” on page 272

19.2.1 Introduction Page

Path: Home > Introduction

Figure 19-1 *Introduction Page*

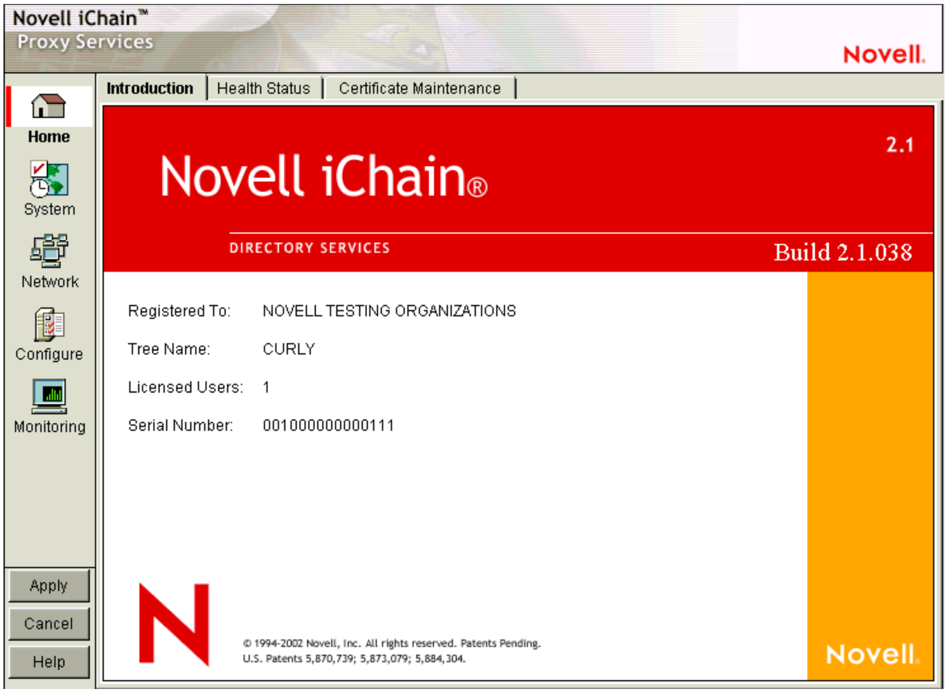

Registered To: NOVELL TESTING ORGANIZATIONS

Tree Name: CURLY

Licensed Users: 1

Serial Number: 001000000000111

© 1994-2002 Novell, Inc. All rights reserved. Patents Pending.
U.S. Patents 5,870,739; 5,873,079; 5,884,304.

The Introduction page displays iChain Proxy Server information, such as version, build number, and licensing information.

19.2.2 Health Status Page

Path: Home > Health Status

Figure 19-2 Health Status Page

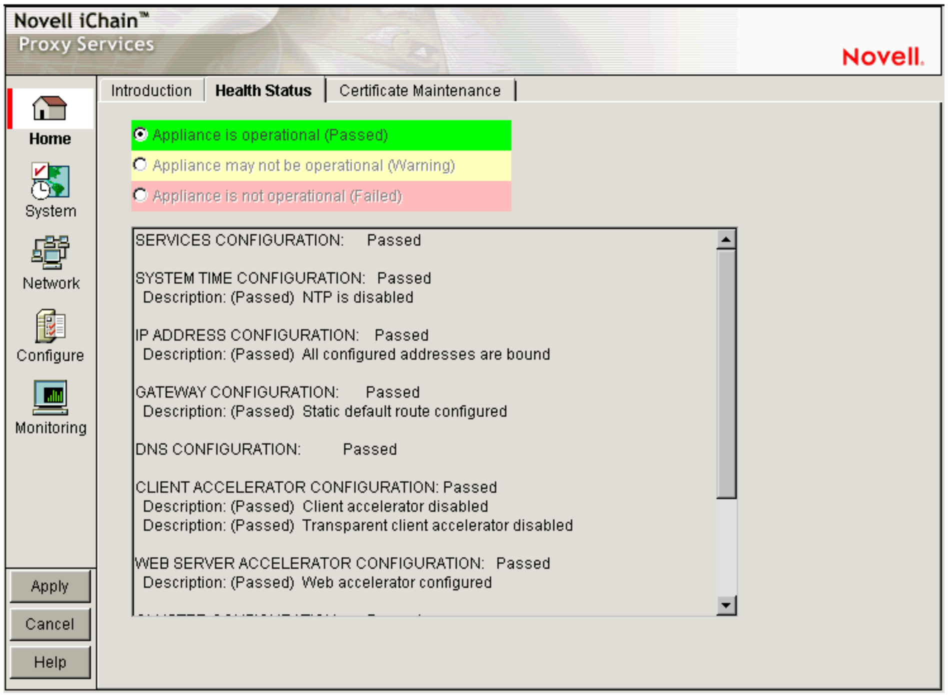

● Appliance is operational (Passed)

○ Appliance may not be operational (Warning)

○ Appliance is not operational (Failed)

```
SERVICES CONFIGURATION:    Passed

SYSTEM TIME CONFIGURATION:  Passed
 Description: (Passed)  NTP is disabled

IP ADDRESS CONFIGURATION:   Passed
 Description: (Passed)  All configured addresses are bound

GATEWAY CONFIGURATION:      Passed
 Description: (Passed)  Static default route configured

DNS CONFIGURATION:          Passed

CLIENT ACCELERATOR CONFIGURATION: Passed
 Description: (Passed)  Client accelerator disabled
 Description: (Passed)  Transparent client accelerator disabled

WEB SERVER ACCELERATOR CONFIGURATION:  Passed
 Description: (Passed)  Web accelerator configured
```

The Health Status page indicates general status of appliance configurations, including which services are currently configured and the operational status of selected services.

A green status indicates that iChain Proxy Services has not detected any configuration discrepancies.

A yellow status indicates that iChain Proxy Services might be functioning sub-optimally because of configuration discrepancies.

A red status indicates that the iChain Proxy Services 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. iChain Proxy Services 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 appliance have Web browsing problems, check the status here and in "Services Page" on page 329.

Web Accelerator Configuration: Reports the 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 Page" on page 329.

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).

19.2.3 Certificate Maintenance Page

Path: Home > Certificate Maintenance

Figure 19-3 Certificate Maintenance Page

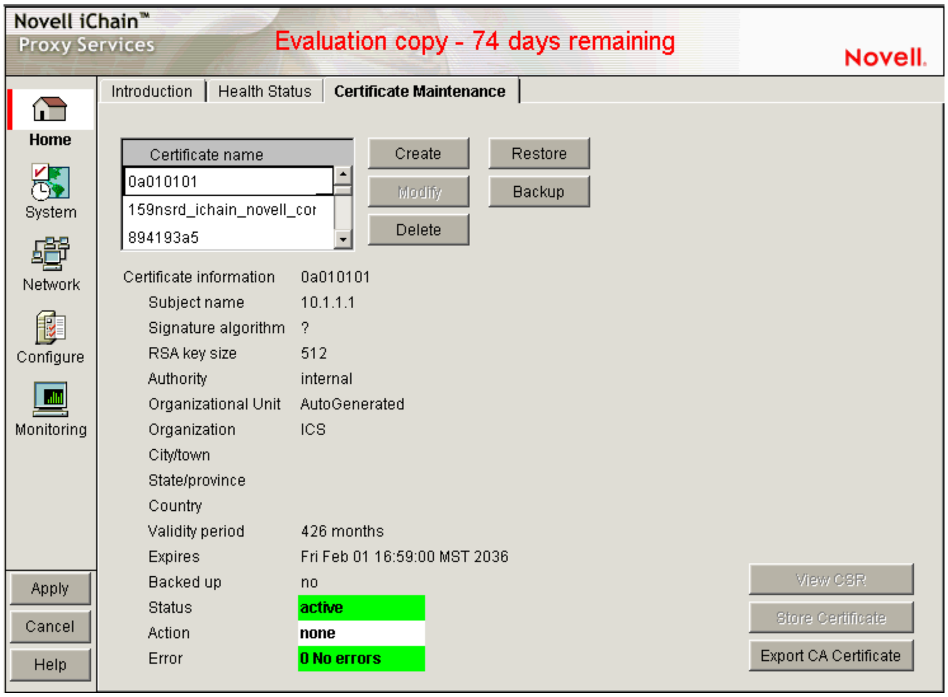

Certificate name
0a010101
159nsrd_ichain_novell_cor
894193a5

Certificate information	0a010101
Subject name	10.1.1.1
Signature algorithm	?
RSA key size	512
Authority	internal
Organizational Unit	AutoGenerated
Organization	ICS
City/town	
State/province	
Country	
Validity period	426 months
Expires	Fri Feb 01 16:59:00 MST 2036
Backed up	no
Status	active
Action	none
Error	0 No errors

The Certificate Maintenance page 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 "Using iChain to Manage Certificates" on page 193.

Certificate Name: A list of certificates created on the appliance.

Certificate Information: Information for the certificate selected.

View CSR: Displays the certificate signing request of a certificate you have created.

This is used to request a certificate from a certificate authority. For more information, see Section 15.3, "Obtaining a Certificate from an External CA," on page 195.

Store Certificate: Stores certificate information received from a certificate authority. For more information, see Section 15.3, "Obtaining a Certificate from an External CA," on page 195.

Export CA Certificate: Displays a certificate authority's certificate. For more information, see Section 15.4, "Viewing (Exporting) a Certificate's CA," on page 198.

19.3 The System Panel Options

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.

This section explains the following System Panel pages:

- "Timezone Page" on page 273
- "Date/Time Page" on page 274
- "Actions Page" on page 275
- "SNMP Page" on page 279
- "Import/Export Page" on page 280
- "Upgrade Page" on page 281
- "Alerts Page" on page 283
- "Admin ACL Page" on page 284

19.3.1 Timezone Page

Path: System > Timezone

Figure 19-4 *Timezone Page*

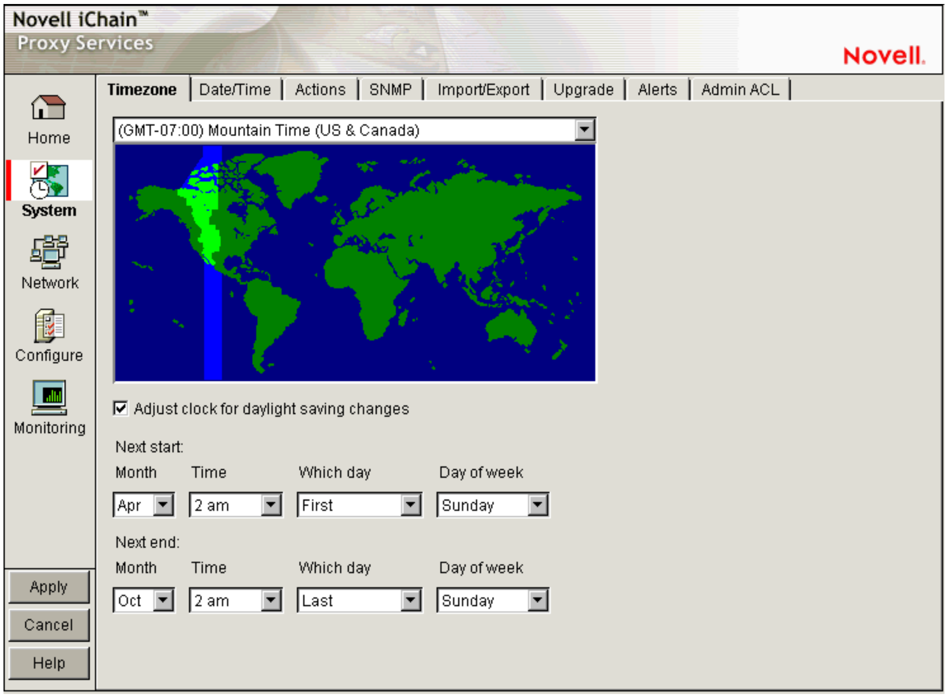

(GMT-07:00) Mountain Time (US & Canada)

☑ Adjust clock for daylight saving changes

Next start:

Month	Time	Which day	Day of week
Apr	2 am	First	Sunday

Next end:

Month	Time	Which day	Day of week
Oct	2 am	Last	Sunday

The Timezone page 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 page 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 Section 19.7, “Using Appliance Commands,” on page 333 for more information.

Adjust Clock for Daylight Saving Changes: If you select 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 page, 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 Section 19.7, “Using Appliance Commands,” on page 333 for more information.

19.3.2 Date/Time Page

Path: System > Date/Time

Figure 19-5 *Date/Time Page*

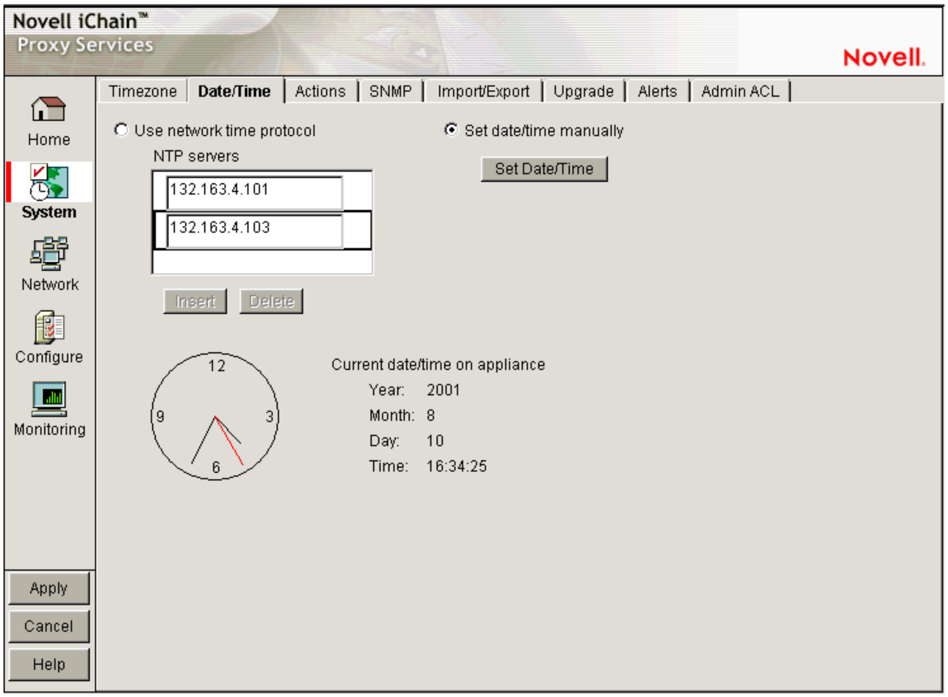

○ Use network time protocol

NTP servers

132.163.4.101

132.163.4.103

● Set date/time manually

Current date/time on appliance

Year: 2001

Month: 8

Day: 10

Time: 16:34:25

The Date/Time page 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: iChain Proxy Services 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, iChain Proxy Services 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: Selecting this option turns the 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 additional 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, iChain Proxy Services slows the appliance clock down 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 is later than the appliance clock, synchronization between the two is usually immediate. However, in certain situations you might observe the appliance clock incrementing by six hundred minute intervals. This is normal system behavior.

The fact that the Apply button changes from Wait back to Apply indicates only that the NTP configuration change has been made, not that the appliance clock is fully synchronized with the NTP server.

If the above features are problematic in your situation, you can set appliance time manually to the target time and then re-enable the NTP feature.

Set Time Manually: The following dialog box appears when you select this option and click Set Time. Set the date and time using the drop-down lists. Clicking OK immediately resets the system clock.

Figure 19-6 *Set Time Dialog Box*

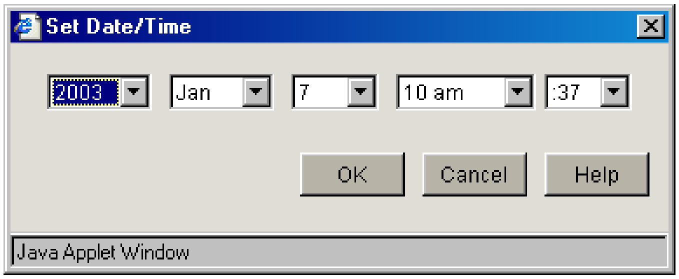

Use this option if NTP is not available to your appliance or you need to set a specific time for some reason.

19.3.3 Actions Page

Path: System > Actions

Figure 19-7 Actions Page

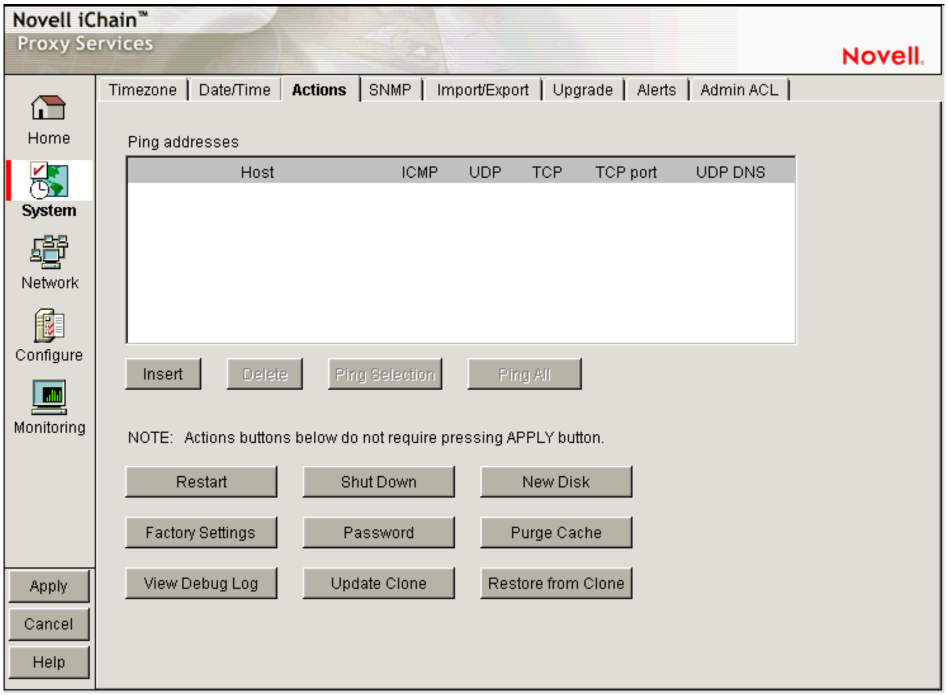

NOTE: Actions buttons below do not require pressing APPLY button.

The Actions page 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. However, changes made in the Actions page are immediately effective.

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), etc.

Restart: Shuts down the caching system and then restarts it. Configuration settings are retained but cached objects are removed.

Shut Down: Shuts down the caching system. The hardware remains turned on until it is manually powered off.

When the appliance has 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 auto-detected.

Factory Settings: Resets the appliance to its original factory configuration as explained in "Restoring Factory Settings" on page 349. Passwords are retained. If you want to preserve other settings for later use on this or another appliance, see "Import/Export Page" on page 280.

Password: See "Password Dialog Box" on page 277.

Purge Cache: See "Purge Cache Dialog Box" on page 278.

View Debug Log: When an appliance experiences an abnormal shutdown because of a configuration error or other problem, iChain Proxy Services 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, or e-mail it to Technical Support.

Update Clones: Each appliance stores a clone image that is initially 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, iChain Proxy Services assumes the current configuration is faulty and automatically replaces it with the clone image.

You can overwrite the default clone image with an alternate 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 also need to reconfigure proxy services on the appliance or use a .nas file to restore these. See "Restoring the Appliance to the Clone Image" on page 349.

Password Dialog Box

Path: System > Actions > Password

Figure 19-8 Password Dialog Box

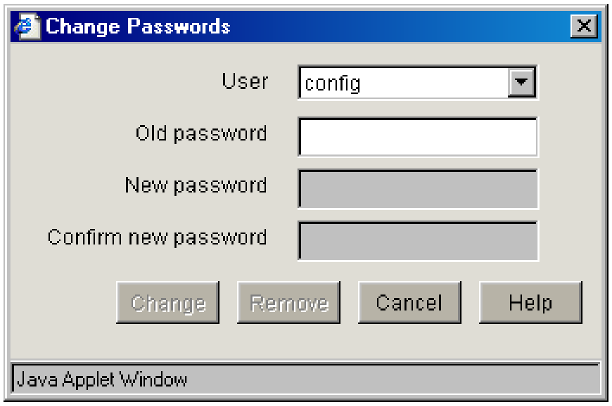

Change Passwords

User config

Old password

New password

Confirm new password

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 and execute get commands from the command line. The Apply function and the set command are not available. The server license information is also not available.

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 user selected.

Remove: Removes (sets to null) the password for the user selected.

Appliance passwords are case-sensitive.

Purge Cache Dialog Box

Path: System > Actions > Purge Cache

Figure 19-9 *Purge Cache Dialog Box*

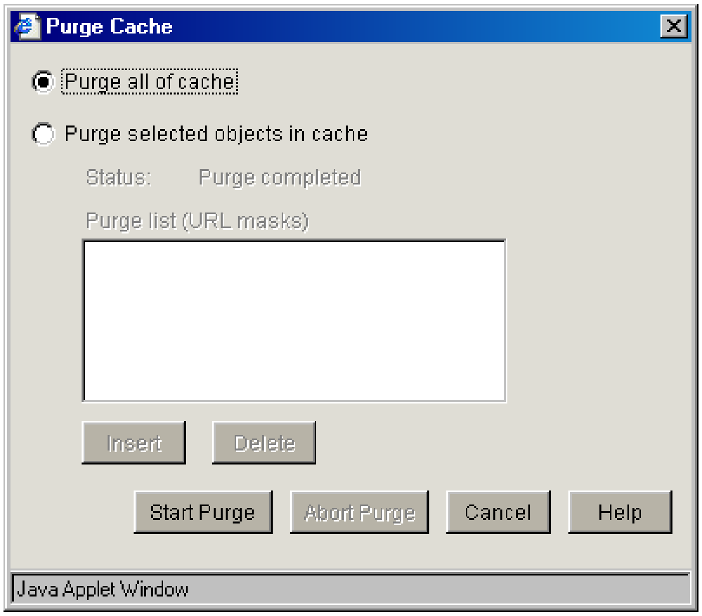

● Purge all of cache

○ Purge selected objects in cache

Status: Purge completed

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: Starting the purge with this option selected will purge everything from the appliance's cache.

Purge Selected Objects in Cache: Selecting 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 purging of 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” or any other combination of uppercase and lowercase letters for “=SPORTS” following the question mark in the URL.

19.3.4 SNMP Page

Path: System > SNMP

Figure 19-10 *SNMP Page*

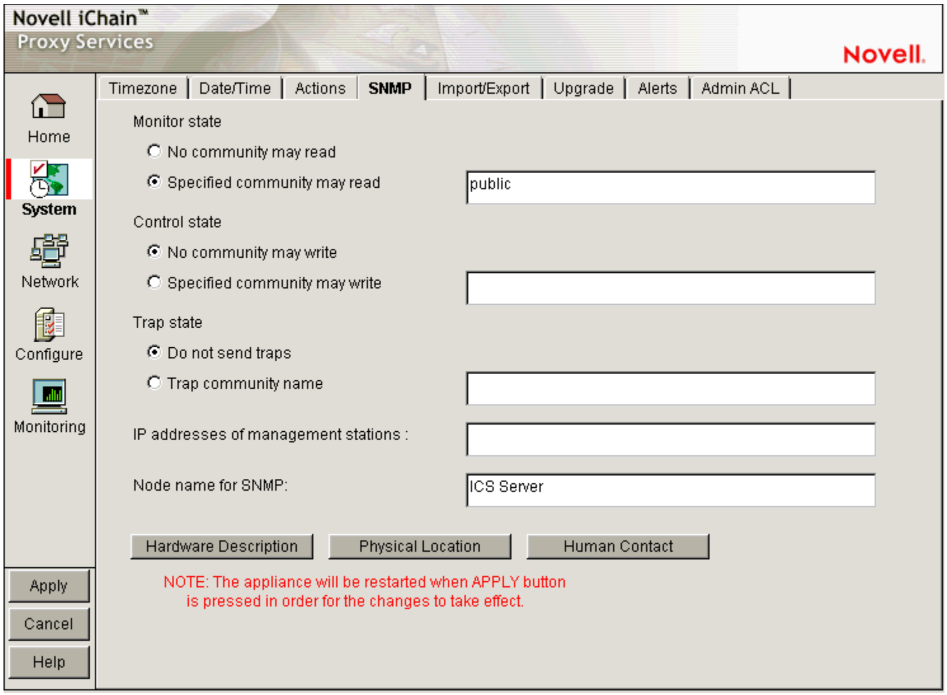

Monitor state

○ No community may read

● Specified community may read: public

Control state

● No community may write

○ Specified community may write

Trap state

● Do not send traps

○ Trap community name

IP addresses of management stations :

Node name for SNMP: ICS Server

NOTE: The appliance will be restarted when APPLY button is pressed in order for the changes to take effect.

The SNMP page 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 1 standard outlined in RFC 1067: A Simple Network Management Protocol (http://www.faqs.org/rfcs/rfc1067.html).

When SNMP-enabled appliance components start, 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.

Each appliance contains an ichain.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 page, all alerts you check in the Alerts Page (see “Alerts Page” on page 283) are automatically sent as SNMP traps even if you have not configured syslog or e-mail alert notification on the Alerts page.

Monitor State: Allows you to specify community Read access and the community name or password to be used. Community names must contain only ASCII characters and must not have spaces.

Control State: Allows you to specify community Write access and the community name or password to be used. Community names must contain only ASCII characters 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 either 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 only ASCII characters and must not have spaces.

IP Addresses of Management Stations: One or more management station IP addresses, separated by semicolons.

Node Name for SNMP: Lets you specify a node name for management of the appliance through SNMP.

The buttons below the node name field let you enter additional information regarding the hardware, the appliance's physical location, and information regarding the person responsible for the appliance.

19.3.5 Import/Export Page

Path: System > Import/Export

Figure 19-11 *Import/Export Page*

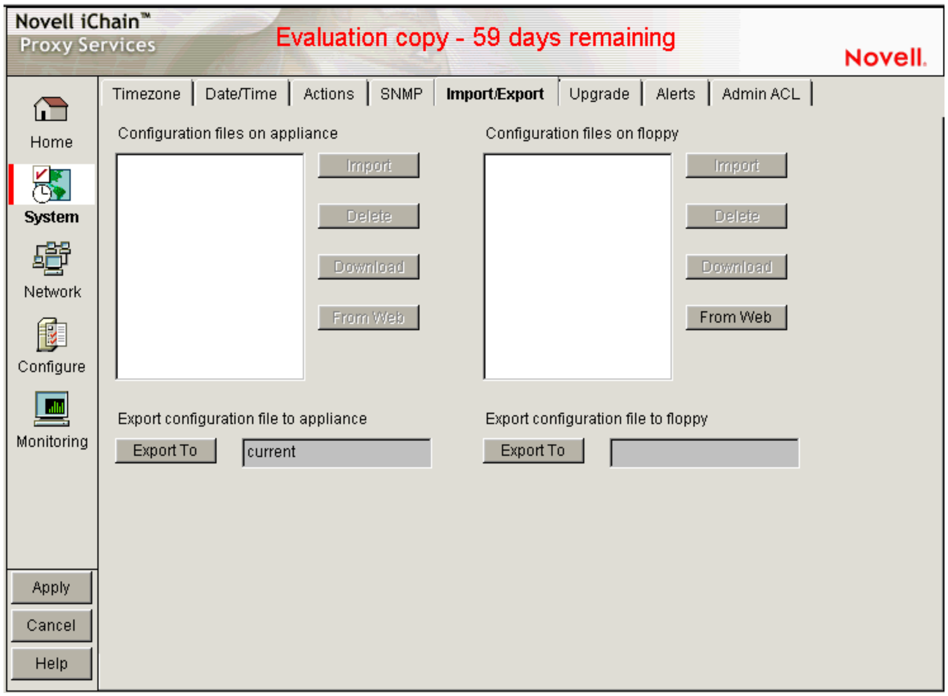

current

The Import/Export page lets you manage appliance configuration files on the appliance and on floppy disk.

Configuration Files on Appliance: Displays a list of all of the configuration files stored on the appliance. These files are used to instantly configure the appliance, 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 iChain Proxy Services. The .nas extension of these files is not shown in this list but 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 drive. You can download, import, and delete any file in this list. You can also copy a configuration file from any URL to a floppy in the appliance's floppy 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 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: Clicking the button under one of these titles creates a configuration file on the appliance or on the diskette in the appliance's floppy drive.

Files saved using the Export feature contain the complete configuration of the appliance at the time of export. The default filename is current.nas. 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.

WARNING: Do not create .nas filenames longer than eight characters, because the system might overwrite a previous file. For example, If you create 000000005.nas and 000000006.nas, the system overwrites the older file.

19.3.6 Upgrade Page

Path: System > Upgrade

Figure 19-12 *Upgrade Page*

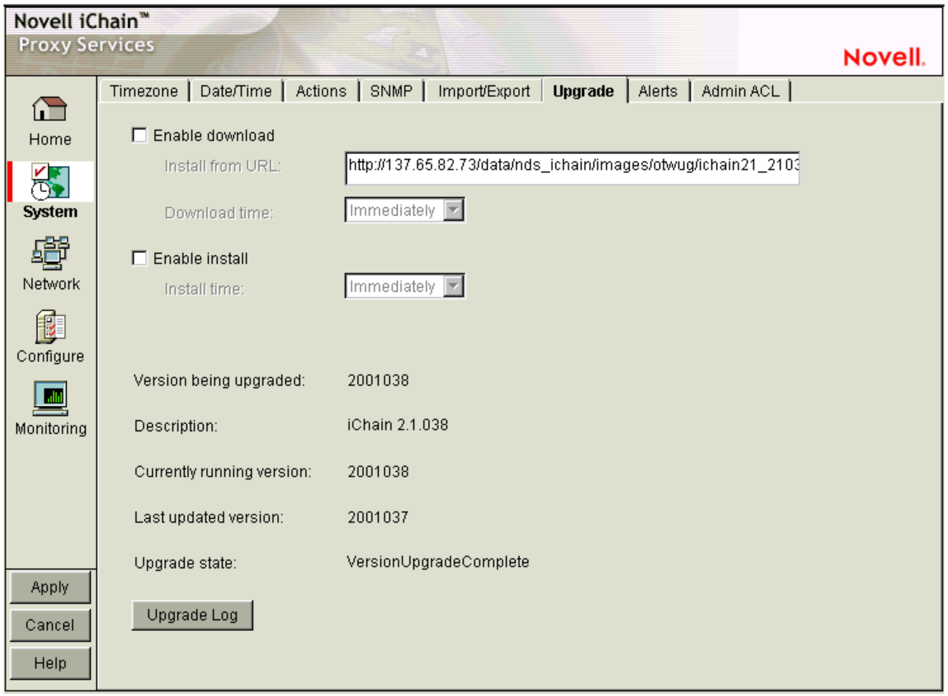

☐ Enable download

Install from URL: http://137.65.82.73/data/nds_ichain/images/otwug/ichain21_2103

Download time: Immediately

☐ Enable install

Install time: Immediately

Version being upgraded: 2001038

Description: iChain 2.1.038

Currently running version: 2001038

Last updated version: 2001037

Upgrade state: VersionUpgradeComplete

The Upgrade page 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.

NOTE: We recommend that you update the appliance's clone image after an upgrade. See "Restoring the Appliance to the Clone Image" on page 349, "Actions Page" on page 275, and Section 19.8, "Performing Patch Upgrades," on page 341 for more information.

Enable Download: Lets you set the appliance to automatically download updates. If you select this option and enter the URL for the patch in the Install from URL field, it 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 automatically install patches. If you select this option, patches downloaded to the appliance are automatically installed as scheduled in the Install Time field.

Version Being Upgraded: Each update has a version number. The version of the current update appears in this field the moment the update process begins. You cannot upgrade the proxy server 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, Version Download Complete, etc. The field is updated each time you click Upgrade.

Upgrade Log: Displays the text messages that have been generated by the upgrade process.

19.3.7 Alerts Page

Path: System > Alerts

Figure 19-13 *Alerts Page*

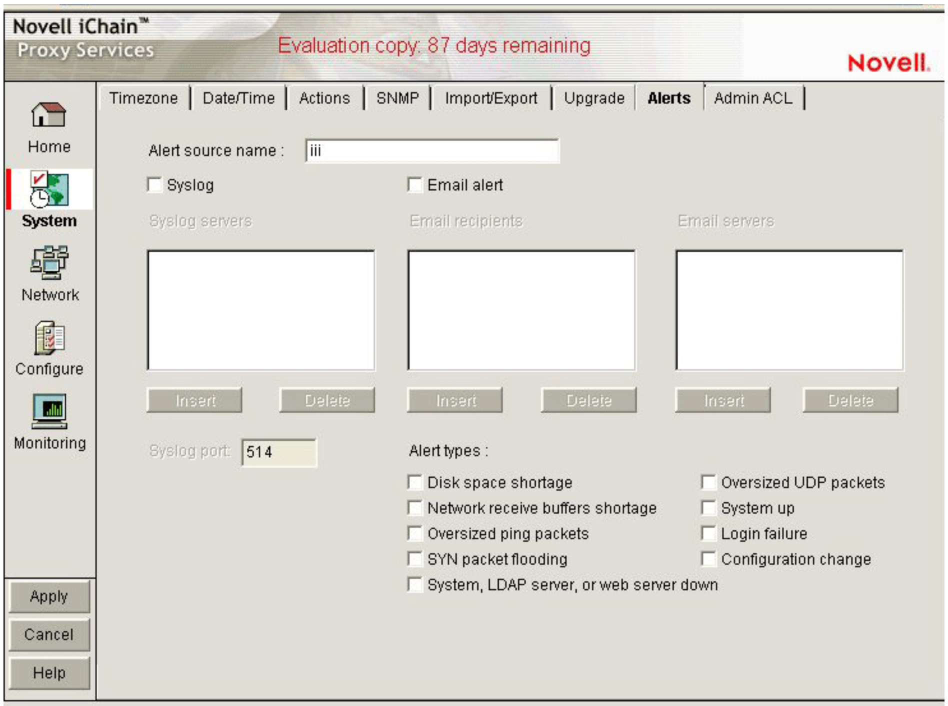

Alert source name : iii

□ Syslog

□ Email alert

Syslog port: 514

Alert types :

□ Disk space shortage

□ Network receive buffers shortage

□ Oversized ping packets

□ SYN packet flooding

□ System, LDAP server, or web server down

□ Oversized UDP packets

□ System up

□ Login failure

□ Configuration change

The Alerts page 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: This identifies the appliance as the source of an alert. The system inserts this in the From field of an e-mail alert and in the Syslog alert message. Be sure to use only those characters that are compatible with the e-mail servers you are specifying. Using spaces and other characters might prevent an e-mail server from accepting the alert message.

Syslog: Selecting this option enables syslog alerts. Alert messages are then sent to one of the syslog servers.

E-mail Alert: Selecting this option enables e-mail alerts. Alert messages are then sent to all of the e-mail recipients. However, in order for this type of alert to function properly, you must not have any spaces in the alert address.

IMPORTANT: For this feature to work, e-mail servers must be able to relay e-mail from the appliance without authentication.

Because of 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 that the e-mail server can relay unauthenticated messages, or you must configure the server to accept mail from the appliance without authentication.

Syslog Servers: This is a list of syslog servers to which the appliance sends alerts. The appliance pings servers in the list, starting with the first server, until it receives an acknowledgement. It then sends a syslog alert using UDP to the responding server.

E-Mail Recipients: This is a list of e-mail recipients to whom the appliance sends alert e-mails. The appliance sends e-mails to all addresses in the list.

E-Mail Servers: This is a list of 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: This is the port the syslog server listens for syslog alerts on. The default port is 514, but this can be changed if required.

Alert Types: Appliance-generated alerts are sent for the following conditions. You enable or disable notification of generated alerts to the configured syslog server, and you e-mail recipients by selecting or deselecting an alert type.

- `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 10 KB) PING packet.
- `SYN Packet Flooding`: The appliance generates this alert when TCP/IP detects a SYN packet flooding attack (half-open connections).
- `System, LDAP server, or web server down`: The appliance generates this alert each time the appliance is shut down properly or restarted manually.
- `Oversized UDP Packets`: The appliance generates this alert when TCP/IP receives an oversized (greater than 16 KB) UDP packet.
- `System Up`: The appliance generates this alert each time the appliance starts.
- `Login Failure`: The appliance generates this alert for all failed login attempts, including failed login attempts to proxy accelerators. 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.

19.3.8 Admin ACL Page

Path: System > Admin ACL

Figure 19-14 Admin ACL Page

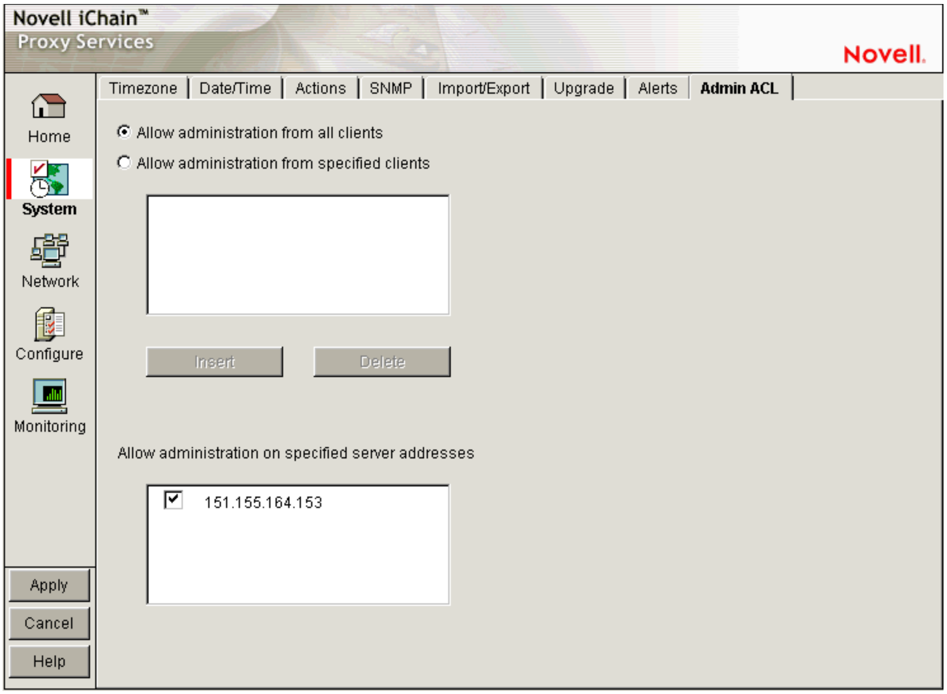

● Allow administration from all clients

○ Allow administration from specified clients

Allow administration on specified server addresses

☑ 151.155.164.153

The Admin ACL page lets you regulate access to appliance administrative functions in the browser-based management tool and the command line interface. You can restrict administrative client access and 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 appliance 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 deselects the option to prevent a global lockout.

NOTE: If you do not include the IP address from which you are specifying client access, and you click Apply, the address is not 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 change administrative access by selecting and deselecting addresses in the list. The system doesn't allow deselecting all addresses. If this is attempted, the system reverts to the default setting by re-selecting all first-assigned addresses.

19.4 The Network Panel Options

The Network panel lets you configure the appliance to function on the network where it is installed.

This section explains the following Network Panel pages:

- Section 19.4.1, “IP Addresses Page,” on page 286

- Section 19.4.2, “DNS Page,” on page 289
- Section 19.4.3, “Gateway/Firewall Page,” on page 291

19.4.1 IP Addresses Page

Path: Network > IP addresses

Figure 19-15 *IP Addresses Page*

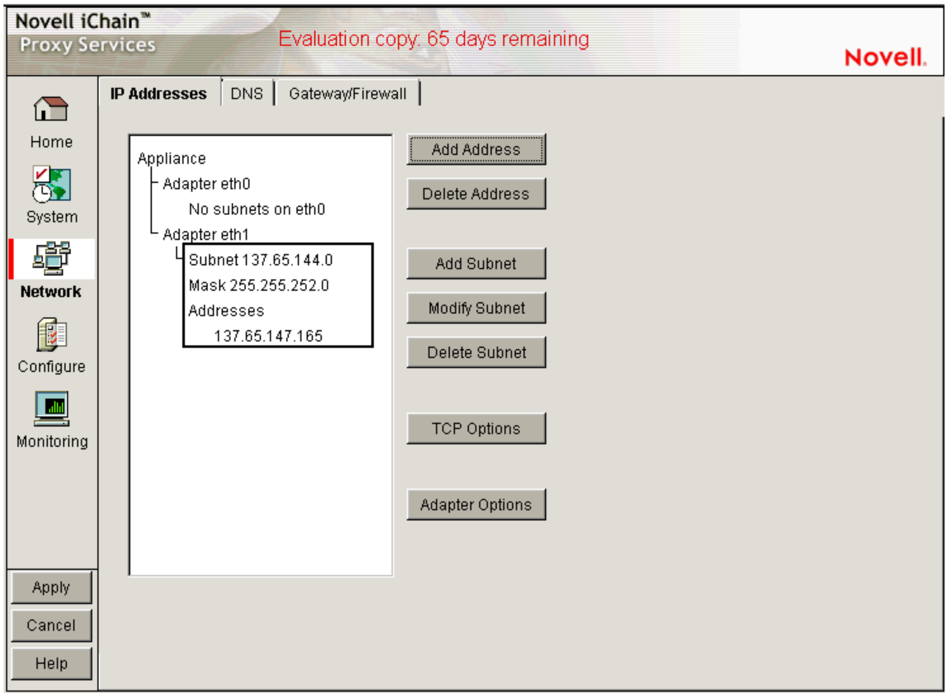

```
Appliance
 Adapter eth0
   No subnets on eth0
 Adapter eth1
   Subnet 137.65.144.0
   Mask 255.255.252.0
   Addresses
     137.65.147.165
```

The IP Addresses page 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 has 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.

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 that spans multiple adaptors.

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 19-16 *TCP Options Dialog Box*

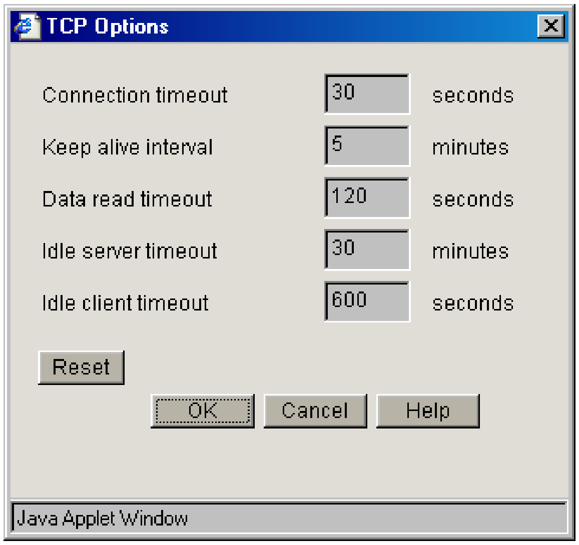

TCP Options

Connection timeout 30 seconds

Keep alive interval 5 minutes

Data read timeout 120 seconds

Idle server timeout 30 minutes

Idle client timeout 600 seconds

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: Keep-alives can be used by the proxy server to verify that the browser at the remote end of a connection is still available. The Keep Alive interval is the number of minutes a connection is idle before the proxy server queries to check if the client is still responding. This Keep Alive parameter applies to the proxy server only, and is not for connections between the proxy server and the back-end Web server (where the proxy is acting as a TCP client).

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 19-17 *Adapter Options Dialog Box*

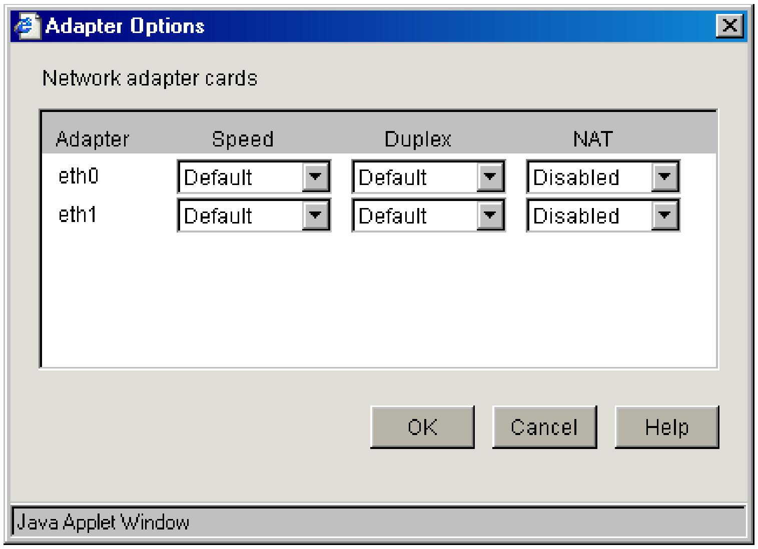

Adapter	Speed	Duplex	NAT
eth0	Default	Default	Disabled
eth1	Default	Default	Disabled

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 correctly detect duplex settings. This is a general industry problem with Fast Ethernet technology.

If your appliance isn't performing as expected, check to ensure 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, provided that 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.

19.4.2 DNS Page

Path: Network > DNS

Figure 19-18 *DNS Page*

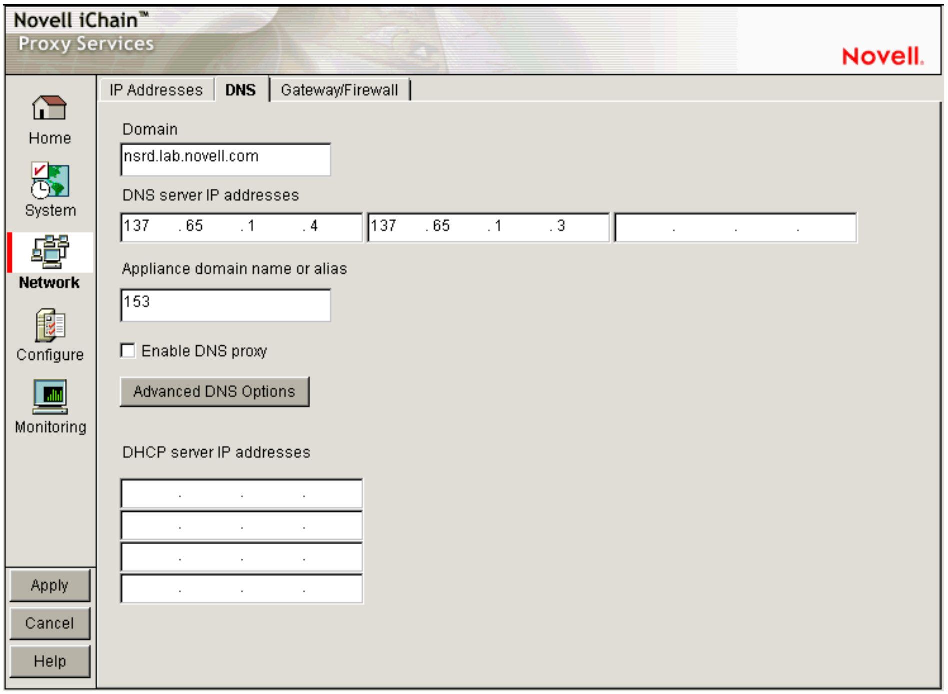

Domain: nsrd.lab.novell.com

DNS server IP addresses: 137 . 65 . 1 . 4 | 137 . 65 . 1 . 3 | . . .

Appliance domain name or alias: 153

□ Enable DNS proxy

DHCP server IP addresses: . . . | . . . | . . . | . . .

The DNS page lets you configure the domain name service that the appliance uses, 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: Specify the domain of your appliance. Valid ranges include all valid domain names.

DNS Server IP Addresses: Specify the IP addresses of the DNS servers you are using. You can enter up to three.

Appliance Domain Name or Alias: (Optional) 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: Because of a potential security risk through the DNS port, the DNS proxy is disabled by default. You can enable the DNS proxy by selecting this option.

Advanced DNS Options: See "Advanced DNS Options Dialog Box" on page 290.

DHCP Server IP Addresses: 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 19-19 *Advanced DNS Options Dialog Box*

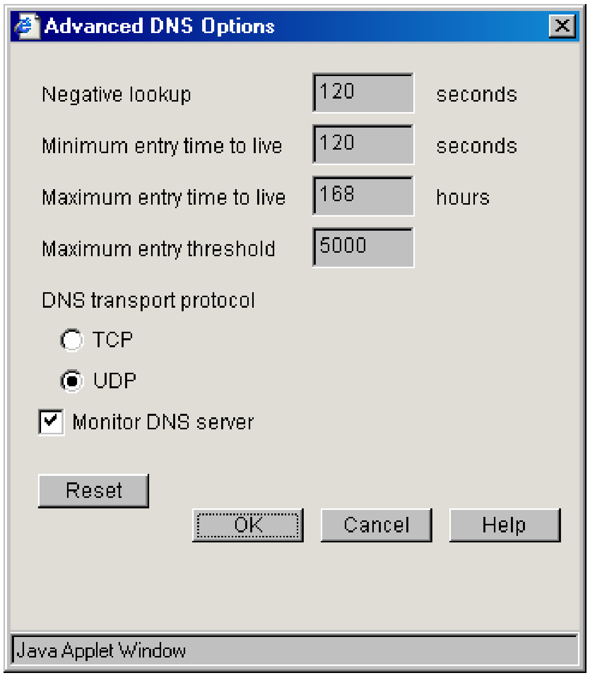

Negative lookup 120 seconds

Minimum entry time to live 120 seconds

Maximum entry time to live 168 hours

Maximum entry threshold 5000

DNS transport protocol

○ TCP

● UDP

☑ Monitor DNS server

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: How long 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 deselect this item if the appliance accesses DNS through a connection that should not be kept continually open, such as a dial-up phone line or ISDN connection. Keep in mind, however, that deselecting the option causes the DNS configuration on the Health Status Page to fail.

Reset: The default settings. Click the Reset button to reset the advanced options to their default values.

19.4.3 Gateway/Firewall Page

Path: Network > Gateway/Firewall

Figure 19-20 *Gateway/Firewall Page*

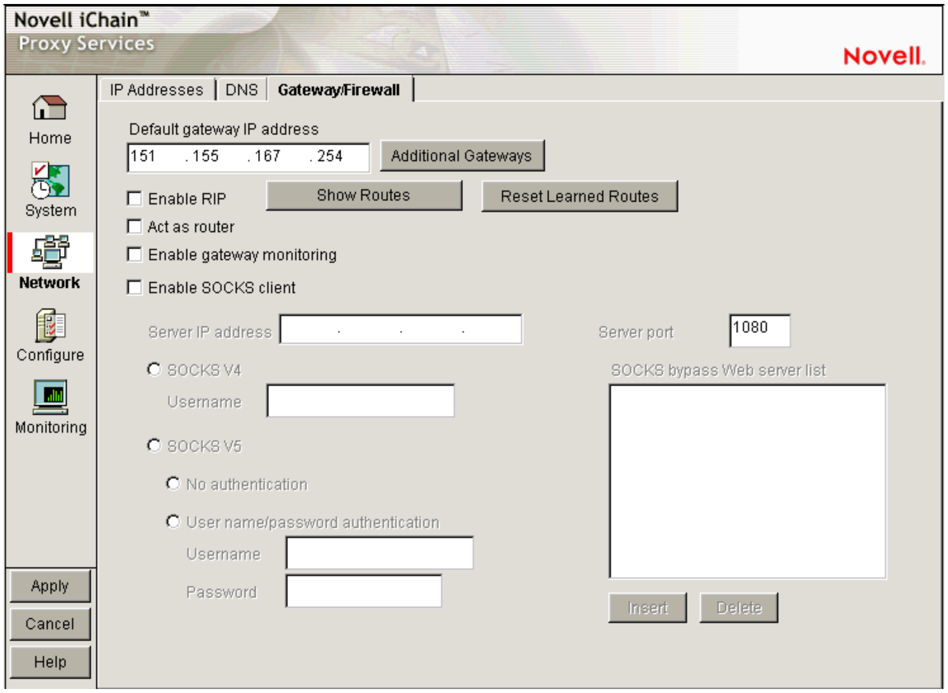

Default gateway IP address: 151 . 155 . 167 . 254

☐ Enable RIP

☐ Act as router

☐ Enable gateway monitoring

☐ Enable SOCKS client

Server IP address: . . .

Server port: 1080

○ SOCKS V4

Username:

○ SOCKS V5

○ No authentication

○ User name/password authentication

Username:

Password:

SOCKS bypass Web server list

The Gateway/Firewall page 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.

In order for the appliance to 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: You can configure static routes under Additional Gateways without having to enable routing. See "Additional Gateways Dialog Box" on page 293.

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 294.

Reset Learned Routes: Throws away all information acquired through RIP. RIP must be turned on for this to have any effect.

Act As Router: Select this option if the appliance functions as the default gateway for clients on the network. If you select this option, you can specify additional gateways. However, you can configure static routes in the Additional Gateways dialog box without enabling routing.

Enable Gateway Monitoring: The appliance normally monitors gateway availability by pinging the configured gateways every minute. You should deselect this option if the appliance accesses its gateways through a connection that should not be kept continually open, such as a dial-up phone line or ISDN connection. Keep in mind, however, that deselecting the option causes the gateway configuration on the Health Status Page to fail.

Enable SOCKS Client: SOCKS is a firewall communication protocol. If there is a firewall preventing 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 option must be selected (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: Specify your SOCKS username.

Password: Specify 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 are not sent to the SOCKS server.

Additional Gateways Dialog Box

Path: Network > Gateway/Firewall > Additional Gateways

Figure 19-21 *Additional Gateways Dialog Box*

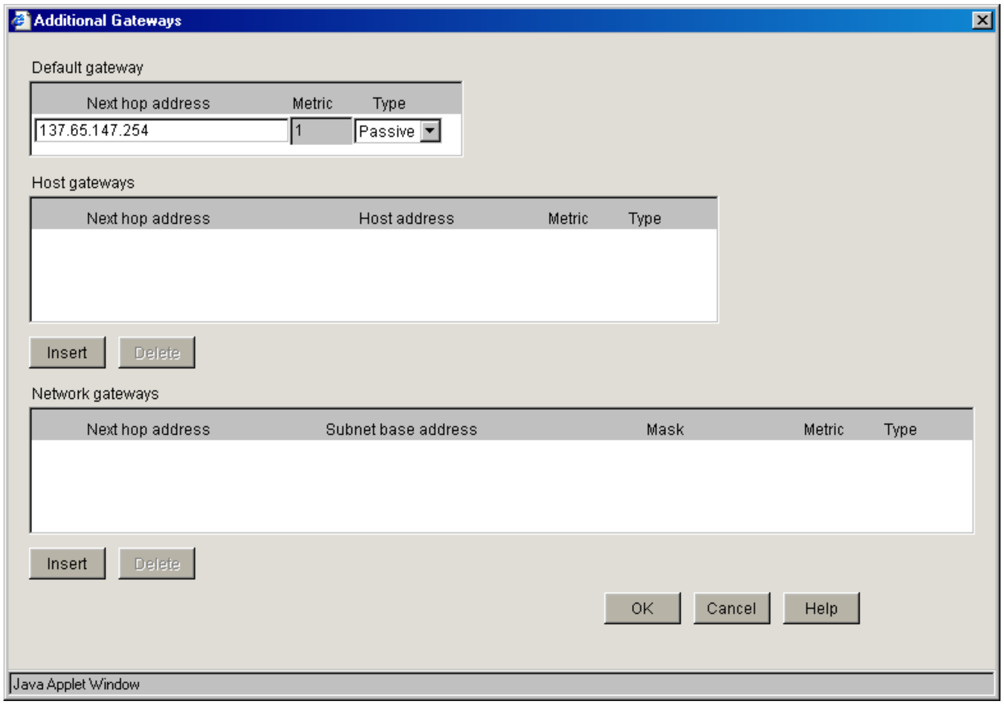

Default gateway

Next hop address	Metric	Type
137.65.147.254	1	Passive

Host gateways

Next hop address	Host address	Metric	Type

Network gateways

Next hop address	Subnet base address	Mask	Metric	Type

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. You can configure static routes under Additional Gateways without enabling routing.

IMPORTANT: The appliance uses additional gateways only when the Act As Router option is selected on the Gateway/Firewall page.

Gateways fall within the following 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 the normal gateway use logic depending on a relative cost factor for using the gateway. 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 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 you can 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*: 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. Valid addresses 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*: 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, 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*: 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 19-22 *Routes Dialog Box*

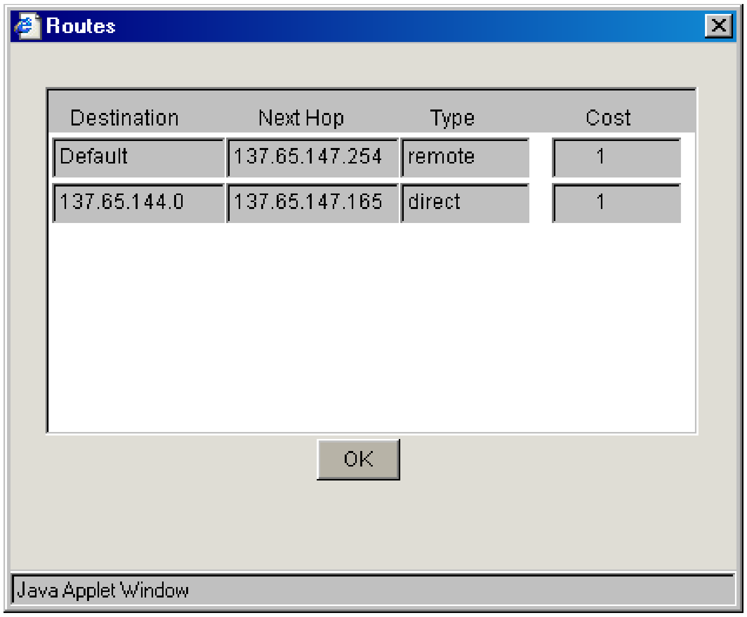

Destination	Next Hop	Type	Cost
Default	137.65.147.254	remote	1
137.65.144.0	137.65.147.165	direct	1

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 option is selected. Clicking Reset Learned Routes clears RIP entries from the list.

Destination: The default route is named and listed first. For other routes, the subnet address is shown.

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 it is a relative cost factor assigned by the RIP function if the Enable RIP option is selected.

19.5 The Configure Panel Options

The Configure panel lets you set up Web acceleration. It also lets you configure authentication profiles and access control information, fine-tune caching services, and specify how error pages are vended to browsers.

This section explains the following Configure Panel pages and options:

- Section 19.5.1, “Web Server Accelerator Page,” on page 296
- Section 19.5.2, “Add Authentication Profiles Dialog Box,” on page 305
- Section 19.5.3, “Using Authentication Profiles,” on page 307
- Section 19.5.4, “LDAP Search Query Authentication,” on page 312
- Section 19.5.5, “Authentication Page,” on page 314

- Section 19.5.6, “Access Control Page,” on page 319
- Section 19.5.7, “FTP Page,” on page 322
- Section 19.5.8, “Management Page,” on page 323
- Section 19.5.9, “Tuning Page,” on page 325
- Section 19.5.10, “Mini Web Page,” on page 327

19.5.1 Web Server Accelerator Page

Path: Configure > Web Server Accelerator

Figure 19-23 *Web Server Accelerator Page*

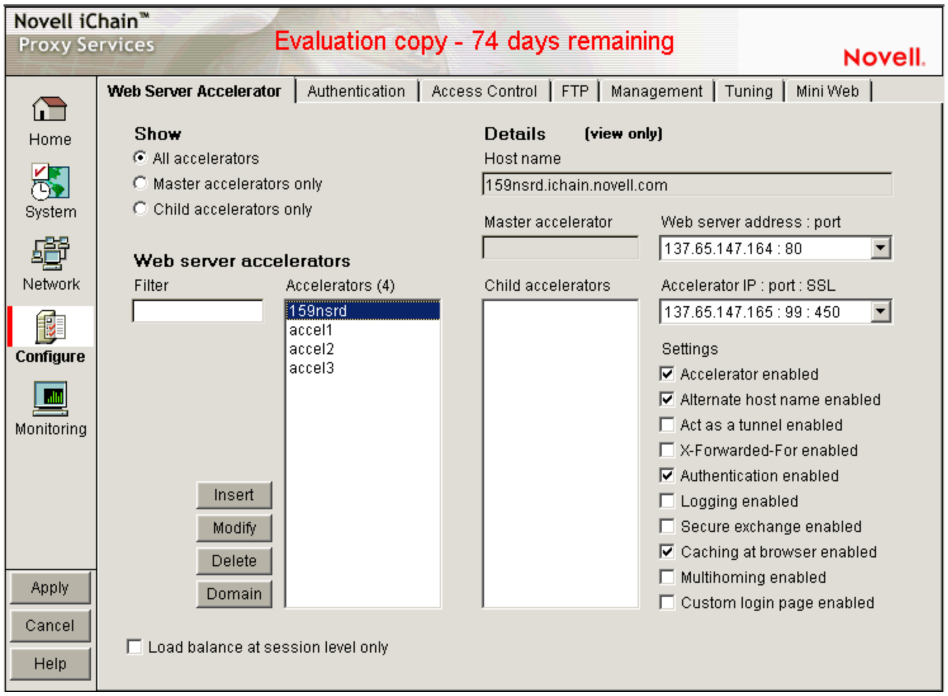

Show

● All accelerators
○ Master accelerators only
○ Child accelerators only

Web server accelerators

Filter:

Accelerators (4)

159nsrd
accel1
accel2
accel3

□ Load balance at session level only

Details (view only)

Host name: 159nsrd.ichain.novell.com

Master accelerator:

Web server address : port: 137.65.147.164 : 80

Child accelerators

Accelerator IP : port : SSL: 137.65.147.165 : 99 : 450

Settings

☑ Accelerator enabled
☑ Alternate host name enabled
□ Act as a tunnel enabled
□ X-Forwarded-For enabled
☑ Authentication enabled
□ Logging enabled
□ Secure exchange enabled
☑ Caching at browser enabled
□ Multihoming enabled
□ Custom login page enabled

The Web Server Accelerator page lets you add one or more Web server (reverse proxy) accelerators. The proxy server 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 55.

Show section: These options allow you to select which accelerators are displayed in the accelerators list.

- All Accelerators: All accelerators are displayed.
- Master Accelerators only: Only accelerators that have child accelerators associated with them (multi-homing relationships) are displayed.

Web Server Accelerators section: This section displays a list of accelerators as specified by the radio buttons discussed above. In addition, a filter field is provided to specify which accelerators are displayed in the accelerators list.

- Filter: By typing characters in the filter field, the list of accelerators is modified according to the filter. Only straight text is used to filter; that is, no wildcards are supported. For example, consider that the list contains four accelerators: Accel1, MyAccel, MyServer, and Zebras. If the letter “a” is typed in the filter list field, only Accel1 would show up in the list. If the letter “m” was then typed, the list would be empty because no accelerators match the pattern “am”. Typing “my” would display MyAccel and MyServer. Typing “mys” would display only MyServer. Typing “z” would display only zebras.
- Accelerators: This field displays the list of accelerators. The number of accelerators in the list is displayed in parenthesis after the heading.

Details section: Information about the selected accelerator is shown in this section. These fields are all viewable only and cannot be changed or modified on this page. To modify any settings for an accelerator, click Modify after selecting an accelerator in the accelerator list.

- Host name: The host name of the accelerator and will include the sub-path if one was specified through multi-homing options.
- Master accelerator: This field will contain the name of the master accelerator of the selected accelerator only if the selected accelerator is the child of that master.
- Child accelerators — Contains a list of accelerators that are the child accelerators of the selected accelerator.
- Web server address: port: Contains the accelerator's Web server IP address or DNS name, followed by the Web server port it is using. This is a multiple value field and if multiple values are present, they can be viewed by clicking the down-arrow.
- Accelerator IP: port: SSL: Shows the accelerator IP address being used, followed by the accelerator port. Also included is the SSL listening port that was specified, whether or not it is actually in use. This is a multiple value field and if multiple values are present, they can be viewed by clicking the down-arrow.
- Settings: This list of check boxes shows various settings for the accelerator.

For a complete explanation of the settings shown on this page, see “Web Server Accelerator Dialog Box” on page 297.

Load Balance at Session Level Only: Selecting this option causes the proxy server to use the same Web server for all fills during a session. This prevents eBusiness users from needing to log in multiple times. This setting affects all Web server accelerators configured on the proxy server.

For more information on how iChain handles load balancing, see the Novell Technical Information Document (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10069756.htm).

Web Server Accelerator Dialog Box

Path: Configure > Web Server Accelerator > Insert

Figure 19-24 *Web Server Accelerator Dialog Box*

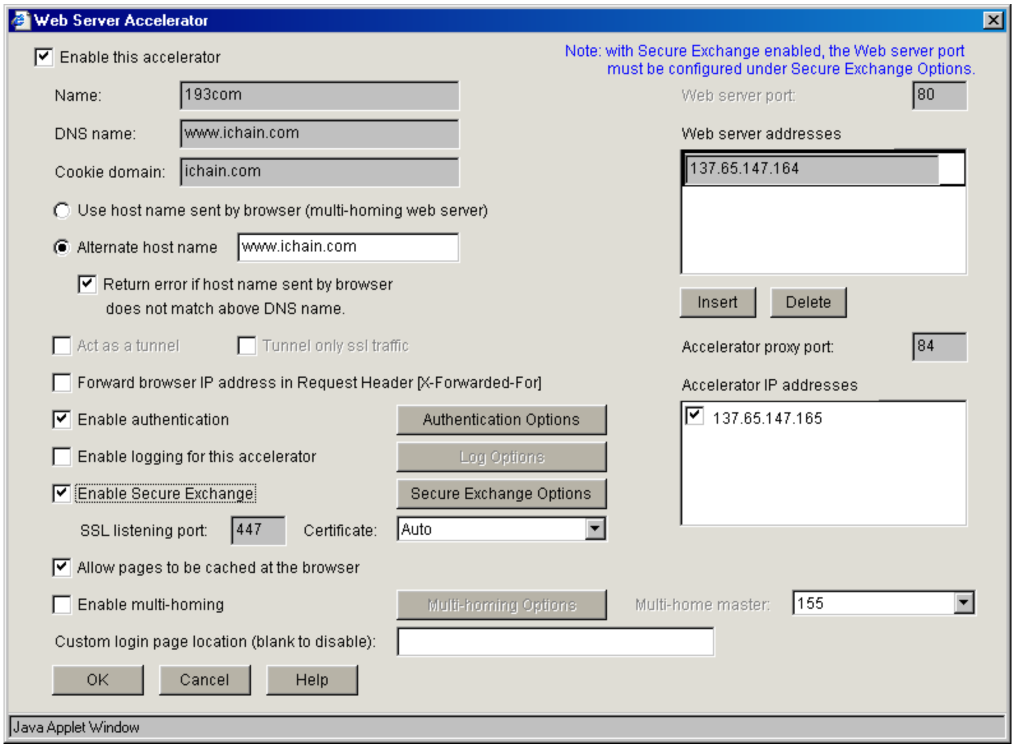

Web Server Accelerator

☑ Enable this accelerator

Note: with Secure Exchange enabled, the Web server port must be configured under Secure Exchange Options.

Name: 193com

DNS name: www.ichain.com

Cookie domain: ichain.com

○ Use host name sent by browser (multi-homing web server)

● Alternate host name www.ichain.com

☑ Return error if host name sent by browser does not match above DNS name.

□ Act as a tunnel □ Tunnel only ssl traffic

□ Forward browser IP address in Request Header [X-Forwarded-For]

☑ Enable authentication

□ Enable logging for this accelerator

☑ Enable Secure Exchange

SSL listening port: 447 Certificate: Auto

☑ Allow pages to be cached at the browser

□ Enable multi-homing

Custom login page location (blank to disable):

Web server port: 80

Web server addresses: 137.65.147.164

Accelerator proxy port: 84

Accelerator IP addresses: ☑ 137.65.147.165

Multi-home master: 155

The Web Server Accelerator dialog box lets you create Web server accelerator services for handling requests to Web servers.

Enable This Accelerator: Specifies whether to enable the defined Web server accelerator after you have configured it. The default is Enabled.

Name: 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 directory name that the log files for the Web server accelerator are kept in.

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. If your infrastructure supports multi-homing, refer to Section 4.7, “Multi-Homing,” on page 65, “Path-Based Multi-Homing” on page 67, and “Path-Based Multi-Homing Example” on page 68 for further information.

If you are accelerating multiple Web servers for a single Web site and plan to use path-based multi-homing, you must use the same DNS name in every accelerator definition.

NOTE: If you are changing the DNS name on an accelerator which has authentication enabled, the existing cookie domain might not be valid with the new DNS name if the DNS name is not a subdomain of the cookie domain (resulting in the browser displaying a 403 Forbidden Error message). You can check the cookie domain at Authentication Options > Cookie Domain.

Cookie Domain: In Configure > Web Server > Modify > Authentication Options, there is a Cookie Domain field. Single authentication across two accelerators within the same iChain Proxy Server is not possible when the Cookie Domain is different on each accelerator. The Cookie Domain field can be used to allow single authentication across multiple accelerators where the cookie domain by default would be different for each accelerator. For example, if Accelerator One has a DNS name of www.support.novell.com and Accelerator Two has a DNS name of www.developernet.novell.com, the cookie domains would be support.novell.com for Accelerator One and developernet.novell.com for Accelerator Two. The cookie domains are different for each accelerator, so users are prompted for authentication when accessing Accelerator Two, even if they have already authenticated to Accelerator One. If the cookie domain is changed to novell.com on each accelerator, then users do not need to authenticate again when accessing Accelerator Two if they have already authenticated through Accelerator One. This is not possible if one of the DNS names ends with acme.com and the other DNS name ends with novell.com. Both accelerators need to have a common subdomain for this to work.

NOTE: A cookie is specific to a single authentication profile, so even though the cookie domains are the same, the cookie set by the Accelerator One is not valid for the Accelerator Two if you use different authentication profiles. See "Add Authentication Profiles Dialog Box" on page 305 and Section 19.5.3, "Using Authentication Profiles," on page 307 for more information about authentication profiles.

Use Host Name Sent by Browser (Multi-homing Web Server): Selecting this option preserves the host name in the HTTP header exactly as it came in the browser request.

Alternate Host Name: Selecting this option causes the specified string to be substituted for the host name in the HTTP header before the request is forwarded to the Web server.

Return Error if Host Name Sent by Browser Does Not Match above DNS Name: Selecting this option causes iChain Proxy Services to match the host name 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, iChain Proxy Services returns an error to the requesting browser.

Act As a Tunnel: 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.

The Act as a Tunnel option 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. When the option is selected, the accelerator sets up a tunnel for all incoming traffic.

When you select the Act as a Tunnel option, you have the additional option of having the accelerator service tunnel only SSL traffic.

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 then verifies that the address and port being accessed are actually an SSL Web site. If verification fails, the service tears down the connection.

NOTE: The SSL port number for the SSL tunnel is specified via the Accelerator proxy port and not the SSL listening port.

Forward Browser IP address in Request Header [X-Forwarded-For]: X-Forwarded-For headers are 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.

Selecting 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 deselected causes the appliance to remove X-Forwarded-For headers from any Web accelerator requests passing through the appliance.

Deciding whether to select the option requires that you 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 Authentication: Selecting this option 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.

Enable Logging for This Accelerator: Causes log files to be kept for this Web server accelerator. Clicking Log Options displays the Log Options dialog box.

Enable Secure Exchange: Selecting this option allows SSL to be used for HTTP requests between the client and the iChain box, and optionally between the iChain box and Web servers. The default is Disabled. Clicking Secure Exchange Options displays the Secure Exchange Options dialog box.

This option must be enabled to use Basic Authentication. See Section 6.2, “Enabling Authentication Through the HTTP Authorization Header,” on page 93.

SSL Listening Port: Specifies the port to be used for Secure Exchange communication.

IMPORTANT: iChain Proxy Services requires that each service (including HTTPS) use a unique IP address and port combination. The default HTTPS port is 443. Attempts to enable HTTPS for more than one service on the same IP address and port results in a TCP bind error.

Certificate: This drop-down list displays certificates you have stored on your appliance. System-generated certificates do not appear in the list.

Use this field after you have stored a certificate you created specifically for the Web server accelerator you are creating. This prevents browsers from receiving certificate confirmation messages each time they access the appliance. For more information, see “Using iChain to Manage Certificates” on page 193.

Allow Pages to Be Cached at the Browser: With this option selected, all pages are marked cacheable on the browser. If this option is not selected, all pages are marked non-cacheable on the browser.

Enable Path-Based Multi-Homing: This option is enabled only when you have created another accelerator definition and have not created a standard multi-homing relationship between previously defined accelerators on the iChain Proxy Services appliance. In other words, you don't have multiple accelerators sharing the same accelerator IP address and port.

Multi-homing lets you configure the system so that a multi-homing master accelerator fills general requests from a site's main Web server and routes specific requests to specialized child accelerators that fill from other specialized Web servers. This option lets you create child accelerators for path-based multi-homing configurations.

When you enable path-based multi-homing for the accelerator you are defining, you must also click the Multi-homing Options button and specify settings that the multi-homing master can use to route traffic to the accelerator you are defining. For more information, see "Multi-Homing Options Dialog Box" on page 304.

If you have created multiple accelerators that can function as multi-homing masters, when you select a name in the Multi-home 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 multi-homing, see "What Is Multi-Homing?" on page 65.

Multi-Home Master: This drop-down list contains the names of accelerators you have defined that can function as multi-homing masters, meaning they are not configured as child accelerators to other multi-homing masters.

Custom Login Page Location (blank to disable): The field in the Web Server Accelerator dialog box for the custom login page is populated with the directory from sys:etc\proxy\data, where the login pages can be found. If a user enters NIKE in this field, the user is specifying that the custom login/logout pages are found in sys:etc\proxy\data\nike.

Web Server Port: The port number that the origin Web server is listening on for incoming connections. The default for HTTP is 80 (1 - 65535).

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 appliance must be able to fill all requests through any of these names or addresses unless path-based multi-homing is being used.

Accelerator Proxy Port: The port number that the proxy server is listening on for incoming connections. The default for HTTP is 80 (1 - 65535).

Accelerator IP Addresses: For normal accelerator situations, non-path-based multi-homing configurations, and accelerators configured as multi-homing 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 multi-homing configurations, this lists the IP address or addresses to which the multi-homing master forwards browser requests that match the specified path rule.

Log Options Dialog Box

Path: Configure > Web Server Accelerator > Insert (Modify) > Log Options

Figure 19-25 *Log Options Dialog Box*

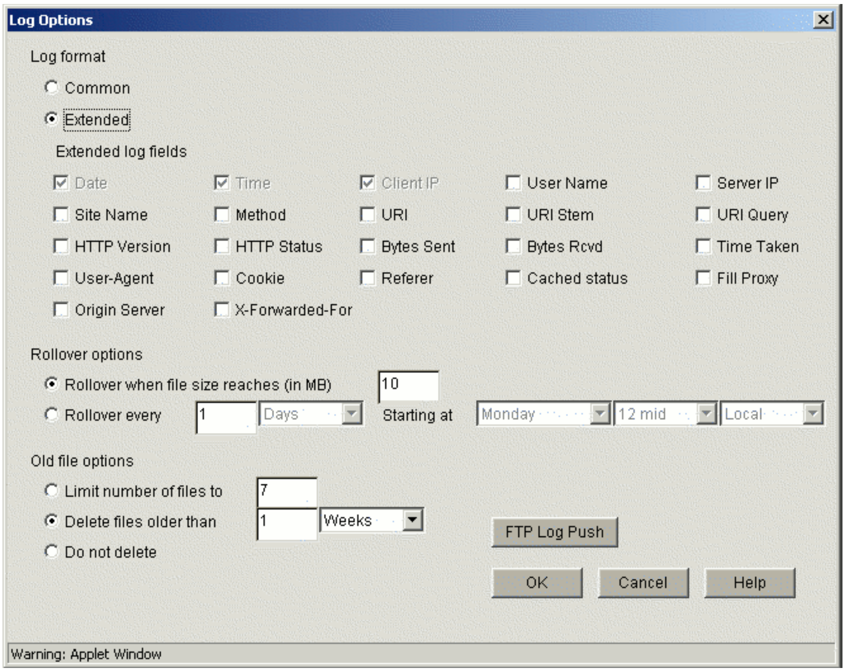

Log Options

Log format

○ Common

● Extended

Extended 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 □ X-Forwarded-For

Rollover options

● Rollover when file size reaches (in MB) 10

○ Rollover every 1 Days Starting at Monday 12 mid Local

Old file options

○ Limit number of files to 7

● Delete files older than 1 Weeks

○ Do not delete

Warning: Applet Window

This dialog box contains numerous settings for logging the operations and errors of the iChain Proxy Server.

Common: If this option is selected, only three parameters (Date, Time, and Client IP) are enabled. No others are selectable.

Extended: If this option is selected, all of the Extended log fields are enabled. The Date, Time, and Client IP options are selected by default and cannot be deselected.

Rollover When File Size Reaches (in MB): If this option is selected, the log file rolls over (empties and starts writing from the beginning) after the file reaches the size specified in the edit field.

Rollover Every: If this option is selected, the log file rolls over automatically at a specified time interval or starting at a certain day.

Limit Number of Files To: If this option is selected, the number of old log files kept on the server is limited to the number entered in the edit field.

Delete Files Older Than: If this option is selected, old log files are deleted on the time interval specified.

Do Not Delete: If this option is selected, none of the log files are deleted.

FTP Log Push Button: Launches the FTP Log Push Configuration dialog box.

FTP Log Push Configuration Dialog Box

Path: Configure > Web Server Accelerator > Insert (Modify) > Log Options > FTP Log Push

Figure 19-26 Log Push Configuration Dialog Box

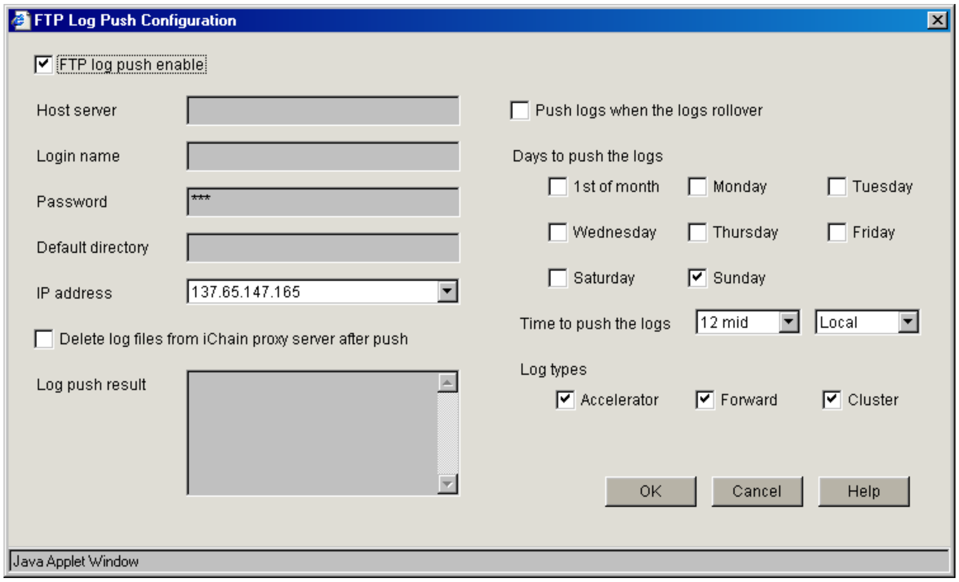

FTP Log Push Configuration

☑ FTP log push enable

Host server

Login name

Password ***

Default directory

IP address 137.65.147.165

☐ Delete log files from iChain proxy server after push

Log push result

☐ Push logs when the logs rollover

Days to push the logs

☐ 1st of month ☐ Monday ☐ Tuesday

☐ Wednesday ☐ Thursday ☐ Friday

☐ Saturday ☑ Sunday

Time to push the logs 12 mid Local

Log types

☑ Accelerator ☑ Forward ☑ Cluster

This dialog box enables you to set FTP log push options.

FTP Log Push Enable: If this option is selected, FTP log push is enabled.

Host Server: The name of the host server where the logs are pushed.

Login Name: The user name on the host server to log in as. This user must have appropriate rights to be able to push logs to the server.

Password: The password for the user on the host server.

Default Directory: The directory on the host server where the logs are pushed.

IP Address: The IP address of the host server.

Delete Log Files from iChain Proxy Server after Push: If this option is selected, the logs are deleted from the iChain Proxy Server after they have been pushed to the specified host server.

Log Push Result: The result of the log push.

Push Logs When the Logs Rollover: If this option is selected, the logs are automatically pushed to the host server when they roll over.

Days to Push the Logs: The Logs are pushed to the host server every day that has been selected.

Time to Push the Logs: The logs are pushed to the host server starting at the time specified by these two fields.

Log Types: Select each type of log to push to the host server.

Secure Exchange Options Dialog Box

Path: Configure > Web Server Accelerator > Insert (Modify) > Secure Exchange Options

Figure 19-27 Secure Exchange Options Dialog Box

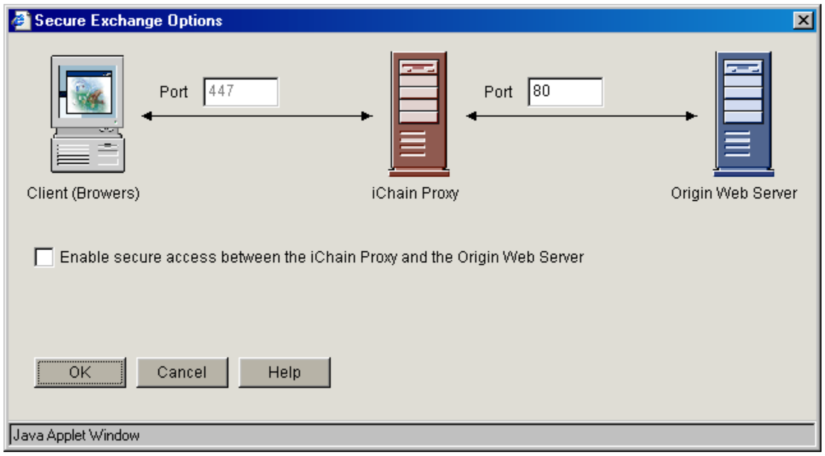

Secure Exchange Options

Port 447

Port 80

□ Enable secure access between the iChain Proxy and the Origin Web Server

This dialog box provides additional Secure Exchange options.

NOTE: When Secure Exchange is enabled, the Web server port must be configured under Secure Exchange Options.

Port (Client to Proxy): Specifies the port to be used for communication between the client (browser) and the proxy server. The default is port 443.

Port (Proxy to Web Server): Specifies the port to be used for communication between the proxy server and the origin Web server. The default is port 80. When Secure Exchange is enabled, this port must be configured here and not at the Web Server Accelerator dialog.

Enable secure access between the iChain Proxy and the Origin Web Server: If this option is selected, the connection between the proxy server and the origin Web server is secured.

Multi-Homing Options Dialog Box

Path: Configure > Web Server Accelerator > Insert > Enable Multi-homing > Multi-homing Options

Figure 19-28 Multi-homing Options Dialog Box

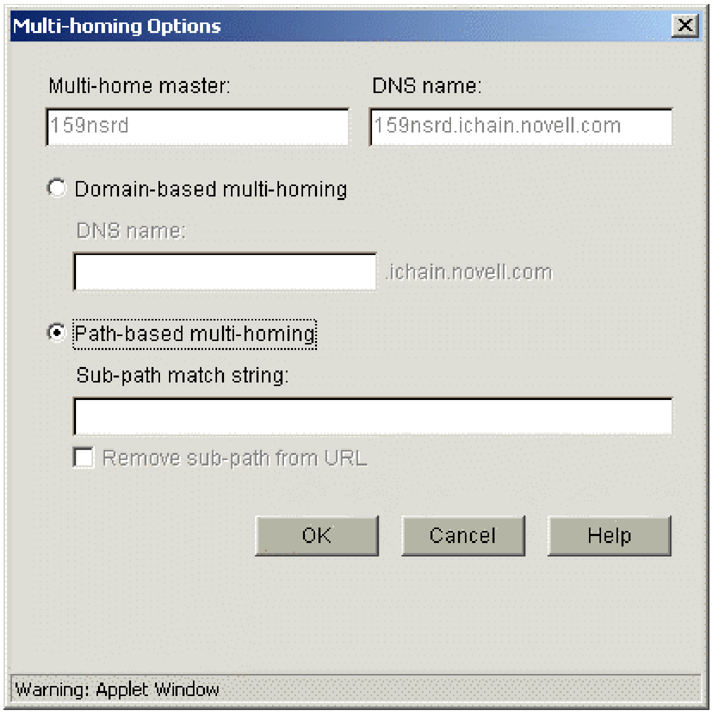

Multi-homing Options

Multi-home master: 159nsrd

DNS name: 159nsrd.ichain.novell.com

○ Domain-based multi-homing

DNS name: .ichain.novell.com

● Path-based multi-homing

Sub-path match string:

□ Remove sub-path from URL

This dialog box lets you specify a string that, if present in the browser request, causes the multi-homing 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).

Sub-Path Match String: The string the multi-homing master compares against the browser request. If the string is not found, the multi-homing master accelerator attempts to fill the request through the Web server addresses in its accelerator definition. If the string is found, the multi-homing master accelerator routes the request to the accelerator with the matching string.

Remove Sub-Path from URL: Select this option if the path string doesn't actually appear at the root of the Web server. If this option is selected, 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 selected, the matched string is retained in the request sent to the Web server.

If you enable this option, JavaScript might not function properly. JavaScript can obscure the URL data that iChain needs to access and modify in order to direct requests to backend servers. Because an absolute path reference can occur anywhere in JavaScript, the iChain parser does not parse through JavaScript.

We recommend that you do not use JavaScript applications with the iChain path-based multi-homing services. All other types of iChain accelerators should function correctly.

19.5.2 Add Authentication Profiles Dialog Box

Path: Configure > Web Server Accelerator > Insert or Modify > Enable Authentication > Authentication Options

Figure 19-29 *Add Authentication Profiles Dialog Box*

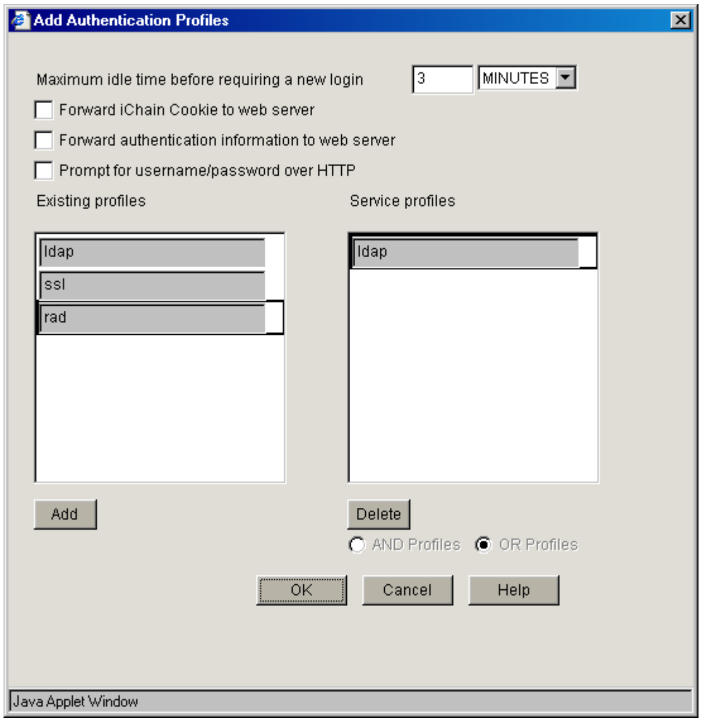

Add Authentication Profiles

Maximum idle time before requiring a new login 3 MINUTES

□ Forward iChain Cookie to web server

□ Forward authentication information to web server

□ Prompt for username/password over HTTP

Existing profiles

- ldap
- ssl
- rad

Service profiles

- ldap

○ AND Profiles ● OR Profiles

The Add Authentication Profiles dialog box lets you select one authentication profile to authenticate the users of the proxy service from which you accessed this dialog box.

The Existing Profiles list shows all the authentication profiles you have created. See "Authentication Dialog Box" on page 315. The Service Profiles list contains the profile that is active for the proxy service from which you accessed the Add Authentication Profiles dialog box.

Maximum Idle Time Before Requiring a New Login: The period of browser inactivity allowed before the proxy server requests a new login.

Forward iChain Cookie to Web Server: If this option is selected, the iChain cookie is sent to the Web server along with the other data being sent.

Forward Authentication Information to Web Server: If this option is selected, the username and password are sent to the Web server. If you have created OLAC parameters for ICHAIN_UID or ICHAIN_PWD, iChain ignores this option and always sends a basic authentication header to the Web server. For more information about OLAC parameters, see Section 13.2, "Using OLAC Custom Header Variables," on page 156.

Prompt for Username/Password over HTTP: If this option is selected, authentication is done over unencrypted HTTP instead of HTTPS.

Existing Profiles: A list of the authentication profiles you create in Cache > Authentication. For more information, see "Authentication Page" on page 314.

Service Profiles: The authentication profile that is active for the proxy service from which you accessed the Add Authentication Profiles dialog box. List content depends on whether And Profiles or Or Profiles is selected. You add a profile to the list by clicking a profile in the Existing Profiles list, then clicking Add. You remove a profile from this list using the Delete button.

AND Profiles: If this option is selected, the profiles in the Service profiles list are used in conjunction with each other to authenticate to users accessing the Web servers specified by this accelerator.

OR Profiles: If this button is selected, either one of the profiles in the Service profiles list are used to authenticate the users accessing the Web servers specified by this accelerator.

19.5.3 Using Authentication Profiles

There are three classes of authentication profiles:

- LDAP
- RADIUS
- Mutual

You cannot have two or more authentication profiles from the same class in combination. If two or more authentication profiles are provided in an “AND” condition, only one user can be represented from a successful login.

Mutual and RADIUS authentications first try to obtain a username or distinguished name (DN), or some unique information about a user. This data is then used to either search for the specific user or bind to the specific user using an LDAP authentication profile. The end result is that iChain will have a full DN.

A user must be found so that OLAC variables can be obtained and ACLCheck can provide access control for the user. Only the ACLCheck authentication profile is used for OLAC and ACLCheck.

The following table shows the categories of authentication profiles that should be tested:

Authentication Class	Sub-class	Description
Mutual	DN	The full DN is stored within the certificate. A potential issue is that the format of the DN might be in reverse order. No formatting is done to the DN.
Mutual	Certificate mapping	Map data in the certificate to a user in an LDAP profile.
RADIUS	Novell RADIUS	The username and token are used to authenticate to the RADIUS server. This can be configured to return the full DN of the user that authenticated. With this DN, a user can be found in an LDAP profile.
RADIUS	Normal	The username and token are used to authenticate to the RADIUS server. The username must be used in some fashion to match to a user in the LDAP profile. This limits the kinds of LDAP profiles that can be used.
LDAP	DN bind with DN input	This configuration lists no LDAP contexts. A full DN is required for input. This profile might not work in conjunction with RADIUS:normal.

Authentication Class	Sub-class	Description
LDAP	DN bind with search field	A DN is built from the search field and username values. A bind () is called for all contexts or a specific context with the DN and password. The default search field value is **cn**.
	Attribute search	A search is performed for equivalence on the defined attribute and the username value. The search is performed for each subtree listed until a match is found. Multiple matches are not allowed. The search is performed using the rights of the LDAP username defined in the LDAP profile. After a match is found, a bind () on the DN of the matched user with the input password is called to validate the password.

The following sections describe the various ways you can configure authentication profiles:

- "The Authentication Process with Mutual Authentication" on page 308
- "The Authentication Process with RADIUS Authentication" on page 308
- "Using LDAP AND Mutual Over HTTP" on page 309
- "Using LDAP AND Mutual Over HTTPS" on page 309
- "Using LDAP OR Mutual, RADIUS OR Mutual" on page 309
- "Using RADIUS AND Mutual" on page 309
- "Using LDAP AND RADIUS" on page 309
- "Using LDAP OR RADIUS" on page 309
- "Using LDAP AND RADIUS AND Mutual" on page 309
- "Using LDAP OR RADIUS OR Mutual" on page 310
- "Setting Timeouts and Pool Limits for LDAP Profiles" on page 310
- "Diagnosing the Need for More Pool Handles" on page 311

The Authentication Process with Mutual Authentication

The data is extracted from the certificate and used to locate a user in an LDAP tree. The LDAP tree that is selected comes from an LDAP authentication profile named ldapcert. If this authentication profile is not found, the LDAP server defined in the aclcheck profile is used. The LDAP authentication profile must have a user and password with rights to read and verify the information within the mutual certificate.

When a mutual authentication profile is only used to authenticate a user, a password is not provided. The username is extracted from searches that are performed with the certificate data and the mapping rules. This means that the password field value is empty if passed to a back-end Web server.

The Authentication Process with RADIUS Authentication

Only a username and token value are needed to log in to a RADIUS server. The LDAP tree that is selected comes from an LDAP authentication profile named ldaprad. If this authentication profile is not found, the LDAP server defined in the aclf profile is used. The LDAP authentication profile needs to have a user and password that has rights to read and verify the username provided in the form. With a properly configured Novell RADIUS server, the full DN is returned with a successful login. This full DN represents the user.

The LDAP password is optional. If a password is provided, the input username and LDAP password are used in a bind () call against the selected LDAP authentication profile. If successful, the user is authenticated. If it is not successful, the user must try again. If the password is not provided, a search is made on the selected LDAP tree to locate the user.

Using LDAP AND Mutual Over HTTP

The user should be prompted for a certificate when a restricted or secure page is hit. The mutual screen is displayed first, followed by the LDAP login screen. Both username values must be the same or the authentication fails.

Using LDAP AND Mutual Over HTTPS

The user should be prompted for a certificate upon initial connection to an accelerator. If required, this could be changed when a restricted or secure page is hit. The LDAP login screen comes up when a restricted or secure page is hit. Both username values must be the same or the authentication fails.

Using LDAP OR Mutual, RADIUS OR Mutual

This configuration raises issues when the user authenticates with mutual authentication and a password is not defined that can be passed to origin Web servers.

Using RADIUS AND Mutual

This configuration is similar to LDAP AND Mutual with the exception that the LDAP password is optional with a RADIUS profile. If a password is provided, then an LDAP bind takes place. The Mutual authentication can use the ldapcert authentication profile and the RADIUS authentication can use the ldaprad authentication profile.

Using LDAP AND RADIUS

This combination forces the user to input the LDAP password. Only one authentication page needs to be presented to the user. The current RADIUS login page lists the LDAP password as optional. The administrator is responsible for configuring it correctly.

Using LDAP OR RADIUS

This combination should give the user the RADIUS login page. The user can either input the username with the token or the username with the LDAP password (either one works). After a RADIUS authentication, the LDAP profile in the combination is used to locate the user within the tree. If the user is not found, the authentication fails.

Using LDAP AND RADIUS AND Mutual

This combination is similar to RADIUS AND Mutual. When the user is located, the LDAP profile is used to validate that the user exists instead of the default aclcheck profile. The main difference is that the LDAP password is required and an LDAP bind takes place.

Using LDAP OR RADIUS OR Mutual

This combination is similar to RADIUS OR Mutual, where once the user is located, the LDAP profile is used to validate that the user exists instead of the default aclcheck profile. The main difference is that the LDAP password is required and an LDAP bind takes place.

Setting Timeouts and Pool Limits for LDAP Profiles

LDAP timeouts and pool limits cannot be set from the Web application; they must be set from the command line. The default limits work in most environments. You can configure the following:

- **Bind Timeout:** By default, a bind request can take up to 10 seconds before it times out and returns an error.
- **Search Timeout:** By default, a search query can take up to 10 seconds before it times out and returns either partial results or an error.
- **Pool Size:** By default, an LDAP server can have 30 open handles between iChain and each LDAP server in the pool. If you have a very busy iChain server and your users are having trouble authenticating or performing LDAP searches, you might want to increase this value. For more information on how to discover whether you need more pool handles, see "Diagnosing the Need for More Pool Handles" on page 311.

These options use the following syntax:

```
SET authentication <name> ldap <option>=<value>
```

Replace <name> with the name of an LDAP profile. Replace <option> and <value> with one of the following:

Option	Description	Example
`bindtimeout=<n>`	Specifies the amount of time a login (LDAP bind) can take. Supported values: ◆ 1 to 255, to specify the number of seconds to wait ◆ 0 to specify 10 seconds Default: 10 seconds	`set authentication corp ldap bindtimeout = 15`
`searchtimeout=<n>`	Specifies the amount of time a search query can take. Supported values: ◆ 1 to 255, to specify the number of seconds to wait ◆ 0 to specify 10 seconds Default: 10 seconds	`set authentication corp ldap searchtimeout = 20`

Option	Description	Example
poolsize=<n>	Specifies the total number of LDAP handles that can be open between iChain and an LDAP server. Supported values: ◆ 1 to 255, to specify the number of handles ◆ 0 to specify 30 handles Default: 30 handles for the LDAP profile, 50 handles for the ACLCHECK profile	`set authentication corp ldap poolsize = 60`

Diagnosing the Need for More Pool Handles

The <n> parameter in the poolsize option specifies the number of handles that each LDAP server in your profile can have. If you have three LDAP servers in a profile and specify poolsize=60, the pool actually has 180 handles for search and bind requests. The search handle gets a bind with the profile user and password, and the bind handle gets a bind with the user identity for that session.

As long as you have two or more LDAP servers in each profile (three is recommended), the default value should be enough. Having only one LDAP server in a profile is not a recommended configuration. Many of the changes in iChain 2.3 SP5 were to improve the failover and load balancing of the LDAP servers. If you use only one LDAP server, iChain has a single point of failure.

If you have only one LDAP server in your profile, 30 handles in the pool might not be enough, especially if you have a busy site. Most authentication requests go through a throttle queue, which is limited to 30 concurrent requests. However, mutual authentication and basic auth authentications do not because of the nature of the protocol. These types of authentications could saturate the pool if the machine is under heavy load. Select one or both of the following methods to determine if your machine is experiencing this problem:

- "Statistical Counter View" on page 311
- "Dynamic View" on page 312

Statistical Counter View

To view whether the pool has been saturated, do the following:

1. At the proxy console screen, select Display LDAP Pool Status.
2. Analyze the data on the new screen. This data should look similar to the following:

```
[ldap] - Some servers down!
   Username: cn=admin,o=novell - Password is valid
   SIMPLE Bind - port 389
   Servers: 10.10.164.176 10.10.167.176(Down) 10.10.164.175
   Pending LDAP Calls: 0
   Counts: Down=5 Timeout=0 Overflow=0
```

3. Check to see if your data has a Counts line.

   The Counts line only appears when there are problems accessing the LDAP servers.The counters appear when the issues start happening, and they cannot be reset.

4 Analyze the counters in the Counts line.

   Down: Indicates the number of times an LDAP server was detected down in the pool.

   Timeout: Indicates the number of times the LDAP server did not respond before the timeout limit was reached.

   Overflow: Indicates the number of times a handle was requested but not obtained.

5 If the Overflow counter is increasing, you need to increase the number of pool handles.

Dynamic View

To view the pool problem as it is happening, do the following:

1 At the proxy console screen, select Debug Options.

2 Select LDAP Trace Options.

3 Change the trace level to 1.

4 At the proxy console screen, select LDAP Pool Messages.

5 View the trace messages looking for BH (NONE FOUND) and SH (NONE FOUND) messages. The messages have the following format:

```
<application name>:(poolname-<address>): Get BH (NONE FOUND) SH(NONE FOUND)
```

6 If you find these messages, you need to increase the number of pool handles.

19.5.4 LDAP Search Query Authentication

There are two LDAP authentication methods for locating a user candidate for login. These methods are very simple to administer and provide common functionality for authentication. The first method builds a list of DN values given a single naming attribute value and a list of contexts. Each of the DN values are tried along with a password using the LDAP bind() operation. The second method performs a search on a single attribute value along with a list of subpaths to search on. If a single user is found, the password is checked against that user using the LDAP bind() operation.

A new feature in iChain 2.3 allows the administrator to set up a search query on more than one single attribute. This feature still searches through a list of subpaths and expects that a single record (user) matches the search criteria. You must know how to express an LDAP search query. See RFC 2254, “The String Representation of LDAP Search Filters,” if you are looking for specifics about building LDAP search strings.

Simple Query Example

The default value for the LDAP search query submits the same query that would be set up for a single attribute search on the cn attribute. This results in a default search query of (&(objectclass=person)(cn=%username%)).

This query has three sections. The first section is the AND logical operator noted as an ampersand (&) value. In LDAP searches, the operators are placed in a prefix notation. In this example, the remaining two sections must return a TRUE result in order for a record to match the search query. The second section of (objectclass=person) is a way to narrow down the search to just person records. Objectclass is an attribute name and person is a specific value.

The last section of (cn=%username%) consists of a search attribute of cn and the value section that expects input from the login form. Any POST variable name can be used. The leading and training % character is used to identify the POST variables that are expected for the search query. In this example, a POST value of username=sam will substitute the value of sam for %username% in the search query.

Figure 19-30 displays the default LDAP search login method:

Figure 19-30 *Default LDAP Search Login Method*

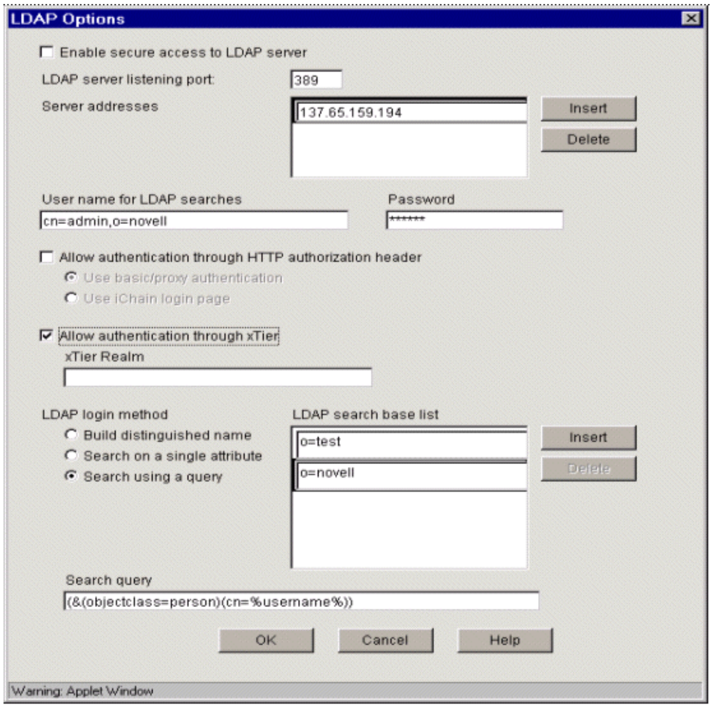

LDAP Options

□ Enable secure access to LDAP server

LDAP server listening port: 389

Server addresses: 137.65.159.194

User name for LDAP searches: cn=admin,o=novell

Password: ******

□ Allow authentication through HTTP authorization header

● Use basic/proxy authentication

○ Use iChain login page

☑ Allow authentication through xTier

xTier Realm:

LDAP login method

○ Build distinguished name

○ Search on a single attribute

● Search using a query

LDAP search base list: o=test, o=novell

Search query: (&(objectclass=person)(cn=%username%))

Warning: Applet Window

Multiple Attribute Query

In this example, the user inputs a name that matches either the cn attribute or the mail attribute. The search query is:

(&(objectclass=person)(|(cn=%username%)(mail=%username%)))

If the POST value submitted during a login is username=sam@company.com&password=test, then the LDAP search query is:

(& (objectclass=person)(|(cn=sam)(mail=sam@company.com)))

The %username variables are replaced with the username POST variables that come from the login page. This provides more flexibility in defining how to locate a specific user in the tree.

At a more complex level, if you want to require the user to match both the cn and mail attribute values of a unique user for all authentication requests, you would configure an LDAP query like the following:

(&(objectclass=person)(cn = %username%)(mail=%emailValue%))

The login page sends a POST with the emailValue tag as well as the username tag. In concept, the login page can have additional POST variables that iChain uses for building the LDAP query. This is a simple search-and-replace functionality. Any %tagnames% that do not have a value or are not in the POST request are replaced with no value (cn=).

For a final example, consider the following LDAP search query:(&(objectclass=person)(| (cn=%username)(ssn=%ssn%)(employeeID=%workID%)))

The login page should prompt the user for a username, social security number, and work ID where one of the three values must match the data within a user record.

This flexibility of a complex search query has a price. You must be well-versed in building LDAP queries. The basics of LDAP search queries are not difficult to learn. Complex search queries might require you to add an index so that searching does not compare each of the records in a subtree. Without an index, some searches might take a very long time to complete.

The search query must be well-formed. A validating LDAP query parser is not part of checking the search query input. If you do not place enough matching parentheses, the LDAP server might appear to be down.

19.5.5 Authentication Page

Path: Configure > Authentication

Figure 19-31 Authentication Page

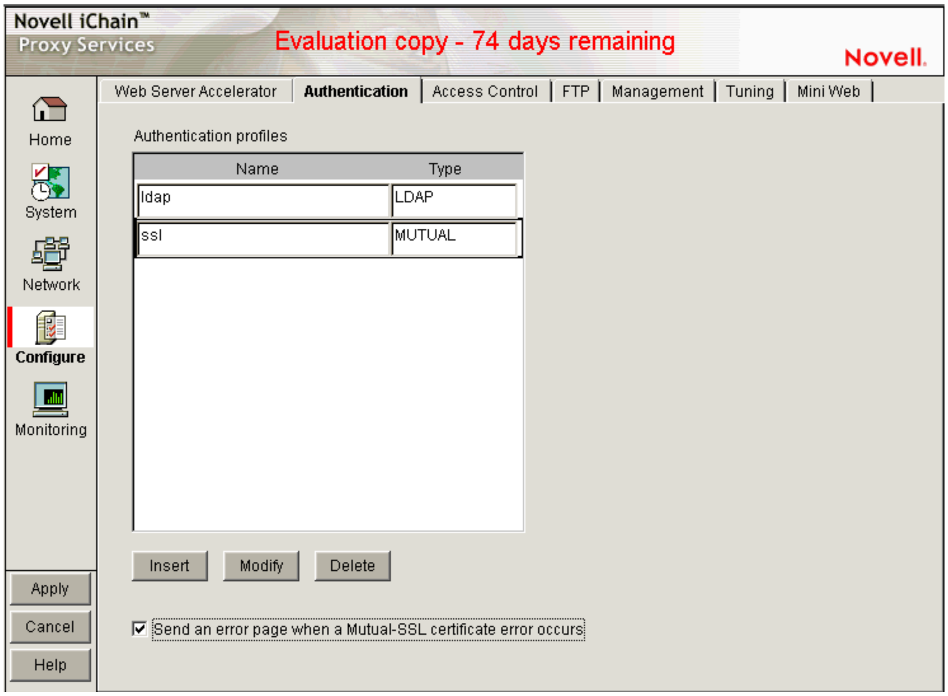

Authentication profiles

Name	Type
ldap	LDAP
ssl	MUTUAL

☑ Send an error page when a Mutual-SSL certificate error occurs

The Authentication page lets you enable the appliance to authenticate users to an LDAP or RADIUS authentication source.

Authentication Profiles List: Lists the authentication profiles you have configured on the appliance using the Authentication dialog box.

Send an Error Page when a Mutual-SSL Certificate Error Occurs: When this option is selected, an error message is sent whenever a Mutual-SSL certificate error occurs during authentication operations.

Authentication Dialog Box

Path: Configure > Authentication > Insert under the Authentication Profiles list

Figure 19-32 *Authentication Dialog Box*

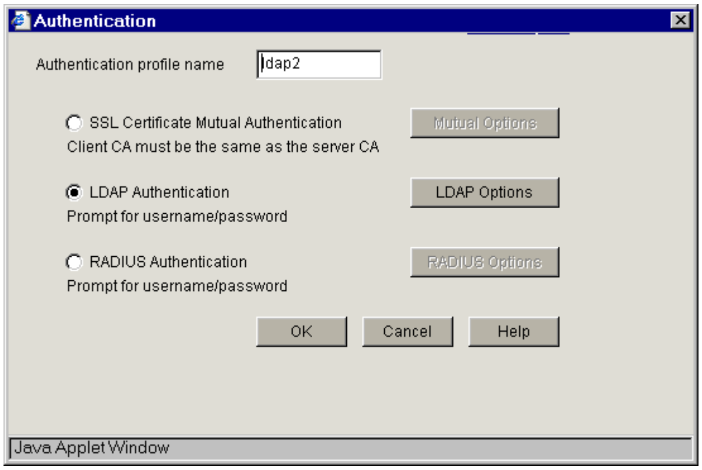

Authentication

Authentication profile name: ldap2

○ SSL Certificate Mutual Authentication
Client CA must be the same as the server CA

● LDAP Authentication
Prompt for username/password

○ RADIUS Authentication
Prompt for username/password

The Authentication dialog box lets you assign an authentication profile name and specify either LDAP or RADIUS as the authentication source.

IMPORTANT: iChain Proxy Services doesn't recognize case differences in profile names. MyProfile and myprofile are, effectively, the same profile name.

Also, iChain Proxy Services 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, iChain Proxy Services removes the first name, concatenates parts of the first profile with the second, and uses 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.

LDAP Options Dialog Box

Path: Configure > Authentication > Insert > LDAP Authentication > LDAP Options

Figure 19-33 *LDAP Options Dialog Box*

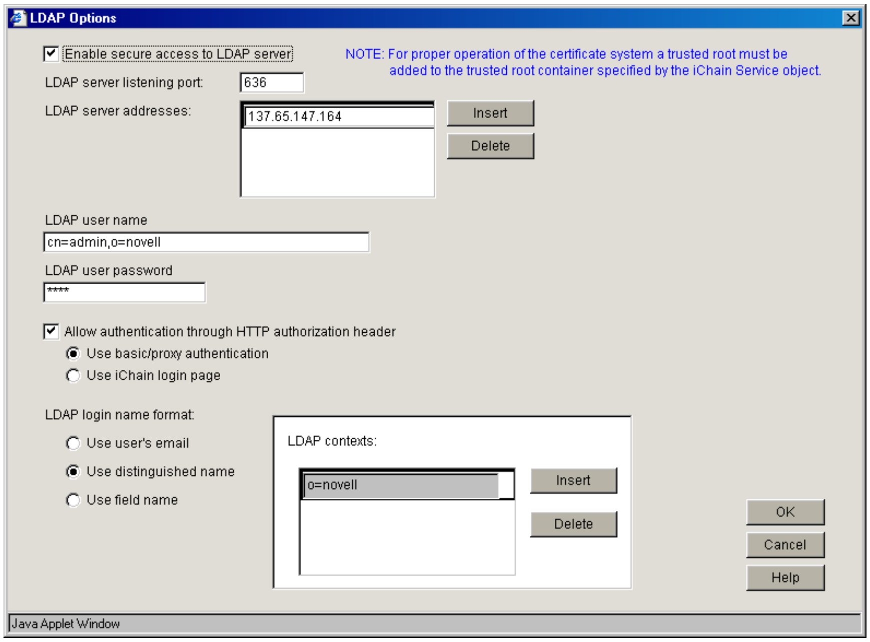

☑ Enable secure access to LDAP server

NOTE: For proper operation of the certificate system a trusted root must be added to the trusted root container specified by the iChain Service object.

LDAP server listening port: 636

LDAP server addresses: 137.65.147.164

LDAP user name: cn=admin,o=novell

LDAP user password: ****

☑ Allow authentication through HTTP authorization header

● Use basic/proxy authentication

○ Use iChain login page

LDAP login name format:

○ Use user's email

● Use distinguished name

○ Use field name

LDAP contexts: o=novell

Use the LDAP Options dialog box to configure the appliance so users can authenticate through an LDAP database.

LDAP Server Addresses: The IP addresses of the LDAP servers. Specifying multiple LDAP servers allows for failover LDAP servers. If the first LDAP server cannot be accessed, the next LDAP server on the list is tried, and so on down the list in the order specified. For more information, see Appendix B, “Using LDAP Server Load Balancing and Failover,” on page 387.

LDAP Server Listening Port: The port number the LDAP server is listening on for requests from LDAP clients. The default is 389 for normal access. Use 636 for secure access.

Enable Secure Access to LDAP Server: Selecting this option causes the data sent between the LDAP client and the LDAP server to be sent using SSL.

Username: Specify the username to use for accessing the LDAP server.

Password: Specify the password to use for accessing the LDAP server.

LDAP Login Name Format Dialog Box

The contents of this box change depending on the option selected.

Use User's E-Mail: 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 searches 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 specify the context of one or more LDAP containers from which the search for the e-mail name should begin.

Use Distinguished Name: Select this option to allow users to authenticate using their LDAP usernames. Users can use either their fully distinguished (full LDAP contexts) LDAP usernames, or you can provide a list of LDAP contexts so users only need to type their usernames.

LDAP Contexts: Contains specific contexts in which the LDAP server looks 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: Select this option to require that users enter a specific LDAP field name.

Field Name: The LDAP field name (such as User ID) through which users can authenticate.

LDAP Search Base: Click Insert to specify the context of one or more LDAP containers. The appliance performs a subtree search in all containers in the list and in their subcontainers.

RADIUS Authentication Dialog Box

Path: Configure > Authentication > Insert > RADIUS Authentication > RADIUS Options

Figure 19-34 *RADIUS Options Dialog Box*

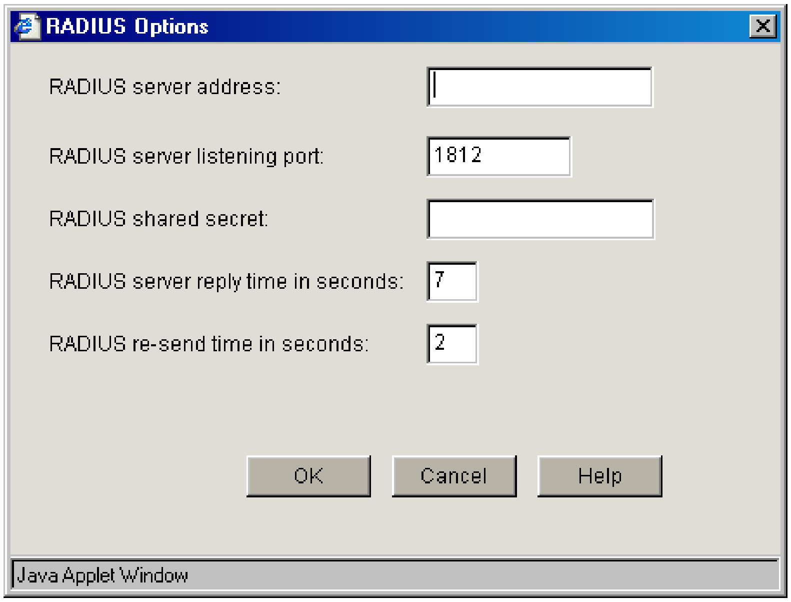

RADIUS Options

RADIUS server address:

RADIUS server listening port: 1812

RADIUS shared secret:

RADIUS server reply time in seconds: 7

RADIUS re-send time in seconds: 2

Use this dialog box to specify a RADIUS server the appliance can use for authentication.

NOTE: Port 1812 is the default RADIUS server port; however, the Novell RADIUS server defaults to port 1645. It is possible for a different default port to be specified when the Novell RADIUS server is loaded.

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 Server Reply Time in Seconds: The total time the appliance waits for a response from the RADIUS server before authentication fails. The default is 7 seconds.

RADIUS Resend Time in Seconds: The interval in seconds between appliance requests to the RADIUS server. The default is two seconds. This means that the appliance could send three requests before the 7-second default limit expires and the authentication request fails.

Third-Party RADIUS Server Support

iChain supports RADIUS authentication via the RADIUS server supplied by Novell or a RADIUS server supplied by third-party vendors. The Novell RADIUS server uses Novell eDirectory as the underlying database for storing user objects.

If the Novell RADIUS server is running on the iChain authorization server, it has the ability to indicate the fully distinguished name of the user object that was authenticated via RADIUS. iChain must use the fully distinguished name of the user in order to perform ACL rights selecting when accessing a protected resource. A RADIUS server provided by a vendor other than Novell stores user objects in its own database.

iChain must be able to map this third-party name to a user object stored in eDirectory on the authorization server in order to determine if the user has rights to the protected resource he or she is attempting to access. This is done by performing an LDAP subtree search on the authorization server for the user name entered on the iChain login dialog box. In order to perform this search, two fields must be changed on the ACLCheck authentication profile.

A search base must be specified that identifies an eDirectory container that defines where in the tree to begin the search and the anonymous bind feature must be set to No. This is done by entering the following commands on the iChain Proxy Services console screen:

- `add authentication aclcheck ldap searchbase =` *`container name`*
- `set authentication aclcheck ldap bindanonymous = no`
- `apply`

If there is already a value specified for the search base, the set command must be used instead of the add command, as shown below:

```
set authentication aclcheck ldap searchbase = container name
```

The container name can specify an eDirectory organization or an organizational unit that appears on the authorization server. For example:

```
add authentication aclcheck ldap searchbase = o=novell
```

The implication here is that an eDirectory user object must be created on the authorization server whose CN is the same as the RADIUS user name and it must appear hierarchically in the tree somewhere below the container specified by the search base.

Also verify that iChain has defined an LDAP server, user, and password. This is needed to perform the search base lookup.

For example:

```
>get authentication aclcheck ldap
authentication aclcheck ldap address = 10.252.3.5
authentication aclcheck ldap bindusername = cn=admin;o=novell
authentication aclcheck ldap bindpassword = password
```

If the above values are empty, they can be input at the iChain browser GUI > Configure > ACL page.

19.5.6 Access Control Page

Path: Configure > Access Control

Figure 19-35 *Access Control Page*

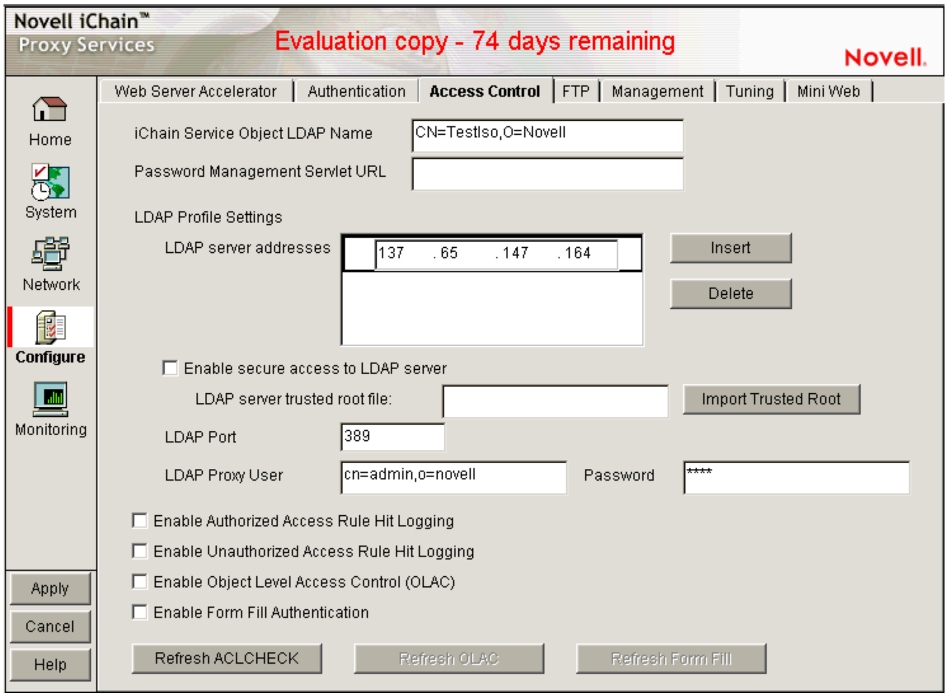

iChain Service Object LDAP Name: CN=TestIso,O=Novell

Password Management Servlet URL:

LDAP Profile Settings

LDAP server addresses: 137 . 65 . 147 . 164

□ Enable secure access to LDAP server

LDAP server trusted root file:

LDAP Port: 389

LDAP Proxy User: cn=admin,o=novell Password: ****

□ Enable Authorized Access Rule Hit Logging

□ Enable Unauthorized Access Rule Hit Logging

□ Enable Object Level Access Control (OLAC)

□ Enable Form Fill Authentication

The Access Control page lets you define the parameters necessary to set up access control. This page includes the following fields:

iChain Service Object LDAP Name: Specifies the name of the iChain Service Object (ISO) containing parameter settings defining your iChain domain or infrastructure.

Password Management Servlet URL: Defines the URL of the password management servlet on the iChain Proxy Server. This servlet enables users within your iChain infrastructure to change their passwords. Use the full http URL. For example, http://ichain.provo.novell.com/servlet/iChainPasswordMgr where /servlet/ is the servlet directory of your servlet engine.

NOTE: Novell Technical Support does not support the password management servlets used in the iChain configuration. These servlets are not developed by the iChain development team; therefore, no bugs related to these servlets will be fixed.

The ConsoleOne Password Policy TAB contains options for writing to the ISO object. This means that the password management servlet from forge.novell.com or a custom developed servlet can read password properties from an iChain object. However, Novell Technical Services does not support servlets from forge.novell.com. For more information, see TID 407040 (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10097123.htm).

LDAP Server Addresses: Specifies the IP addresses of the iChain LDAP access control servers. Specifying multiple LDAP servers allows for load-balancing of the LDAP servers. Requests to the LDAP servers for access control are done in a round-robin fashion, meaning, the requests are distributed among multiple LDAP servers to balance the load. For more information, see Appendix B, "Using LDAP Server Load Balancing and Failover," on page 387.

Enable Secure Access to LDAP Server: Selecting this option causes the data sent between the LDAP client and the LDAP server to be sent using SSL.

LDAP Server Trusted Root File: The path on the appliance to a trusted root file that contains the Certificate Authority (CA) used by the LDAP server in the profile you are creating. The system fills this field with information for the trusted root file you create using the Import Trusted Root button. 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: Clicking this button opens the Import Trusted Root dialog box.

Import Trusted Root Dialog Box

Path: Configure > Access Control > Import Trusted Root

Figure 19-36 *Import Trusted Root Dialog Box*

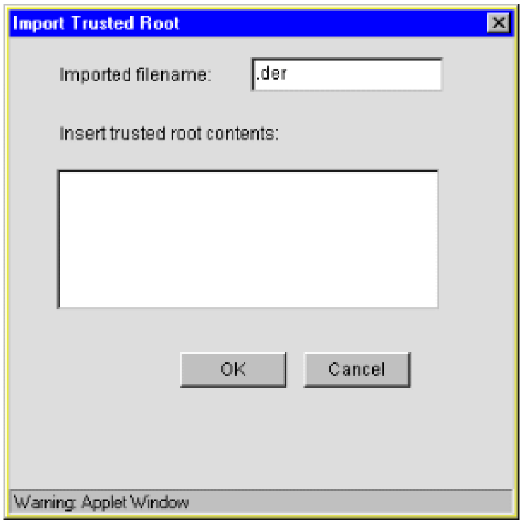

Import Trusted Root

Imported filename: .der

The Import Trusted Root dialog box lets you create a trusted root file with a .der extension that contains information identifying the Certificate Authority used by the LDAP server.

To create a trusted root file:

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. 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 iChain Proxy Services 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 "Limitations of the Appliance's Mini FTP Server" on page 358.

If you don't include a path with the filename, the proxy server creates the file at the root of the sys: volume. You cannot see the root of the sys: volume using FTP.

2 Export the Certificate Authority in .b64 format.

2a Using ConsoleOne, open the properties of the Tree Organizational CA.

2b Select Self Signed under the certificate tab in the Security Container.

2c Click Export.

Do not export the private key.

2d Select File in Base64 Format and then click Next.

2e Using a text editor on your configuration workstation, open the .b64 file for the Certificate Authority, select the file contents, then paste the contents to the Clipboard.

The content should be similar to the following example:

```
-----BEGIN CERTIFICATE-----
MIIFOjCCBCKgAwIBAgIhAhwR+pxWWF78DPWCaZJz8K245MzcSvS0EwzRbBJwAgEKMA0GCSqGS
Ib3DQEBBQUAMDcxGjAYBgNVBAsTEU9yZ2FuaXphdGlvbmFsIENBMRkwFwYDVQQKExBKU0hPUl
QtTlc2NS1UUkVFMB4XDTA0MDUzMTE0MDY0OFoXDTE0MDUzMBgwEAIBAAIIf////////
8BAQACBBH6nFaiTjBMAgECAgIA/wIBAAMNAID//////////////wMJAID/////////
MBIwEAIBAAIIf/////////8BAf8wEjAQAgEAAgh//////////wEB/
zANBgkqhkiG9w0BAQUFAAOCAQEAUo9ffBgSGaYYDA5s6nqXdPxAb9XTVMz14mUT0JPOKtFgzd
KLxyzS9GZAZrNTIz5izgrrAI7NoG5aNciohoCpEroUBPFCrwzGowqaixug/
82DSjxhxljrVWyY7RH24N3uLbtVmLTrlUj+b1wnx3+wPxpxLpxQR84nWwA5JtsNG+E6/
MT64oUYPqSHvBmrcmSl3+X2+pbTxJB/
S+Znh7NDiIGvwAgM+N2JaCfUYWKzHeio1PhHyaqqSugZiYQ6DPXLXg7QXU+zVNjTQJDFx6BPC
rjFbUHGsyVdJUiqDVtyHupb6qDeGplzssp/C/nYY5wpPf0P5HGk8lKV8Uv20UX/fw==-----
END CERTIFICATE-----
```

3 Return to the Import Trusted Root dialog box, then paste the clipboard contents into the Insert Trusted Root contents text box.

4 Click OK.

5 Apply the changes.

6 At the iChain command line interface, initiate Set Authentication refreshaclcheck=1 to renew LDAP connections.

LDAP Port: Specifies the port on which the LDAP server listens for access control requests. The default is 389 for normal access. Uses 636 for secure access.

LDAP Proxy User: Specifies the LDAP username for the iChain Proxy Server to use when making requests for access control information from the iChain LDAP access control server.

Password: Specifies the LDAP password for the iChain Proxy Server to use when making requests for access control information from the iChain LDAP access control server.

Enable Authorized Access Rule Hit Logging: Specifies whether to enable logging for authorized access rule hits. The default is Disabled.

Enable Unauthorized Access Rule Hit Logging: Specifies whether to enable logging for unauthorized access rule hits. The default is Disabled.

Enable Object Level Access Control: Specifies whether to enable object level access control.

IMPORTANT: When the access control profile is first generated, object level access control must be disabled. Apply the ACL profile, mark Enable, then click Apply.

Enable Form Fill Authentication: Specifies whether to enable single sign-on with login HTML forms from origin servers.

Refresh ACLCheck: Refreshes the ACLCheck parameter settings.

Refresh OLAC: If OLAC has been enabled, use Refresh to force an immediate reread of OLAC parameters.

Refresh Form Fill: Refreshes the connection to LDAP servers and Form Fill policy rules.

19.5.7 FTP Page

Path: Configure > FTP

Figure 19-37 *FTP Page*

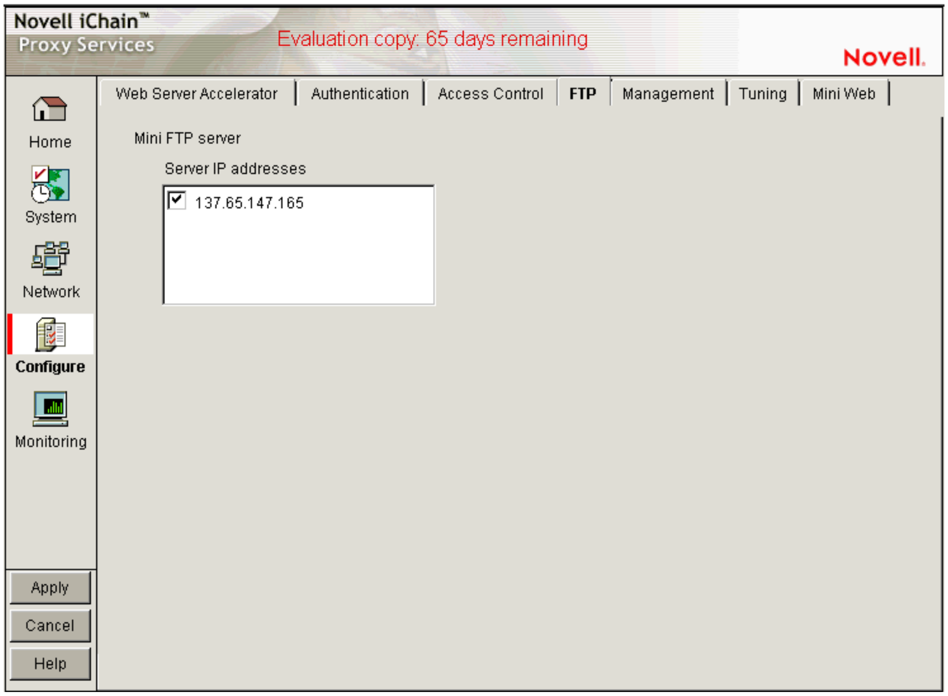

Mini FTP server

Server IP addresses

☑ 137.65.147.165

The FTP page lets you configure the appliance to provide an FTP listening address (Mini FTP Server) for appliance management.

Mini FTP Server: Select an address in the Server IP Addresses list to enable it for FTP listening. If an address is not selected, you cannot upload or download files using FTP.

19.5.8 Management Page

Path: Configure > Management

Figure 19-38 *Management Page*

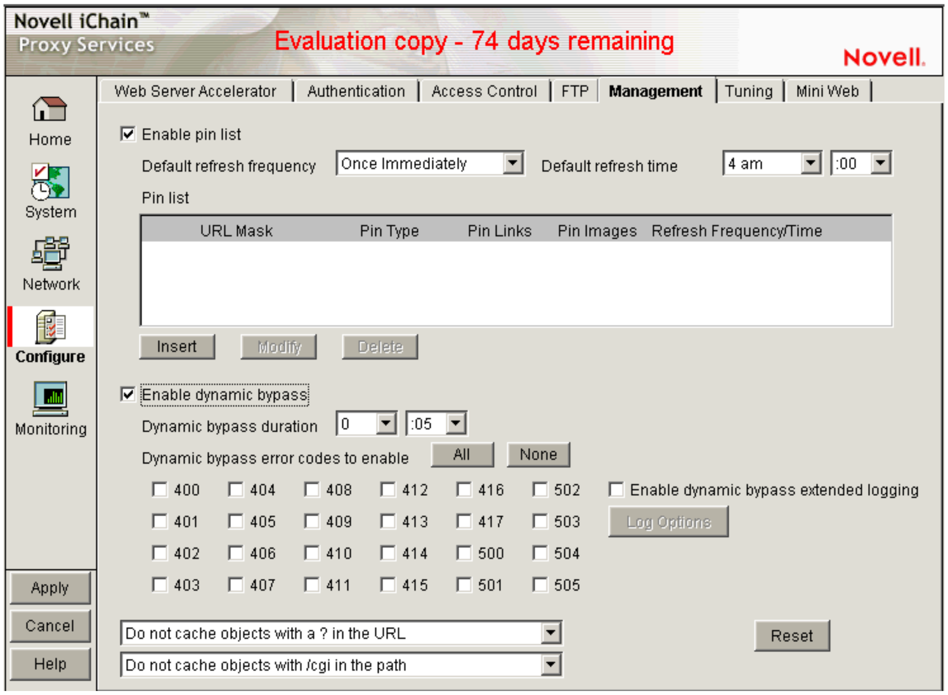

☑ Enable pin list

Default refresh frequency: Once Immediately   Default refresh time: 4 am :00

Pin list

URL Mask	Pin Type	Pin Links	Pin Images	Refresh Frequency/Time

☑ Enable dynamic bypass

Dynamic bypass duration: 0 :05

Dynamic bypass error codes to enable

□ 400 □ 404 □ 408 □ 412 □ 416 □ 502

□ 401 □ 405 □ 409 □ 413 □ 417 □ 503

□ 402 □ 406 □ 410 □ 414 □ 500 □ 504

□ 403 □ 407 □ 411 □ 415 □ 501 □ 505

□ Enable dynamic bypass extended logging

Do not cache objects with a ? in the URL

Do not cache objects with /cgi in the path

The Management page lets you identify objects that are 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: Select this option to activate pinning on the appliance. For information regarding what pin lists are and how they function, see "The Pin List" on page 234 and "Pin List Examples" on page 238.

NOTE: Selecting 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: Lets you 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: Lets you specify the time of day or the time interval when pinned objects are 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: This 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 238.

Pin Type: This specifies whether and how the appliance will cache objects that match the URL mask.

- **Normal:** iChain Proxy Services 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:** iChain Proxy Services keeps the pinned objects in cache as long as possible, although they might be written to the appliance's hard disk.
- **Memory:** iChain Proxy Services 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:** iChain Proxy Services does not cache the objects. In other words, you can use this option to prevent objects from being cached.

Pin Links: Specifies how many link levels iChain Proxy Services follows the pin type rule you've established.

Pin Images: Used to pin image files that reside on a different host than the page requested.

Refresh Frequency/Time: Lets you 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.

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 cachable 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 proxy server appliance does not cache objects that Web server administrators have marked non-cachable.

By the same token, the Do Not Cache option is sometimes misinterpreted to imply that objects meeting the criteria are never cached. That is also not the case. Objects containing question marks and /cgi in the path might meet other criteria that cause them to be cached. This option only causes the proxy server to ignore question marks and /cgi in determining whether to cache objects.

Resetting the Defaults

Clicking the Reset button resets the two drop-down lists back to their default (do not cache) settings.

19.5.9 Tuning Page

Path: Configure > Tuning

Figure 19-39 *Tuning Page*

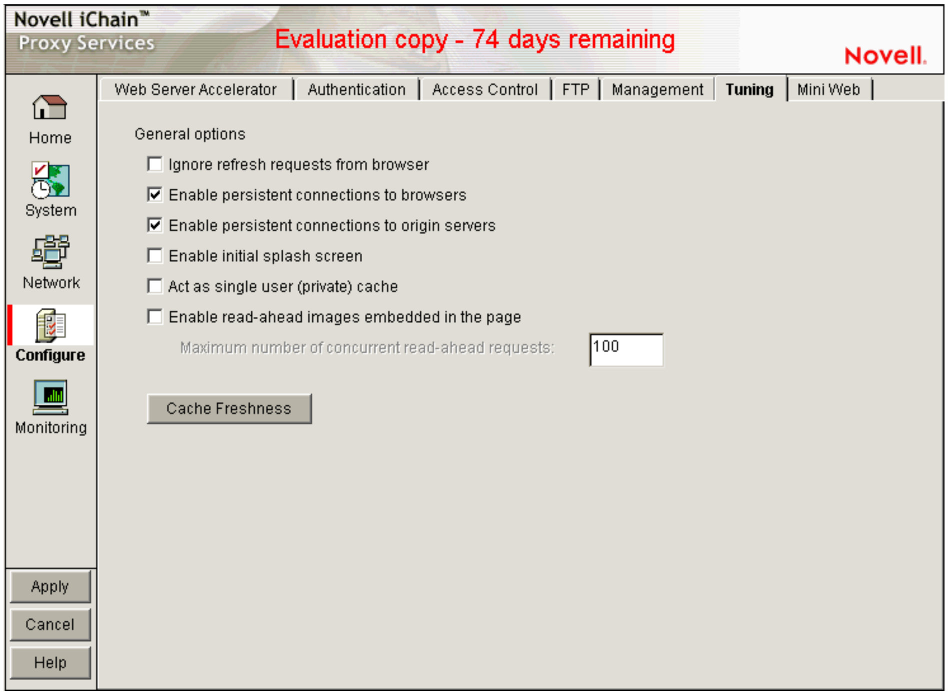

General options

☐ Ignore refresh requests from browser

☑ Enable persistent connections to browsers

☑ Enable persistent connections to origin servers

☐ Enable initial splash screen

☐ Act as single user (private) cache

☐ Enable read-ahead images embedded in the page

Maximum number of concurrent read-ahead requests: 100

The Tuning page lets you restrict and enable functionality that affects all appliance operations. The implications for each option are explained below.

Ignore Refresh Requests from Browser: When a user clicks Refresh or Reload in the browser, the default system action is to send a new request to the Web server. Selecting this option causes refresh or reload requests to be filled from the cache until cache freshness parameters are met. See "Cache Freshness Dialog Box" on page 326.

Enable Persistent Connections to Browsers: If enabled, 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: If enabled, 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: If enabled, 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 "Using FTP to Customize the Appliance Splash Screen" on page 360. The splash screen is disabled (turned off) by default.

Act as a Single User (Private) Cache: If enabled, iChain Proxy Services caches objects that have been flagged for private caches only.

Enable Read-Ahead Images Embedded in the Page: If selected, iChain Proxy Services retrieves and caches objects that have been flagged Read-Ahead. You specify the maximum number of Read-Ahead objects the proxy server retrieves in the Maximum Number of Concurrent Read-Ahead Requests field.

Cache Freshness Dialog Box

Path: Configure > Tuning > Cache Freshness

Figure 19-40 *Cache Freshness Dialog Box*

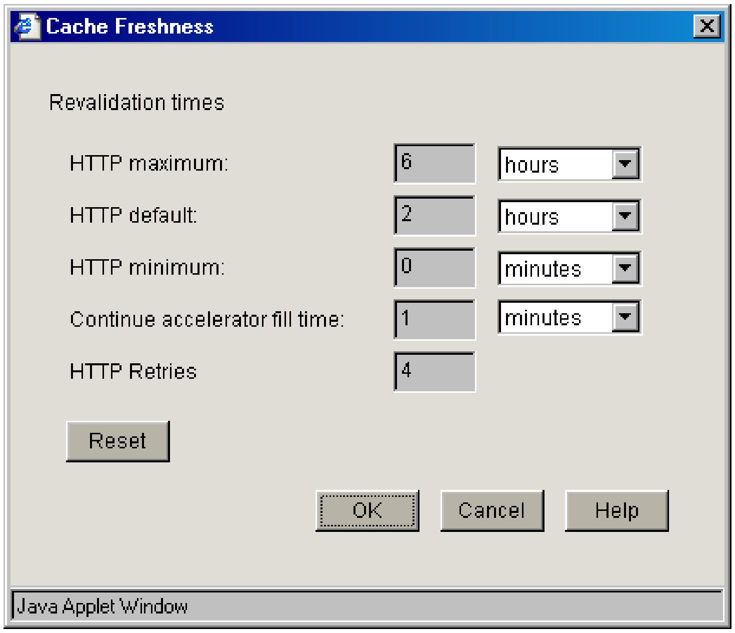

Cache Freshness

Revalidation times

HTTP maximum:	6	hours
HTTP default:	2	hours
HTTP minimum:	0	minutes
Continue accelerator fill time:	1	minutes
HTTP Retries	4	

The Cache Freshness dialog box lets you set time values governing when iChain Proxy Services revalidates requested cached objects against those on their respective origin Web servers. If requested objects have changed, iChain Proxy Services re-caches them. Default field values are shown in Figure 19-40.

iChain Proxy Services does not automatically recache objects when they expire. Expired objects are revalidated (and recached 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 Section 4.9, “Cache Freshness,” on page 73.

HTTP Maximum: The maximum number of hours or days iChain Proxy Services serves HTTP data from cache before revalidating it against content on the origin Web server. No object is served from cache after this value expires without being revalidated.

This overrides a freshness or Time to Expire header value specified by the Webmaster if he or she specified a longer time.

You use this value to reduce the maximum time iChain Proxy Services waits before checking whether requested objects need to be refreshed.

HTTP Default: The value iChain Proxy Services 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 iChain Proxy Services serves HTTP data from cache before revalidating it against content on the origin Web server. No requested object is 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 iChain Proxy Services waits before checking whether requested objects need to be refreshed. This parameter does not override No Cache or Must Revalidate directives from the origin Web server.

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.

HTTP Retries: The number of retry requests to be issued.

19.5.10 Mini Web Page

Path: Configure > Mini Web (to see this page, click the upper-right corner of the Tuning page.)

Figure 19-41 *Mini Web Page*

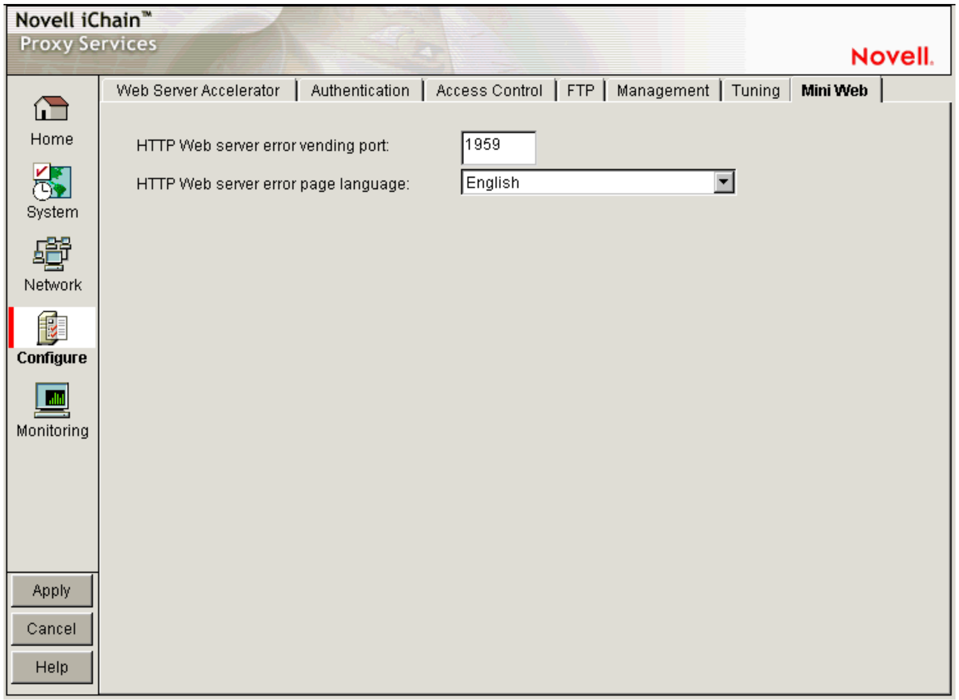

HTTP Web server error vending port: 1959

HTTP Web server error page language: English

The Mini Web page lets you configure how appliance-generated error pages are vended to browsers.

HTTP Web Server Error Vending Port: The port the browser uses 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: From the drop-down list, you can select a language for appliance-generated error messages to browsers.

19.6 The Monitoring Panel Options

You can monitor various appliance activities and statistics in the browser-based tool. This section explains the following pages:

- "Summary Page" on page 328
- "Services Page" on page 329
- "Performance Page" on page 330
- "Cache Page" on page 331
- "Cache Logs Page" on page 332
- "Top Ten Sites Page" on page 332

19.6.1 Summary Page

Path: Monitoring > Summary

Figure 19-42 *Summary Page*

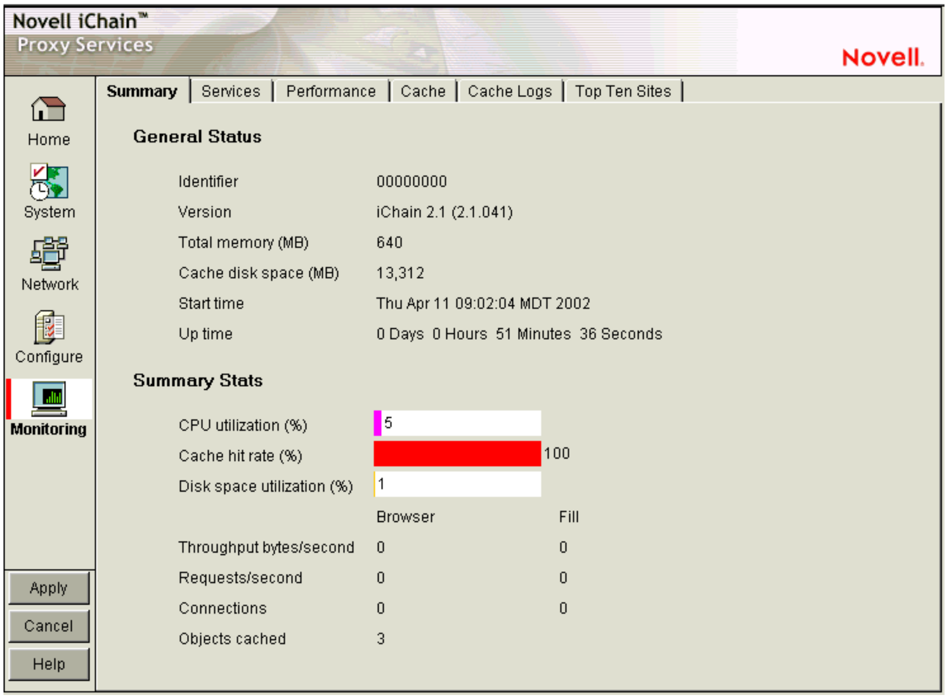

General Status

Identifier	00000000
Version	iChain 2.1 (2.1.041)
Total memory (MB)	640
Cache disk space (MB)	13,312
Start time	Thu Apr 11 09:02:04 MDT 2002
Up time	0 Days 0 Hours 51 Minutes 36 Seconds

Summary Stats

CPU utilization (%)	5	
Cache hit rate (%)	100	
Disk space utilization (%)	1	
	Browser	Fill
Throughput bytes/second	0	0
Requests/second	0	0
Connections	0	0
Objects cached	3	

The Summary page 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. Use this field to verify whether the proxy server 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.

19.6.2 Services Page

Path: Monitoring > Services

Figure 19-43 *Services Page*

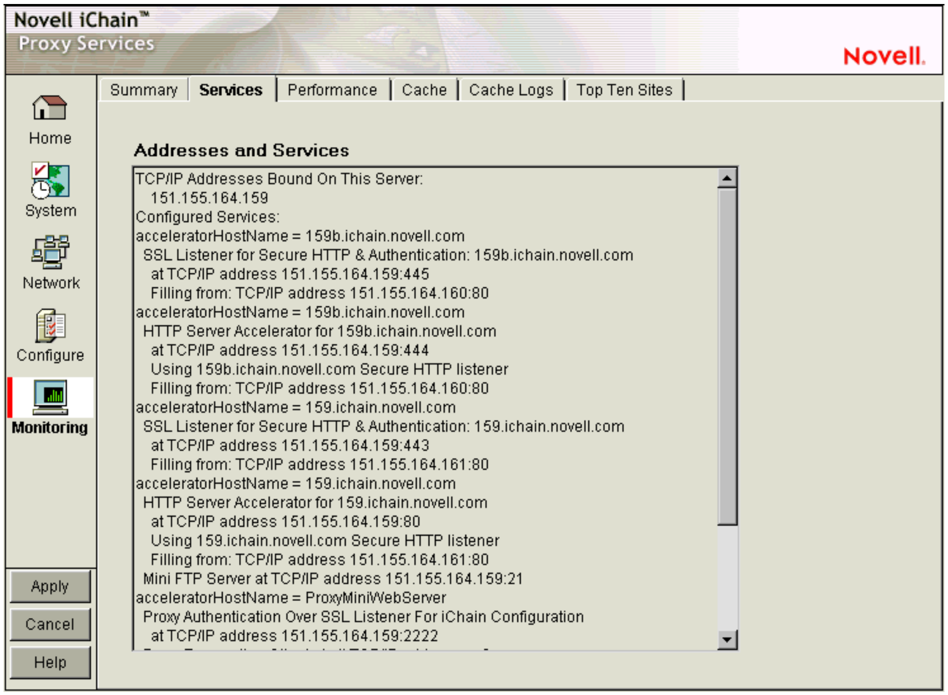

Addresses and Services

```
TCP/IP Addresses Bound On This Server:
  151.155.164.159
Configured Services:
acceleratorHostName = 159b.ichain.novell.com
 SSL Listener for Secure HTTP & Authentication: 159b.ichain.novell.com
  at TCP/IP address 151.155.164.159:445
  Filling from: TCP/IP address 151.155.164.160:80
acceleratorHostName = 159b.ichain.novell.com
 HTTP Server Accelerator for 159b.ichain.novell.com
  at TCP/IP address 151.155.164.159:444
  Using 159b.ichain.novell.com Secure HTTP listener
  Filling from: TCP/IP address 151.155.164.160:80
acceleratorHostName = 159.ichain.novell.com
 SSL Listener for Secure HTTP & Authentication: 159.ichain.novell.com
  at TCP/IP address 151.155.164.159:443
  Filling from: TCP/IP address 151.155.164.161:80
acceleratorHostName = 159.ichain.novell.com
 HTTP Server Accelerator for 159.ichain.novell.com
  at TCP/IP address 151.155.164.159:80
  Using 159.ichain.novell.com Secure HTTP listener
  Filling from: TCP/IP address 151.155.164.161:80
 Mini FTP Server at TCP/IP address 151.155.164.159:21
acceleratorHostName = ProxyMiniWebServer
 Proxy Authentication Over SSL Listener For iChain Configuration
  at TCP/IP address 151.155.164.159:2222
```

The Services page 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: Use this list for troubleshooting problems with configured services. It shows 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.

19.6.3 Performance Page

Path: Monitoring > Performance

Figure 19-44 *Performance Page*

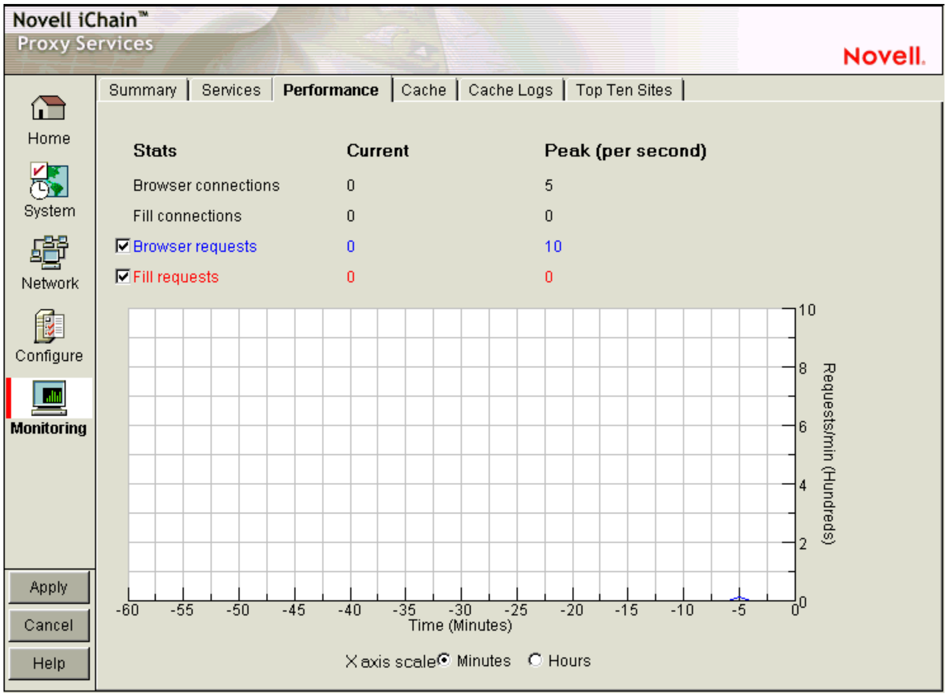

Stats	Current	Peak (per second)
Browser connections	0	5
Fill connections	0	0
☑ Browser requests	0	10
☑ Fill requests	0	0

X axis scale ● Minutes ○ Hours

The Performance page shows current and peak levels of usage in terms of TCP connections and HTTP requests. The page 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. Select this option to enable graphing of browser requests.

Fill Requests: The current and peak number of appliance requests per second to origin Web servers. Select this option 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.

19.6.4 Cache Page

Path: Monitoring > Cache

Figure 19-45 *Cache Page*

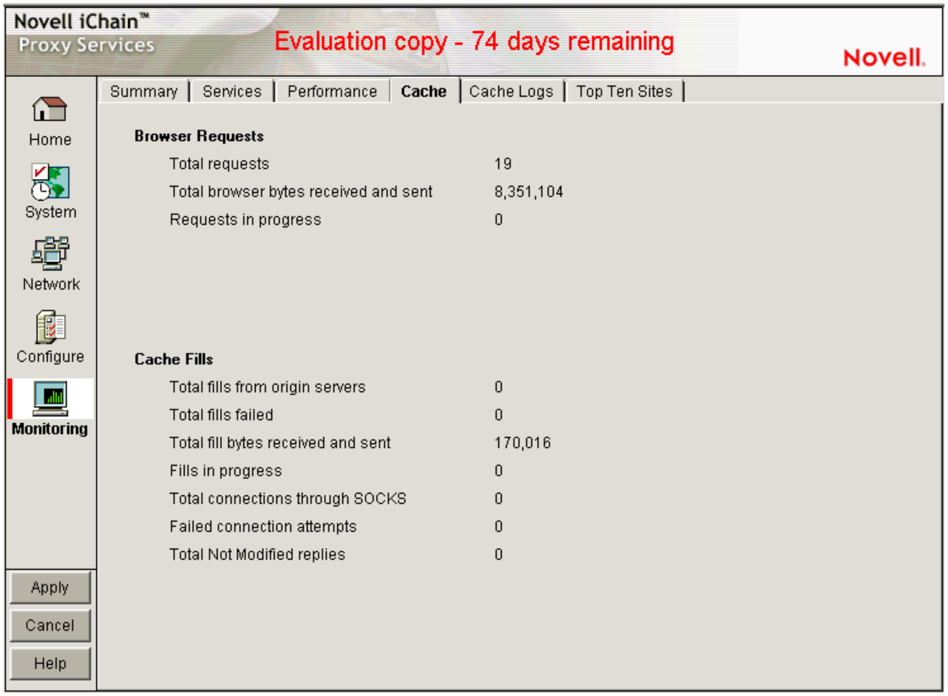

Browser Requests

Total requests	19
Total browser bytes received and sent	8,351,104
Requests in progress	0

Cache Fills

Total fills from origin servers	0
Total fills failed	0
Total fill bytes received and sent	170,016
Fills in progress	0
Total connections through SOCKS	0
Failed connection attempts	0
Total Not Modified replies	0

The Cache page 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 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 Fills from Origin Servers: The total number of fill requests the appliance has made to origin Web servers.

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 Not Modified Replies: The total number of 304 Not Modified replies received for all fill requests to origin servers.

19.6.5 Cache Logs Page

Path: Monitoring > Cache Logs

Figure 19-46 Cache Logs Page

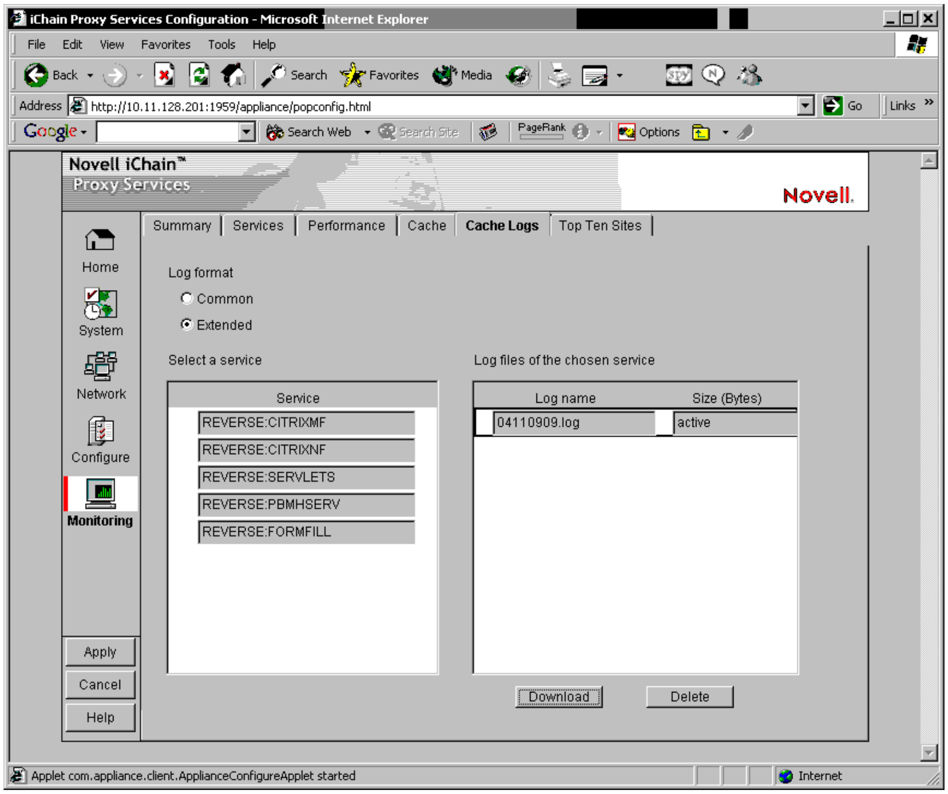
iChain Proxy Services Configuration - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://10.11.128.201:1959/appliance/popconfig.html
Go Links
Google Search Web Search Site PageRank Options
Novell iChain™
Proxy Services
Novell.
Home
System
Network
Configure
Monitoring
Apply
Cancel
Help
Summary Services Performance Cache Cache Logs Top Ten Sites
Download
Delete
Internet

Log format

○ Common

● Extended

Select a service

Service
REVERSE:CITRIXMF
REVERSE:CITRIXNF
REVERSE:SERVLETS
REVERSE:PBMHSERV
REVERSE:FORMFILL

Log files of the chosen service

Log name	Size (Bytes)
04110909.log	active

Applet com.appliance.client.ApplianceConfigureApplet started

The Cache Logs page 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.

19.6.6 Top Ten Sites Page

Path: Monitoring > Top Ten Sites

Figure 19-47 *Top Ten Sites Page*

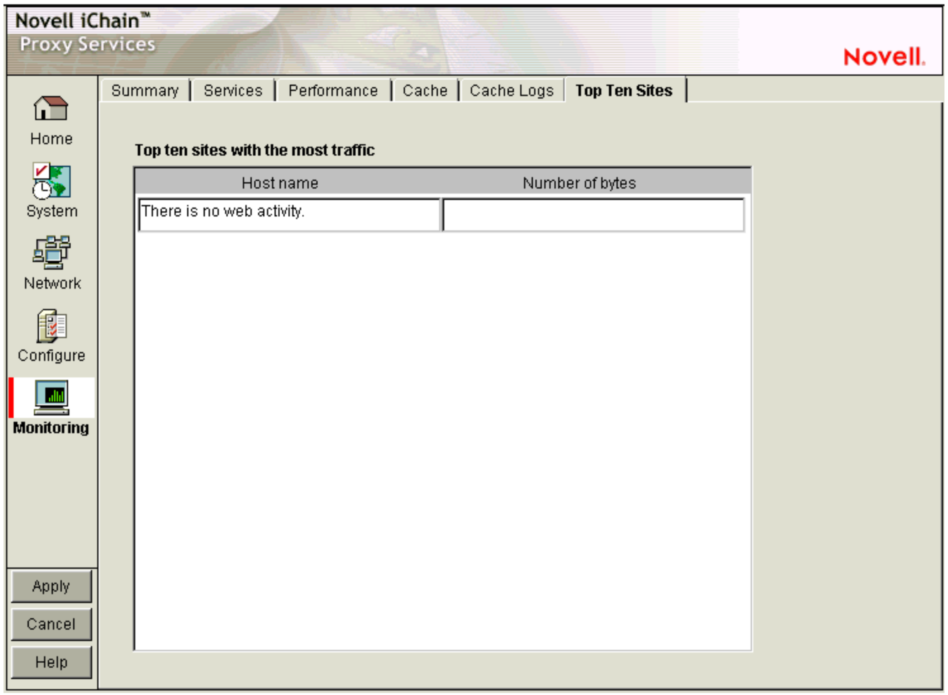

Top ten sites with the most traffic

Host name	Number of bytes
There is no web activity.	

The Top Ten Sites page 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.

19.7 Using Appliance Commands

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.

iChain Proxy Services 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

Command	Function	Requires Arguments?
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
createsessionbrokerkey	Creates an encryption key for Session Broker communication	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
installsessionbrokerkey	Installs a Session Broker encryption key from a floppy	No
lock	Re-locks the iChain Proxy Server console	No
ping	Sends a ping request to the addresses specified. Ports are optional.	Yes
purgecache	Purges the cache buffers.	No
remove	Deletes the specified value.	Yes
resetlearnedroutes	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
restore	Restores a configuration file to its directory	Yes
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

19.7.1 Troubleshooting the Command Line

Commands entered return an error

- Make sure you use the equal (=) sign when setting or adding, for 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 show when I use the GET command to display them

- You must conclude with the apply command to make the values take effect.

19.7.2 Connecting through Telnet

Telnet is disabled by default. At the command line interface, you can use the following command to enable Telnet:

```
set listener telnet enable=yes
```

IMPORTANT: Telnet access is not secure unless a password is set. We strongly recommend that you set system passwords as part of the appliance initialization process. For more information, see "Password Dialog Box" on page 277. Remember that when you are using Telnet, the login username and password are sent in clear text.

You can manage the iChain Proxy Services by using commands from a workstation with a Telnet connection to the appliance.

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, then run telnet.exe.

3 From the Telnet screen, click Terminal > Preferences.

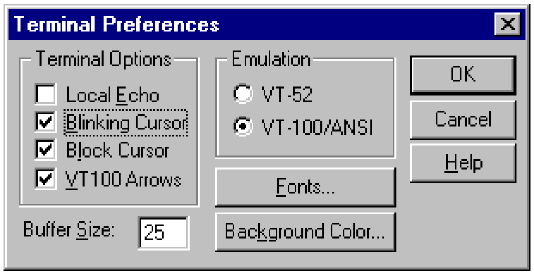

Terminal Preferences

Terminal Options
☐ Local Echo
☑ Blinking Cursor
☑ Block Cursor
☑ VT100 Arrows

Emulation
○ VT-52
● VT-100/ANSI

Buffer Size: 25

4. Under Terminal Options, select VT 100 Arrows.
5. Under Emulation, select 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, then click Remote System.

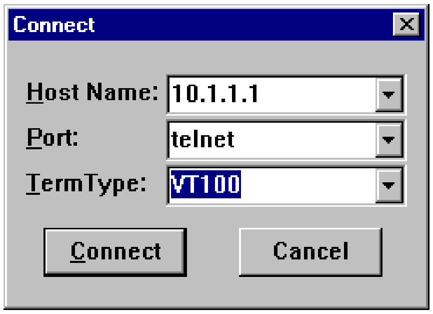

Host Name: 10.1.1.1

Port: telnet

TermType: VT100

10. Enter the appliance IP address in the Host Name field, then 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.

Additional Telnet Information

If the appliance has a monitor attached, 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 336, and ensure that you have selected the VT100 terminal type in Step 10 on page 336.

Setting Up an Appliance Through Telnet

Chapter 2, “Installing iChain Components,” on page 23 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 335.

   The starting address for Telnet should be 172.16.0.1, 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 that you set system passwords as part of the initialization process. For more information, see "Password Dialog Box" on page 277.

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:

```
System Console

>
```

2 At the System Console 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 are *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 System Console 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 Chapter 2, "Installing iChain Components," on page 23.

Troubleshooting Telnet

- "Telnet never starts" on page 337
- "Telnet starts after a long time" on page 337
- "Commands at the bottom of the screen look strange or don't make sense" on page 338

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.

19.7.3 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:

1. Connect a null-modem cable to the serial port on the workstation and the appliance.
2. Click Start, then run hypertrm.exe.

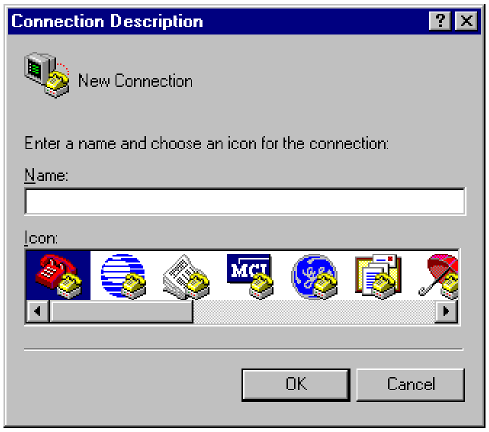

Connection Description

New Connection

Enter a name and choose an icon for the connection:

Name:

Icon:

3. Specify a name for the connection, select an icon as instructed, then click OK.

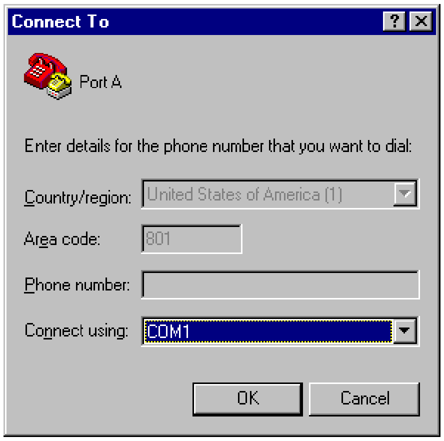

Connect To

Port A

Enter details for the phone number that you want to dial:

Country/region: United States of America (1)

Area code: 801

Phone number:

Connect using: COM1

4. Click the Connect Using drop-down list, select a Direct to Com option corresponding to the serial port connection on your workstation, then click OK.

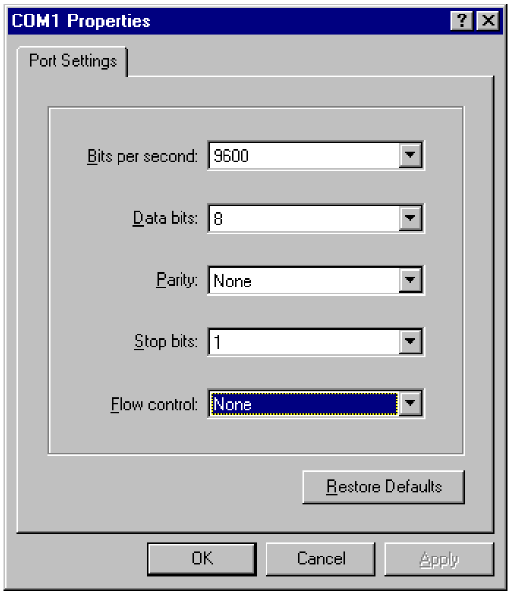

Bits per second: 9600

Data bits: 8

Parity: None

Stop bits: 1

Flow control: None

5 Set the properties according to the following table, 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:

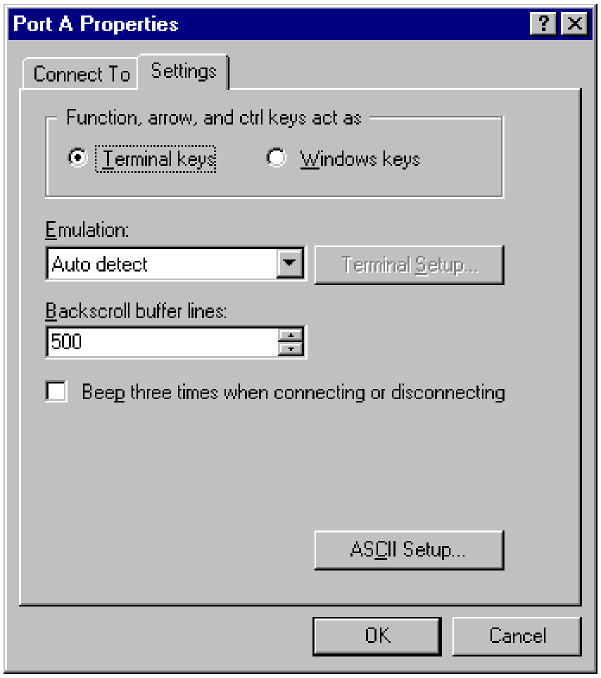

Port A Properties

Function, arrow, and ctrl keys act as

● Terminal keys ○ Windows keys

Emulation: Auto detect

Backscroll buffer lines: 500

□ Beep three times when connecting or disconnecting

7 Click Terminal Keys.

8 Click the Emulation drop-down list, then select ANSI.

9 (Optional) Specify cursor behavior by clicking Terminal Setup.

10 Click ASCII Setup, then ensure that only Wrap Lines That Exceed Terminal Width is selected.

11 Click OK twice.

12 Press Enter.

A command line prompt appears in Hyper Terminal.

You can now use all console commands to manage the appliance.

Additional Null-Modem 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 Left-arrow acts like a backspace or erase.
- The Right-arrow acts like a space.
- 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.

19.8 Performing Patch Upgrades

You can apply over-the-wire upgrades to the appliance as they become available. These over-the-wire upgrades provide a quick and easy method to install patches and fixes from Novell. For more information, see "Upgrade Page" on page 281.

19.8.1 Making Sure You Update the Clone Image before and after Upgrading

Both before and after an upgrade or support pack installation, you must update the appliance's clone image to avoid having the upgraded system overwritten by the older clone image. For more information, see Step 8 on page 342.

19.8.2 Preserving Configuration Settings

The system upgrade retains all appliance configuration settings. As a precaution, we recommend that you also export your appliance's configuration settings to a DOS-formatted floppy diskette prior to the upgrade.

After the upgrade is completed, if you need to import the configuration file, refer to the "Apply a named configuration file from a floppy" task in "Using Telnet or the Command Line" on page 348.

19.8.3 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 three ways:

1. First, the over-the-wire upgrade checks whether 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 checks whether the appliance is configured as a forward proxy with access through the firewall. If it is, the appliance tries the following two methods, in 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.

19.8.4 Downloading and Installing the Upgrade

1. In the browser-based management tool, click System > Actions > Update Clone > Update Clone.
2. When the update is complete, click the Upgrade page.
3. Select Enable Download, then type the URL for the upgrade.

   The URL is available from your appliance vendor.
4. Click the Download Time drop-down list, then select the time you want the download to occur.
5. Select Enable Install.
6. Click the Install Time drop-down list, then 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 that could make the proxy server unstable.

Using Logging Tools

20

Novell® iChain® includes a variety of logging tools that allow you to view information about access rules, news servers, mail servers, and Web server activity to help you manage your infrastructure.

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

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

The following topics are covered:

- Section 20.1, "Using Strong Cryptography," on page 343
- Section 20.2, "Using Step-Up Cryptography," on page 345
- Section 20.3, "Automatic Configuration Mechanisms," on page 345
- Section 20.4, "DNS Name Resolution," on page 351
- Section 20.5, "Appliance Error Messages," on page 353
- Section 20.6, "FTP Services," on page 357
- Section 20.7, "Shutting Down and Restarting," on page 361
- Section 20.8, "Using the SOCKS Client Service," on page 362
- Section 20.9, "Time Synchronization," on page 362

20.1 Using Strong Cryptography

The strong cryptography settings allow the server to be configured to force strong encryption to be used in SSL sessions (as in https). Client mode (when the proxy server initiates the SSL session) and server mode (when the proxy server accepts an SSL session from another machine) can be configured separately. The default is to not force the use of strong cryptography in either mode.

The configuration can be done from the iChain Proxy Server system console using the following commands:

```
set authentication strongserverenable = (yes/no)
```

No: Clients can initiate an SSL session with the proxy server using weak or strong cryptography.

Yes: Clients must initiate an SSL session with the proxy server using strong cryptography, or the session fails.

```
set authentication strongclientenable = (yes/no)
```

No: The proxy server initiates an SSL session with another server using any cryptography that server supports (strong or weak).

Yes: The proxy server only initiates an SSL session with another server using strong cryptography; if unsupported by the other server, it will fail.

Applying these settings stores them in the ISO object and creates a nile.cfg file. This file is read by nile.nlm at startup, so the server must be restarted for these settings to take effect.

20.1.1 Cryptography Settings

- **Weak cryptography:** Encryption is done with key sizes less than 128 bits.
- **Strong cryptography:** Encryption is done with key sizes of 128 bits or larger.

20.1.2 Configuring Federal Information Processing Standards in iChain

This section discusses the Federal Information Processing Standards (FIPS) option, including how to turn this option on or off, and the cipher options that go with it.

Turning the FIPS Option On/Off

To turn on the FIPS option in iChain, add the following load command line before “load proxy” in the appstart.ncf:

Syntax: load nile {-|/}{F|f}

The following are examples of this syntax:

```
load nile -F


load nile -f
load nile /F
load nile /f
```

The original appstart.ncf:

```
. . . . .
load dbypass
load proxy
load caconfig
. . . . .
```

The updated appstart.ncf:

```
. . . . .
load dbypass
load nile /F
load proxy
load caconfig
. . . . .
```

To turn this option off, users can either delete the load nile /F line or remove the /F option.

After updating the appstart.ncf when turning the option on or off, you need to restart the iChain server for the update to become effective.

Cipher Options For FIPS

iChain supports the following cipher options for FIPS:

For the server side (from the viewpoint of the browser):

- SSL_RSA_WITH_DES_CBC_SHA (weak)
- SSL_RSA_WITH_3DES_EDE_CBC_SHA (strong)

For the client side (from the viewpoint of the Web server):

- SSL_RSA_WITH_DES_CBC_SHA (weak)
- SSL_RSA_WITH_3DES_EDE_CBC_SHA (strong)

For information on how to configure weak or strong cryptography, see Section 20.1, "Using Strong Cryptography," on page 343.

20.2 Using Step-Up Cryptography

Step-Up Cryptography is a variation of SSL that provides a way for weaker clients to detect the need for strong cryptography. This feature is referred to as Server Gated Cryptography (SGC) by Microsoft, and Step-Up Cryptography by Netscape*. iChain supports Netscape's Step-Up Cryptography. This feature is especially applicable for users running on Windows 98, Windows NT, users with older browsers (Internet Explorer 5.0, 5.5, and Netscape 4.7*x*), and machines that are used outside the United States.

Step-Up Cryptography depends on a server to have a special certificate that permits it to participate in strong cryptography with the client. This certificate must be issued by a trustworthy CA (currently only Verisign and Thawte), and must contain an extension that indicates it is step-up capable. This means that the clients that have step-up capability must contain the code for strong cryptography. Step-Up Cryptography is conducted by using SSL's handshake feature. After the client views the server's certificate and verifies the appropriateness of Step-Up Cryptography, it initiates a second handshake, and the messages for the second handshake are transmitted over the current protected session.

To use Step-Up Cryptography:

1. Obtain a Verisign Secure Site Pro certificate or a Thawte 128-bit SuperCert certificate by going to the applicable Web site.
2. Select Novell as the vendor. (If Novell is not available, select Silverstream. Do not select Microsoft.)
3. Follow the normal import process.

20.3 Automatic Configuration Mechanisms

To	See
Learn about appliance configuration files	"About Appliance Configuration Files" on page 346
Learn about the three methods of managing configuration files	"Managing Configuration Files" on page 347
Save appliance configurations	"Using Customized Configuration Files to Change the System Configuration" on page 347

To	See
Change the current appliance configuration	"Using Customized Configuration Files to Change the System Configuration" on page 347
Configure multiple appliances	"Creating Appliance Configuration Shortcuts" on page 348
Restore the original factory settings	"Restoring Factory Settings" on page 349
Change the clone image. The appliance uses this to restore the system if it senses the system has become unstable.	"Restoring the Appliance to the Clone Image" on page 349
Reimage the appliance	"Reimaging and Restoring the Appliance System" on page 350

20.3.1 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 sample 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=172.16.0.2,mask=255.255.255.0
set eth0 name=eth0
set eth0 speed=default
set eth0 duplex=default
clear eth0 address
add eth0 address=172.16.0.1,mask=255.255.255.0
set floppy poll=no
set floppy interval=120
set floppy saveonapply=no
. . .
apply
```

20.3.2 System-Generated Configuration Files

The iChain Proxy Server employs two configuration files: factory.nas and current.nas.

factory.nas

This file contains the appliance configuration as it came from the factory. This is a system file that is never modified.

current.nas

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 of the commands used to create the current appliance configuration. To view the file, click System, click Import/Export, then select Current under Configuration Files on Appliance, then click Download.

20.3.3 Using Customized Configuration Files to Change the System Configuration

In addition to using system-level configuration files, iChain Proxy Services lets you save the appliance's current configuration to arbitrarily named .NAS files and apply them to the system later.

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 347.

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 that you have the files if the clone image is ever applied or the appliance is ever reimaged.

For a summary where having customized configuration files is an advantage, see "Creating Appliance Configuration Shortcuts" on page 348.

20.3.4 Managing Configuration Files

You can manage appliance configuration files using the browser-based management tool, using Telnet or the command line interface, and using the appliance's FTP functionality. The following three sections briefly explain how to use each of these management options:

- "Using the Browser-Based Management Tool" on page 347
- "Using Telnet or the Command Line" on page 348
- "Using FTP" on page 348

Using the Browser-Based Management Tool

You can export and import configurations and manage the creation of the autoload configuration from the browser-based management tool.

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	Enter	Notes
Apply an autoload file from a floppy	`import floppy`	First verify that the disk is inserted into the appliance.
Export a named configuration file to the appliance's hard drive	`export` *`filename`*	*Filename* is the name of the configuration file without the .nas extension specified.
Export a named configuration file to a floppy	`export` *`filename`* `floppy`	The file is saved on the DOS-formatted floppy disk inserted into the appliance.
Apply a named configuration file from the appliance's hard drive	`import` *`filename`*	*Filename* is the name of the configuration file without the .nas extension.
Apply a named configuration file from a floppy	`import` *`filename`* `floppy`	The file is 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 359), use one of the following commands, where *filename* is the name of your configuration file:

Command	Description
`get` *`filename`*`.nas`	Downloads the specified configuration file to your FTP local directory on your client workstation
`put` *`filename`*`.nas`	Uploads the specified configuration file from the FTP local directory to the appliance
`put` *`filename`*`.nas,execute`	Uploads the configuration file and applies it to the proxy server

20.3.5 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.

20.3.6 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 for the appliance as described in the Initial Installation Guide or use a previously created .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 172.16.0.1 on subnet 72.16.0 with a subnet mask of 255.255.255.0.
- Other network adapters have no addresses bound.
- No caching, proxy cache, 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.

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 Chapter 2, "Installing iChain Components," on page 23 or use a previously created .nas file on a floppy diskette to restore the appliance's configuration settings.

20.3.7 Restoring the Appliance to the Clone Image

Each appliance stores a clone image that is initially 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 Chapter 2, "Installing iChain Components," on page 23, or, if you have saved the appliance configuration, you can use a .nas file to restore the configuration.

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.

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.

20.3.8 Reimaging and Restoring the Appliance System

The appliance comes with a CD that can be used to reimage the system. It reformats the hard disks and reinstalls the appliance system. After reimaging an appliance, you must either reinitialize it as described in the Chapter 2, "Installing iChain Components," on page 23 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 346.

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, then turn the appliance's power switch off and wait a few seconds before restarting.

2. If the appliance configuration has not been previously saved, insert a formatted, blank floppy disk into the appliance's floppy disk drive.

   If you have previously saved the appliance configuration, skip to Step 7.

3. If you have access to the appliance through the browser-based management tool, click System > Import/Export; otherwise, continue to Step 5.

   If there is an AUTOLOAD file on the floppy disk, skip to Step 6.

4. 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 6.

5. 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.

6. Remove the floppy disk from the appliance.

7. Insert the appliance system CD into the CD drive.

8. Turn the appliance's power switch off, wait a few seconds, then turn the appliance back on.

   The CD automatically launches and the appliance reinitializes.

9. After the initialization process starts, insert the configuration diskette with the AUTOLOAD.NAS file into the appliance.
10. After all disk activity ceases and the system prompt appears, remove the appliance system CD and the floppy disk.
11. 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.

20.4 DNS Name Resolution

As iChain Proxy Services 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 straightforward, the proxy server tries various permutations to 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.

20.4.1 How the Appliance Resolves DNS Names and Formulates Subsequent DNS Queries

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.

If the DNS name server can’t resolve the query, the appliance formulates subsequent DNS queries based on the following:

- The appliance’s domain name
- The appliance’s r_append.cfg file

For example, assume the following:

- The browser request URL is webserver.
- The appliance’s domain name is acme.com.
- The appliance’s r_append.cfg file has the following content:

```
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. The appliance formulates a second query by appending the appliance's domain name to the URL as follows:

   webserver.acme.com
2. If this query fails, the appliance appends the appliance's subdomain name to the URL as follows:

   webserver.com
3. If this query fails, the appliance appends each entry in the r_append.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 lines of the r_append.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 r_append.cfg file.

Modifying the R_APPEND.CFG File

1 Start an FTP client on a workstation with access to the appliance.

   For help, see "Starting an FTP Session with the Appliance" on page 359.

2 Point the FTP client to one of the appliance's IP addresses.

3 Enter the following command:

```
get /etc/proxy/appliance/config/user/r_append.cfg
```

   The file is transferred to the FTP client's default directory.

4 Referring to the example in "How the Appliance Resolves DNS Names and Formulates Subsequent DNS Queries" on page 351, modify the r_append.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 r_append.cfg file back in \etc\proxy\appliance\config\user on the appliance.
6. Restart the appliance.

20.5 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 on how to modify error message text and create support for additional languages.

The appliance provides error messages to browsers through a set of language-specific directories, each of which contains two 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 two files the appliance uses are:

- Message.cfg, which contains the text of all appliance error messages
- Pxyerr.htm, which is the HTML template file that applies a format to the applicable error text and other error information and is sent to the receiving browsers

The language-specific directories that contain these two 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 Configure > Mini Web.

20.5.1 Checking the Language Directories on Your Appliance

To see a list of the directories on your appliance:

1. In the browser-based management tool, click Configure > Mini Web, then open the drop-down list.

20.5.2 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, as described in the following topics:

- "Creating a New Language" on page 354
- "Customizing the Error Message Text of an Existing Language" on page 355

- "Customizing the Error Message Format of an Existing Language" on page 355
- "Certificate Error Handling" on page 355
- "Using Certificate Error Handling" on page 355
- "Customizing Error Messages" on page 356
- "Customizing the Error Page" on page 357
- "Localizing Error Messages" on page 357
- "Localizing Error Pages" on page 357

Creating a New Language

1 Start FTP and log in as explained in "Starting an FTP Session with the Appliance" on page 359.

2 Enter the following:

```
get /etc/proxy/data/errpage/nls/english/pxyerr.htm

get /etc/proxy/data/errpage/nls/english/message.cfg
```

3 Modify the message.cfg file.

There are explicit instructions in the file that clearly indicate which parts can be translated.

The http status messages contain a number at the beginning of the Translated Message. You should not modify the number followed by a space at the beginning of the message. Anything after the space can be modified.

You cannot delete or add error messages, nor can you change the number or order of the messages.

IMPORTANT: When customizing the messages, do not use a semi-colon (;) in the message. If you do, the message truncates at the semi-colon and only the portion of the message before the semi-colon is displayed.

4 Modify the pxyerr.htm file.

Because this is an HTML file, you can customize it 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 Save the file when modifications are complete.

6 Using FTP, create a new language directory in the path given in Section 20.5, "Appliance Error Messages," on page 353.

7 Enter the following, where *language* is the new language directory you created:

```
put pxyerr.htm /etc/proxy/data/errpage/nls/language/message.cfg

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 354 for details, complete the following basic steps:

1. Open the message.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 354 for details.

3. Save the file when modifications are completed.

Customizing the Error Message Format of an Existing Language

Refer to the procedure in "Creating a New Language" on page 354 for details, then complete the following basic steps:

1. Open the pxyerr.htm 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 354 for details.

3. Save the file when modifications are completed.

Certificate Error Handling

Currently if accelerators have mutual authentication (this can include mutual and other authentication) enabled, when users present bad certificates (expired or revoked) to access these accelerators, the browsers display a "page not found" error. The certificate error handling feature enables administrators to configure the error messages so that they define what the problem is with a particular certificate. For example, if a certificate is expired, the administrator can configure an error message to let the user know that the certificate has expired. This feature helps administrators more effectively troubleshoot certificate problems. It also helps the user better understand why certificates might not be working.

NOTE: If users use mutual authentication or other authentication, and they cancel a certificate or have a bad certificate, the authentication fails at mutual authentication and they are not prompted for other authentication. If you turn on the certificate error handling feature, the system prompts you for other authentication. There is no error page for failure of mutual authentication.

Using Certificate Error Handling

To use the certificate error handling feature, you must go to the Authentication page and select the Send an Error Page When a Mutual-SSL Certificate Error Occurs option, then select the language from the drop-down list (see Figure 20-1).

NOTE: The certificate error handling feature applies to all accelerators. After this feature is enabled or disabled, you must restart the iChain server. Also, if you change the error messages or error pages, the iChain server must be restarted.

Figure 20-1 *Certificate Error Handling*

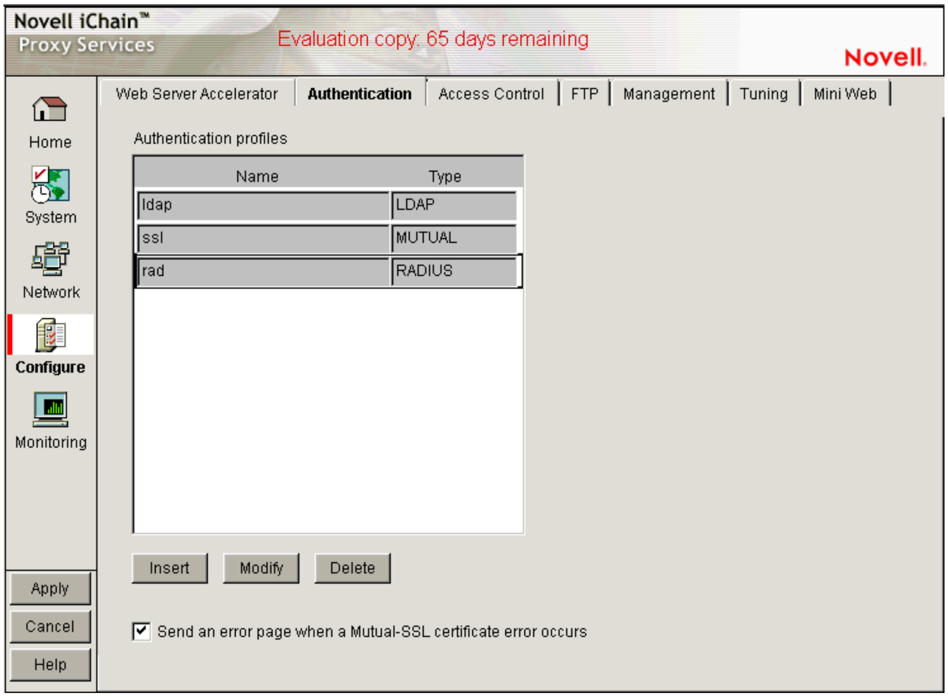

Novell iChain™ Proxy Services

Evaluation copy: 65 days remaining

Novell.

Web Server Accelerator | **Authentication** | Access Control | FTP | Management | Tuning | Mini Web

Home
System
Network
Configure
Monitoring

Authentication profiles

Name	Type
ldap	LDAP
ssl	MUTUAL
rad	RADIUS

Insert | Modify | Delete

Apply
Cancel
Help

☑ Send an error page when a Mutual-SSL certificate error occurs

Customizing Error Messages

There are two error message files, message.cfg (message file) and crfterrpg.htm (error page). These files are located at sys:\etc\proxy\data\errpage\nls\english.

The message.cfg file is for error status and description in the cererrpg.htm.

There are 60 certificate messages (101 to 160) in the message.cfg file. You can change the content after Translated Message=. For messages 109 and 113, uses must keep %d for error code.

The messages are paired for status and description fields in the error page. For example, message 111 (status) and message 112 (description) are shown here:

```
[Message 11]
Message ID=LOC_MSG_CERT_NOT_YET_STATUS
English Message=Certificate Not Valid Yet
Translated Message=Certificate Not Valid Yet

[Message 12]
Message ID=LOC_MSG_CERT_NOT_YET_DESC
English Message=Your certificate's start date is in the future. Please try later or
contact your system administrator.
Translated Message=Your certificate's start date is in the future. Please try later
or contact your system administrator.
```

When customized, message 111 (status) and 112 (description) might look like this:

```
[Message 11]
Message ID=LOC_MSG_CERT_NOT_YET_STATUS
English Message=Certificate Not Valid Yet
Translated Message=Certificate Not For Now

[Message 12]
Message ID=LOC_MSG_CERT_NOT_YET_DESC
English Message=Your certificate's start date is in the future. Please try later or
contact your system administrator.
Translated Message=Your certificate is not for now. Please try later.
```

NOTE: Only the content following Translated Message= is changed.

Customizing the Error Page

To customize an error page, change the static messages in the crterrpg.htm. Users can redesign their own pages as long as they use the crterrpg.htm as the file name and they have <ERROR_STATUS> and <ERROR_DESCRIPTION> fields in the HTML page.

Localizing Error Messages

The current default language is English; however, you can translate error messages and the error page into other languages.

To localize error messages:

1. After modifying the MESSAGE.CFG file, make sure you save it in UTF-8 format.
2. Translate the content of the Translated Message for every message (messages 101 to 160). The *_STATUS and *_DESC messages in the message.cfg file are used to replace the <ERROR_STATUS> and <ERROR_DESCPTION> fields in the crterrpg.htm file.

Localizing Error Pages

To localize an error page, translate the static messages in the CRTERRPG.HTM. Users can redesign their own pages as long as they use the CRTERRPG.HTM as the file name and they have <ERROR_STATUS> and <ERROR_DESCRIPTION> fields in the HTML page.

20.6 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).

The following sections contain important information about using iChain Proxy Services FTP services:

- Section 20.6.1, “Accessing Active and Passive FTP,” on page 358
- “Setting Up Appliance FTP Services” on page 358
- “Starting an FTP Session with the Appliance” on page 359
- “Changing the FTP Working Directory” on page 359
- “Managing Configuration Files with FTP” on page 360

- “Using FTP to Download a Configuration File to Your Workstation” on page 360
- “Using FTP to Move a Configuration File to the appliance from a Workstation” on page 360
- “Using FTP to Move a Configuration File to the Appliance and Execute It” on page 360
- “Using FTP to Customize the Appliance Splash Screen” on page 360

20.6.1 Accessing Active and 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: 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 cannot usually access it through a DOS window on a client outside the firewall.

FTP access through a browser outside the firewall: You cannot usually access an appliance inside a firewill through a browser that is outside the firewall.

FTP access through a browser inside the firewall: Because Netscape browsers use passive FTP, you can usually 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.

1. In the browser, click Tools > Internet Options > Advanced.
2. Under Browsing, select Use Web-Based FTP.

   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.

Appliance routing and transparent proxy limitations on FTP: When using the built-in appliance router capabilities in a transparent proxy situation, users cannot be able to perform browser-based FTP using ftp:// as the protocol. Browser-based FTP works normally with forward proxy.

20.6.2 Setting Up Appliance FTP Services

Before using FTP services to manage the appliance, ensure that the appliance's Mini FTP Server is properly configured.

1. Start the browser-based management tool > click Configure > FTP.
2. Ensure that at least one of the IP addresses in the Server IP Addresses list is selected.

   You use the selected address for your FTP session. IP address 172.16.0.1 is selected by default.

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 and currently supports the following commands:

```
CDDELELS and DIRPWDQUITSYSTTYPEUSER
```

The Mini FTP Server has no support for wildcard characters.

The Mini FTP Server does not work with directory or file names that contain spaces.

Mini FTP server access is limited to the following directories and their subdirectories:

- sys:etc\proxy\mftp
- sys:etc\proxy\appliance
- sys:etc\proxy\appliance
- sys:etc\proxy\data
- sys:etc\appliance
- sys:etc\applianc
- 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.

20.6.3 Starting an FTP Session with the Appliance

1. Launch your FTP application and enter a valid appliance IP address, for example:

   ```
   ftp 172.16.0.1
   ```

2. Log in to the appliance using the Config username and the password you have set.

20.6.4 Changing the FTP Working Directory

The default working volume and directory for FTP sessions is sys:\etc\proxy\appliance\config\user. This means that 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 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.

20.6.5 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 quickly apply different scenarios, 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 on a floppy disk, it is updated with each change. From the browser-based management tool or the Telnet/command line interface, you can export other configuration files.

20.6.6 Using FTP to Download a Configuration File to Your Workstation

1 Start FTP and log in as explained in "Starting an FTP Session with the Appliance" on page 359.

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.

20.6.7 Using FTP to Move 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 359.

2 Enter the following, where *filename* is the name of your configuration file:

```
put filename.nas
```

The file is transferred to the appliance.

20.6.8 Using FTP to Move a Configuration File to the Appliance and Execute It

1 Start FTP and log in as explained in "Starting an FTP Session with the Appliance" on page 359.

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
```

20.6.9 Using FTP to Customize the Appliance Splash Screen

The appliance has an optional splash screen that displays before pages are vended.

The splash screen has the root filename bmsplash. A three-letter extension indicates whether the splash screen is disabled (.off) or enabled (.htm). The screen is disabled by default.

To alter the appearance of the splash screen:

1 Start FTP and log in as explained in "Starting an FTP Session with the Appliance" on page 359.

2. To download the splash screen to you can modify it, do one of the following:
   - If the splash screen is enabled, enter:

     ```
     get /etc/proxy/data/bmsplash.htm
     ```
   - If the splash screen is not enabled, enter:

     ```
     get /etc/proxy/data/bmsplash.off
     ```
3. Modify the file.
4. To upload the splash screen after you have modified it, do one of the following:
   - If the splash screen is enabled, enter:

     ```
     put bmsplash.htm /etc/proxy/data/bmsplash.htm
     ```
   - If the splash screen is not enabled, enter:

     ```
     put bmsplash.off /etc/proxy/data/bmsplash.off
     ```

20.7 Shutting Down and Restarting

If you need to shut down or restart an appliance, you should do it properly to protect the data in memory and ensure that it is written to disk. There are several ways to properly shut down or restart the appliance.

20.7.1 Restarting from the Browser-Based Management Tool

1. Start the browser-based management tool.
2. Click System > Actions.
3. Shut down iChain Proxy Services 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 down iChain Proxy Services, you hear a three-beep sequence that repeats until the appliance is turned off or restarted.

If you choose to restart iChain Proxy Services, 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.

20.7.2 Shutting Down and Restarting from the Command Line

You can shut down or restart and appliance from the command line.

NOTE: Both actions break the connection. If you *restart* the appliance from a remote connection, you can reconnect after the appliance restarts. However, If you *shut down* the appliance, someone needs physical access to the appliance to restart it.

1. 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.

2 To shut down the appliance from the command line, enter

```
Shutdown
```

If you are near the appliance, you will hear a continuous three-beep sequence after the system is down. You can restart the system from an attached keyboard by pressing Ctrl+Alt+Delete, or you can shut off the power after you hear the three-beep sequence.

20.8 Using the SOCKS Client Service

For sites that have deployed SOCKS firewalls, the appliance provides a SOCKS Client service to redirect all forward proxy traffic through the firewall. Redirecting appliance traffic to the SOCKS firewall might significantly reduce appliance performance.

The appliance currently supports both SOCKS4 and SOCKS5 protocols.

20.9 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 "Using the Browser-Based Management Tool to Set the Time" on page 362.

The following additional flexibility in setting system time, including changing the GMT offset and daylight saving parameters, is available through the command line interface:

- "Synchronizing Time" on page 362
- "Using the Browser-Based Management Tool to Set the Time" on page 362
- "Using the Command Line to Set NTP Time Sources" on page 363

20.9.1 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 pre-configured with two servers:

132.163.4.101
132.163.4.103

You can add additional servers or delete them using the browser-based management tool and the command line interface. For more information regarding NTP functionality, see "Using the Command Line to Set NTP Time Sources" on page 363.

20.9.2 Using the Browser-Based Management Tool to Set the Time

To add or delete an NTP Server:

1 Start the browser-based management tool > click System > Date/Time.

2. Select Use Network Time Protocol.
3. Do one of the following:
   - To add a server, click Insert, then type the URL or IP address of the server.
   - To delete a server, select the server, then click Delete.

   Changes in the Date/Time page are immediately effective.

To set the time manually:

1. Start the browser-based management tool > click System > Date/Time.
2. Select Set Time Manually.
3. Click Set Time.
4. Using the drop-down lists, select the correct time and date.
5. Click OK.

20.9.3 Using the Command Line to Set NTP Time Sources

1. Do one of the following:
   - Add an NTP server address by entering

     ```
     add ntp server=128.115.14.97
     ```
   - 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 and time setting than the appliance clock, the system slows 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 is later than the appliance clock, synchronization between the two is generally be immediate. However, in certain situations you might observe the appliance clock incrementing by 600-minute intervals. This is normal system behavior.

The fact that the Apply button changes from Wait back to Apply indicates only that the NTP configuration change has been made, not that the appliance clock is fully synchronized with the NTP server.

If the above features are problematic in your situation, you can manually set appliance time to the target time and then re-enable the NTP feature.S

Upgrading Your iChain System 21

This chapter provides suggestions and steps on how to upgrade your current Novell® iChain® installation to iChain 2.3 Proxy Server and Authorization Server software.

The following topics are included:

- Section 21.1, "Upgrading from iChain 2.0, 2.1, and 2.2," on page 365
- Section 21.2, "Upgrading from iChain 1.5," on page 368

21.1 Upgrading from iChain 2.0, 2.1, and 2.2

This section discusses upgrading from the iChain 2.0, 2.1, and 2.2 software versions. The following steps should be considered:

- 1. Prepare the Current iChain Platform
- 2. Back Up the Existing iChain Configuration
- 3. Upgrade eDirectory with the iChain Schema Using the Install CD

  NOTE: This step is only required if upgrading from iChain 2.0.

- 4. Install ConsoleOne 1.34 and the iChain Snap-Ins
- 5. Convert and Modify Existing ACL/ISO Definitions

  NOTE: This step is only required if upgrading from iChain 2.0, and is done by default if the ConsoleOne® snap-ins detect that the iChain attributes are in an older format.

- 6. Upgrade the Proxy Server to iChain 2.3
- 7. Test the System
- 8. Implement New Features

This section also addresses:

- Schema Differences Between 2.0, 2.1, 2.2, and 2.3

21.1.1 1. Prepare the Current iChain Platform

1 Prepare a test scenario with the customer (for each app, identify key profiles). Be aware of what each application requires as input (for example, simple authentication header, parameters passed using the command line).

2 Test the scenario on the running iChain 2.0/2.1/2.2 system and confirm that it is working.

21.1.2 2. Back Up the Existing iChain Configuration

You should back up both the Authorization Server and the iChain Proxy Server.

To Back Up the Authorization Server

1. Back up eDirectory™.
2. Do an export to LDIF of the iChain objects (Access Control List, iChain Service Object, Communities).
3. Back up any custom tools or modules that might have been running on the Authorization server.
4. Rename the ConsoleOne directory if ConsoleOne version 1.2x is installed the (iChain 2.3 Authorization Server CD ships with version 1.34). If this is not an option, rename the iChain snap-in and lib directories.

To Back Up the iChain Proxy Server

1. Do an export to a NAS file of the Proxy Server configuration and a screen shot of all configuration screens.
2. Export or back up certificates that are being used by the proxy server.
3. Back up the following files:
   - /etc/hosts — contains host mappings to IP addresses
   - /etc/proxy/data — contains custom login pages (ca*.html)
   - /ichain/oac/oac.properties — contains advanced OLAC configuration settings
   - /etc/proxy/r_append.cfg — if any DNS search types changed
   - /system/appstart.ncf and /system/tune.ncf
4. Copy any tools or modules that you might have used from the server (for example, lsearch.nlm for LDAP testing, netmon.nlm for taking traces).

NOTE: If you want to save your configuration so you can quickly revert back to it, the fastest way to do this is to pull SCSI drive 0 (if you have a multi-drive system) and replace it with another drive, such as the highest numbered drive from your SCSI sub-system, or better yet, use a spare you have on a shelf. Label the original “Drive 0," and set it aside. Then proceed with the install as normal.

Make sure you have everything needed to restore a valid iChain 2.0/2.1/2.2 Proxy Server image.

The 2.1/2.2/2.3 schema is compatible with 2.0, meaning that if you leave your 2.0 iChain Service Object (ISO) untouched, you could have one proxy server running 2.0 while a second one is being upgraded. (This could help in doing a seamless migration and an easy rollback.)

21.1.3 3. Upgrade eDirectory with the iChain Schema Using the Install CD

NOTE: This step is only required if upgrading from iChain 2.0. No schema changes exist between iChain 2.1, 2.2, and 2.3.

The install script generates many BURP errors during this phase. They can be ignored. These errors are generated because many of the modifications to the schema that the install script is trying to perform are already in place.

NOTE: If the tree you are upgrading also contains Novell BorderManager® schema extensions, you will need to manually re-link the brdsrvsOutgoingAcl attribute with the object class named brdsrvsACLRule. This is done easily in ConsoleOne schema manager, after applying the new schema and reloading ConsoleOne.

21.1.4 4. Install ConsoleOne 1.34 and the iChain Snap-Ins

If it isn't already installed, install Console 1.34 and also install the iChain snap-ins from the Authorization Server CD. This is required for any RADIUS or token-based authentication setup.

21.1.5 5. Convert and Modify Existing ACL/ISO Definitions

NOTE: This step is only required if upgrading from iChain 2.0.

Convert and modify existing Access Control List (ACL)/iChain Service Object (ISO) definitions to match specifications in iChain 2.1, 2.2, and 2.3.

The ConsoleOne snap-ins that ship with iChain 2.1 and 2.2 can detect iChain 2.0-formatted objects. After upgrading the Authorization Server from 2.0 to 2.2 and selecting properties of the original 2.0 ISO with the 2.2 snap-ins, the ISO is automatically extended with the new required attributes.

NOTE: If administrators are creating completely new objects, the following should be considered:

1. The ISO has many new attributes in 2.0. The most important of these involves ACLCHECK dynamic LDAP search attributes.

2. If you decide to re-create the ISO, the corresponding Rule Objects referencing the old ISO's protected resources must be re-created. If this is not done, ACLCHECK reports “old version" errors.

21.1.6 6. Upgrade the Proxy Server to iChain 2.3

1 Image the proxy server with iChain 2.3.

   WARNING: When installing from a CD, both the original drive and the clone drive are overwritten. You cannot restore from the clone in this case, unless you first remove the clone drive from the system before installation.

2 Unlock the Proxy Server system console by entering `unlock` at the prompt. You do not need to specify a password.

3 Import the NAS file by placing the floppy containing the current.nas file into the proxy server. Enter `import floppy`. (If autoload does not exist, enter `import current floppy`.

   Wait until the system displays “completed execution of current” at the server console.

4 Import the server certificates that were backed up from the 2.0/2.1/2.2 server.

   If problems exist accessing the proxy server GUI, do the following from the Internet Caching System console:

   4a Run the _kill application to kill the java ServerApplication thread and all support modules.

   4b Unload the cert.nlm file at the system console.

4c Reload cert.nlm.

4d Execute appstart.ncf at the system console.

5 Restore the files backed up in "2. Back Up the Existing iChain Configuration" on page 365.

Do not copy appstart.ncf and tune.ncf from your old 2.0 or 2.1 server. Make a note of the changes, and edit the appstart.ncf and tune.ncf on your new 2.3 iChain server.

Some default settings have been changed and we recommend that you do not overwrite the existing 2.3 files.

NOTE: The oac.properties file is not needed unless some non-default parameters were required for functionality in 1.5 (for example, increasing worker threads, synchronization interval).

6 Using the proxy server GUI, run the health check to make sure that all services are up and running.

7 Verify if the eDirectory server still has community objects (which shipped with 1.5, but not with 2.*x*) and rules based on community objects. If this is the case, modify the APPSTART.NCF to load ACLCHECK with the /M option.

8 Verify that you can access the iChain protected resource from the browser.

21.1.7 7. Test the System

1 Complete an offline test using your defined scenario.

2 Complete a production test.

21.1.8 8. Implement New Features

Only after you have confirmed that the old features are working should you enable any of the new iChain 2.3 features.

21.1.9 Schema Differences Between 2.0, 2.1, 2.2, and 2.3

The iChain 2.3 schema file is found on the Authorization Server CD in the \schema subdirectory. This file documents all iChain attributes and lists the new attributes that have been added.

21.2 Upgrading from iChain 1.5

Upgrading from iChain 1.5 to iChain 2.3 requires the same steps as described in Section 21.1, "Upgrading from iChain 2.0, 2.1, and 2.2," on page 365, with a primary difference from what is described in "5. Convert and Modify Existing ACL/ISO Definitions" on page 367. You should replace this section with the instructions below:

21.2.1 5. Converting and Modifying Existing ACL/ISO Definitions When Upgrading from iChain 1.5

The ConsoleOne snap-ins that ship with iChain 2.3 can detect iChain 2.0-formatted objects. After upgrading the Authorization Server from 1.5 to 2.3 and selecting properties of the original 1.5 ISO with the new 2.3 snap-ins, you see the following message, asking whether the 1.5 ISOs should be upgraded:

Figure 21-1 *iChain 1.5 Protected Resources Found*

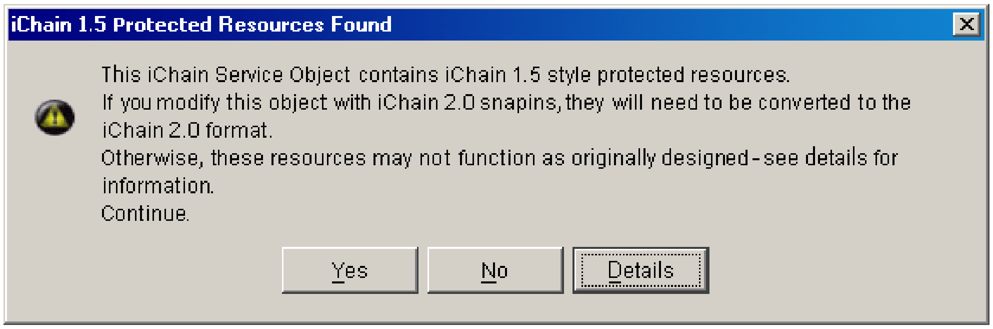

iChain 1.5 Protected Resources Found

This iChain Service Object contains iChain 1.5 style protected resources.
If you modify this object with iChain 2.0 snapins, they will need to be converted to the iChain 2.0 format.
Otherwise, these resources may not function as originally designed - see details for information.
Continue.

If you select Yes, the ISO attributes are converted to the 2.3 format, which means the ISO access mode defaults to Secure (requiring users to authenticate before authorization to access the protected resource at the origin server is granted). If needed, modify the resources when the conversion is completed. The modify dialog box is shown here:

Figure 21-2 *Modify Existing Protected Resource*

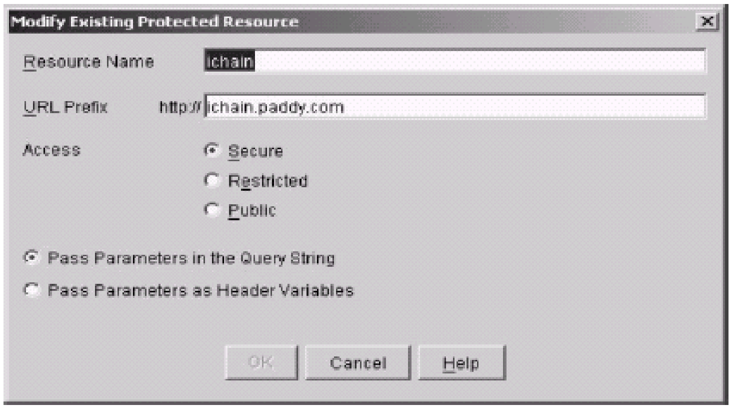

Modify Existing Protected Resource

Resource Name: ichain

URL Prefix: http:// ichain.paddy.com

Access: ● Secure ○ Restricted ○ Public

● Pass Parameters in the Query String
○ Pass Parameters as Header Variables

If you select No at the iChain 1.5 Protected Resources Found dialog box, the object remains as an iChain 1.5 protected resource and the resources on this ISO continues to function as they did in iChain 1.5. For backward compatibility, you should make all changes to the protected resources with the iChain 1.5 snap-ins.

NOTE: If administrators are creating completely new objects, the following should be considered:

1. The ISO has many new attributes. The most important of these involves the protected resource mode (public, secured, or protected). The defined protected resource needs a /* at the end for accessing resources in all subdirectories.

2. If you decide to re-create the ISO, the corresponding Rule Objects referencing the old ISO's protected resources must be re-created. If this is not done, ACLCHECK reports "old version" errors.

Known Issues A

The following issues have been discovered in previous releases of iChain.

- Section A.1, "Installation and Upgrade Issues," on page 371
- Section A.2, "Known Issues With iChain 2.3 and NetWare 6.5," on page 373
- Section A.3, "Known Issues With iChain Form Fill," on page 374
- Section A.4, "Known Authentication Issues," on page 376
- Section A.5, "Known Proxy Administration Tool Issues," on page 379
- Section A.6, "Known Multi-Homing and Path-Based Issues," on page 380
- Section A.7, "Known Session Broker Issues," on page 381
- Section A.8, "Known Certificate Issues," on page 382
- Section A.9, "Known Rewriter Issues," on page 382
- Section A.10, "Miscellaneous Issues," on page 383

A.1 Installation and Upgrade Issues

- Section A.1.1, "Over-the-Wire Upgrade (OTWUG) Requirements," on page 371
- Section A.1.2, "OTWUG vs. CD Upgrade," on page 371
- Section A.1.3, "Upgrading From iChain 2.2 to iChain 2.3," on page 372
- Section A.1.4, "OTWUG Might Have to Use IP Addresses," on page 372
- Section A.1.5, "Additional RADIUS Configuration Required," on page 372
- Section A.1.6, "Importing a .NAS File, Changing to the Factory Settings Can Cause a Reboot," on page 372
- Section A.1.7, "Password Management Information Lost When Importing .NAS File," on page 372
- Section A.1.8, "The Servlet Directory Has Been Removed in iChain 2.3," on page 373
- Section A.1.9, "Accessing the SecretStore Client Utilities," on page 373

A.1.1 Over-the-Wire Upgrade (OTWUG) Requirements

You must have iChain 2.3 SP1 or later to complete the OTWUG.

A.1.2 OTWUG vs. CD Upgrade

The CD upgrade provides large DOS and Sys volume sizes. The advantage of the CD upgrade are the ability to save core dumps out to the DOS volume for debugging.

A.1.3 Upgrading From iChain 2.2 to iChain 2.3

When upgrading from iChain 2.2 to iChain 2.3, you are prompted with the EULA. This is a one-way process and you must confirm it in order to continue. If you are upgrading from iChain 2.2, enable telneton.nas from the GUI. It is disabled by default.

IMPORTANT: You cannot upgrade from iChain 2.2 to iChain 2.3 SP3 using the OTWUG. You must upgrade to iChain 2.3 SP1 or SP2 first.

A.1.4 OTWUG Might Have to Use IP Addresses

If a DNS services is not available, the OTWUG must use IP addresses.

A.1.5 Additional RADIUS Configuration Required

With the added functionality in iChain 2.3 of being able to combine RADIUS authentication with LDAP authentication, additional configuration is needed to map the RADIUS user common name to a distinguished name in the LDAP authentication tree.

The recommended way to do this is to create an LDAP authentication profile named ldaprad to be used to find the distinguished name of the user in the LDAP authentication tree. See the Novell *iChain 2.3 Administration Guide* online for details.

To do this using the aclcheck profile, as done in previous versions of iChain, the ldap logintype of the aclcheck profile needs to be modified, as well as the ldap searchbase and ldap bindanonymous settings.

Set the ldap logintype to FieldName using the following command on the iChain server command line interface:

```
set authentication aclcheck ldap logintype = FieldName
```

A.1.6 Importing a .NAS File, Changing to the Factory Settings Can Cause a Reboot

The default settings for SNMP have changed to:

```
set snmp monitor=no
set snmp name=iChain
```

These settings changes can cause a server reboot if you import a .nas file from a previous version or change to the factory settings.

A.1.7 Password Management Information Lost When Importing .NAS File

If you import an iChain 2.2 .nas file that contains password management servlet information to an iChain 2.3 server, the password information might be lost because the SNMP files change. If this occurs, you need to reboot. Importing iChain 2.3 .nas files with the same information works properly.

A.1.8 The Servlet Directory Has Been Removed in iChain 2.3

The servlet directory is no longer available in iChain 2.3. This includes the java, class, and jar files related to the servlets that were previously available on the iChain Authorization CD. To access these servlets, see the Novell Cool Solutions Web site (http://www.novell.com/coolsolutions/icmag/).

A.1.9 Accessing the SecretStore Client Utilities

For the iChain 2.3 release, the SecretStore client utilities have been removed from the authorization server CD. To get the latest version of these utilities, go to the Novell NDK Web site (http://developer.novell.com/ndk/downloadaz.htm).

A.2 Known Issues With iChain 2.3 and NetWare 6.5

- Section A.2.1, "Install the Latest Support Pack for NetWare 6.5 Before Testing iChain in Your Environment," on page 373
- Section A.2.2, "Using the iFolder Client Through iChain," on page 373
- Section A.2.3, "The iPrint Client for Novell Open Enterprise Server No Longer Supports iChain," on page 374
- Section A.2.4, "Accessing iFolder Through an Accelerator Requiring Authentication," on page 374
- Section A.2.5, "Leaving the Challenge/Response Blank in RADIUS Causes Error," on page 374
- Section A.2.6, "Issue With RADIUS Loading On a NetWare 6.5 Server," on page 374

A.2.1 Install the Latest Support Pack for NetWare 6.5 Before Testing iChain in Your Environment

Because of potential compatibility issues, we recommend that you install NetWare 6.5 with the latest support pack before testing iChain in your environment.

A.2.2 Using the iFolder Client Through iChain

The latest Novell iFolder client and server components are required for use through iChain. The iFolder client only works through an accelerator using an LDAP profile with basic authentication enabled. iFolder can be accessed from a browser through iChain only by using NetStorage. The iFolder applet for browser access to iFolder is not supported through iChain.

Also, if you are using the iFolder client through iChain, you must use the Use host name sent by browser option. The Alternate host name option should not be used.

Do not use path-based multi-homing when configuring an accelerator for iFolder client access.

A.2.3 The iPrint Client for Novell Open Enterprise Server No Longer Supports iChain

For information about how to accelerate NetWare 6.5 iPrint with iChain 2.3, see TID 10092295 (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10092295.htm).

A.2.4 Accessing iFolder Through an Accelerator Requiring Authentication

When you access iFolder through an accelerator requiring authentication, if the authentication fails for any reason, the iFolder client login dialog box cannot be closed with the Cancel or close buttons. Instead, it re-prompts for login.

A workaround for closing the login dialog box is to choose Exit from the right-click menu of the iFolder tool-tray icon.

A.2.5 Leaving the Challenge/Response Blank in RADIUS Causes Error

If you leave the Response field blank in the Challenge/Response, a 400 Bad Request error displays.

A.2.6 Issue With RADIUS Loading On a NetWare 6.5 Server

Because of an issue with the ConsoleOne snap-in installation or an issue in ConsoleOne on the NetWare server, if the eDirectory RADIUS objects (the DAS object, etc.), are created from ConsoleOne running on the server, RADIUS will not load. This issue is under investigation.

To work around this issue, install ConsoleOne and the NMAS/RADIUS snap-ins on a workstation and then create all RADIUS configuration objects from the workstation.

A.3 Known Issues With iChain Form Fill

- Section A.3.1, "New Version of the SecretStore Plug-in," on page 374
- Section A.3.2, "ConsoleOne Crashes When Trying to Pass Large or Numerous Form Fill Polices," on page 375
- Section A.3.3, "Form Fill Does Not Use Modified LDAP Options Without Being Restarted," on page 375
- Section A.3.4, "Form Fill Fails to Remember Attribute When Spaces Exist in the <Name> Field," on page 375
- Section A.3.5, "Using <JavaScript> to Keep Specified Functions Does Not Work," on page 376

A.3.1 New Version of the SecretStore Plug-in

A new version of the SecretStore Plug-in for iManager is now available. To download this new version, see SecretStore 3.4 Plug-in for iManager2.5 & 2.6 (http://download.novell.com/SummaryFree.jsp?buildid=l3PWqDrLhiw~).

You should update both the iChain server and eDirectory servers (all platforms) with this new version.

A.3.2 ConsoleOne Crashes When Trying to Pass Large or Numerous Form Fill Polices

If you have too large or too many (over 16KB) Form Fill files, ConsoleOne displays an error when trying to submit the data.

To work around this, Form Fill supports a file stored on the local file system: <LocalPolicy>*filename*</LocalPolicy>. If the filename does not contain a "/" "\" or ":" symbol, the file goes into the following path:

sys:etc\proxy\appliance\config\user\formfill

Otherwise, the file is taken as an absolute path. You can use multiple <LocalPolicy> tags, but they are limited to 1 MB.

In a policy, you can use the <debugPost/> tag, which requires <post/> or <maskedPost/>. This tag lets SSO modify the HTML page with the changes need for SilentPost, but you can look at the source before you post the information. This also lets you debug without needing a sniffer.

Filterpoint 3 (determining Cutthrough) supports the following extension exclusion list:

```
"gif", "jpg", "jpeg", "pdf", "png", "zip", "jar", "bmp", "iso", "ico", "exe",
"dll", "doc", "mov", "mp3", "mpeg", "ppt", "rpm", "tar", "wav", "sxi", "xls","wmf",
"wpd", "sxw", "gz"
```

This speeds up SSO processing.

A.3.3 Form Fill Does Not Use Modified LDAP Options Without Being Restarted

After changing the LDAP options on the Access Control page (for example, port, server IP address, LDAP user), Form Fill does not communicate with the LDAP server using the new settings. To work around this problem:

1. Disable Form Fill and click Apply.
2. Enable Form Fill and click Apply.

A.3.4 Form Fill Fails to Remember Attribute When Spaces Exist in the <Name> Field

If the <name> attribute for the policy has a value with a space in it (i.e. <name>APACHE1 Login</name>), any <input> that has a value="~" does not get properly stored in the directory after a user enters it and submits the form. Attributes being called from LDAP (i.e. value="~cn") still work and get filled in the form.

A.3.5 Using <JavaScript> to Keep Specified Functions Does Not Work

When you create a customized <JavaScript> Form Fill policy that allows only certain functions, the policy keeps all <JavaScript> functions instead of only the ones you specify.

A.4 Known Authentication Issues

- Section A.4.1, "iChain Single Sign-on is Not Compatible With iManager 2.5," on page 376
- Section A.4.2, "Users Prompted to Authenticate Twice When Accessing Password Management Servlet," on page 376
- Section A.4.3, "iChain Prompts for SSL Mutual Authentication Even After Disabling Authentication on an Accelerator," on page 377
- Section A.4.4, "Time Restriction, Intruder Lockout, and Login Disabled are not Checked During Radius Token Authentication," on page 377
- Section A.4.5, "Certificate Revocation Checking (revocationcheckmethod =)," on page 377
- Section A.4.6, "OSCP 'Unknown' Response Lets Certificate Authenticate User," on page 377
- Section A.4.7, "Certificate Authentication Problems When the CRL Is Invalid," on page 378
- Section A.4.8, "Session Times Out During the 0 TTL State," on page 378
- Section A.4.9, "NMAS RADIUS Update Required for iChain RADIUS Token Authentication," on page 378
- Section A.4.10, "Use the Latest NetIdentity Client With iChain 2.3," on page 378
- Section A.4.11, "Using an iChain Accelerator for WebDAV Connections to NetStorage," on page 378
- Section A.4.12, "Do Not Use NetIdentity with LDAP/RADIUS ANDing," on page 379
- Section A.4.13, "iChain is Unable to And LDAP with XTier Based Profiles," on page 379
- Section A.4.14, "OLAC Authentication Profile Parameters Still Passed When OLAC Is Disabled," on page 379

A.4.1 iChain Single Sign-on is Not Compatible With iManager 2.5

iChain's single sign-on functionality (including forward authentication, OLAC, and Form Fill) is not compatible with iManager 2.5. To log in, iManager 2.5 requires a username, password, and treename. This prevents forward authentication and OLAC from working. With Form Fill, the Exit button in iManager directs you back to the initial login form.

A.4.2 Users Prompted to Authenticate Twice When Accessing Password Management Servlet

When you set up an additional accelerator, your users might be required to authenticate again to the second accelerator. This occurs when the password management server is a separate accelerator in a CDA setup that requires authentication. And when the authentication profile names for the accelerator you are authenticating to and the password management accelerator have different names.

A.4.3 iChain Prompts for SSL Mutual Authentication Even After Disabling Authentication on an Accelerator

If you disable authentication on an accelerator, the user is still prompted to supply their user certificate. This continues until a purge cache is done. You can cancel the certificate until the cache can be purged.

A.4.4 Time Restriction, Intruder Lockout, and Login Disabled are not Checked During Radius Token Authentication

Time restriction, intruder lockout, and login disabled are checked only if you are using Novell's NMAS Radius server in the same tree as the iChain Authorization tree. Also, the LDAP does the checks if the Radius Token authentication is ANDed with an LDAP authentication.

A.4.5 Certificate Revocation Checking (revocationcheckmethod =)

Certificate revocation is checked in the URL of the certificate. The certificate revocation method is set using the following setting:

revocationcheckmethod=*method*

The *method* is the type of certificate revocation checking performed during mutual authentication. The following Certificate Revocation checks are available depending on the value of the revocationcheckmethod parameter:

- OCSP only. This method checks only the online certificate status protocol (OCSP). The CRL is checked in the URL and overrides the client certificate. The certificate does not pass if it is revoked or there is a miscommunication.
- CRL only. This checks only the Certificate Revocation list.
- OCSP- CRL. This method checks the OCSP first, then checks the CRL. The CRL protocols include HTTP, LDAP, and X.500 directory name (provides only the directory name).

The Certificate Revocation Lists (CRL) is checked when the following conditions exist:

- The client certificate contains a CRL Distribution point (CDP).
- The client certificate contains a CRL CDP but not an Authority Information Access (AIA) extension for OSCP.
- The OCSP configured sources is disabled.

A.4.6 OSCP ‘Unknown' Response Lets Certificate Authenticate User

If you have your certificate revocation method set to OCSP-CRL, the certificate is allowed to authenticate the user if an unknown response occurs and the client certificate does not contain a CRL distribution point. If you do not want the certificate to authenticate the user under these circumstances, you need to set your revocation method to OCSP only.

A.4.7 Certificate Authentication Problems When the CRL Is Invalid

An invalid Certificate Revocation List (CRL) prevents mutual or certificate authentication from working properly. The CRL includes a dated time stamp indicating when the CRL is invalid. The Certificate Authority (CA) needs to update the CRL periodically with a new expiration date and time. If the CA does not update the CRL, perhaps because the CA is down or for any other reason, the CRL becomes invalid. During certificate or mutual authentication, the iChain Proxy Server compares the time stamp of the CRL with its own time and if the CRL time stamp has expired, then the authentication fails.

A.4.8 Session Times Out During the 0 TTL State

If you are using Mutual SSL Authentication, a certificate error might occur if the user attempts to access the site while their user ID is in the 0 TTL state. During this state, the user's session times out. However, there is a 60 second window where the user ID is still registered with the IAGENT database.

A.4.9 NMAS RADIUS Update Required for iChain RADIUS Token Authentication

If you are using NetWare 5 with Support Pack 6 or higher, or NetWare 6 with Support Pack 3 or higher, you need to update these versions in order for the NMAS RADIUS Server to do iChain RADIUS token authentication. See the Novell Technical Information Novell Technical Information document (http://support.novell.com/cgi-bin/search/searchtid.cgi?/2965335.htm) for the download and details.

A.4.10 Use the Latest NetIdentity Client With iChain 2.3

If you are using NetIdentity for iChain authentication, you must use NetIdentity client 1.2.1 or later. Also, NetIdentity-aware server components that released with NetWare 6.5 that are accessed through iChain, such as NetStorage, Virtual Office, and iManager, must be updated with NetWare 6.5 SP1 or later for full functionality.

A.4.11 Using an iChain Accelerator for WebDAV Connections to NetStorage

If you are using an iChain accelerator for WebDAV connections to NetStorage, the Allow authentication through HTTP authorization header and Use basic/proxy authentication options must be enabled in the Authentication Profile used by the accelerator. Users are required to provide login credentials both for iChain and for NetStorage authentications; however, if the workstation has the NetIdentity client installed and configured to trust the NetStorage CA, the user will not be prompted for the additional login to NetStorage.

The Authentication profile option, Allow authentication through NetIdentity, might give unexpected results if you enable it in addition to basic/proxy authentication. If you enable the Allow authentication through NetIdentity option, the workstations with the NetIdentity client installed are unable to make WebDAV connections through the accelerator. This problem occurs because the WebDAV OPTIONS header is a non-redirectable request. In this case, iChain returns a 409 Conflict error, resulting in a failed authentication and connection.

A.4.12 Do Not Use NetIdentity with LDAP/RADIUS ANDing

NetIdentity authentication does not work if you have configured an LDAP profile, created a RADIUS profile, and have ANDed the two profiles together.

A.4.13 iChain is Unable to And LDAP with XTier Based Profiles

If you attempt to And LDAP with XTier, you lose the auth profile that was added last.After you apply the settings to the server, only one of the profiles appears.

A.4.14 OLAC Authentication Profile Parameters Still Passed When OLAC Is Disabled

You can configure an accelerator and initially enable OLAC with authentication profile parameters using a query string. If you then disable OLAC, the authentication profile parameters are still passed in the query string for existing users who are logged in.

A.5 Known Proxy Administration Tool Issues

- Section A.5.1, “Command Line and Proxy Administration Tool Sluggishness During Startup,” on page 379
- Section A.5.2, “Daylight Saving Time Adjustment,” on page 380
- Section A.5.3, “Speed and Duplex Settings Might Not Be Synchronized Between the Proxy Administration Tool and INETCFG,” on page 380
- Section A.5.4, “Proxy Administration Tool Slowness With Sun JRE,” on page 380
- Section A.5.5, “The PIN List Might Not Display Properly with SUN JRE,” on page 380

A.5.1 Command Line and Proxy Administration Tool Sluggishness During Startup

Appliance with more than one disk drive execute mirroring and cloning processes when the system starts the first time. These one-time processes are required for system fault tolerance and must run to completion.

While the processes are running, the console and the Proxy Administration Tool might seem sluggish for a couple of minutes.

Cache performance is also somewhat affected by the processes

IMPORTANT: Do not restart the appliance. This causes the mirroring and cloning processes to restart and delays the arrival of normal system response times.

A.5.2 Daylight Saving Time Adjustment

The default settings for Adjust Clock for Daylight Saving Changes in the iChain Proxy Administration Tool and the actual iChain Proxy Server do not match. To get the iChain Proxy Server to match the Proxy Administration Tool:

1. In the Proxy Administration Tool, go to System > Timezone > then deselect Adjust Clock for Daylight Saving Changes.
2. Apply the changes.
3. Select the Adjust Clock for Daylight Saving Changes box again.

A.5.3 Speed and Duplex Settings Might Not Be Synchronized Between the Proxy Administration Tool and INETCFG

iChain 2.3 currently does not support additional settings set by Gigabit cards. Use the INETCFG command line utility to adjust speed and duplex settings on Gigabit cards. Do not use the Proxy Administration Tool to change these settings. Also, do not use INETCFG if you have an open session to the server with the Proxy Administration Tool.

A CLI command was added that lets you add additional load line parameters without needed to go through the INETCFG. This should allow the additional settings needed on the newer cards. The syntax is as follows:

```
set eth0 loadlineparameters=IOMAPMODE=1
```

A.5.4 Proxy Administration Tool Slowness With Sun JRE

If you are running iChain on a Windows XP machine with Sun JRE, you might experience a delay when selecting Configure > Authentication in the Proxy Administration Tool. While the page is loading, you cannot access the Proxy Administration Tool until all of the authentication profiles are added.

A possible workaround is to go to the control panel and double-click the Java Plug-in, then select the Browser tab and deselect Microsoft Internet Explorer so that the Sun JRE is not used.

If you need to use the Sun JRE, we recommend that you use the Sun JRE 1.5 SP4 version to avoid the slowness issue on Windows platforms. You can use the Microsoft* JVM, but you need to already have it downloaded because it is no longer available for download.

A.5.5 The PIN List Might Not Display Properly with SUN JRE

If you use a large PIN list, the entire list might not display. Currently, Java versions 1.5.0_02 and 1.4.2_08 can handle only approximately 762 entries.

A.6 Known Multi-Homing and Path-Based Issues

- Section A.6.1, "Disabling a Multi-Homing Master Accelerator," on page 381
- Section A.6.2, "Cannot Use Double-byte Path for Path Based Multi-Homing," on page 381

- Section A.6.3, “Multi-Homing Display in the iChain GUI,” on page 381
- Section A.6.4, “Path-Based Multi-Homing Might Not Function Properly with JavaScript,” on page 381

A.6.1 Disabling a Multi-Homing Master Accelerator

Disabling a multi-homing master accelerator also disables all of its children. The iChain Proxy Administration Tool, however, doesn't reflect that the children are disabled.

A.6.2 Cannot Use Double-byte Path for Path Based Multi-Homing

All data in the NAS file is not yet double-byte international ready.

A.6.3 Multi-Homing Display in the iChain GUI

The multi-homing displays correctly in the individual accelerator screen. It also displays correctly in the Monitoring->Services screen. It is also configured and works correctly. However, at times the multi-homing parent/child relationship is not displayed correctly in the Configure->Web Server Accelerator screen. This is a display problem only.

A.6.4 Path-Based Multi-Homing Might Not Function Properly with JavaScript

This is particularly important if you have the option Remove sub-path from URL enabled. JavaScript might obscure the URL data that iChain needs to access and modify to direct requests to backend servers. Because an absolute path reference can occur anywhere in JavaScript, the iChain word parse does not know how to assemble or parse through JavaScript.

A.7 Known Session Broker Issues

- Section A.7.1, “Configuration Via the ConsoleOne Wizard,” on page 381
- Section A.7.2, “Enabling Basic Authentication for LDAP and Disabling Secure Exchange Causes Session Broker to Fail,” on page 382

A.7.1 Configuration Via the ConsoleOne Wizard

When you configure the session broker IP address on the first page of the iChain Web Server Accelerator, the change is only recognized by the proxy server if one (or both) of two requirements are met:

1. A subsequent page of the wizard is accessed before clicking Finish,

   or

2. Apply is either clicked from the iChain Proxy Administration Tool or entered from the iChain console.

Until at least one of these requirements is met, the iChain Proxy server does not communicate with the session broker.

A.7.2 Enabling Basic Authentication for LDAP and Disabling Secure Exchange Causes Session Broker to Fail

Setting basic authentication for LDAP and turning secure exchange off causes SB to fail and prompt for the user to log in on both servers.

A.8 Known Certificate Issues

- Section A.8.1, "Creating a Certificate With an Unspecified Purpose Breaks XTier," on page 382
- Section A.8.2, "Cannot Use UTF8 Encoded Server Certificates with iChain 2.3," on page 382
- Section A.8.3, "Netscape Browsers and Certificate Database Passwords," on page 382

A.8.1 Creating a Certificate With an Unspecified Purpose Breaks XTier

If you create certificate in ConsoleOne without specifying a purpose, XTier does not function properly. You need to specify a purpose for the authentication to work.

A.8.2 Cannot Use UTF8 Encoded Server Certificates with iChain 2.3

For more information about this issue, see TID 10096804 (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10096804.htm).

A.8.3 Netscape Browsers and Certificate Database Passwords

An issue occurs with users setting up their Netscape browsers to prompt for the certificate database password each time they want to select a certificate. After they enter the password, the browser appears to hang and in some cases, it eventually times out. This is because of a defect in Netscape. If users enter the URL again without closing the browser, they are prompted to select their certificates again and re-enter their passwords. After the second time, they are given access.

A.9 Known Rewriter Issues

- Section A.9.1, "Changes to the Rewriter," on page 382
- Section A.9.2, "Rewriting Fails on a Page with Numerous HREFs," on page 383

A.9.1 Changes to the Rewriter

- "Viewing Changes to the Rewriter" on page 383
- "Default Mime Type Changes" on page 383
- "The Internal Rewriter Command Has Been Removed" on page 383

Viewing Changes to the Rewriter

For a listing of new features in the rewriter, see the rewriter.sam file located at sys:/etc/proxy/rewriter.sam.

Default Mime Type Changes

Text/plain is no longer a default mime type. If you want to have this mime type rewritten, you need to add it to the sys:etc\proxy\rewriter.cfg file in the [Mime Content-type] section.

The Internal Rewriter Command Has Been Removed

The [Internal Rewriter] command has been removed. In its place is a new command:

```
set accelerator <name> DisableRewriter=Yes|No
```

This setting is exported to the .nas file.

A.9.2 Rewriting Fails on a Page with Numerous HREFs

Although the rewriting failure occurs when downloading large amounts of data from the accelerated Web server, it's not the size or timeout of the page which is the issue. It is the number of links to be rewritten. There is a data size limit to the number of references that the rewriter can rewrite on a page.

The solution is to reduce the number of HREFs that need to be rewritten. If the problem is occurring because the rewriter is rewriting HTTP to HTTPS, you can solve this problem by disabling multi-homing for the Web server and by rewriting the Web page to use relative links. This reduces the number of links that need to be rewritten. For more information, see TID 10101106, iChain 2.3 not rewriting Web data correctly (http://support.novell.com/docs/Tids/Solutions/10101106.html).

A.10 Miscellaneous Issues

- Section A.10.1, “Sending the Via Header to the Web Server,” on page 384
- Section A.10.2, “Cannot Delete Email Messages in Outlook Web Access,” on page 384
- Section A.10.3, “License Agreement Does Not Appear in GUI When Logged In as the View User,” on page 385
- Section A.10.4, “Trustedroot Configured .nas File Fails on Import,” on page 385
- Section A.10.5, “Server Might Abend When Restart Server is Entered on the iChain Server,” on page 385
- Section A.10.6, “Creating Custom Login Pages without a UTF-8 Charset,” on page 385
- Section A.10.7, “Issue With Compressed Data Not Being Sent From Web Servers,” on page 385
- Section A.10.8, “Nsure Audit Logging Issue When Using iChain to Log Events,” on page 386
- Section A.10.9, “Cannot Add a Second Server to an Accelerator,” on page 386

A.10.1 Sending the Via Header to the Web Server

The [HTTP Headers] section of the proxy.cfg file is responsible for setting the Via host header. If the section is missing from the proxy.cfg file, a default Via host header is sent to the Web server, which consists of the host server name of the proxy (ICS_Server) and the proxy build number.

Use the following steps to edit this file, add the [HTTP Headers] section, and insert your values for the Via header:

1 At the command line, unlock the cole.

2 To enter debug mode, enter the following command.

```
debug
```

3 To stop the proxy, enter the following command:

```
_kill
```

4 Wait for the proxy to unload.

5 When the proxy has unloaded, edit the proxy.cfg file. Enter the following command:

```
edit sys:\etc\proxy\proxy.cfg
```

6 Add the [http headers] section with the Via host header values you need. Use the following fomat:

```
[HTTP Headers]
ViaHeaderHostName = <@name>
ViaHeaderBuildVersion = <version>
```

Replace <@name> with the host name, for example www.novell.com.

Replace <version> with the build version of the header, for example 3.0.223.

7 Save the changes.

8 Reboot.

A.10.2 Cannot Delete Email Messages in Outlook Web Access

You cannot delete email messages from the Microsoft* Outlook Web Access (OWA) Exchange server when secure exchange is enabled unless the Alternate hostname is the same as the Accelerator hostname. If these names are different, the MOVE method has a host header that matches the alternate hostname but the WebDAV destination header matches the accelerator hostname. This causes a 502 error when processing the request. If the Alternate hostname and the Accelerator hostname match, this error does not occur.

For more information about this issue, see TID 10091523 (http://support.novell.com/cgi-bin/search/searchtid.cgi?/10091523.htm).

A.10.3 License Agreement Does Not Appear in GUI When Logged In as the View User

The iChain GUI homepage shows that the server is not licensed when a user logs in using the View user account. If the user logs in using the Config user, the server shows that it is licensed. The issue is that when a user logs in with the View user account, the GUI does not make any LDAP queries to the ISO object, such as a query to obtain license information.

A.10.4 Trustedroot Configured .nas File Fails on Import

If you export a .nas file from a machine that has a trusted root configured on the Access Control tab, the .nas environment does not work when you restore the file. To prevent this issue, do not export the current.nas file unless you set the export trustedroot=yes and make it the default setting.

A.10.5 Server Might Abend When Restart Server is Entered on the iChain Server

If you enter the Restart Server command at the NetWare command line, the server might abend and display the following error message:

```
Error 'Abend on p00:  KiClearKernelSetJmp:  passed on setjmp structure in at head
of list.
OS version:  Novell NetWare 5.60.06.
Running Process:  Interrupt Services Routing (nested count 1)
interrupt process:  Console Command Process.
```

The correct way to shut down a proxy server is to use the shutdown or restart commands from the iChain command line interface or GUI.

This issue will not be resolved.

A.10.6 Creating Custom Login Pages without a UTF-8 Charset

When you create or modify a custom login page, you must save it in the UTF-8 format. You can use <FORM...ACCEPT-CHARSET= “UTF-8”...> as the form statement on the login page. This is important because the decoder for the iChain login pages now assumes that the post data coming back from the browser is UTF-8. The HTTP/HTML specifications have no other mechanism in place to ensure that this is the case. This is necessary to support any usernames not into the 7-bit ASCII set.

For more information about this issue, see TID 10099466 (http://support.novell.com/cgi-bin/search/searchtid.cgi?10099466.htm).

A.10.7 Issue With Compressed Data Not Being Sent From Web Servers

There are noted instances where a Web server fails to send compressed data to iChain, yet sends the compressed data to a browser. Some Web servers, including Microsoft Internet Information Services (IIS), do not respond with compressed (GZIP) data if a Via HTTP header is sent from a proxy to the Web server.

A.10.8 Nsure Audit Logging Issue When Using iChain to Log Events

If you have iChain configured to use Nsure Audit to log events, the events that correspond with mutual authentication using revoked certificates might not be logged. This is because when certificate error pages are enabled, nothing is logged. When certificate error pages are disabled, a log entry is created but it uses the information from a previous successful login and not the current data.

A.10.9 Cannot Add a Second Server to an Accelerator

If you add a second back-end Web server to an accelerator, the accelerator does not initialize. Instead, an error message appears. This might occur if you use a DNS name on the server instead of an IP address. If you try a second time to apply the configuration, the accelerator initializes properly and is functional.

Using LDAP Server Load Balancing and Failover

B

You can add multiple LDAP servers to an authentication profile and to an access control pool:

- **Authentication profile pool.** Each profile specifies the LDAP servers that can be contacted for authentication information. If you have LDAP servers that contain identical user information, you can provide failover for your users by adding them to the same authentication profile. Then, when one server goes down, the user can still be authenticated through one of the other servers in the pool. For information on adding LDAP servers to an authentication profile, see "LDAP Options Dialog Box" on page 315.
- **Access control pool.** The LDAP servers in this pool are used for ACL checking, single sign-on, and OLAC. If you have multiple LDAP servers that are in the same directory tree, you can provide failover for your users by adding them to the access control pool. Then, when one server goes down, the user can still be authorized through one of the other servers in the pool. For information on adding LDAP servers to the access control pool, see "Access Control Page" on page 319.

iChain® 2.3 SP4 IR3 has modified the code that performs such tasks as monitoring the health of each LDAP server and re-enabling servers that have come online. All pools now use the same algorithms for these tasks. For load balancing, they use a modified round robin algorithm. The next server in the list is used for the next request unless the request is an authentication request. When a user requests authentication, the initial request derives persistence from the username. This allows the user to use the same LDAP server for subsequent requests. This solves a problem with Form Fill when deleteRemembered is enabled. The user returns to the same server where the secrets have been deleted rather than being sent to the next LDAP server in the list, which might not have been synchronized with the LDAP server that deleted the outdated secrets.

The following screens allow you to monitor the health of your LDAP configuration:

- **iChain Console.** Two console screens display information about the LDAP servers. From the main list of screens, you can access the LDAP Pool Messages screen. From the Proxy Console, you can access the Proxy LDAP Pool Information screen to view the status of each configured LDAP server.
- **Services Page.** In the Web application, you can view the status of the LDAP servers. Click Monitoring > Services and scroll to the LDAP section. Each authentication profile is listed by name, for example [ldap], and displays the status of each LDAP server configured for the profile. The [aclcheck] section lists the status of the LDAP servers that have been added to the access control pool.

Documentation Updates

C

This section lists updates to the *Novell® iChain® 2.3 Administration* guide that have been made since the initial release of iChain. The information will help you to keep current on documentation updates and, in some cases, software updates (such as a Support Pack release).

The information is grouped according to the date when the *Novell iChain Administration* guide was republished. Within each dated section, the updates are listed by the names of the main table of contents sections.

The *Novell iChain Administration* guide has been updated on the following dates:

- Section C.1, "February 2009 (SP6)," on page 389
- Section C.2, "March 2007 (SP5)," on page 390
- Section C.3, "March 16, 2006 (SP4)," on page 391
- Section C.4, "October 24, 2005," on page 391
- Section C.5, "August 8, 2005 (SP3)," on page 391
- Section C.6, "June 24, 2005 (SP3)," on page 392
- Section C.7, "April 12, 2005 (SP2)," on page 392
- Section C.8, "March 3, 2005 (SP2)," on page 392
- Section C.9, "January 26, 2005 (SP2)," on page 393

C.1 February 2009 (SP6)

Location	Change
"Advanced DNS Options Dialog Box" on page 290	Added information about the Reset button.
Section 16.1.1, "Secure Access to Citrix Thin Clients," on page 204	Added information on how to use the Citrix Secure Gateway with iChain and the Citrix thin client.
Section A.9.2, "Rewriting Fails on a Page with Numerous HREFs," on page 383	Added information about a known rewriter issue when a page has numerous HREFs.
"Enabling the Custom Rewriter" on page 255	Fixed the error in the rewriter unload filter command.
Section A.10.1, "Sending the Via Header to the Web Server," on page 384	Fixed the problem with the parameter names for the [HTTP Header} section in the proxy.cfg file.
Section 18.2.2, "New Commands," on page 262	Corrected the certificate commands in the .nas file.
"Add Authentication Profiles Dialog Box" on page 305 and Section 13.2, "Using OLAC Custom Header Variables," on page 156	Added information to clarify that if you create OLAC parameters for ICHAIN_UID or ICHAIN_PWD, the Forward Authentication Information to Web Server option, the setting is ignored and a basic authentication header is always sent.
"Internal Rewriter Summary" on page 252	Added information that rewriter does not support nested double quotes.

Location	Change
Section 15.9, “Renewing a Third-Party Certificate,” on page 201	Fixed the link to the AppNote on how to renew a certificate issued by an external CA.
Section A.10.1, “Sending the Via Header to the Web Server,” on page 384	Explains how to add this section to the file so that the via host header can be sent.
“Pin Type” on page 234	Removed the information claiming that URL masks work differently with the Bypass option. They work the same.
“Setting Timeouts and Pool Limits for LDAP Profiles” on page 310 and “Diagnosing the Need for More Pool Handles” on page 311	Added information on how you determine whether your system needs more pool handles for LDAP profiles.

C.2 March 2007 (SP5)

A number of minor technical corrections were made to the *Novell iChain Administration* guide to clarify configuration issues. The following sections are new or were revised:

Location	Change
Section 5.2, “Disabling Mutual SSL,” on page 86	Added instructions on the need to purge cache to enable the new configuration.
Section 13.2, “Using OLAC Custom Header Variables,” on page 156	Added that an <olac-param name> cannot contain any extended characters.
“Rewriter Support” on page 243	The section has been revised to clarify the differences between the internal and custom rewriter. Add information on starting the custom rewriter on reboot (see “Configuring the Custom Rewriter to Start on Reboot” on page 256).
“Using the Proxy Services Interface” on page 267	The section has been revised to make access to the sections covering the main configuration pages easier.
“Using LDAP Server Load Balancing and Failover” on page 387	iChain 2.3 SP5 modified the methods used for LDAP load balancing and failover.
“Using ACLCHECK options” on page 150	The default for dynamic group option (/g) has been changed from enabled to disabled. Two other options are no longer valid (/c and /l), and one option (/z) has been added.
“Configuring OLAC for Shared Secrets” on page 158	The section was revised because the SecretStore processor no longer uses the Provider URL option in the configuration file.
“Setting Timeouts and Pool Limits for LDAP Profiles” on page 310	Added for the new commands you can use to set time limits for LDAP login requests and LDAP search queries.

C.3 March 16, 2006 (SP4)

A number of minor technical corrections were made to the *Novell iChain Administration* guide in March 14, 2006, for Support Pack 4. The following table lists the major updates:

Location	Change
"Using ACLCHECK options" on page 150	Added documentation about new command line options.
"Form Fill" on page 161	Reorganized the information and revised the description of many of the tags.
"Known Issues" on page 371	Created an appendix which describes the issues that have been discovered in previous releases of iChain.

C.4 October 24, 2005

The following section was updated in the *Novell iChain Administration* guide:

Location	Change
"Pin Type" on page 234	Clarified the cache bypass list option and added a TID reference (10097536).
Section 15.9, "Renewing a Third-Party Certificate," on page 201	Improved the accuracy of the information and included a TID reference for converting a certificate to .p7b.

C.5 August 8, 2005 (SP3)

The following table lists the updates that were made to the *Novell iChain Administration* guide in August 2, 2005, for Support Pack 3:

Location	Change
"Using ACLCHECK options" on page 150	Added documentation about a new command line option that lets you disable or enable the processing of Dynamic groups. By default, this option is disabled.
"Define the Location of the Trusted Root Container for all Trusted Roots" on page 88	Added the section and command line option.
Section 16.1.2, "Configuring iChain to Accelerate Citrix MetaFrame Servers with Nfuse," on page 204	Corrected the command to enable the ICA thin client tunnel. It needs to be set accelerator *accelerator name* tunnelauthforica = yes.

Location	Change
"Using ACLCHECK options" on page 150	Added information about the following ACLCHECK options: - /g - /l - /m - /p - /t - /v

C.6 June 24, 2005 (SP3)

The following table lists the updates that were made to the *Novell iChain Administration* guide in June 24, 2005, for Support Pack 3:

Location	Change
Chapter 16, "Using Advanced Accelerator Features," on page 203	Removed Dynamic Bypass information. This feature is no longer available.
Section 8.1.3, "RADIUS Authentication Load Balance and Failover," on page 105	Added a section that explains RADIUS load balance and failover.

Removed the Dynamic Bypass information from Chapter 16, "Using Advanced Accelerator Features," on page 203. That feature is no longer available.

C.7 April 12, 2005 (SP2)

The following table lists the updates that were made to the *Novell iChain Administration* guide in April, 2005, for Support Pack 2:

Location	Change
"Create a New Authentication Profile Via the Proxy Server Administration Tool" on page 79	Updated the steps and note.
Section 10.3, "Configuring Session Broker without a Floppy Drive," on page 125	Added the section.
"iChain Server: Secure Exchange Between the Browser and iChain" on page 80	Updated the steps and notes.

C.8 March 3, 2005 (SP2)

The following table lists the updates that were made to the *Novell iChain Administration* guide in March, 2005, for Support Pack 2:

Location	Change
Section 16.6, “Concurrent Login Restriction,” on page 240	New section.
Section 4.6, “Setting Up Secure Exchange,” on page 63	Revised procedures.
Section 4.6.1, “Creating a Trusted Root Container and Trusted Root Objects,” on page 64	Revised the configuration information.
Section 4.6.2, “Configuring iChain to Use the Trusted Root Container and Objects,” on page 65	Revised the configuration information.

C.9 January 26, 2005 (SP2)

In addition to style and consistency changes throughout, the following sections list the updates that were made to the *Novell iChain Administration* guide in January, 2005, for Support Pack 2:

Location	Change
Form Fill (page 161)	Revised entire section.
Setting Up Enhanced Security Within the Authentication Cookie (page 265)	Added a command line parameter for increasing the number of bits of random data that must be matched for authentication.
Importing a CSR Signed by Intermediates (page 197)	Added instructions for importing a CSR that is signed by intermediates.
Renewing a Third-Party Certificate (page 201)	Added instructions for renewing a third-party certificate.