

Novell Access Manager

3.0

March 7, 2007

ADMINISTRATION GUIDE

www.novell.com



Novell®

Legal Notices

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

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

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. Please refer to 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.

Copyright © 2006-2007 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

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 at <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 online documentation for this and other Novell products, and to get updates, see www.novell.com/documentation.

Novell Trademarks

Novell is a registered trademark of Novell, Inc., in the United States and other countries.

SUSE is a registered trademark of Novell, Inc., in the United States and other countries.

Third-Party Materials

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

Contents

About This Guide	17
Part I System Management	19
1 Security Considerations	21
1.1 Certificates	21
1.2 Access Manager Administration Console	21
1.3 Configuration Store	22
1.4 Auditing and Event Notification	22
1.5 Identity Server	23
1.5.1 Federation Concerns	23
1.5.2 Authentication Contracts	23
1.6 NetWare Access Gateway	24
1.7 Linux Access Gateway	25
1.8 SSL VPN	25
1.9 J2EE Agent	25
2 Administration Console	27
2.1 Starting and Stopping Access Manager Components	27
2.1.1 Updating an Identity Server Configuration	27
2.1.2 Restarting the Identity Server	28
2.1.3 Restarting the Access Gateway Service Provider	28
2.1.4 Starting the Access Gateway Service Provider	29
2.1.5 Stopping the Access Gateway Service Provider	29
2.1.6 Restarting the Access Gateway	29
2.1.7 Scheduling a Restart of the Access Gateway	29
2.1.8 Stopping the Access Gateway	30
2.1.9 Scheduling the Shutdown of the Access Gateway	30
2.2 Backing Up and Restoring the Configuration Store	31
2.2.1 How The Backup and Restore Process Works	31
2.2.2 Running the Backup	32
2.2.3 Restoring a Configuration Store	32
2.2.4 Running the Diagnostic Configuration Export	33
2.3 Installing Secondary Versions of the Administration Console	34
2.4 Converting your Secondary Console into your Primary Console	35
2.4.1 Changing the Master Replica	35
2.4.2 Restoring CA Certificates	36
2.4.3 Deleting Objects from the eDirectory Configuration Store	36
2.4.4 Additional Procedures	36
2.5 Multiple Sessions, Multiple Administrators	39
3 Changing the IP Address of Access Manager Devices	41
3.1 Changing the IP Address of the Administration Console	41
3.2 Changing the IP Address of an Identity Server	41
3.3 Changing the IP Address of the Access Gateway	42

4	Maintaining an Identity Server	45
4.1	Managing in Identity Server	45
4.2	Editing Server Details	46
4.3	Security	46
	Part II Novell Identity Server Configuration	51
5	Configuring an Identity Server	53
5.1	Creating an Identity Server Configuration (Advanced Options)	53
5.2	Assigning an Identity Server to a Configuration	59
5.3	Removing a Server from a Configuration	59
5.4	Modifying the Base URL	59
5.5	Enabling Role-Based Access Control	60
5.6	Clustering Identity Servers	60
5.7	Configuring Secure Communication on the Identity Server	63
6	Defining Reusable Configuration Settings	65
6.1	Configuring Attribute Sets	65
6.2	Editing Attribute Sets	67
6.3	Configuring User Matching Expressions	67
6.4	Adding Custom Attributes	68
6.4.1	Creating Shared Secret Names	68
6.4.2	Creating LDAP Attribute Names	69
7	Configuring Local Authentication	71
7.1	Configuring Identity User Stores	71
7.2	Creating Authentication Classes	75
7.2.1	Creating Basic, Form-Based, or NMAS Authentication Classes	75
7.2.2	Creating an X.509 Authentication Class	77
7.2.3	Creating a RADIUS Authentication Class	80
7.3	Configuring Authentication Methods	81
7.4	Configuring Authentication Contracts	82
7.4.1	Creating a Local Contract	82
7.4.2	Creating an External Contract	83
7.5	Specifying Authentication Defaults	84
7.6	Creating Custom Login Pages	85
7.6.1	Modifying the User Portal Page	85
7.6.2	Creating Your Own Login or Error Page	86
8	Configuring Trusted Providers	87
8.1	Understanding the Trust Model	87
8.2	Creating a Trusted Provider Reference	89
8.3	Reimporting a Trusted Provider's Metadata	91
8.4	Editing a SAML 1.1 Trusted Identity Provider's Metadata	92
8.5	Editing a SAML 1.1 Trusted Service Provider's Metadata	93
8.6	Configuring Common Access Settings for a Trusted Provider	94
8.6.1	Configuring Display and Access Settings	94
8.6.2	Configuring Communication Security Settings	95
8.6.3	Specifying the Intersite Transfer Service URL	96

8.7	Selecting Attributes for a Trusted Provider	97
9	Configuring User Authentication and Federation	99
9.1	Configuring Authentication for a Trusted Identity Provider	99
9.2	Configuring User Identification Methods	101
9.2.1	Selecting a User Identification Method	101
9.2.2	Configuring the User Matching Method	102
9.2.3	Defining the User Provisioning Method	103
9.2.4	User Provisioning Error Messages	107
9.3	Configuring Authentication for a Trusted Service Provider	108
9.4	Configuring User Identification Methods for SAML 1.1 Trusted Identity Providers	109
9.5	Specifying a SAML Audience URI	111
10	Configuring Communication Profiles	113
11	Configuring Liberty Web Services	115
11.1	Configuring the Web Services Framework	115
11.2	Enabling Web Services and Profiles	116
11.3	Editing Web Service Descriptions	117
11.4	Configuring Credential Profile Security and Display Settings	118
11.5	Configuring Service and Profile Details	120
11.6	Customizing Attribute Names	122
11.7	Editing Web Service Policies	123
11.8	Configuring the Web Service Consumer	126
11.9	Mapping LDAP and Liberty Attributes	127
11.9.1	Configuring One-to-One Attribute Maps	127
11.9.2	Configuring Employee Type Attribute Maps	131
11.9.3	Configuring Employee Status Attribute Maps	132
11.9.4	Configuring Postal Address Attribute Maps	133
11.9.5	Configuring Contact Method Attribute Maps	134
11.9.6	Configuring Gender Attribute Maps	136
11.9.7	Configuring Marital Status Attribute Maps	136
Part III	Access Gateway Configuration	139
12	Configuring the Access Gateway to Protect Resources	141
12.1	Creating a Reverse Proxy and Proxy Service	142
12.2	Configuring a Proxy Service	146
12.3	Configuring the Web Servers of a Proxy Service	148
12.4	Configuring Protected Resources	150
12.4.1	Setting Up a Protected Resource	150
12.4.2	Assigning an Authorization Policy to a Protected Resource	152
12.4.3	Assigning an Identity Injection Policy to a Protected Resource	153
12.4.4	Assigning a Form Fill Policy to a Protected Resource	154
12.5	Configuring HTML Rewriting	155
12.5.1	Understanding the Rewriting Process	156
12.5.2	Configuring the HTML Rewriter	162
12.6	Configuring Connection Limits	166
12.6.1	Configuring TCP Listen Options for Clients	167
12.6.2	Configuring TCP Connect Options for Web Servers	168
12.6.3	Configuring Connection and Session Persistence	170

13	Configuring the Access Gateway for SSL	171
13.1	Using SSL on the Access Gateway Communication Channels	172
13.2	Prerequisites for SSL	173
13.2.1	Prerequisites for SSL Communication between the Access Gateway and the Web Servers	173
13.2.2	Prerequisite for SSL Communication between the Identity Server and the Access Gateway	173
13.3	Configuring SSL Communication with the Browsers and the Identity Server	174
13.4	Configuring SSL between the Proxy Service and the Web Servers	176
13.5	Managing the Certificates of the Embedded Service Provider	178
13.6	Configuring the Encryption Key	179
14	Additional Configuration Settings	181
14.1	Changing the Name of an Access Gateway and Modifying Other Descriptive Details	181
14.2	Modifying the Base URL of the Identity Server	182
14.3	Setting Date and Time	182
14.4	Setting Up a Tunnel	184
14.5	Configuring Error Page Presentation	185
14.6	Configuring Console Access	186
14.6.1	Setting Up an FTP Listening Address	186
14.6.2	Enabling Console Access with SSH and Telnet Sessions	187
14.6.3	Setting the Password for the admin and config Users	188
14.7	Configuring Network Settings	188
14.7.1	Viewing and Modifying Adapter Settings	189
14.7.2	Viewing and Modifying Gateway Settings	191
14.7.3	Viewing and Modifying DNS Settings	194
14.7.4	Configuring Hosts	196
14.7.5	Configuring IPQoS	197
14.8	Customizing Log Out	198
14.9	Configuring X-Forwarded-For Headers	198
14.10	Upgrading the Access Gateway Software	199
15	Configuring the Cache Settings	201
15.1	Configuring Global Caching Options	202
15.2	Controlling Browser Caching	204
15.3	Configuring Custom Cache Control Headers	205
15.3.1	Understanding How Custom Cache Control Headers Work	205
15.3.2	Enabling Custom Cache Control Headers	206
15.4	Configuring a Pin List	207
15.4.1	URL Mask	209
15.4.2	Pin Type	210
15.4.3	Follow Links	211
15.5	Configuring a Purge List	211
15.6	Purging Cached Content	212
15.7	Preventing a Web Site from Being Cached	212
16	Protecting Multiple Resources	215
16.1	Setting Up a Group of Web Servers	216
16.2	Using Multi-Homing to Access Multiple Resources	216
16.2.1	Domain-Based Multi-Homing	217
16.2.2	Path-Based Multi-Homing	218

16.2.3	Virtual Multi-Homing	221
16.2.4	Creating a Second Proxy Service	221
16.2.5	Configuring a Path-Based Multi-Homing Proxy Service	223
16.3	Managing Multiple Reverse Proxies	224
16.3.1	Managing Entries in the Reverse Proxy List	225
16.3.2	Changing the Authentication Proxy Service	226
16.4	Managing a Group of Access Gateways	226
16.4.1	Applying Configuration Changes to a Group	227
16.4.2	Issuing Commands to a Group of Access Gateways	228
16.4.3	Managing the Servers in the Group	229
16.4.4	Adding Servers to an Existing Group	229
16.4.5	Changing the Primary Cluster Server	230
16.4.6	Deleting Members from a Group	230
16.4.7	Viewing the Health of the Group	230
16.4.8	Viewing Group Statistics	231
16.4.9	Viewing Group Command Status	231
16.4.10	Viewing the Status of Group Alerts	232
17	Configuring Access Manager for Citrix Clients	233
17.1	Prerequisites	233
17.2	Configuring the Access Gateway for Citrix Clients	234
18	Virtualization on the Linux Access Gateway	237
18.1	Virtualization Overview	237
18.2	Configuring VMs	237
18.2.1	Creating a Novell COS Partition	238
18.2.2	Configuring a VM	238
18.3	Modifying a VM	239
18.4	Deleting a VM	240
18.5	Setting Maximum Memory Limit	240
18.6	Scan	241
Part IV	Configuring the SSL VPN Gateway	243
19	Overview of SSL VPN	245
19.1	Understanding How SSL VPN Works	245
19.2	SSL VPN Servers	246
20	Managing SSL VPN Servers	247
20.1	Viewing SSL VPN Servers	247
20.2	Viewing SSL VPN Server Details	248
20.3	Monitoring Alerts	250
21	Configuring the SSL VPN Servers	251
21.1	Viewing and Modifying Server Configuration	252
21.2	Configuring the DNS Server List	253
21.3	Viewing and Modifying Traffic Policies	254
21.4	Configuring Full Tunneling	256
21.5	Viewing Client Integrity Check Policies	257

21.5.1	Configuring Category	257
21.5.2	Configuring Application Names	258
21.5.3	Configuring Application Attributes and Details	259
21.6	Creating New Client Integrity Check Policies	260
21.7	Configuring Alerts Settings	261
21.8	Adding Certificates to the SSL VPN Keystore	262
21.9	Adding Trusted Roots for SSL VPN	263
21.10	Configuring SSL VPN to Download Applet on Internet Explorer	263
21.11	Accelerating SSL VPN	264
21.11.1	Injecting the SSL VPN Header	264
21.12	Moving SSL VPN Server to a Different Administration Console	267
21.13	Accelerating SSL VPN with a New Access Gateway	268
22	Load Balancing and Fault Tolerance	269
22.1	Installing SSL VPN Servlet on a Separate Machine	269
22.2	Load Balancing	270
22.3	Fault Tolerance	270
23	Accessing the SSL VPN User Portal	273
23.1	Accessing SSL VPN on Linux	273
23.2	Accessing SSL VPN on Macintosh	275
23.3	Accessing SSL VPN on Windows	276
24	Understanding the SSL VPN User Interface	279
24.1	Viewing the Home Page	280
24.2	Viewing the Statistics Page	280
24.3	Viewing the Policies Page	281
24.4	Viewing the Log Entries Page	281
24.4.1	Saving Log Entries	282
24.5	Viewing the Applications Page	283
24.5.1	SSLizing Applications on Linux	283
24.5.2	SSLizing Applications on Macintosh	284
24.6	Monitoring the Connection Status	284
Part V	Security and Certificate Management	285
25	Understanding How Access Manager Uses Certificates	287
26	Managing Certificates	289
26.1	Creating Certificates	289
26.1.1	Creating a Locally Signed Certificate	290
26.1.2	Generating a Certificate Signing Request	294
26.1.3	Importing a Signed Certificate	295
26.2	Auto-Importing Certificates from Servers	295
26.3	Importing a Private/Public Key Pair	296
26.4	Exporting a Private/Public Key Pair	296
26.5	Importing Public Key Certificates (Trusted Roots)	297
26.6	Renewing a Certificate	297
26.7	Exporting a Public Certificate	298

26.8	Enabling 4096k Keys	299
26.8.1	Jurisdiction Policy Files	299
27	Assigning Certificates to Access Manager Devices	301
27.1	Importing a Trusted Root to the LDAP User Store	301
27.2	Replacing Identity Server SSL Certificates	303
27.3	Changing a Trust Store Password	304
27.4	Assigning Certificates to an Access Gateway	305
27.5	Assigning Certificates to J2EE Agents	305
27.6	Configuring SSL for Authentication between the Identity Server and Access Gateway	305
27.7	Changing a Non-Secure (HTTP) Environment to a Secure (HTTPS) Environment	306
27.8	Creating Keystores and Trust Stores	306
Part VI	Policy Management	309
28	Managing Policies	311
28.1	Selecting a Policy Type	311
28.2	Managing Policy Containers	312
28.3	Managing Policies	312
28.3.1	Creating Policies	312
28.3.2	Deleting Policies	313
28.3.3	Importing and Exporting Policies	313
28.4	Managing a Rule List	313
28.4.1	Rule Evaluation for Role Policies	314
28.4.2	Rule Evaluation for Authorization Policies	314
28.4.3	Rule Evaluation for Identity Injection and Form Fill Policies	314
28.5	Enabling Policy Logging	315
29	Creating Role Policies	317
29.1	Understanding RBAC in Access Manager	317
29.2	Creating Roles	320
29.2.1	Selecting Conditions	321
29.2.2	Selecting an Action	323
29.2.3	Reviewing the Rules	323
29.2.4	Example Role Policies	323
29.3	Creating Access Manager Roles from an Existing Role-Based Policy System	327
29.3.1	Creating a Role Using an LDAP Attribute	328
29.3.2	Creating a Role Using the Location of the User Objects	330
29.3.3	Creating a Role Using Role or Group Membership	332
29.4	Mapping Roles between Trusted Providers	335
29.5	Enabling and Disabling Role Policies	336
29.6	Importing and Exporting Role Policies	337
30	Creating Authorization Policies	339
30.1	Designing an Authorization Policy	339
30.1.1	Many Rules or Many Conditions	340
30.1.2	Controlling Access with a Deny Rule	340
30.1.3	Controlling Access with Multiple Conditions	341
30.1.4	Using Permit Rules with a Deny Rule	342
30.1.5	Public Policies	344

30.1.6	General Design Principles	344
30.1.7	Assigning Policies to Resources	344
30.2	Creating Access Gateway Authorization Policies	344
30.2.1	The Process	345
30.2.2	Sample Policies	346
30.3	Creating Web Authorization Policies for J2EE Agents	350
30.4	Creating Enterprise JavaBean Authorization Policies for J2EE Agents	351
30.5	Conditions	352
30.5.1	Authentication Contract Condition	353
30.5.2	Client IP Condition	355
30.5.3	Credential Profile Condition	356
30.5.4	Current Date Condition	357
30.5.5	Current Day of Week Condition	358
30.5.6	Current Day of Month Condition	359
30.5.7	Current Time of Day Condition	359
30.5.8	HTTP Request Method Condition	360
30.5.9	LDAP Attribute Condition	361
30.5.10	Liberty User Profile Condition	362
30.5.11	Roles for Current User Condition	363
30.5.12	URL Condition	363
30.5.13	URL Scheme Condition	364
30.5.14	URL Host Condition	365
30.5.15	URL Path Condition	366
30.5.16	URL File Name Condition	367
30.5.17	URL File Extension Condition	368
30.5.18	X-Forward-For IP Condition	369
30.6	Using Multiple Conditions	370
30.7	Importing and Exporting Authorization Policies	371

31 Creating Identity Injection Policies 373

31.1	Designing an Identity Injection Policy	373
31.2	Configuring an Identity Injection Policy	374
31.3	Configuring an Authentication Header Policy	375
31.4	Configuring a Custom Header Policy	378
31.5	Configuring a Custom Header with Tags	379
31.6	Specifying a Query String for Injection	381
31.7	Injecting into the Cookie Header	382
31.8	Importing and Exporting Identity Injection Policies	382

32 Creating Form Fill Policies 385

32.1	Understanding an HTML Form	385
32.2	Creating a Form Fill Policy for the Sample Form	388
32.3	Implementing Form Fill Policies	391
32.3.1	Creating a Form Fill Policy	391
32.3.2	Creating a Login Failure Policy	395
32.3.3	Troubleshooting a Form Fill Policy	396
32.4	Creating and Managing Shared Secrets	396
32.4.1	Local Shared Secret	397
32.4.2	Remote Shared Secret	397
32.5	Importing and Exporting Form Fill Policies	398

Part VII Monitoring Access Manager Components	399
33 Enabling Auditing	401
33.1 Configuring Access Manager for Novell Auditing	402
33.1.1 Specifying the Logging Server and Events	402
33.1.2 Generating Queries	403
33.2 Enabling Identity Server Audit Events	403
33.3 Enabling Access Gateway Audit Events	405
33.4 Enabling SSL VPN Audit Events	406
34 Configuring Logging	409
34.1 Understanding the Types of Logging	409
34.1.1 Component Logging for Troubleshooting Configuration or Network Problems	409
34.1.2 Debug Trace Logging to Discover Software Problems	410
34.1.3 HTTP Transaction Logging for Proxy Services	410
34.2 Configuring Component Logging	410
34.2.1 Enabling Component Logging	411
34.2.2 Downloading the Log Files	412
34.2.3 Understanding the Log File Format	414
34.3 Configuring Debug Trace Logging	415
34.4 Configuring Access Gateway Logging	416
34.4.1 Determining Logging Requirements	416
34.4.2 Calculating Roll Over Requirements	417
34.4.3 Enabling Logging	420
34.4.4 Configuring Common Log Options	420
34.4.5 Configuring Extended Log Options	422
34.4.6 Configuring Log Pushing	424
34.4.7 Configuring the Size of the Log Partition	426
35 Viewing Statistics	427
35.1 Monitoring Identity Server Statistics	427
35.2 Monitoring Access Gateway Statistics	429
35.2.1 Server Activity	429
35.2.2 Server Benefits	433
35.2.3 Server SSL Activity	434
35.2.4 Top 20 Sites	434
35.2.5 Service Provider Activity	435
35.2.6 Configured Services	438
35.3 Viewing SSL VPN Statistics	438
35.3.1 Viewing the Bytes Graphs	440
36 Managing Server Health	443
36.1 Health States	443
36.2 Monitoring the Health of an Identity Server	444
36.3 Monitoring the Health of an Access Gateway	446
36.4 Monitoring the Health of an SSL VPN Server	448
37 Reviewing Command Status	449
37.1 Viewing the Command Status of the Identity Server	449
37.2 Viewing the Command Status of the Access Gateway	449

37.2.1	Viewing the Status of Current Commands	450
37.2.2	Viewing Detailed Command Information	450
37.2.3	Issuing Commands	451
37.2.4	Scheduling a Command	452
37.3	Viewing Command Status of the SSL VPN Server	453
37.3.1	Viewing Command Information	453
38	Reviewing Alerts	455
38.1	Monitoring Identity Server Alerts	455
38.2	Monitoring Access Gateway Alerts	455
38.2.1	Reviewing Java Alerts	455
38.2.2	Configuring Access Gateway Alerts	456
38.2.3	Enabling SNMP	461
Part VIII	Troubleshooting	465
39	Troubleshooting the Administration Console	467
39.1	Session Conflicts	467
39.2	Unable to Login to the Administration Console	467
39.3	Exception Processing IdentityService_ServerPage.JSP	467
39.4	Logging	467
40	Troubleshooting for the Identity Server and Authentication	469
40.1	Authentication	469
40.1.1	Browser Hangs in an Authentication Redirect	469
40.1.2	Prompting Users to Refresh a Session before Expiration	469
40.2	Translating the Identity Server Configuration Port	470
40.3	Problems Reading Keystores after Identity Server Re-installation	471
41	Troubleshooting Access Manager Policies	473
41.1	Roles and Authorization Policies	473
41.2	Turning on Logging for Policy Evaluation	474
41.3	Understanding Policy Evaluation Traces	475
41.3.1	Format	475
41.3.2	Policy Result Values	480
41.3.3	Sample Policies	481
41.4	Policy Page Timeout	483
41.5	Policy Changes Are Not Saved	483
41.6	Policy Creation and Storage	483
41.7	Policy Distribution	484
41.8	Policy Evaluation: Access Gateway Devices	485
41.8.1	Successful Policy Configuration Example	486
41.8.2	No Policy Defined Configuration Example	487
41.8.3	Deny Access Configuration/Evaluation Example	487
42	Troubleshooting the Access Gateway	491
42.1	Fixing Problems Common to Both Platforms	491
42.1.1	Mismatch Certificates in a Group of Access Gateways	491
42.2	Troubleshooting the Linux Access Gateway	492

42.2.1	Linux Access Gateway Logs	492
42.2.2	Troubleshooting a Linux Access Gateway Crash	494
42.2.3	Connection Details	498
42.2.4	Network Socket Issues	498
42.2.5	Authentication Issues	498
42.2.6	Rewriter Issues	501
42.2.7	Useful Linux Access Gateway Commands	502
42.2.8	Using the Linux Access Gateway Console	503
42.2.9	COS Related Issues	503
42.2.10	Memory Related Issues	503
42.2.11	Authorization and Identity Injection Issues	504
42.2.12	Form Fill Issues	505
42.3	Troubleshooting the NetWare Access Gateway	507
42.3.1	Additional Options during the Boot Process	507
42.3.2	Setting the Date and Time at the Console	508
42.3.3	Unlocking the NetWare Access Gateway Console	509
42.3.4	Command Line Options	509
42.3.5	Telnet Fails after Performing an Upgrade	509
42.3.6	SSL Certificate Error with X.509 Authentication from NetWare Access Gateway	510
43	Troubleshooting SSL VPN	511
43.1	Connecting Successfully to the Server	511
43.1.1	Connection Problems with Mozilla Firefox	512
43.1.2	Connection Problems with Internet Explorer	513
43.2	SSL VPN Not Reporting	513
43.3	Verifying SSL VPN Components	514
43.4	Issues With Keep Alive	514
43.5	Idle Connection	515
43.6	Mozilla Firefox Browser Displays "X" Mark	515
43.7	Unable to Contact the SSL VPN Server	515
43.8	Unable to Get Authentication Headers	515
43.9	The SSL VPN Connection Is Successful But There Is No Data Transfer	516
43.10	Applications Are Not Getting SSLized from Terminal After Running the <code>su</code> Command	516
43.11	Unable to Connect to SSL VPN Gateway	516
Part IX	Appendixes	517
A	About Liberty and SAML 2.0	519
A.1	Liberty Alliance Resources	519
A.2	What's New in SAML 2.0?	520
B	Certificates Terminology	521
C	Data Model Extension XML	523
C.1	Elements	523
C.2	Writing Data Model Extension XML	526
D	Logging: Using the Custom Content Filter	529
D.1	Custom Content Filter XML Syntax	529
D.2	Examples of Custom Content Filter XML	530

D.2.1	Example One	530
D.2.2	Example Two	530
D.2.3	Example Three.....	532
D.3	Custom Content Filter Thread Identifiers	533

E Authentication Classes and Duplicate Common Names **535**

About This Guide

This guide describes all of the interface features of Novell® Access Manager, including:

- ◆ Part I, “System Management,” on page 19
- ◆ Part II, “Novell Identity Server Configuration,” on page 51
- ◆ Part III, “Access Gateway Configuration,” on page 139
- ◆ Part IV, “Configuring the SSL VPN Gateway,” on page 243
- ◆ Part V, “Security and Certificate Management,” on page 285
- ◆ Part VI, “Policy Management,” on page 309
- ◆ Part VII, “Monitoring Access Manager Components,” on page 399
- ◆ Part VIII, “Troubleshooting,” on page 465
- ◆ Part IX, “Appendixes,” on page 517

This administration guide is intended to help you understand and configure all of the features provided by Access Manager, and includes advanced topics. It is recommended that you first become familiar with the information in the *Novell Access Manager 3.0 Setup Guide*.

The basic setup guide can help you understand how to perform a basic Identity Server configuration, set up a resource protected by an Access Gateway, and configure SSL. The basic setup and the administration guides are designed to work together, and important information and setup steps are not necessarily repeated in both places.

Audience

This guide is intended for Access Manager administrators. It is assumed that you have knowledge of evolving Internet protocols, such as:

- ◆ Extensible Markup Language (XML)
- ◆ Simple Object Access Protocol (SOAP)
- ◆ Security Assertion Markup Language (SAML)
- ◆ Public Key Infrastructure (PKI) digital signature concepts and Internet security
- ◆ Secure Socket Layer/Transport Layer Security (SSL/TSL)
- ◆ Hypertext Transfer Protocol (HTTP and HTTPS)
- ◆ Uniform Resource Identifiers (URIs)
- ◆ Domain Name System (DNS)
- ◆ Web Services Description Language (WSDL)

Feedback

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

Additional Documentation

Before proceeding, you should be familiar with the *Novell Access Manager 3.0 Installation Guide* and the *Novell Access Manager 3.0 Setup Guide*, which provides information about setting up the Access Manager system.

If you are unfamiliar with SAML 1.1, see [SAML Overview \(http://www.novell.com/documentation/saml/saml/data/ag8qdk7.html\)](http://www.novell.com/documentation/saml/saml/data/ag8qdk7.html) on the [Documentation \(http://www.novell.com/documentation/a-z.html\)](http://www.novell.com/documentation/a-z.html) Web site.

For conceptual information about Liberty, and to learn about what is new for SAML 2.0, see [Appendix A, “About Liberty and SAML 2.0,” on page 519.](#)

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.

When a single pathname can be written with a backslash for some platforms or a forward slash for other platforms, the pathname is presented with a backslash. Users of platforms that require a forward slash, such as Linux* or UNIX*, should use forward slashes as required by your software.

System Management

This section discusses system and server management topics that apply to the Novell® Access Manager Administration Console.

- ♦ [Chapter 1, “Security Considerations,” on page 21](#)
- ♦ [Chapter 2, “Administration Console,” on page 27](#)
- ♦ [Chapter 3, “Changing the IP Address of Access Manager Devices,” on page 41](#)
- ♦ [Chapter 4, “Maintaining an Identity Server,” on page 45](#)

Security Considerations

1

This section describes some security checks that you can use to help verify the security of your Novell® Access Manager configuration.

- ♦ [Section 1.1, “Certificates,” on page 21](#)
- ♦ [Section 1.2, “Access Manager Administration Console,” on page 21](#)
- ♦ [Section 1.3, “Configuration Store,” on page 22](#)
- ♦ [Section 1.4, “Auditing and Event Notification,” on page 22](#)
- ♦ [Section 1.5, “Identity Server,” on page 23](#)
- ♦ [Section 1.6, “NetWare Access Gateway,” on page 24](#)
- ♦ [Section 1.7, “Linux Access Gateway,” on page 25](#)
- ♦ [Section 1.8, “SSL VPN,” on page 25](#)
- ♦ [Section 1.9, “J2EE Agent,” on page 25](#)

For firewall information, see “[Setting Up Firewalls](#)” in the *Novell Access Manager 3.0 Setup Guide*.

1.1 Certificates

Your security deployment plan should contain policies for the following:

- ♦ **Key size for certificates.** The Access Manager product ships with a CA that can create certificates with a key size of 2048, which is the maximum size supported by older software. For information about increasing the key size to 4096, see [Section 26.8, “Enabling 4096k Keys,” on page 299](#).
- ♦ **Certificate renewal dates.** We recommend that certificates should be renewed every two years. Your security needs might allow for a longer or shorter period.
- ♦ **Trusted Certificate Authorities.** The Access Manager ships with a CA, and during installation of the various components, creates and distributes certificates. If this CA is not on your list of trusted CAs, you need to add certificates created by your trusted CAs. See [Chapter 27, “Assigning Certificates to Access Manager Devices,” on page 301](#).

1.2 Access Manager Administration Console

The admin user you create when you install the Administration Console has all rights to the Access Manager components. We recommend that you protect this account by configuring the following features:

- ♦ **Password Restrictions:** When the admin user is created, no password restrictions are set. To ensure that the password meets your minimum security requirements, you should configure the standard eDirectory password restrictions for this account. Go to the Administration Console and click *Users*. Browse to the admin user (found in the novell container), then click *Restrictions*. For configuration help, use the *Help* button.
- ♦ **Intruder Detection:** The admin user is created in the novell policy container. To modify the intruder detection policy for this container, go to the Administration Console and click

Directory Administration > Modify Object. Select *novell*, then click *OK*. Click *Intruder Detection*. For configuration help, use the Help button.

You also need to protect the Administration Console from Internet attacks. It should be installed behind your firewall.

If you install secondary consoles for redundancy, these secondary consoles should be on the same network. For a secure system, they should not be required to cross routers to communicate with each other.

Also, if you are installing the Administration Console on a separate machine, ensure that the DNS names resolve between the Identity Server and the Administration Console. This ensures SSL security functions correctly between the Identity Server and the configuration store in the Administration Console.

1.3 Configuration Store

The configuration store is an embedded, modified version of eDirectory™. It can only be backed up and restored with command line options. The backup file is not encrypted, so it should not be used to back up user accounts with their passwords. Because of this limitation, it should not be used for a user store.

You should back up the configuration store on a regular schedule, and the files created (an LDIF and ZIP file) should be stored in a secure place. See [Section 2.2, “Backing Up and Restoring the Configuration Store,” on page 31](#).

In addition to backing up the configuration store, you should also install at least two Administration Consoles (a primary console and a secondary replica). This ensures that if the primary console goes down, the secondary console can keep the communication channels open between the various components. You can install up to three Administration Consoles.

1.4 Auditing and Event Notification

For a secure system, you need to set up either auditing or syslogging to notify the system administrator when certain events occur. The most important audit events to monitor are the following:

- ◆ Configuration changes
- ◆ System shutdowns and startups
- ◆ Server imports and deletes
- ◆ Intruder lockout detection (available only for eDirectory user stores)
- ◆ User account provisioning

Audit events are device specific. To select auditing events, use the following:

- ◆ **Administration Console:** In the Administration Console, click *Access Manager > Auditing*
- ◆ **Identity Server:** In the Administration Console, click *Access Manager > Identity Servers > Configuration Assignment > Logging*.
- ◆ **Access Gateway:** In the Administration Console, click *Access Manager > Access Gateways > Edit > Novell Audit*.
- ◆ **J2EE Agent:** In the Administration Console, click *Access Manager > J2EE Agents > Edit*.

- ◆ **SSL VPN:** In the Administration Console, click *Access Manager > SSL VPNs > Edit > Novell Audit Settings*.

In addition to the selectable events, device-generated alerts are automatically sent to the audit server. These Management Communication Channel events have an ID of 002e0605. All Access Manager events begin with 002e. SSL VPN starts with 0031. You can set up Novell Auditing to send e-mail whenever these events or your selected audit events occur. See [Configuring System Channels \(http://www.novell.com/documentation/novellaudit20/novellaudit20/data/al6t4sd.html\)](http://www.novell.com/documentation/novellaudit20/novellaudit20/data/al6t4sd.html).

The Access Gateway also supports a syslog and allows you to send e-mail notification to system administrators. To configure this system:

- ◆ **Linux Access Gateway:** In the Administration Console, click *Access Manager > Access Gateways > Edit > Alerts*.
- ◆ **NetWare Access Gateway:** In the Administration Console, click *Access Manager > Access Gateways > Edit > Legacy Alerts*.

1.5 Identity Server

By default, all Access Manager components (Identity Server, Access Gateway, SSL VPN, and J2EE agents) trust the certificates signed by the local CA. We recommend that you configure the Identity Server to use an SSL certificate signed externally, and that you configure the trusted store of the service provider for each component to trust this new CA. See [Chapter 27, “Assigning Certificates to Access Manager Devices,” on page 301](#).

- ◆ [Section 1.5.1, “Federation Concerns,” on page 23](#)
- ◆ [Section 1.5.2, “Authentication Contracts,” on page 23](#)

1.5.1 Federation Concerns

When you set up federation between an identity provider and a service provider, you can select either to exchange assertions with a post method or to exchange artifacts. An artifact is a randomly generated ID, it contains no sensitive data, and only the intended receiver can use it to retrieve assertion data. Assertions might contain the user’s password or other sensitive data, which can make them less secure than an artifact when the assertion is sent to the browser. It is possible for a virus on the browser machine to access the memory where the browser decrypts the assertion. If both providers support artifacts, you should select this method because it is more secure. For more details, see the Response Protocol Binding option in [Section 9.1, “Configuring Authentication for a Trusted Identity Provider,” on page 99](#).

1.5.2 Authentication Contracts

By default, the Administration Console allows you to select from the following contracts and options when specifying whether a resource requires an authentication contract:

- ◆ **None:** Allows public access to the resource and does not require authentication contract.
- ◆ **Name/Password - Basic:** Requires that the user enter a name and password that matches an entry in an LDAP user store. The credentials do not need to be sent over a secure port. This uses the unprotected BasicClass, which is not recommended for a production environment.

- ◆ **Name/Password - Form:** Requires that the user enter a name and password that matches an entry in an LDAP user store. The credentials do not need to be sent over a secure port. This uses the unprotected PasswordClass, which is not recommended for a production environment
- ◆ **Secure Name/Password - Basic:** Requires that the user enter the name and password from a secure (SSL) connection. This uses the ProtectedBasicClass, which is recommended for a production environment. If your Web servers are using basic authentication, this contract provides the credentials for this type of authentication.
- ◆ **Secure Name/Password - Form:** Requires that the user enter the name and password from a secure (SSL) connection. This uses the ProtectedPasswordClass, which is recommended for a production environment.
- ◆ **Any Contract:** Allows the user to use any contract defined for the Identity Server configuration.

If you have set up the Access Manager to require SSL connections among all of its components, you should delete the Name/Password - Form and the Name/Password - Basic contracts. This removes them from the list of available contracts when configuring protected resources and prevents them from being assigned as the contract for a protected resource. If these contracts are assigned, the user's password goes across the wire in clear text format. At some future date, if your system needs this type of contract, you can re-create it from the method. To delete these contracts, go to the Administration Console and click *Identity Servers > Setup > [Configuration] > Local > Contracts*.

1.6 NetWare Access Gateway

The NetWare Access Gateway is installed with two user accounts: config and admin. The config user has no assigned password and the admin user is given the password of novell.

IMPORTANT: Before your Access Gateways is placed in a production environment, you need to assign a password for the config user, and you need to change the password for the admin user. For instructions, see [Section 14.6.3, "Setting the Password for the admin and config Users," on page 188](#).

Intruder detection lockout has been setup for these accounts. The config and admin users are allowed 5 attempts to log in successfully. If the user fails on the 5th attempt, the account is locked for 15 minutes.

Before you enable any of the following protocols, you need to be aware of their security issues:

- ◆ **Telnet:** Opens a clear text communication channel and sends passwords in clear text.
- ◆ **FTP:** Opens a clear text communication channel and sends passwords in clear text.
- ◆ **SSH:** Requires a LDAPS listener on port 636, on all IP addresses configured for the NetWare[®] Access Gateway. It cannot be restricted to a single IP address.
- ◆ **SFTP:** Requires the NCPIP.NLM to be loaded with a listener on port 524.

If you enable any of these protocols, the NetWare Access Gateway needs to be installed behind a firewall appliance, and the firewall needs to block the following ports:

- ◆ 21 for FTP
- ◆ 23 for Telnet
- ◆ 524 for SFTP

- ♦ 636 for SSH

For more information about installing the Access Gateway behind a firewall, see “[Setting Up Firewalls](#)” in the *Novell Access Manager 3.0 Setup Guide*.

1.7 Linux Access Gateway

The Linux Access Gateway is installed with two user accounts: `config` and `root` with the password as `novell`.

IMPORTANT: Before your Access Gateways is placed in a production environment, you need to change the passwords for both `root` and `config`. See the last steps in “[Using a Standard Linux Installation with the Default Settings](#)” in the *Novell Access Manager 3.0 Installation Guide*.

Before you enable the SSH protocol, it requires a LDAPS listener on port 636, on all IP addresses configured for the Linux Access Gateway. It cannot be restricted to a single IP address:

If SSH is enabled, the Linux Access Gateway needs to be installed behind a firewall appliance, and the firewall needs to block port 636 for SSH.

For more information about installing the Access Gateway behind a firewall, see “[Setting Up Firewalls](#)” in the *Novell Access Manager 3.0 Setup Guide*.

1.8 SSL VPN

Use of AES 265 mode of encryption is recommended.

1.9 J2EE Agent

All communication should be sent over a secure channel.

Administration Console

2

This section discusses the following Administration Console topics:

- ◆ [Section 2.1, “Starting and Stopping Access Manager Components,” on page 27](#)
- ◆ [Section 2.2, “Backing Up and Restoring the Configuration Store,” on page 31](#)
- ◆ [Section 2.3, “Installing Secondary Versions of the Administration Console,” on page 34](#)
- ◆ [Section 2.4, “Converting your Secondary Console into your Primary Console,” on page 35](#)
- ◆ [Section 2.5, “Multiple Sessions, Multiple Administrators,” on page 39](#)

2.1 Starting and Stopping Access Manager Components

Access Manager has three services that can be stopped and started: the Identity Server, the Access Gateway, and the embedded service provider within the Access Gateway. Normally, you do not need to stop and start these services. However, if you have change certain configuration options, you can be prompted to update the Identity Server or to restart the embedded service provider.

The following sections explain how to stop, start, and schedule a restart of the various Access Manager components:

- ◆ [Section 2.1.1, “Updating an Identity Server Configuration,” on page 27](#)
- ◆ [Section 2.1.2, “Restarting the Identity Server,” on page 28](#)
- ◆ [Section 2.1.3, “Restarting the Access Gateway Service Provider,” on page 28](#)
- ◆ [Section 2.1.4, “Starting the Access Gateway Service Provider,” on page 29](#)
- ◆ [Section 2.1.5, “Stopping the Access Gateway Service Provider,” on page 29](#)
- ◆ [Section 2.1.6, “Restarting the Access Gateway,” on page 29](#)
- ◆ [Section 2.1.7, “Scheduling a Restart of the Access Gateway,” on page 29](#)
- ◆ [Section 2.1.8, “Stopping the Access Gateway,” on page 30](#)
- ◆ [Section 2.1.9, “Scheduling the Shutdown of the Access Gateway,” on page 30](#)

2.1.1 Updating an Identity Server Configuration

Whenever you change an Identity Server configuration, the system prompts you to update the configuration. An *Update Servers* status is displayed under the *Status* column on the Setup page. You must click *Update Servers* to update the configuration so that your changes take effect.

When clicked, this link sends a reconfigure command to all servers that use the configuration. The servers then begin the reconfiguration process. This process occurs without interruption of service to users who are currently logged in.

When you update a configuration, the system blocks inbound requests until the update is complete. The server checks for any current requests being processed. If there are such requests in process, the server waits five seconds and tests again. This process is repeated three times, thus waiting up to

fifteen seconds for these requests to be serviced and cleared out. After this period of time, the update process begins. Any remaining requests have the potential for errors.

During the update process, all settings are reloaded with the exception of the base URL. In most cases, user authentications are preserved; however, there are conditions during which some sessions are automatically timed out. These conditions are:

- ◆ A user logged in via an authentication contract that is no longer valid. This occurs if an administrator removes a contract or changes the URI that is used to identify it.
- ◆ A user logged in to a user store that is no longer valid. This occurs if you remove a user store or change its type. Changing the LDAP address to a different directory is not recommended, because the system will not detect the change.
- ◆ A user received authentication from an identity provider that is no longer trusted. This occurs if you remove a trusted identity provider or if the metadata for the provider changed.

Additionally, if you remove a service provider from an identity provider, the identity provider removes the provided authentication to that service provider. This does not cause a timeout of the session to occur.

Changes to the SAML and Liberty protocol profiles can result in the trusted provider having outdated metadata for the Identity Server being reconfigured. This necessitates an update at the other provider and may cause unexpected behavior until that occurs.

- 1 In the Administration Console, click *Access Manager > Identity Servers*, then click the *Setup* tab.
- 2 Select the Identity Server configuration, then click *Update Servers*.
This link is available only when you have made changes that require a server update.

2.1.2 Restarting the Identity Server

Starting and stopping an Identity Server terminates active user sessions. These users receive a prompt to log in again.

- 1 In the Administration Console, click *Access Manager > Identity Servers* and select the Identity Server to stop.
- 2 Click *Stop*.
- 3 Wait for the *Command Status* to change from *Pending* to *Complete*.
- 4 Select the Identity Server, then click *Start*.
- 5 When the *Command Status* changes to *Complete*, click *Refresh*.
The status icon of the Identity Server should turn green.

2.1.3 Restarting the Access Gateway Service Provider

The service provider should be restarted whenever you enable or modify logging on the Identity Server. Otherwise, it is automatically restarted when a configuration change requires a restart.

To stop and start the Access Gateway service provider:

- 1 In the Administration Console, click *Access Manager > Access Gateways > [Name of Access Gateway]*.

- 2 Click *Actions > Restart Service Provider*, then click *OK*.

In a few seconds, the status icon of the Access Gateway should turn green.

2.1.4 Starting the Access Gateway Service Provider

When an Access Gateway is removed from a group configuration, the embedded service provider is stopped. It should remain stopped until you have reconfigured the Access Gateway. When you have finished the reconfiguration, you should start the embedded service provider.

The service provider should be restarted whenever you enable or modify logging on the Identity Server. Otherwise, it is automatically restarted when a configuration change requires a restart.

To start the Access Gateway service provider:

- 1 In the Administration Console, click *Access Manager > Access Gateways > [Name of Access Gateway]*.
- 2 Click *Actions > Start Service Provider*, then click *OK*.

In a few seconds, the status icon of the Access Gateway should turn green.

2.1.5 Stopping the Access Gateway Service Provider

To stop the Access Gateway service provider:

- 1 In the Administration Console, click *Access Manager > Access Gateways > [Name of Access Gateway]*.
- 2 Click *Actions > Stop Service Provider*, then click *OK*.

In a few seconds, the status icon of the Access Gateway should turn red.

2.1.6 Restarting the Access Gateway

Restarting the Access Gateway makes all protected resources unavailable until the Access Gateway returns to a server status of green.

To restart the Access Gateway:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Name of Access Gateway*.
- 2 Click *Restart*.

In a few minutes, the status icon of the Access Gateway should turn green.

2.1.7 Scheduling a Restart of the Access Gateway

Restarting the Access Gateway makes all protected resources unavailable until the Access Gateway returns to a server status of green. Scheduling this event allows you to pick the best time for your resources to be momentarily unavailable.

- 1 In the Administration Console, click *Access Manager > Access Gateways > [Name of Access Gateway] > Actions > Schedule Restart*.
- 2 Fill in the following fields:

Name Scheduled Command: (Required) Specifies a name for this scheduled command. This name is used in log and trace files.

Type: Displays the type of command that is being scheduled, such as *Access Gateway Shutdown*, *Access Gateway Restart*, *Access Gateway Upgrade*, *Device Configuration*.

Group: If you are scheduling a command for a group of Access Gateways, this option displays the name of the group.

Description: (Optional) Provides a field to describe the reason for the command.

Date & Time: The drop-down menus allow you to select the day, month, year, hour, and minute when the command should execute.

- 3 Click *OK*.

2.1.8 Stopping the Access Gateway

You should stop the Access Gateway only when you plan on turning off the power or to configure boot options for troubleshooting. After you have stopped the Access Gateway, you must have physical access to the machine to start it.

In the Administration Console, click *Access Manager > Access Gateways > [Name of Access Gateway] > Shutdown*.

The machine is physically turned off. Before you start the Access Gateway again, you can modify the boot options on a NetWare Access Gateway. For information about these boot options, see [Section 42.3.1, “Additional Options during the Boot Process,” on page 507](#).

2.1.9 Scheduling the Shutdown of the Access Gateway

You should stop the Access Gateway only when you plan on turning off the power or to configure boot options for troubleshooting. After you have stopped the Access Gateway, you must have physical access to the machine to start it. Scheduling this event allows you to pick the best time for the Access Gateway to be unavailable.

- 1 In the Administration Console, click *Access Manager > Access Gateways > [Name of Access Gateway] > Actions > Schedule Shutdown*.

- 2 Fill in the following fields:

Name Scheduled Command: (Required) Specifies a name for this scheduled command. This name is used in log and trace files.

Type: Displays the type of command that is being scheduled, such as *Access Gateway Shutdown*, *Access Gateway Restart*, *Access Gateway Upgrade*, *Device Configuration*.

Group: If you are scheduling a command for a group of Access Gateways, displays the name of the group.

Description: (Optional) Provides a field to describe the reason for the command.

Date & Time: The drop-down menus allow you to select the day, month, year, hour, and minute when the command should execute.

- 3 Click *OK*.

The machine is turned off when the scheduled command executes.

Before you start the Access Gateway again, you can modify the boot options on a NetWare Access Gateway. For information about these boot options, see [Section 42.3.1, “Additional Options during the Boot Process,”](#) on page 507.

2.2 Backing Up and Restoring the Configuration Store

In addition to installing secondary consoles to protect the configuration store, you can also back it up so you can restore it. You cannot restore data from a previous version of Access Manager to a new version.

Backup and restore utilities are scripts that are run from the command line, and they allow you to back up and restore your Access Manager configuration. An additional script allows you to export your configuration so Novell® Technical Support can help diagnose possible configuration problems.

Before running these scripts, verify the following:

- ◆ You have `root` access.
- ◆ You have changed the directory to `/opt/novell/devman/bin`.

This section discusses the following topics:

- ◆ [Section 2.2.1, “How The Backup and Restore Process Works,”](#) on page 31
- ◆ [Section 2.2.2, “Running the Backup,”](#) on page 32
- ◆ [Section 2.2.3, “Restoring a Configuration Store,”](#) on page 32
- ◆ [Section 2.2.4, “Running the Diagnostic Configuration Export,”](#) on page 33

2.2.1 How The Backup and Restore Process Works

Default Parameters

All of the scripts call the `getparams.sh` script to request the parameters from the user. The `defbkparm.sh` script is created by the Access Manager installation. It contains default parameters for several of options required by the underlying backup and restore utilities. If the entries in this file are commented out, the user is prompted for the additional parameters.

Certtool.jar

`Certtool.jar` is a key certificate utility for eDirectory built on top of the same interfaces that the Access Manager certificate management features use. It provides some features similar to the Java `keytool` utility. It must run on a computer that had eDirectory installed, or at least NPki. The basic command line to invoke the tool is:

```
java -Djava.library.path=/opt/novell/lib -jar certtool.jar -h
```

The `-h` option produces help listing of command line options.

2.2.2 Running the Backup

- 1 Change to the `/opt/novell/devman/bin` directory.
- 2 Run the following command from root: `./ambkup.sh`.
- 3 Enter the Access Manager administration user ID.
- 4 Enter the Access Manager administration password.
- 5 Re-enter the password for verification.
- 6 Enter a password for encrypting and decrypting private keys, then re-enter for verification.
You must use the same password for both backup and restore.
- 7 Press Enter.

During backup, the system creates two files (`.ldif` and `.zip`) that you can identify by the date, time, and server name. For example:

- ◆ Configuration information: `server3_20060602_0729.ldif`
- ◆ Certificates and trusted roots: `server3_20060602_0729.zip`

The backup produces two files: LDIF and ZIP.

The LDIF File: The LDIF file is created by the eDirectory ICE utility. Everything in the `OU=accessManagerContainer,O=novell` container is exported to the LDIF file.

When ICE creates the LDIF file, it records all the objects inside a container, and then records the container object. This can be a problem when importing the LDIF file, because records cannot be added to a container before the container is created. Therefore, when the backup script ICE runs, it creates a file that has `_pre` appended to the name of the file. After ICE creates this file, the backup script runs a program that reverses the order of the records in the LDIF file. This program is `ldifReverse` and it exists in the same directory where the backup and restore scripts are located. When `ldifReverse` runs, it reads the records from the file ending with `_pre` produced by ICE, and creates the final LDIF file which has the `.ldif` extension.

The ZIP File: The backup script uses `certtool.jar` to create a ZIP file of the certificate information (see “[Certtool.jar](#)” on page 31). This jar file uses the same API that Access Manager’s certificate management features use. This `.zip` file contains the following:

- ◆ The configurations store’s CA key.
- ◆ The certificates contained in the configuration store.
- ◆ The trusted roots in the `trustedRoots` container in `accessManagerContainer`. The trusted roots are backed up in both the LDIF file and the ZIP file. They are added to the ZIP file so that the ZIP file has the complete certificate-related configuration.

IMPORTANT: The backup utility prompts you for a location to store the backup files, so that they are not erased if you uninstall the product. The default location is the logged-in user’s home directory.

2.2.3 Restoring a Configuration Store

Restore adds the records produced by the backup to the configuration store.

- 1 Ensure that the two files created during the backup process are accessible.

- 2 Change to the `/opt/novell/devman/bin` directory.
- 3 Run the following command from root: `./amrestore.sh`.
- 4 Enter the Access Manager administration user ID.
- 5 Enter the Access Manager administration password.
- 6 Enter the name of the backup file. Do not include the extension `.ldif` or `.zip`.
- 7 Press Enter.

In order to accurately restore the configuration store to the state it was in when the backup occurred, the `accessManagerContainer` is removed from the configuration store. Then it is restored from the backup file.

Because it first removes the `accessManagerContainer` container, the restore process ensures a successful restore before it proceeds.

- ◆ Restore checks to make sure that the LDIF file looks like a valid backup of Access Manager. This is done by calling the program `ldifReverse` with the `-t` option. This causes `ldifReverse` to test the LDIF file to see that it is an LDIF file, and that it begins with the `accessManagerContainer` record that will be restored.
- ◆ Creates a temporary backup file called `__recover__.ldif`. This file can be used as a backup file to restore the configuration store back to the state it was in before the restore was attempted.

After taking these precautionary steps, the restore process imports the backed up LDIF file using the `eDirectory ICE` tool, then it uses `certtool.jar` restore the certificates.

If you have secondary Administration Console installed, you must restart Tomcat (`/etc/init.d/novell-tomcat4 restart`) in order to re-establish LDAP connections to the primary Administration Console.

2.2.4 Running the Diagnostic Configuration Export

To create an `.ldif` file that you can export for diagnostic purposes:

- 1 Change to the `/opt/novell/devman/bin` directory.
- 2 Run the following command from root: `./amdiagcfg.sh`.
- 3 Enter the Access Manager administration user ID.
- 4 Enter the Access Manager password.
- 5 Re-enter the password for verification.
- 6 Press Enter.

The diagnostic configuration export utility is almost identical to the backup utility with two differences: the ZIP file is not created, and the final LDIF file is scanned to have passwords removed. Passwords are blanked out by a program called `Strippasswd`.

`Strippasswd` removes occurrences of passwords in the LDIF file, replacing them with empty strings. If you look at the LDIF file, you will see that password strings are blank. You might see occurrences within the file or text that looks similar to `password="String"`. These are not instances of passwords, but rather definitions that describe passwords as string types.

The LDIF file can then be sent to Novell Support for help in diagnosing configuration problems.

2.3 Installing Secondary Versions of the Administration Console

The Administration Console contains an embedded version of eDirectory, which contains all the configuration information for the Access Manager. It also contains a server communications module, which is in constant communication with the Access Manager modules. If the Administration Console goes down and you have not installed any secondary consoles, your Access Manager components also go down and your protected resources become unavailable.

You can create fault tolerance by installing up to two secondary consoles. We highly recommend that you install at least one secondary console. For a secure configuration, secondary consoles must be installed on the same network as the primary console. The administration consoles should not be required to use a router to communicate with each other.

As long as the primary console is running, you can use the secondary consoles to manage certificates and component configuration changes. When the primary console goes down, you can use the secondary consoles to make configuration changes, but certificate management must wait until the primary console is running. When the primary console is down, the Access Manager components use the secondary consoles to access to their configuration information. As soon as the primary console comes back online, the components revert to using the primary machine, but they continue to accept commands from the secondary consoles.

After installing a secondary console, you might have to wait from 30 to 60 minutes before using it. The components query the primary console hourly for information about available consoles, and they reject commands from a console that is not in their approved list. You can force the components to recognize the secondary console by restarting the Integration Agent on each Identity Server, Linux Access Gateway, and Linux J2EE Agent with the following command:

```
/etc/init.d/novell-jcc restart
```

For the NetWare Access Gateway, you need to wait until the primary console informs it of the new secondary console.

When installing the Administration Console, you must declare whether the installation is for a primary console or a secondary console. To install a secondary console, answer No to the following prompt:

```
Is this the primary administration server in a failover group?
```

When prompted, enter the IP address of the primary console. For installation instructions, see [“Installing the Access Manager Administration Console”](#) in the *Novell Access Manager 3.0 Installation Guide*.

The primary console must be used for the following tasks:

- ♦ **New Device Installation:** The primary console must be running when you install new devices such as another Access Gateway or SSL VPN server.
- ♦ **Backup and Restore:** Backup and restore must be run on the primary console. Once the restore has completed, you must restart Tomcat on all secondary consoles. Use the following command:

```
/etc/init.d/novell-tomcat4 restart
```

For more information about backup and restore, see [Section 2.2, “Backing Up and Restoring the Configuration Store,”](#) on page 31.

2.4 Converting your Secondary Console into your Primary Console

In order for the secondary administration console to be converted into the primary administration console, a recent backup of the administration console must be available. This is necessary in order to restore the certificate authority (CA). This procedure involves changing the master replica and restoring the CA certificates.

WARNING: Perform these steps only if the primary administration console cannot be restored.

If the failed server holds a master replica of any partition, you must use DSRepair to designate a new master replica on a different server in the replica list.

This conversion includes the following tasks:

- ♦ [Section 2.4.1, “Changing the Master Replica,” on page 35](#)
- ♦ [Section 2.4.2, “Restoring CA Certificates,” on page 36](#)
- ♦ [Section 2.4.3, “Deleting Objects from the eDirectory Configuration Store,” on page 36](#)
- ♦ [Section 2.4.4, “Additional Procedures,” on page 36](#)

2.4.1 Changing the Master Replica

- 1 At the console of one of the servers that shared a replica with the failed server, load DSRepair with the switch that lets you access the advanced options.
 - ♦ Linux: Use the `-Ad` switch (`/opt/novell/eDirectory/bin/ndsrepair -P -Ad`)

For more information on how to run DSRepair with advanced options using the `-a` or `-Ad` switches, see [Advanced DSRepair Options \(http://www.novell.com/documentation/edir88/edir88/data/aflm3p7.html#aflm3p7\)](http://www.novell.com/documentation/edir88/edir88/data/aflm3p7.html#aflm3p7) on the [eDirectory 8.8 Documentation Web site \(http://www.novell.com/documentation/edir88/treetitl.html\)](http://www.novell.com/documentation/edir88/treetitl.html)

WARNING: If you use DSRepair with `-a` or `-Ad`, some of the advanced options can cause damage to your tree. For more information on these options, refer to the [Novell Support Web site, Solution 2938493 \(http://support.novell.com/servlet/tidfinder/2938493\)](http://support.novell.com/servlet/tidfinder/2938493).

- 2 Select *Replica and Partition Operations*.
- 3 Select the partition you want to edit, so you can remove the failed server from the replica ring of that partition.
- 4 Select *View Replica Ring* to see a list of servers that have replicas of the partition.
- 5 (Conditional) If the failed server held the master replica, select another server to hold the master by selecting *Designate This Server As the New Master Replica*.

The replica ring now has a new master replica. All replicas participating in the ring are notified that there is a new master.
- 6 Wait for the master replica to be established. Make sure the other servers in the ring acknowledge the change before proceeding.
- 7 Go back to *View Replica Ring*. Select the name of the failed server, then select *Remove This Server from the Replica Ring*.

If you have not loaded DSRepair with `-a` or `-Ad` (depending on the platform) for advanced options, you will not see this option in the list.

WARNING: Make sure you do not do this if the failed server is designated as the master replica. You can see this information in the list of servers in the ring. If it is the master, designate a different server as the master as noted in [Step 5](#). Then, come back to this step and remove the failed server from the replica ring.

- 8 Log in as Admin.
- 9 After reading the explanation message, enter your agreement to continue.
- 10 Exit DSRepair.
All servers participating in that replica ring are notified.

2.4.2 Restoring CA Certificates

Run `aminst-certs.sh` from the `/opt/novell/devman/bin` directory in Linux.

2.4.3 Deleting Objects from the eDirectory Configuration Store

Several objects representing the failed primary administration console in the configuration store must be manually deleted.

- 1 Log in to the new Administration Console and click the *View Objects* icon at the top of the page.
- 2 Expand the *novell* container.
A list of objects is displayed. There should be 13 objects to delete. Delete only the objects that have the name of the failed primary administration console in them.
- 3 To delete an object, click on the object and click *Delete Object*.
- 4 In the *novell* container, expand the *accessManagerContainer* object.
- 5 Expand *VCDN_Root > PartitionsContainer > Partition > ROMAServerContainer*.
- 6 Delete the object that has the name of the failed primary Administration Console.
The object's name is the hostname of your Linux server during installation. The default name is *linux*.
- 7 Click the *Roles and Tasks* icon at the top of the page to return to the Access Manager menu.

2.4.4 Additional Procedures

You might need to perform additional steps for the following components:

- ♦ [“Third Administration Console” on page 37](#)
- ♦ [“NetWare Access Gateways” on page 37](#)
- ♦ [“Linux Access Gateways” on page 37](#)
- ♦ [“SSL VPN” on page 38](#)
- ♦ [“Identity Servers” on page 39](#)
- ♦ [“J2EE Agents” on page 39](#)

Third Administration Console

If you installed a third Administration Console used for failover, you must manually perform the following steps on that server:

- 1 Edit the `vcdn.conf` file at `/opt/novell/devman/share/conf/`.

In the file you will find a section of XML that looks like:

```
<vcdnPrimaryAddress>10.1.1.1</vcdnPrimaryAddress>
```

(Where 10.1.1.1 represents the failed primary Administration Console IP address.) You must change this IP address to the new Administration Console's address that you are converting to the primary server.

- 2 Restart the Administration Console by typing `/etc/init.d/novell-tomcat4 restart` on the command line interface.

NetWare Access Gateways

For each NetWare Access Gateway server imported in the Administration Console, you must perform the following steps:

- 1 Enter debug mode on the server by typing `debug` and inputting the password `proxydebug`.
- 2 Go to the NetWare prompt by typing `Ctrl-Esc` if you are using the keyboard. If you are remote via ssh, type `Ctrl-Z`, then select screen 1.
- 3 Type `java -show`, and note the process ID next to `JCCServerImpl`.
- 4 Type `java -kill###`, where `###` represents the process ID.
- 5 Edit the `ecc.cfg` file by typing `edit sys:\etc\proxy\ecc.cfg`.
Find the section labeled `[jccsettings]` and change the IP address of the line labeled `serveraddress` to the new primary Administration Console.
- 6 Edit the `settings.properties` file by typing: `edit sys:\jcc\conf\settings.properties`.
Change the IP address list labeled `remotemgmtip`, removing the IP address of the failed Administration Console. Ensure that the address of the new primary server is listed.
- 7 Restart the server by typing `appboot` at the NetWare prompt.

Linux Access Gateways

For each Linux Access Gateway in the failed Administration Console, you must perform the following steps:

- 1 Log in as the root user.
- 2 Open a terminal window and shut down all services by typing the following commands:
 - ♦ `/etc/init.d/novell-jcc stop`
 - ♦ `/etc/init.d/novell-tomcat4 stop`
 - ♦ `/etc/init.d/novell-sslvpn stop`
- 3 Edit the `config.xml` file by typing: `vi /var/novell/cfgbd/.current/config.xml`.
Type `/Remote`, then press Enter.

In the *IPv4Address* field, change the IP address from the failed Administration Console to the new primary Administration Console address.

- 4 Type `:wq!` to save and exit.
- 5 Edit the `settings.properties` file by typing `vi /opt/novell/devman/jcc/conf/settings.properties`.
Change IP address in the *remotemgmtip* list by removing the IP address of the failed Administration Console. Ensure that the address of the new primary server is listed.
- 6 Type `:wq!` to save and exit.
- 7 Start all services by typing the following commands:
 - ♦ `/etc/init.d/novell-jcc start`
 - ♦ `/etc/init.d/novell-tomcat4 start`
 - ♦ `/etc/init.d/novell-vmc start`

SSL VPN

For each SSL VPN component in the Administration Console, you must perform the following steps:

- 1 Log in as the root user.
 - 2 Open a terminal window and shut down all services by typing the following commands:
 - ♦ `/etc/init.d/novell-jcc stop`
 - ♦ `/etc/init.d/novell-tomcat4 stop`
 - ♦ `/etc/init.d/novell-sslvpn stop`
 - 3 Edit the `config.xml` file by typing: `vi /etc/opt/novell/sslvpn/config.xml`.
 - 4 Type `/DeviceManagerAddress`, then press Enter.
 - 5 Change the IP address to that of the new primary Administration Console.
 - 6 Type `:wq!` to save and exit.
 - 7 Edit the `settings.properties` file by typing:
`vi /opt/novell/devman/jcc/conf/settings.properties`
Change the IP address list labeled *remotemgmtip* by removing the IP address of the failed Administration Console. Ensure that the address of the new primary Administration Console is listed.
 - 8 Type `:wq!` to save and exit.
 - 9 Start all services by typing the following commands:
 - ♦ `/etc/init.d/novell-jcc start`
 - ♦ `/etc/init.d/novell-tomcat4 start`
 - ♦ `/etc/init.d/novell-sslvpn start`
- If the SSLVPN is no longer functioning, restart the Linux server by typing `reboot`.

Identity Servers

For each Identity Server in the Administration Console, you must perform the following steps:

- 1 Log in as the root user.
- 2 Open a terminal window and shut down all services by typing: `/etc/init.d/novell-jcc stop`, then `/etc/init.d/novell-tomcat4 stop`.
- 3 Edit the `settings.properties` file by typing:
`vi /opt/novell/devman/jcc/conf/settings.properties`
Change the IP address list labeled *remotemgmtip* by removing the IP address of the failed Administration Console. Ensure that the address of the new primary Administration Console is listed.
- 4 Type `:wq!` to save and exit.
- 5 Start the services by typing `/etc/init.d/novell-jcc start`, then `/etc/init.d/novell-tomcat4 start`

J2EE Agents

For each J2EE agent in the Administration Console, you must perform the following steps:

- 1 Log in as the root user.
- 2 Open a terminal window and shut down all services by typing: `/etc/init.d/novell-jcc stop`.
- 3 Edit the `settings.properties` file by typing:
`vi /opt/novell/devman/jcc/conf/settings.properties`
Change the IP address list labeled *remotemgmtip* by removing the IP address of the failed Administration Console. Ensure that the address of the new primary Administration Console is listed.
- 4 Type `:wq!` to save and exit.
- 5 Start the services by typing `/etc/init.d/novell-jcc start`.

2.5 Multiple Sessions, Multiple Administrators

The Administration Console has been designed to warn you when another administrator is making changes to an Access Manager device (Identity Server, Access Gateway, SSL VPN, or J2EE Agent). The policy options in the interface have no such protection. If you have more than one administrator that is configuring and managing policies, we suggest that you create separate policy containers for each administrator. This allows them to configure policies at the same time without session conflicts. See [Section 28.2, “Managing Policy Containers,” on page 312](#).

You should not start multiple sessions to the Administration Console with the same browser on a workstation. Browser sessions share settings which can result in problems when you apply changes to configuration settings. However, if you are using two different brands of browsers simultaneously, such as Internet Explorer and Firefox, it is possible to avoid the session conflicts.

Changing the IP Address of Access Manager Devices

3

The following sections explain how to change the IP address on the following devices:

- ◆ [Section 3.1, “Changing the IP Address of the Administration Console,” on page 41](#)
- ◆ [Section 3.2, “Changing the IP Address of an Identity Server,” on page 41](#)
- ◆ [Section 3.3, “Changing the IP Address of the Access Gateway,” on page 42](#)

NOTE: Changing the IP address of an SSL VPN component is not recommended.

3.1 Changing the IP Address of the Administration Console

We recommend that you install the Administration Console with the IP address that it will always use because all of the devices that import into the Administration Console use this address to establish secure communication with the Administration Console.

The only tested method of changing the IP address so that all other devices trust the Administration Console is to install a secondary console with the new IP address and then promote the secondary console to be the primary console. Remember to change the IP addresses of all components pointing to the new Administration Console.

See the following sections:

- ◆ [Section 2.3, “Installing Secondary Versions of the Administration Console,” on page 34](#)
- ◆ [Section 2.4, “Converting your Secondary Console into your Primary Console,” on page 35](#)

Converting a secondary console into a primary console is not a simple task. The task was designed as a disaster recovery solution when the primary console is no longer available. It is not a simple configuration change.

3.2 Changing the IP Address of an Identity Server

These instructions assume that your Identity Server and Administration Console are not on the same machine. If they are on the same machine, see [Section 3.1, “Changing the IP Address of the Administration Console,” on page 41](#).

To move a machine or change the IP address for the Identity Server:

- 1 In the Administration Console, click *Access Manager > Identity Servers*.
- 2 Click the server name.
- 3 On the General page, click *Edit*.
- 4 Specify the new IP address in the *Management IP Address* field and, if necessary, a port.

- 5 Click *OK*, then click *Close*.
- 6 In Linux, open the console shell and stop the server communication service by using the following command:

```
/etc/init.d/novell-jcc stop
```
- 7 Using YaST, change the IP address on the physical Linux server hosting the Identity Server.
- 8 At the console shell, access the `/opt/novell/devman/jcc` directory, then enter the following command:

```
conf/Configure.sh
```
- 9 When you are prompted for the local listener IP address, enter the new IP.
- 10 When you are prompted for the administration server IP, enter the same IP that you used during the initial installation.
- 11 Follow the prompts and accept the defaults for ports and admin user.
- 12 At the console shell, start the server communication service using the command:

```
/etc/init.d/novell-jcc start
```
- 13 Restart the Identity Server application on the Servers page.

3.3 Changing the IP Address of the Access Gateway

If you need to change the IP address of the Access Gateway machine, you need to configure the Access Gateway for this change. This is especially significant when the Access Gateway machine has only one IP address.

IMPORTANT: The new IP address must be configured in the Administration Console before you change it on the Access Gateway. If you change in on the Access Gateway first, the Administration Console will not trust the Access Gateway and re-establish communication.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Adapter List*.
- 2 (Conditional) If the machine belongs to a group, select the Access Gateway from the *Group Member* list.
- 3 In the *Adapter eth0* section, select the subnet mask that contains the IP address.
If the new IP address is in a different subnet mask, see “[Moving Staged Components](#)” in the *Novell Access Manager 3.0 Installation Guide*.

When you select the subnet mask, the Adapter page appears.

Adapter eth0

Subnet: 10.10.159.206

Subnet Mask: * 255.255.0.0

IP Address List *

New... | Delete | Change IP Address...

<input type="checkbox"/>	IP Addresses
<input type="checkbox"/>	10.10.159.206

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

OK Cancel

- 4** Select the old IP address, click *Change IP Address*, specify the new IP address, then click *OK*
This option changes all configuration instances of the old IP address to the new IP address. For example, any reverse proxies that have been assigned the old IP address as a listening address are modified to use the new IP address as the listening address.
- 5** To save these changes, click *Configuration Panel*, then click *Apply Changes*.
The configuration changes are applied to the Access Gateway machine.
- 6** (Conditional) If you are physically moving the machine, move it before completing the rest of these steps.
- 7** Check the IP address that the Administration Console uses for managing the Access Gateway. Click *Access Manager > Access Gateways > [Name of Access Gateway] > Edit*.
- 8** If the old IP address is listed as the *Management IP Address*, select the new IP address. If your Access Gateway has multiple IP addresses, select the one that you want the Administration Console to use for communication with the Access Gateway.
The port should only be modified if there is another device on the Access Gateway that is using the default port of 1443.
- 9** If the name of the Access Gateway is the old IP address, modify the *Name* option.
- 10** Click *OK*.
The Administration Console uses the configured IP address to find the Access Gateway.
- 11** Reboot the Access Gateway. Click *Access Gateways > [Name of Access Gateway] > Restart*.

Maintaining an Identity Server

4

Server maintenance involves tasks that you perform after you have configured the server. Maintenance includes monitoring server and statistics, configuring Identity Server logging, replacing certificates, and so on.

- ◆ [Section 4.1, “Managing in Identity Server,” on page 45](#)
- ◆ [Section 4.2, “Editing Server Details,” on page 46](#)
- ◆ [Section 4.3, “Security,” on page 46](#)

For information about server health, see [Section 36.2, “Monitoring the Health of an Identity Server,” on page 444](#).

For information about configuring the Identity Server, see [Part II, “Novell Identity Server Configuration,” on page 51](#).

4.1 Managing in Identity Server

The Identity Servers page is the starting point for managing Identity Servers. Most often, you use this page to stop and start servers, and to assign servers to Identity Server configurations. An Identity Server cannot operate until you have assigned it to an Identity Server configuration.

- 1 In the Administration Console, click *Access Manager > Identity Servers*.

Server	Server Status	Alerts	Command Status	Statistics	Configuration Assignment
<input type="checkbox"/> 10.10.157.30		0	Complete	View	IDP_A

- 2 On the Servers tab, you can perform the following functions by clicking the server’s check box, then clicking any of the following options:

Refresh: Refreshes the server list.

Start: Starts the selected server. (See [Section 2.1, “Starting and Stopping Access Manager Components,” on page 27](#).)

Stop: Stops the selected server.

Actions: Enables you to perform the following tasks:

- ◆ **Add to configuration:** Enables you to assign a server to a configuration. See [Assigning an Identity Server to a Configuration](#) for more information.
- ◆ **Remove from configuration:** Enables you to remove one or more servers from a configuration. See [Removing a Server from a Configuration](#) for more information.
- ◆ **Delete:** Deletes the selected server.

IMPORTANT: The system does not allow you to delete an Identity Server that is started. All other servers must be manually uninstalled by running the uninstall script on the server.

This page also displays links in the following columns:

Column	Description
Server	Lists the Identity Server name.
Server Status	Lists the status of each configuration. If <i>Update Servers</i> is displayed, click the link to perform the update. Updating an Identity Server configuration reloads the Identity Server configuration data without stopping the server.
Alerts	Displays the Alerts page where you can monitor and acknowledge server alerts.
Command Status	Displays the Command Status page.
Statistics	Displays the Server Statistics page and allows you to view the server statistics. See Section 35.1, "Monitoring Identity Server Statistics," on page 427.
Configuration Assignment	Lists the Identity Server configuration to which this server belongs. An identity server can belong to multiple configurations.

4.2 Editing Server Details

You can edit server details, such as the server name and port. You can also access the other server management tabs from this page.

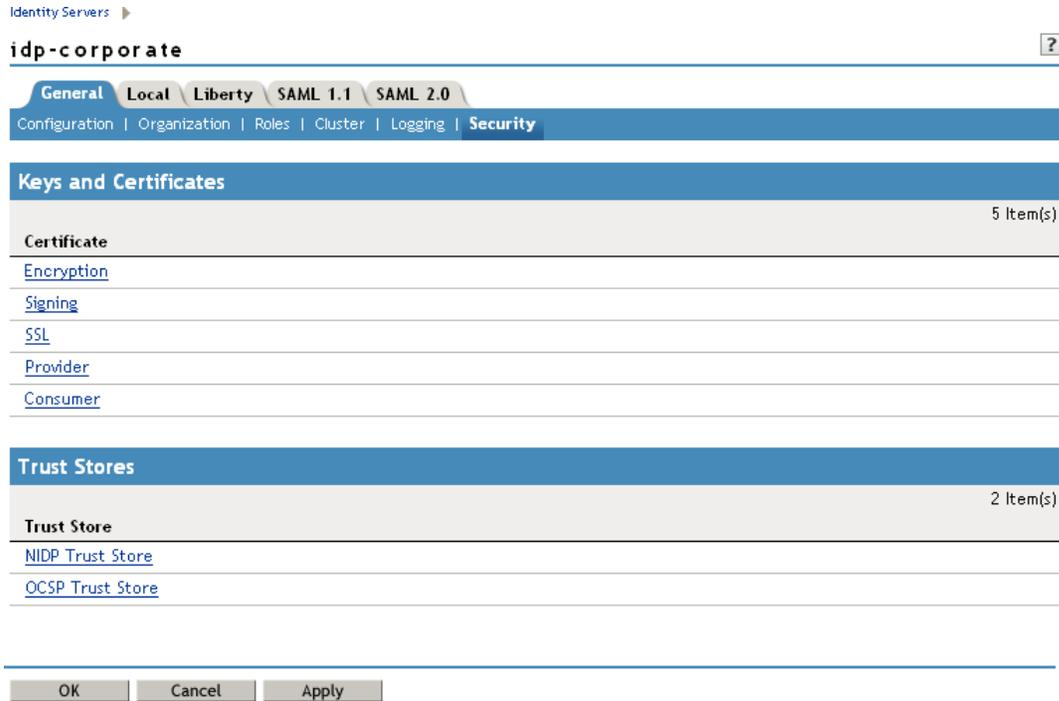
- 1 In the Administration Console, click *Access Manager > Identity Servers*, then click the server name.
- 2 Click *Edit*.
- 3 Fill in the following fields as necessary:
 - Edit:** Displays the Server Details Edit page.
 - Name:** The name of the Identity Server. Names must be alphanumeric and can include spaces, hyphens, and underscores.
 - Management IP Address:** The IP address of the Identity Server. Changing server IP addresses is not recommended and causes the server to stop reporting. See [Section 3.2, "Changing the IP Address of an Identity Server," on page 41.](#)
 - Port:** The Identity Server port.
 - Location:** The location of the Identity Server.
 - Description:** A description of the Identity Server.
- 4 Click *OK*.

4.3 Security

You can view the private keys, CA certificates, and certificate containers associated with the Identity Server configuration. Primarily, you use the Security page to add and replace CA certificates as

necessary and to perform certificate management tasks, such as adding trusted root certificates to a trust store, or changing passwords for the containers.

- 1 In the Administration Console, click *Access Manager > Identity Server > Setup > [Configuration] > Security*.



- 2 Click any of the following links:

Encryption: Displays the NIDP-encryption certificate keystore. The encryption certificate is used to encrypt specific fields or data in the assertions. Click *Replace* the test-encryption certificate.

Signing: Displays the NIDP-signing certificate keystore. The signing certificate is used to sign the assertion or specific parts of the assertion. Click *Replace* to replace the test-signing certificate.

SSL: (Required) Displays the NIDP-connector keystore. Click this link to access the keystore and replace the test-connector certificate.

Provider: Displays the NIDP-provider keystore. Click this link to access the keystore and replace the test-provider certificate used by the identity provider.

Consumer: Displays the NIDP-consumer keystore. Click this link to access the keystore and replace the test-consumer certificate used by the identity consumer (service provider).

Keystore: -connector ?

Keystore name: -connector
Keystore type: Java
NIDP configuration name: TradingCo-ids

NIDP Configuration Members' Keystores			
Change Password...			
<input type="checkbox"/>	Keystore Name	Type	Device
<input type="checkbox"/>	Connector	Java	10.10.159.206

Certificates

1 item(s)

Replace
Certificate: <input type="text"/>
Alias(es): tomcat
<input type="button" value="OK"/> <input type="button" value="Cancel"/>

2a Browse to locate the certificate, then click *OK*.

- 3** To manage trust stores associated with the Identity Server, click either of the following links on the Security page:

NIDP Trust Store: The trusted root certificate container for CA certificates associated with the Identity Server. Click this link to access the trust store, where you can change the password or add trusted roots to the container. Liberty and SAML 2.0 protocol messages that are exchanged between identity and service providers often need to be digitally signed. A provider uses the signing certificate included with the metadata of a trusted provider to validate signed messages from the trusted provider. For protocol messages to be exchanged between providers using SSL, each provider must trust the SSL Certificate Authority (CA) of the other provider. Well-known CAs should already be trustable, but for those that are not, you must import the CA for the other provider. Failure to do so causes numerous system errors.

OCSP Trust Store: The trust store for OCSP certificates. Online Certificate Status Protocol is a method used for checking the revocation status of a certificate. For this feature, you must set up an OCSP server. The Identity Server sends an OCSP request to the OCSP server to determine if a certain certificate has been revoked. The OCSP server replies with the revocation status. If this revocation checking protocol is used, the Identity Server does not cache or store the information in the reply, but sends a request every time it needs to check the revocation status of a certificate. The OCSP reply is signed by the OCSP server. To verify that it was signed by the correct OCSP server, the OCSP server certificate needs to be added to this trust store. The OCSP server certificate itself is added to the trust store, not the CA certificate.

Trust Store: -truststore ?

Trust store name: -truststore
Trust store type: Java
NIDP configuration name: TradingCo-ids

NIDP Configuration Members' Trust Stores

[Change Password...](#)

<input type="checkbox"/>	Trust Store Name	Type	Device
<input type="checkbox"/>	NIDP Trust Store	Java	10.10.159.206

Trusted Roots

[Add...](#) | [Remove](#) |

Trusted Root

[configCA](#)

Auto-Import From Server X

Server IP Address:

Server Port:

1 item(s)

3a Specify the server IP address and port.

The auto-import displays the certificate chain, which you can select for import.

3b Click *OK*, then click *Close*.

4 Restart Tomcat.

The system prompts you with a dialog box to restart Tomcat. This is necessary whenever security changes are made to the Identity Server.

For more information about enabling security for a basic Access Manager configuration, see [“Enabling SSL Communication”](#) in the *Novell Access Manager 3.0 Setup Guide*.

For additional information about managing certificates, see [Part V, “Security and Certificate Management,”](#) on page 285.

Novell Identity Server Configuration



In Access Manager, the Identity Server is responsible for authenticating users, building the user's role information and distributing it to the various components. It also serves as the central point for components that request identity information.

This section of the Administration Guide describes the following topics:

- ◆ [Chapter 5, “Configuring an Identity Server,” on page 53](#)
- ◆ [Chapter 6, “Defining Reusable Configuration Settings,” on page 65](#)
- ◆ [Chapter 7, “Configuring Local Authentication,” on page 71](#)
- ◆ [Chapter 8, “Configuring Trusted Providers,” on page 87](#)
- ◆ [Chapter 9, “Configuring User Authentication and Federation,” on page 99](#)
- ◆ [Chapter 10, “Configuring Communication Profiles,” on page 113](#)
- ◆ [Chapter 11, “Configuring Liberty Web Services,” on page 115](#)

For information about Identity Server maintenance tasks, such as auditing, logging, and health monitoring, see [Part I, “System Management,” on page 19](#).

This section of the administration guide is intended to help you understand and configure the Identity Server for authentication, and includes advanced topics. It is recommended that you first become familiar with the information in the *Novell Access Manager 3.0 Setup Guide*.

The *Novell Access Manager 3.0 Setup Guide* is intended to familiarize you with Access Manager and helps you understand how to perform a basic Identity Server configuration, set up a resource protected by an Access Gateway, and configure SSL. The basic setup and the administration guides are designed to work together, and important information and setup steps are not necessarily repeated in both places.

Configuring an Identity Server

5

After you log in to the Administration Console, click *Access Manager > Identity Servers*. The system displays the installed server.



The screenshot shows the 'Identity Servers' page in an administration console. At the top, there are tabs for 'Servers' and 'Setup'. Below the tabs, there are links for 'Refresh', 'Start', 'Stop', and 'Actions'. On the right, it says '1 Item(s)'. Below this is a table with columns: 'Server', 'Server Status', 'Alerts', 'Command Status', 'Statistics', and 'Configuration Assignment'. There is one row with the following data: '10.10.157.30', a yellow question mark icon, '0', 'Complete', 'View', and 'Not Configured'.

Server	Server Status	Alerts	Command Status	Statistics	Configuration Assignment
<input type="checkbox"/> 10.10.157.30	?	0	Complete	View	Not Configured

At this point, the Identity Server is in an unconfigured state and is halted. It remains in this state and cannot function until you create an Identity Server configuration and assign the Identity Server to the new configuration. The configuration defines how the Identity Server functions in an Access Manager configuration. In an Identity Server cluster, multiple servers must use the same configuration.

- ◆ [Section 5.1, “Creating an Identity Server Configuration \(Advanced Options\),” on page 53](#)
- ◆ [Section 5.2, “Assigning an Identity Server to a Configuration,” on page 59](#)
- ◆ [Section 5.3, “Removing a Server from a Configuration,” on page 59](#)
- ◆ [Section 5.4, “Modifying the Base URL,” on page 59](#)
- ◆ [Section 5.5, “Enabling Role-Based Access Control,” on page 60](#)
- ◆ [Section 5.6, “Clustering Identity Servers,” on page 60](#)
- ◆ [Section 5.7, “Configuring Secure Communication on the Identity Server,” on page 63](#)

Additional Identity Server configuration topics for authentication include:

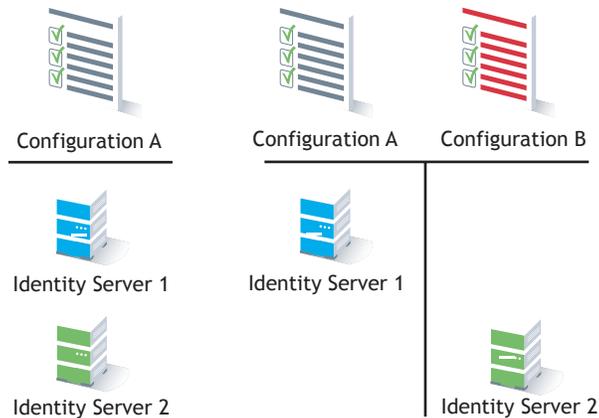
- ◆ [Chapter 6, “Defining Reusable Configuration Settings,” on page 65](#)
- ◆ [Chapter 7, “Configuring Local Authentication,” on page 71](#)
- ◆ [Chapter 11, “Configuring Liberty Web Services,” on page 115](#)

5.1 Creating an Identity Server Configuration (Advanced Options)

This section discusses advanced settings an Identity Server configuration, such as importing SSL certificates, enabling introductions, and configuring identity consumer settings. It is recommended that you are familiar with “[Creating a Basic Identity Server Configuration](#)” in the *Novell Access Manager 3.0 Setup Guide* before proceeding.

After you install an Identity Server, you create an Identity Server configuration. The new configuration is then available for you to assign to one or more Identity Servers. As shown in [Figure 5-1](#), you can also create multiple configurations and assign different Identity Servers to them.

Figure 5-1 Identity Server Configurations



An Identity Server always operates as an identity provider and can optionally be configured to run as an identity consumer (also known as a service provider), using either Liberty, SAML 1.1, or SAML 2.0 protocols. In an Identity Server cluster, multiple servers must use the same configuration.

In an Identity Server configuration, you specify the following information:

- ◆ The base URL for the server site.
- ◆ Certificates for the Identity Server, identity provider, and identity consumer.
- ◆ Authentication settings, such as whether the identity provider requires signed authentications from service providers.
- ◆ The service domains used for publishing and discovering authentications.
- ◆ Organizational and contact information for the server, which is published in the metadata of the Liberty 1.2 and SAML protocols.
- ◆ The LDAP directories (user stores) used to authenticate users, and the trusted root for secure communication between the Identity Server and the user store.

After you create a configuration, you can configure initial settings, such as local authentication, to meet your needs.

To create an Identity Server configuration:

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > New*.

The screenshot shows the 'Create Identity Server Configuration' dialog box, Step 1 of 3: Specify Name and Base URL. The dialog is titled 'Identity Servers' and has a blue header bar with the title and a help icon. Below the header, the current step is 'Step 1 of 3: Specify Name and Base URL'. The main area contains several fields and options:

- Name:** A text field containing 'ids-corporate'. Below it is a hint: '(protocol :// domain : port / application)'. Below that is the 'Base URL' field, which is split into four parts: a protocol dropdown (set to 'http'), a domain field (containing 'ids-corporate.com'), a port field (containing '8080'), and an application field (containing 'nidp'). A 'Select SSL Certificate' link is below the Base URL field.
- LDAP Access:** A field with '20' and a dropdown arrow, followed by the text 'connections'.
- Session timeout:** A field with '15' and a dropdown arrow, followed by the text 'minutes'.
- Allow multiple browser session logout:** An unchecked checkbox.
- Identity Provider:** A section with three checkboxes: 'Show logged out providers', 'Require Signed Authentication Requests', and 'Use Introductions (Publish Authentications)'. Below these is a 'Service domain' field with three sub-fields: 'Local:', 'Common:', and 'Port:'. A 'Select SSL Certificate' link is below the Service domain field.
- Identity Consumer:** A section with a checked 'Enable' checkbox and two unchecked checkboxes: 'Require Signed Assertions' and 'Sign Authentication Requests'.

At the bottom of the dialog are three buttons: '<< Back', 'Next >>', and 'Cancel'.

- 2 Fill in the following fields to specify the properties for your Identity Server configuration:

Name: A name by which you want to refer to the configuration.

IMPORTANT: Carefully determine your settings for the base URL, protocol, and domain. After you have configured trust relationships between providers, changing these settings invalidates the trust model and requires a reimport of the provider's metadata. This is also true for the domain name on the Access Gateway ESP.

Base URL: The application path for the Identity Server. The Identity Server protocols (Liberty 1.2, SAML 1.1, and SAML 2.0) rely on this base URL to generate URL endpoints for each protocol.

- ♦ **Protocol:** The communication protocol. Specify HTTPS in order to run securely (in SSL mode) and for provisioning. Use HTTP only if you do not require security.
- ♦ **Domain:** The domain name used to access the Identity Server. Using an IP address is not recommended for HTTP and is not allowed with HTTPS.
- ♦ **Port:** The port values for the protocol. Default ports are 8080 for HTTP or 8443 for HTTPS. If you want to use a different port, see [Section 40.2, "Translating the Identity Server Configuration Port,"](#) on page 470.
- ♦ **Application:** The Identity Server application. Leave the default value *nidp*.

Select SSL Certificate: Displays the Keystore page that you use to locate and replace the test-connector SSL certificate for this configuration.

The Identity Server comes with a test-connector certificate that you must replace for your production environment. You can replace the test certificate now or after you configure the Identity Server. If you create the certificate and replace the test-connector now, you can save some time by restarting Tomcat only once. Tomcat must be restarted whenever you assign an Identity Server to a configuration and whenever you update a certificate key store. See [Section 4.3, “Security,” on page 46](#).

LDAP Access: The maximum number of LDAP connections allowed to the configuration store. You might adjust this amount for system performance.

Session Timeout: The session inactivity time allowed before timing out. When using Basic authentication and SSL mutual authentication, the browser must be closed to terminate the session. See [Section 40.1.2, “Prompting Users to Refresh a Session before Expiration,” on page 469](#).

Allow multiple browser session logout: Specifies whether a user with more than one session to the server is presented with an option to log out of all sessions. If you do not select this option, only the current session can be logged out. You deselect this option in instances where multiple users log in as guests. Then, when one user logs out, none of the other guests are logged out.

After you enable this option and click OK, you are prompted to apply the changes using *Update Servers* on the Setup page. You must also restart any ESPs in an Access Gateway or J2EE Agent configuration that use this Identity Server configuration.

- 3 To specify identity provider settings, fill in the following fields:

Show logged out providers: Displays logged-out providers on the identity provider’s log-out confirmation page.

Require Signed Authentication Requests: Specifies that for the Liberty 1.2 and SAML 2.0 protocols, authentication requests from service providers must be signed. When you enable this option for the identity provider, you must also enable the *Sign Authentication Requests* option under the *Identity Consumer* heading on this page for the external trusted service provider. (It is possible, however, to configure an identity provider that requires signed requests to function as an identity consumer that does not sign requests.)

Use Introductions (Publish Authentications): Enables single sign-on from the service provider to the identity provider. The service provider determines which identity providers that users are already logged into, and then selectively and automatically asks for authentication from one of the identity providers. Introductions are enabled only between service and identity providers that have agreed to a circle of trust, which means that they have agreed upon a common domain name for this purpose.

After authenticating a user, the identity provider accesses a service at the service domain and writes a cookie to the common part of the service domain, publishing that the authentication has occurred.

- ♦ **Service Domain (Local and Common):** Enables a service provider to access a service at the service domain prior to authenticating a user. This service reads cookies obtained at this domain and discovers if any identity providers have provided authentication to the user. The service provider determines whether any of these identity providers can authenticate a user without credentials. The service domain must resolve to the same IP address as the base URL domain.

For example, if an agreed-upon common domain is *xyz.com*, the service provider can specify a service domain of *sp.xyz.com*, and the identity provider can specify a service domain of *idp.xyz.com*. For the identity provider, *xyz.com* is the common value entered, and *idp* is the local value.

- ♦ **Port:** The port to use for identity provider introductions. Port 8445 for HTTPS is the default and must be opened on your firewall. If you specify a different port, you must edit the Tomcat server XML.

Select SSL Certificate: Displays the Keystore page that you use to locate and replace the test-provider SSL certificate for this configuration.

The Identity Server comes with a test-provider certificate that you must replace for your production environment. This certificate is used for identity provider introductions. You can replace the test certificate now or after you have configured the Identity Server. If you create the certificate and replace the test-connector now, you can save some time by restarting Tomcat only once. Tomcat must be restarted whenever you assign an Identity Server to a configuration and whenever you update a certificate key store. See [Section 4.3, “Security,” on page 46](#).

- 4 Optionally, specify whether the Identity Server also runs as an identity consumer.

If configured to run as an identity consumer, the Identity Server can receive (consume) authentication assertions from other identity providers.

Enable: Enables this site to function as service provider. This setting is enabled by default.

Require Signed Assertions: Specifies that all SAML assertions received by the service provider must be signed by the issuing SAML authority. The signing authority uses a key pair to sign SAML data sent to this trusted provider.

Sign Authentication Requests: Specifies that the service provider signs authentication requests to an identity provider for the Liberty 1.2 and SAML 2.0 protocols.

Use Introductions (Discover IDP Authentications): Enables a service provider to discover whether a user has authenticated to a trusted identity provider, so the user can use single sign-on without requiring authentication credentials.

- ♦ **Service domain:** The shared, common domain for all providers in the circle of trust. This domain must resolve to the same IP address as the base URL domain. You must enable the *Identity Consumer* option to enable this field.
- ♦ **Port:** The port to use for identity consumer introductions. Port 8446 for HTTPS is the default and must be opened on your firewall. If you specify a different port, you must edit the Tomcat server XML.

Select SSL Certificate: Displays the Keystore page that you use to locate and replace the test-consumer SSL certificate for this configuration.

The Identity Server comes with a test-consumer certificate that you must replace for your production environment. This certificate is used for identity consumer introductions. You can replace the test certificate now or after you have configured the Identity Server. If you create the certificate and replace the test-connector now, you can save some time by restarting Tomcat only once. Tomcat must be restarted whenever you assign an Identity Server to a configuration and whenever you update a certificate key store. See [Section 4.3, “Security,” on page 46](#).

- 5 To continue creating the Identity Server configuration, click *Next*.

The system displays the Organization page.

Identity Servers ▶

Create Identity Server Configuration

Step 2 of 3: Specify Organization

Name: *

Display name: *

URL: *

Principal Contact

Company:

First Name:

Last Name:

Email Address:

Telephone Number:

Contact Type: ▼

Use this page to specify organization information for the Identity Server configuration. The information you specify on this page is published in the metadata for the Liberty 1.2 and SAML protocols. The metadata is traded with federation partners and supplies various information regarding contact and organization information located at the Identity Server.

The following fields require information:

- ♦ **Name:** The name of the organization.
- ♦ **Display Name:** The display name for the organization.
- ♦ **URL:** The organization's URL for contact purposes.

Optional fields include Company, First Name, Last Name, Email, Telephone, and Contact Type.

6 Click *Next* to configure the user store.

You must reference your own user store and auto-import the SSL certificate. See [Section 7.1, "Configuring Identity User Stores," on page 71](#) for information about this procedure.

7 After you configure the user store, click *Finish* to save the server configuration.

The system displays the new configuration on the Configurations page.

Identity Servers ?

Servers Setup

Configurations | Attribute Sets | User Matching Expressions | Custom Attributes

New | Delete | Refresh 2 Item(s)

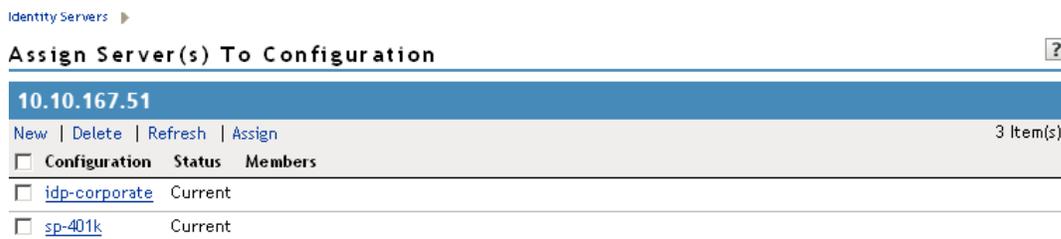
<input type="checkbox"/>	Configuration	Status	Members
<input type="checkbox"/>	ids-a	Current	
<input type="checkbox"/>	spd	Update Servers	Servers

5.2 Assigning an Identity Server to a Configuration

After you create a configuration, you must assign the Identity Server to it. For clustering, you can assign more than one Identity Server to the configuration (see [Section 5.6, “Clustering Identity Servers,”](#) on page 60). A configuration uses any global settings you have specified, such as attribute sets, user matching expressions, and custom attributes that are defined for the server.

- 1 In the Administration Console, click *Access Manager > Identity Servers*.
- 2 On the Servers page, select the server’s check box, then choose *Actions > Assign to configuration*.

You can also select all displayed servers by selecting the top-level Server check box.



- 3 On the Assign Server(s) to Configurations page, select the server configuration check box, then click *OK*.

You are prompted to restart Tomcat. The status icon for the Identity Server should turn green. It might take several seconds for the identity provider to start and for the system to display the green light.

5.3 Removing a Server from a Configuration

Removing an Identity Server from a configuration disassociates the Identity Server from the configuration. The configuration, however, remains intact and can be reassigned later or assigned to another server.

- 1 In the Administration Console, click *Access Manager > Identity Servers*.
- 2 Select the server, then click *Stop*. Wait for the Health indicator to turn red.
- 3 Select the server, then choose *Actions > Remove from configuration*.

5.4 Modifying the Base URL

When configuring an Identity Server, you must carefully determine your settings for the base URL, protocol, and domain. After you have configured trust relationships between providers, changing these settings invalidates the trust model and requires a reimport of the provider’s metadata. See [“Reimporting a Trusted Provider’s Metadata.”](#)

- 1 In the Administration Console, click *Access Manager > Identity Servers > [Configuration]*.
- 2 Change the protocol, domain, port, and application settings, as necessary.
- 3 Click *OK*.
- 4 Click the *Setup* tab, then click *Update Servers*.

- 5 For each device configured to trust the configuration of this modified base URL, you must reconfigure the trust relationship between the device’s embedded service provider and the Identity Server.
 - ♦ For an Access Gateway, see [Section 14.2, “Modifying the Base URL of the Identity Server,” on page 182.](#)
 - ♦ For a J2EE Agent, see “[Modifying the Base URL of the Identity Server](#)” in the *Novell Access Manager 3.0 J2EE Agent Guide*.

For information about setting up SSL and changing an Identity Server from HTTP to HTTPS, see “[Enabling SSL Communication](#)” in the *Novell Access Manager 3.0 Setup Guide*.

5.5 Enabling Role-Based Access Control

Role-based access control (RBAC) is used to provide a convenient way assign a user to a particular job function or set of permissions within an enterprise, in order to control access. In Access Manager, you assign users to roles, based on attributes of their identity, and then associate authorization policies to the role.

For a complete discussion on creating and configuring role policies, see [Chapter 29, “Creating Role Policies,” on page 317](#), in [Part VI, “Policy Management,” on page 309](#).

In order for a role to be assigned to users at authentication, you must enable it for the Identity Server configuration.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Roles*.
- 2 Click the role policy’s check box, then click *Enable*.
- 3 To disable the role policy, click the role policy’s check box, then click *Disable*.
- 4 After enabling or disabling role policies, update the Identity Server configuration on the Setup tab.

5.6 Clustering Identity Servers

To add capacity and for failover, you can cluster a group of Identity Servers and configure them to act as a single server. A cluster of Identity Servers should reside behind an L4 switch. Clients access the virtual IP (VIP) address of the cluster presented on the L4 switch, and the L4 switch alleviates server load by balancing traffic across the cluster.

Whenever a user accesses the virtual IP address (port 8080) assigned to the L4, the system routes the user to one of the Identity Servers in the cluster, as traffic necessitates.

Prerequisites

- An L4 server installed. You can use the same server for Identity Server clustering and Access Gateway clustering, provided that you use different virtual IPs. The LB algorithm can be anything (hash/sticky bit), defined at the Real server level.
- Persistence (sticky) sessions enabled on the L4 server. Normally you define this at the virtual server level.
- An Identity Server configuration created for the cluster. You assign all the Identity Servers to this configuration. See [Section 5.1, “Creating an Identity Server Configuration \(Advanced](#)

Options),” on page 53 for information about creating an Identity Server configuration. See Section 5.2, “Assigning an Identity Server to a Configuration,” on page 59 for information about assigning identity servers to configurations.

The base URL DNS name of this configuration must resolve via DNS to the IP address of the L4 virtual IP address. The L4 balances the load between the identity servers in the cluster.

- ❑ Ensure that the L4 administration server using port 8080 has the following ports open:
 - ◆ 8443 (secure Administration Console)
 - ◆ 7801 (TCP)
 - ◆ 636 (for secure LDAP)
 - ◆ 389 (for clear LDAP, loopback address)
 - ◆ 524 (network control protocol on the L4 machine for server communication)

The identity provider ports must also be open:

- ◆ 8080 (nonsecure login)
- ◆ 8443 (secure login)
- ◆ 1443 (server communication)

If you are using introductions (see Section 5.1, “Creating an Identity Server Configuration (Advanced Options),” on page 53), you must configure the L4 switch to load balance on ports 8445 (identity provider) and 8446 (identity consumer).

Setup

- 1 Install the additional Identity Servers.

During installation, choose option 2, *Install Novell Identity Server*, from CD 1 of the Access Manager installation discs. You run the installation for each new Identity Server you want to add. Specify the IP address and administration credentials of each additional Identity Server. If you are installing on a machine without the Administration Console, the installation asks you for the Administration Console’s IP address. After you install the Identity Servers, the servers are displayed on the Servers page in Identity Servers.

- 2 Assign the Identity Servers to the same configuration.
- 3 Click the configuration name you created for the cluster under *Configuration Assignment*.
- 4 Click *Cluster*.

The screenshot shows the configuration interface for a cluster. At the top, there are tabs for 'General', 'Local', 'Liberty', 'SAML 1.1', and 'SAML 2.0'. Below the tabs is a navigation bar with links for 'Configuration', 'Organization', 'Roles', 'Cluster', 'Logging', and 'Security'. The main content area is divided into two sections:

- Cluster communication backchannel:** This section contains a 'Port:' field with the value '7801' and an 'Encrypt' checkbox which is currently unchecked.
- Level four switch port translation:** This section contains a checkbox for 'Port translation is enabled on switch' which is unchecked, and a 'Cluster member translated port:' field which is empty.

Below these sections is a 'Cluster members' section with a blue header. It includes a 'Refresh' button, a 'Start | Stop | Actions' dropdown menu, and a '1 Item(s)' indicator. A table below this header shows a single entry for a 'Server' with columns for 'Server Status', 'Alerts', 'Command Status', and 'Statistics'.

5 Fill in the following fields as required:

Cluster Communication Backchannel: Provides a communications channel over which the cluster members maintain the integrity of the cluster. For example, this TCP channel is used to detect new cluster members as they join the cluster, and to detect members that leave the cluster. A small percentage of this TCP traffic is used to help cluster members determine which cluster member would best handle a given request. This back channel should not be confused with the IP address/port over which cluster members provide proxy requests to peer cluster members.

- ◆ **Port:** Specifies the TCP port of the cluster back channel on all of the Identity Servers in the cluster. 7801 is the default TCP port.

Because the cluster back channel uses TCP, you can use cluster members on different networks. However, firewalls must allow the port specified here to pass through. To do so use the port number plus 1, for additional devices in the cluster. For example, if you use four devices, your port numbers would be 7801, 7802, 7803, and 7804.

- ◆ **Encrypt:** Encrypts the content of the messages that are sent between cluster members.

Level Four Switch Port Translation: Configures the L4 switch to translate the port of the incoming request to a new port when the request is sent to a cluster member. Because the cluster members communicate with each other over the same IP address/port as the L4 switch, the cluster implementation needs to know what that port is. The translated port is the port on the cluster members where other cluster members can contact it. This is the IP address and port where cluster members provide proxy requests to other cluster members.

- ◆ **Port translation is enabled on switch:** Specifies whether the port of the L4 switch is different from the port of the cluster member.
- ◆ **Cluster member translated port:** Specifies the port of the cluster member.

Under *Cluster Members*, you can refresh, start, stop, and assign servers to Identity Server configurations.

6 Click *OK*, then update the Identity Server as prompted.

Real Server Settings Example

Current real servers settings:

```
1: 149.44.171.116, enabled, name l52, weight 1, timeout 10 mins, maxcon 200000
  backup none, inter 2, retry 4, restr 8
  remote disabled, proxy enabled, subnac disabled
  cookie assignment server: disabled
  exclusionary string matching: disabled
  service ports: 8443 8080
  real ports:
    8443: uport 8443, group 1, pbind clientip
          virtual server: 1, 149.44.174.220, enabled
    8080: uport 8080, group 1, pbind clientip
          virtual server: 1, 149.44.174.220, enabled
2: 149.44.174.51, enabled, name brie, weight 1, timeout 10 mins, maxcon 200000
  backup none, inter 2, retry 4, restr 8
  remote disabled, proxy enabled, subnac disabled
  cookie assignment server: disabled
  exclusionary string matching: disabled
  service ports: 8443 8080
  real ports:
    8443: uport 8443, group 1, pbind clientip
          virtual server: 1, 149.44.174.220, enabled
    8080: uport 8080, group 1, pbind clientip
          virtual server: 1, 149.44.174.220, enabled
```

Virtual Server Settings Example

Current virtual servers settings:

```
1: 149.44.174.220, enabled, dname idp
  virtual ports:
    8443: rport 8443, group 1, pbind clientip, frags
          real servers:
            1: 149.44.171.116, weight 1, enabled, backup none
            2: 149.44.174.51, weight 1, enabled, backup none
    8080: rport 8080, group 1, pbind clientip, frags
          real servers:
            1: 149.44.171.116, weight 1, enabled, backup none
            2: 149.44.174.51, weight 1, enabled, backup none
```

5.7 Configuring Secure Communication on the Identity Server

The Identity Server comes with an SSL test-connector certificate that you must replace. You can also import the trusted root certificate to enable secure communication between the user store and the Identity Server.

This task is part of basic setup. See “[Enabling SSL Communication](#)” in the *Novell Access Manager 3.0 Setup Guide*.

Defining Reusable Configuration Settings

6

You can define reusable settings globally so that they can be available in any Identity Server configuration. The settings include:

- ♦ **Attribute sets:** Sets of attributes that are exchangeable between identity and service providers.
- ♦ **User matching expressions:** The logic of the query to the user store for identification when an assertion is received from an identity provider.
- ♦ **SharedSecret names:** Custom shared secret names that you want to be available when configuring policies.
- ♦ **LDAP attributes:** Custom LDAP attribute names that you want to be available when configuring policies.

This section describes the settings that can apply to any configuration.

- ♦ [Section 6.1, “Configuring Attribute Sets,” on page 65](#)
- ♦ [Section 6.2, “Editing Attribute Sets,” on page 67](#)
- ♦ [Section 6.3, “Configuring User Matching Expressions,” on page 67](#)
- ♦ [Section 6.4, “Adding Custom Attributes,” on page 68](#)

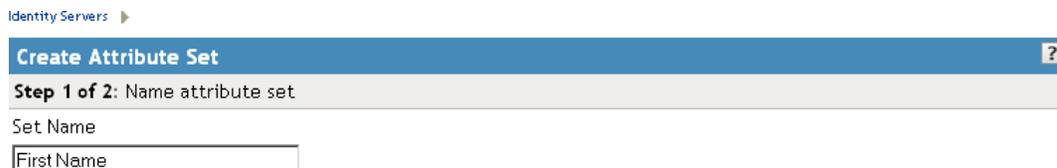
6.1 Configuring Attribute Sets

Attributes you specify on the Identity Server are used in attribute requests and responses, depending on whether you are configuring a service provider (request) or identity provider (response). Attribute sets provide a common naming scheme used in the exchange. For example, an attribute set can map the Liberty attribute FN (first name) to the equivalent remote name used at the service provider, which might be Name.

Attributes also can be defined and used in policy enforcement. They can be attributes defined by the Web Service Profiles, or customized attributes that can be mapped into SAML attributes. You also map user attributes so that the Identity Server can accept them from SAML.

To create and configure an attribute set:

- 1 In the Administration Console, click *Access Manager > Identity Server > Setup > Attribute Sets > New*.



Identity Servers ▶

Create Attribute Set ?

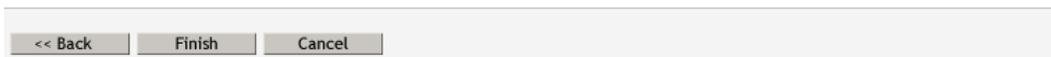
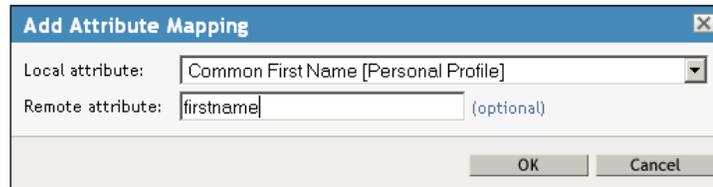
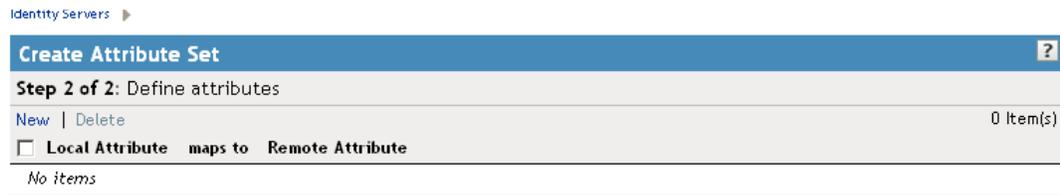
Step 1 of 2: Name attribute set

Set Name

First Name

- 2 Specify a name for identifying the attribute set, then click *Next*.

3 To create a set, click *New*.



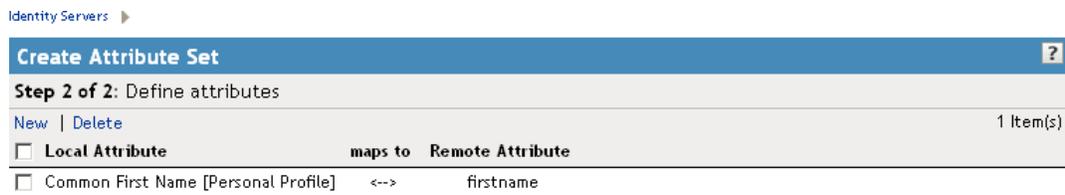
Local Attribute: A drop-down list of all server profile and LDAP attributes. As an example, you can select *All Roles* to use in role policies, which enables trusted providers to send role information in authentication assertions. Customizable attributes can be created and displayed in this list.

Remote Attribute: The name of the attribute defined at the external provider. The text for this field is case sensitive. If you leave this field blank, the system sends an internal value that is recognized between Identity Servers.

For SAML 1.1 identity consumer (service provider), a name identifier received in an assertion is automatically given a remote attribute name of *saml:NameIdentifier*. This allows the name identifier to be mapped to a profile attribute that can then be used in policy definitions.

4 Click *OK*.

The system displays the map settings on the Define Attributes page, as shown below:



You can continue adding as many attributes as you need.

5 Click *Finish* after you created the map.

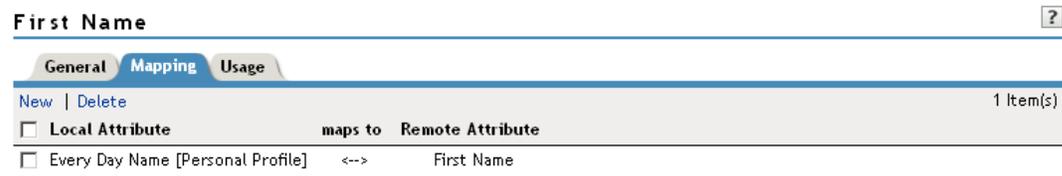
The system displays the map on the Attribute Sets page, as well as indicating whether it is in use by a provider. (See [Section 8.7, “Selecting Attributes for a Trusted Provider,”](#) on page 97.)



6.2 Editing Attribute Sets

You can edit attribute sets that have been created in the system. (See [Section 6.1, “Configuring Attribute Sets,”](#) on page 65.)

- 1 In the Administration Console, click *Access Manager > Identity Server > Setup > Attribute Sets*.
- 2 Click the name of the attribute set that you want to edit.



- 3 The system displays an attribute set page with the following tabs:
 - General:** Click to edit the name of the attribute set.
 - Mapping:** Click to edit the attribute map.
 - Usage:** Displays where the attribute set is used. Informational only.
- 4 Click *OK*, then click *Close*.

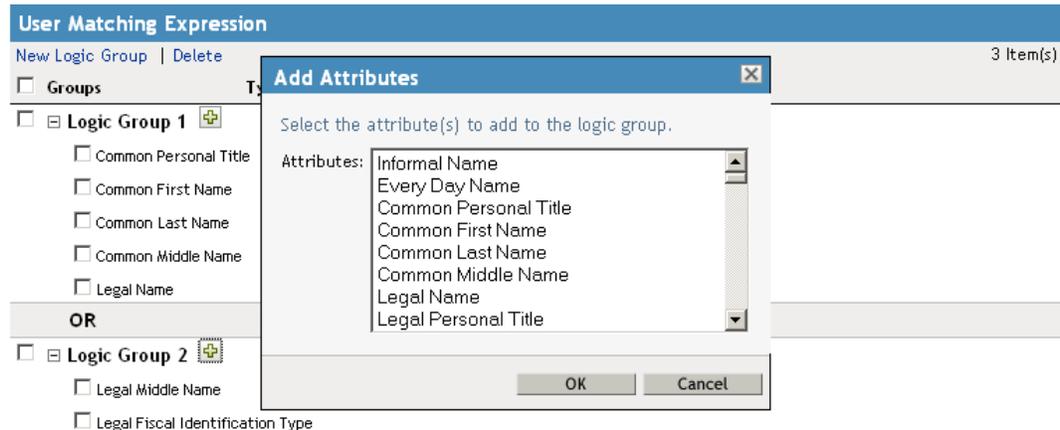
6.3 Configuring User Matching Expressions

One of the user identification methods the Identity Server uses when an assertion is received is to query the user store based on attributes received in the assertion from the identity provider. You configure user matching expressions to define the logic of the query. You must know the LDAP attributes that are used to name the users in the user store and create the user’s distinguished name.

In order to use user matching, you must enable the Personal Profile on the identity provider and the service provider. See [Section 11.2, “Enabling Web Services and Profiles,”](#) on page 116.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > User Matching Expressions*.
- 2 Click *New*, or click the name of an existing user matching expression.
 - Name:** The name of the user lookup expression.

- 3 Click the *Add Attributes* icon (plus sign), then select attributes to add to the logic group. (Use the Shift key to select several attributes.)



- 4 Click *OK*.
- 5 To add logic groups, click *New Logic Group*.
The *Type* drop-down (AND or OR) applies only between groups. Attributes within a group are always the opposite of the type selection. For example, if the *Type* value is AND, the attributes within the group are OR.
- 6 Click the *Add Attributes* icon (plus sign) to add attributes to the next logic group, then click *OK*.
- 7 Click *Finish*.

6.4 Adding Custom Attributes

You can add custom shared secret names or LDAP attribute names that you want to make available for selection when setting up policies.

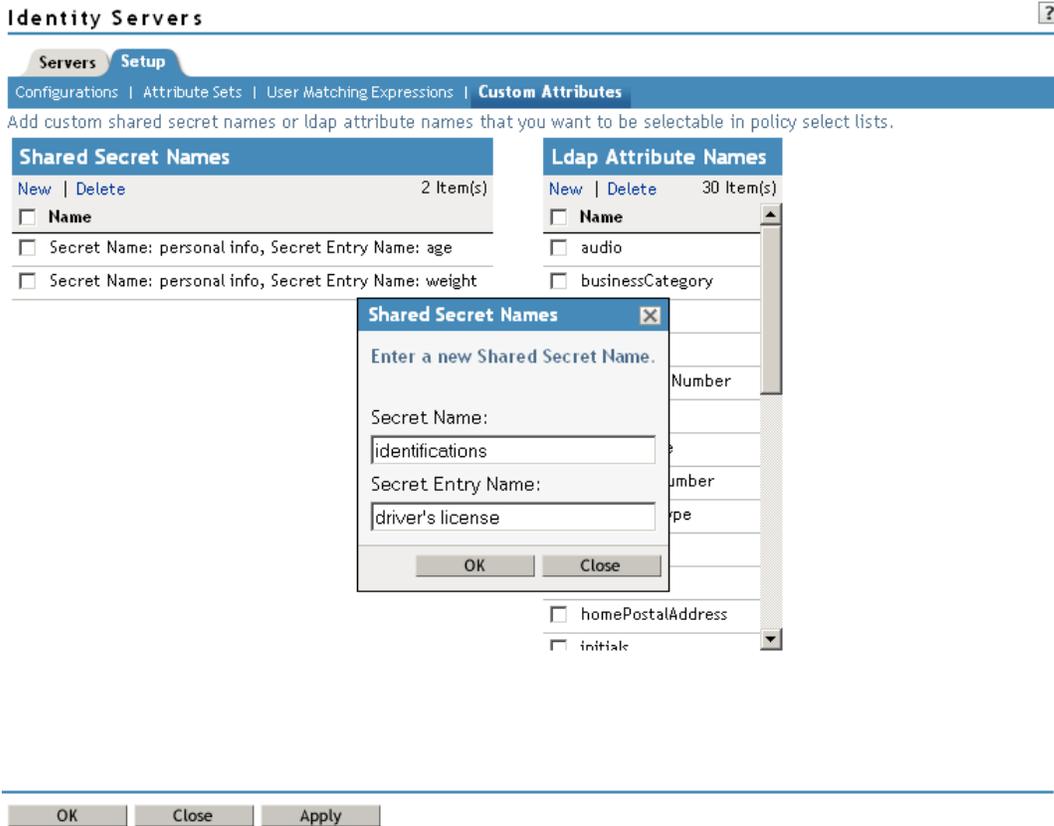
- ♦ [Section 6.4.1, “Creating Shared Secret Names,” on page 68](#)
- ♦ [Section 6.4.2, “Creating LDAP Attribute Names,” on page 69](#)

6.4.1 Creating Shared Secret Names

The shared secret consists of a secret name and one or more secret entry names. You can create a secret name only, or a secret name and an entry name. Shared secret names can be created either on this page or in the associated policy that will consume them.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > Custom Attributes*.

2 To create shared secret names, click *New*.



3 Enter a new shared secret name and, optionally, a secret entry name.

4 Click *OK*.

6.4.2 Creating LDAP Attribute Names

LDAP attributes are available for all policies. You can add available attributes here, as well as on the Policies page. LDAP attribute names can be created either on this page or in the associated policy that will consume them.

1 In the Administration Console, click *Access Manager > Identity Servers > Setup > Custom Attributes*.

2 Click *New* to add a name. This list is customizable. Examples of predefined LDAP attributes include:

- ♦ **audio:** Uses a u-law encoded sound file, stored in the directory.
- ♦ **businessCategory:** Describes the kind of business performed by an organization.
- ♦ **carLicense:** Vehicle license or registration plate.
- ♦ **cn:** The X.500 commonName attribute, which contains a name of an object. If the object corresponds to a person, it is typically the person's full name.
- ♦ **departmentNumber:** Identifies a department within an organization.

- ♦ **displayName:** The preferred name of a person to be used when displaying entries. Identifies a name to be used. When displaying an entry, especially within a one-line summary list, it is useful to use this value. Because other attribute types such as `cn` are multivalued, an additional attribute type is needed.
- ♦ **employeeNumber:** Numerically identifies a person within an organization.
- ♦ **employeeType:** Identifies the type of employee.
- ♦ **givenName:** Identifies the person's name that is not his or her surname or middle name.
- ♦ **homePhone:** Identifies a person by home phone.
- ♦ **homePostalAddress:** Identifies a person by home address.
- ♦ **initials:** Identifies a person by his or her initials. This attribute contains the initials of an individual, but not the surname.
- ♦ **jpegPhoto:** Stores one or more images of a person, in JPEG format.
- ♦ **labeledURI:** Uniform Resource Identifier with an optional label. The label describes the resource to which the URI points.
- ♦ **mail:** A user's e-mail address.
- ♦ **manager:** Identifies a person by manager title.
- ♦ **mobile:** Specifies a mobile telephone number associated with a person.
- ♦ **o:** The name of an organization.
- ♦ **pager:** The pager telephone number for an object.
- ♦ **photo:** Specifies a photograph for an object.
- ♦ **preferredLanguage:** Indicates an individual's preferred written or spoken language.
- ♦ **roomNumber:** The room number of an object.
- ♦ **secretary:** Specifies the secretary of a person.
- ♦ **sn:** The X.500 surname attribute, which contains the family name of a person.
- ♦ **uid:** User ID.
- ♦ **userCertificate:** An attribute stored and requested in the binary form.
- ♦ **userPKCS12:** A format to exchange personal identity information. Use this attribute when information is stored in a directory service.
- ♦ **userSMIMECertificate:** PKCS#7 SignedData used to support S/MIME. This value indicates that the content that is signed is ignored by consumers of `userSMIMECertificate` values.
- ♦ **x500uniqueIdentifier:** Distinguishes between objects when a distinguished name has been reused. This is a different attribute type from both the `uid` and the `uniqueIdentifier` type.

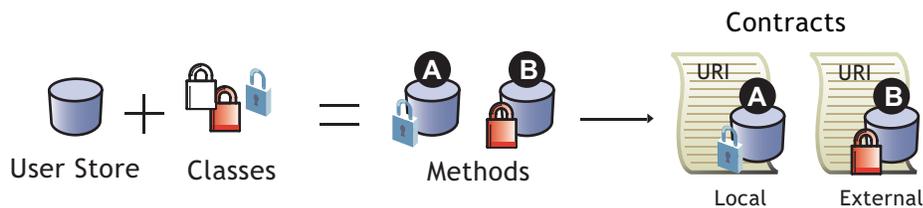
Configuring Local Authentication

7

To guard against unauthorized access, Access Manager supports a number of ways for users to authenticate. These include name/password, RADIUS token-based authentication, and X.509 digital certificates. You configure authentication at the Identity Server by creating authentication contracts that the components of Access Manager (such as an Access Gateway) can use to protect a resource.

Figure 7-1 illustrates the components of a contract:

Figure 7-1 Local Authentication



- ◆ **User stores:** The user directories to which users authenticate. You set up your user store when creating the Identity Server configuration.
- ◆ **Classes:** The code (a Java class) that implements a particular authentication type (name/password, RADIUS, and X.509) or means of obtaining credentials.
- ◆ **Methods:** The pairing of an authentication class with one or more user stores.
- ◆ **Contracts:** The basic unit of authentication. Contracts can be local (executed at the server) or external (executed by another Identity Server). Contracts are identified by a unique URI that can be used by Access Gateways and agents to protect resources. Local contracts are comprised of one or more authentication methods used to uniquely identify a user. You can associate multiple methods with one contract.

You can also use the properties of a class to create custom login pages.

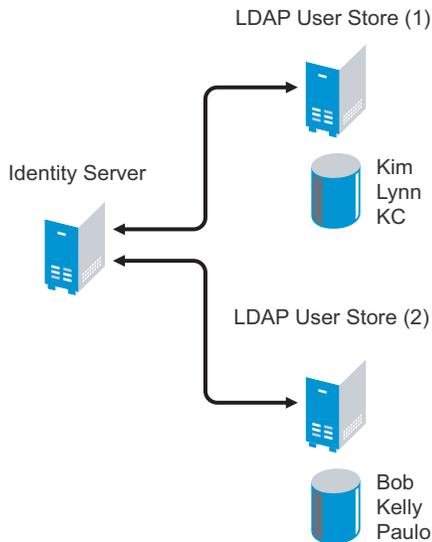
- ◆ [Section 7.1, “Configuring Identity User Stores,” on page 71](#)
- ◆ [Section 7.2, “Creating Authentication Classes,” on page 75](#)
- ◆ [Section 7.3, “Configuring Authentication Methods,” on page 81](#)
- ◆ [Section 7.4, “Configuring Authentication Contracts,” on page 82](#)
- ◆ [Section 7.5, “Specifying Authentication Defaults,” on page 84](#)
- ◆ [Section 7.6, “Creating Custom Login Pages,” on page 85](#)

7.1 Configuring Identity User Stores

User stores are LDAP directory servers to which end users authenticate. You must specify an initial user store when creating an Identity Server configuration. This procedure describes how to add an additional user store to provide load balancing and failover capability. However, you use the same pages for setting up the initial user store or adding a user store.

You can also configure the Identity Server to search more than one user store during authentication. **Figure 7-2** illustrates this type of configuration.

Figure 7-2 Multiple LDAP Directories



It is assumed that each LDAP directory contains different users. You should make sure the users have unique names across all LDAP directories. If both directories contain a user with an identical name, the name and password information discovered in the search of the first directory is always used for authentication. You select the user store and specify the search order when configuring the authentication method.

If you add a secondary Administration Console and you have added replicas to the user store of the primary Administration Console, ensure that you also add the replicas to the secondary Administration Console.

All user stores that you add are included in health checks. If health problems are found, the system displays the user store on the Health page and in the trace log file.

- 1** In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Local*.
- 2** In the User Stores list, click *New*.

If you are creating an Identity Server configuration, this is Step 3 of the wizard.

The screenshot displays the 'Create Identity Server Configuration' wizard, Step 3 of 3: Specify initial User Store. The main window contains the following fields:

- Name: HR User Store
- Admin name: cn=admin,o=novell (Ex: cn=admin,o=novell)
- Admin password: [Redacted]
- Confirm password: [Redacted]
- Directory type: [Redacted]

A 'Specify server replica information' dialog box is open, showing the following fields:

- Name: HR Replica
- IP Address: 10.10.167.50
- Port: 636
- Use secure LDAP connections
- Auto import trusted root (link)
- Connection limit: 20

The dialog box has 'OK' and 'Cancel' buttons. The background shows a table for 'Server replicas' with one entry: HR Replica, 10.10.167.50. Below that is a 'Search Contexts' section with 'No items'.

3 Fill in the following fields:

Name: The name of the user store for reference.

Admin Name: The distinguished name of the admin user of the LDAP directory. Administrator-level rights are required for setting up a user store. This ensures read/write access to all objects used by Access Manager.

Admin Password and Confirm Password: The password for the admin user and confirm the password.

Directory Type: The type of LDAP directory. You can select eDirectory, Active Directory, or Sun ONE

4 Specify a server replica.

For an eDirectory server, it is recommended that you use a replica of the partition where the users reside. Ensure that each LDAP server in the cluster has a valid read/write replica. One option is to create a users partition (a partition that points to the OU containing the user accounts) and reference this server replica.

Name: The display name for the LDAP directory server. If your LDAP directory is replicated on multiple servers, use this name to identify a specific replica.

IP Address: The IP address of the LDAP directory server.

Port: The port of the LDAP directory server.

Use secure LDAP connections: (Required) Specifies that the LDAP directory server requires secure (SSL) connections with the Identity Server.

This option must be enabled if you use this user store as a SecretStore[®] User Store Reference in the Credential Profile details. (See [Section 11.4, “Configuring Credential Profile Security and](#)

Display Settings,” on page 118.) If you have specified that this user store is a SecretStore User Store Reference, this option is enabled but not editable.

Connection limit: The maximum number of pooled simultaneous connections allowed to the LDAP server. Valid values are between 5 and 100.

- 5 Click *Auto import trusted root*.
- 6 Click *OK* to confirm the import.

Select Certificate to Trust

Alias:

Server Certificate

Subject: O=.SPB_UNSTABLE_TREE., CN=spb-unstable.provo.novell.com
Issuer: O=SPB_UNSTABLE_TREE, OU=Organizational CA
Valid starting date: 30 May 2006 17:10:44 GMT
Valid ending date: 29 May 2008 17:10:44 GMT
Signature algorithm: SHA1withRSA
Finger print (MD5): 3C:7A:99:81:05:2F:40:23:0E:94:14:68:A5:D3:29:3D
Finger print (SHA1): 74:86:DD:23:F4:23:5B:95:8C:78:F7:86:6B:05:91:8C:8C:98:0D:99

Root CA Certificate

Subject: O=SPB_UNSTABLE_TREE, OU=Organizational CA
Issuer: O=SPB_UNSTABLE_TREE, OU=Organizational CA
Valid starting date: 28 May 2006 19:10:40 GMT
Valid ending date: 27 May 2016 19:10:40 GMT
Signature algorithm: SHA1withRSA
Finger print (MD5): F4:D9:FE:A5:F9:93:01:02:62:85:29:44:53:D4:5B:90
Finger print (SHA1): AF:EC:A7:1C:22:10:B7:35:91:FE:B9:6E:51:92:B8:9A:6C:0E:A1:5F

OK

Cancel

- 7 Select one of the certificates in the list.
You are prompted to choose either a server certificate or a root CA certificate. To trust one certificate, choose *Server Certificate*. Choose *Root CA Certificate* to trust any certificate signed by that certificate authority.
- 8 Specify an alias, then click *OK*.
- 9 Click *OK* in the *Specify server replica information* dialog box.
- 10 Add a Search Context.

The search context is used to locate users in the directory. If a user exists outside of the specified search context (object, subtree, one level) then the Identity Server cannot find the user, and the user cannot log in.

This is required for Active Directory or Sun ONE; it is optional for eDirectory because the entire tree is searched from the root if a search context is not specified.

NOTE: For Active Directory, do not set the search context at the root level using the Subtree scope. This setting can cause serious performance problems. It is recommended that you set multiple search contexts, one for each top-level organizational unit.

- 11 Click *Finish*.

- 12 Add the new user store to the authentication method. See [Section 7.3, “Configuring Authentication Methods,”](#) on page 81.

7.2 Creating Authentication Classes

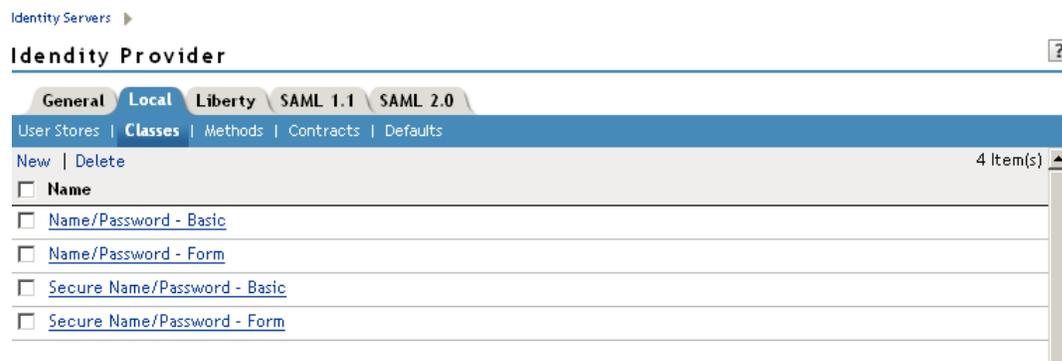
Authentication classes let you define ways of obtaining end user credentials. You specify the code (java class) and properties to be executed to implement a particular authentication type.

Several authentication classes are included with Access Manager to provide a variety of ways to authenticate end users. Custom authentication classes provided by other vendors can also be configured to run in the system.

- ◆ [Section 7.2.1, “Creating Basic, Form-Based, or NMAS Authentication Classes,”](#) on page 75
- ◆ [Section 7.2.2, “Creating an X.509 Authentication Class,”](#) on page 77
- ◆ [Section 7.2.3, “Creating a RADIUS Authentication Class,”](#) on page 80

7.2.1 Creating Basic, Form-Based, or NMAS Authentication Classes

- 1 In the Administration Console, click *Access Manager > Identity Server > Setup > [Configuration] > Local > Classes*.



The following classes are predefined for Access Manager:

Name/Password - Basic: Basic authentication over HTTP using a standard login pop-up screen provided by the Web browser.

Name/Password - Form: Form based authentication over HTTP.

Secure Name/Password - Basic: Basic authentication over HTTPS using a standard login screen provided by the Web browser.

Secure Name/Password - Form: Form based authentication over HTTPS.

- Click *New* to launch the Create Authentication Class Wizard.

Identity Servers ► spb-unstable ►

Create Authentication Class
?

Step 1 of 2: Specify name and java class.

Display name:

Java class:

Java class path: com.novell.nidp.authentication.local.PasswordClass

- Specify a display name, then select one of the following classes from the *Java class* drop-down menu. The following classes are recommended only for testing purposes:

BasicClass: Uses basic HTTP authentication.

PasswordClass: Passes the user name and password over HTTP in readable text, and uses a form-based login to collect the name and password.

RadiusClass: For a production environment, use ProtectedRadiusClass. See [Section 7.2.3, “Creating a RADIUS Authentication Class,” on page 80](#) for configuration steps.

- For a production environment, select one of the following protected classes:

X509Class: See [Section 7.2.2, “Creating an X.509 Authentication Class,” on page 77](#).

ProtectedBasicClass: The BasicClass, protected by HTTPS.

ProtectedPasswordClass: The PasswordClass, protected by HTTPS (form-based).

ProtectedRadiusClass: The RadiusClass, protected by HTTPS.

NMASAuthClass: The authentication class used for Novell® Modular Authentication Service (NMAS™), which uses fingerprint and other technology as a means to authenticate a user. See [NMAS 3.1.1 \(http://www.novell.com/documentation/nmas311/\)](http://www.novell.com/documentation/nmas311/) for information about NMAS.

Other: Used for third-party authorization classes or if you have written your own Java class.

- Click *Next* to configure the properties for each class. The values you enter are case sensitive.

Class	Property Name	Property Value
BasicClass	Query	As an example, if you specify the property value of <i>(&(objectclass=person)(email=%EMail Value%))</i> , the %EMail Value% is replaced with the name entered in the basic authentication login.
PasswordClass	Query	The property value of the Query must be a valid LDAP query string.

Class	Property Name	Property Value
	JSP This property name must be the name of a new JSP file that includes all the needed fields for the Query property. The property value of this attribute should not include the <code>.jsp</code> extension of the file. For example, if you create a new JSP file named <code>login2.jsp</code> then the value of the JSP property would be <code>login2</code> .	The property value for JSP is the name of the JSP page you customized. For example, if you use <code>(&(objectclass=person)(cn=%Ecom_User_ID%)(mail=%Ecom_Email%))</code> as the property value, the system queries for an object of type <code>person</code> that contains a <code>cn</code> equal to <code>Ecom_User_ID</code> from the specified <code>.jsp</code> file, and <code>mail</code> equal to <code>Ecom_Email</code> from the same <code>.jsp</code> file.
NMASAuthClass	Type <code>NMAS_LOGIN_SEQUENCE</code> in the <i>Property Name</i> field.	Specify the name of the NMAS Login Sequence to be used for this type of authentication.

See [Section 7.6, “Creating Custom Login Pages,” on page 85](#) for more information.

7.2.2 Creating an X.509 Authentication Class

The X509 authentication class lets you authenticate users using X509 based certificates. It also identifies the user in user-stores, employing various user-mapping mechanisms.

- 1 In the Administration Console, click *Access Manager > Identity Server > Setup > [Configuration] > Local > Classes*.
- 2 Click *New*.
- 3 Enter a display name, then choose X509Class from the drop-down menu.

Create Authentication Class
?

Step 1 of 2: Specify name and java class.

Display name:

Java class:

Java class path: `com.novell.nidp.authentication.local.X509Class`

4 Click *Next*.

The screenshot shows the 'Create Authentication Class' configuration page, Step 2 of 3: Specify validations. The 'Validations' dropdown is set to 'OCSP-CRL'. Below this, there are three sections:

- CRL Validation:** Includes a checkbox for 'Map X500 CRL to LDAP' and an 'LDAP URL:' text input field.
- OCSP Validation:** Includes checkboxes for 'Sign OCSP request' and 'Use configured OCSP responder URL', and a 'URL:' text input field. There is also a checkbox for 'Disable Root CA revocation check'.
- Trust Stores:** A table with 2 items. The items are 'NIDP Trust Store' and 'OCSP Trust Store', both with blue underlined text.

At the bottom of the form are three buttons: '<< Back', 'Next >>', and 'Cancel'.

5 Configure the following options:

Validations: The validation type. Trust validation occurs if the certificate chain is trusted by verified in the *NIDP Trust Store*. In addition to usual certificate validations, the Identity Server supports CRL (certificate revocation list) and OCSP (Online Certificate Status Protocol) validations for each authentication request.

CRL: Checks the CRL. If you enable CRL validations, the CRL distribution point extension is read out of the user's X.509 certificate. The CRL distribution point contains URL where the complete CRL can be found, as published by the certificate authority. The system performs sanity checks on the CRL itself and then checks to see if the user certificate is in the revoked list. The system can get the CRL over HTTP and LDAP. If you are not expecting the distribution point in user certificates, you can specify a value in the *LDAP URL* to get the CRL.

OCSP: If OCSP validation is enabled, the Authority Info Access point (AIA) is read out of the user certificate, which contains the URL for the OCSP responder. A signed OCSP request for the user certificate is sent to OCSP responder. A signed OCSP response is received from the responder which has the revoked status for the user certificate. Alternately, if you are not expecting AIA in user certificate, you can specify a value in the OCSP responder *URL* field. The value you enter here overrides any OCSP responder URLs in a certificate.

You can specify either CRL or OCSP validations or both in order as defined. These validations can also be turned off. Configuration for user-mappings along with CRL Validations will be in AuthN Class page in Admin. The default setting is to carry out OCSP first, then CRL.

6 Click *Next*.

7 Configure attribute mappings.

Create Authentication Class [?]

Step 3 of 3: Specify attribute mappings.

Show certificate errors

Attributes: [Subject name]

Available attributes: [Directory name, Email, Serial number and issuer name]

Attribute Mappings:

Directory name:

Email:

Serial number and issuer name:

Subject name:

Use this page to specify attribute mappings for the X.509 authentication class. *Subject name* is the default map.

Show certificate errors: Displays an error page when a certificate error occurs.

Attributes: The list of attributes currently in use. You can arrange the order in which the system processes attributes.

Available attributes: The available X.509 attributes. You can select attributes and move them to the *Attributes* field to specify if they are used.

Directory name: Uses *sasAllowableSubjectNames* when using a directory name for mapping.

Email: The LDAP attribute to use for e-mail mapping. By default, this is the LDAP name *mail*.

Serial number and issuer name: Lets you map a user's certificate using the serial number and issuer name. The issuer name and the serial number must be put into the same LDAP attribute of the user.

When using a Case Ignore String attribute, both the issuer name and the serial number must be in the same attribute separated by a dollar sign (\$) character. The issuer name must 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...*) 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*.)

Subject name: Maps the subject name of the client certificate to the *sasAllowableSubjectName*.

The LDAP attribute can be any Case Ignore List or Case Ignore String attribute of the user. If you are configuring your own attribute, ensure that the attribute is added to the Person class. When using a Case Ignore List attribute, both the issuer name and the serial number must 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.

The certificate number is displayed in Internet Explorer with a space after every fourth digit. However, you should enter the certificate number without using spaces.

8 Click *Finish*.

7.2.3 Creating a RADIUS Authentication Class

RADIUS enables communication between remote access servers and a central server. Secure token authentication through RADIUS is possible because Access Manager works with Novell Modular Authentication Service (NMAS) RADIUS software that can run on an existing NetWare® server. Access Manager 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.

- 1 In the Administration Console, click *Access Manager > Identity Server > Setup > [Configuration] > Local > Classes*.
- 2 Click *New*.
- 3 Enter a display name, then select *RadiusClass* from the drop-down menu.
- 4 Click *Next*.

Create Authentication Class ?

Step 2 of 2: Specify properties.

Servers

New | Delete | ↑ | ↓ 0 Item(s)

Server

No items

Port: 1812

Shared secret:

Reply time: 7000 milliseconds

Resend time: 2000 milliseconds

Failed server retry: 5 minutes

JSP:

Require password

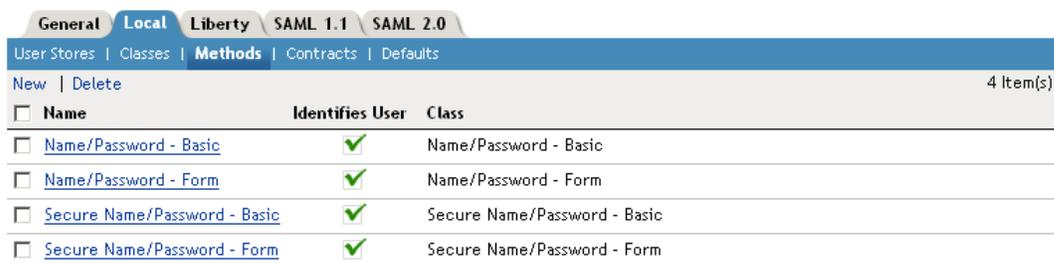
- 5 Click *New* to add an IP address of the RADIUS server. You can add additional servers for failover purposes.
- 6 Click *OK*.
- 7 Fill in the following fields:
 - Port:** The port of the RADIUS server.
 - Shared Secret:** The RADIUS shared secret.
 - Reply Time:** The total time to wait for a reply in milliseconds
 - Resend Time:** The time to wait in milliseconds between requests.
 - Server Failure Retry:** The time in milliseconds which must elapse before a failed server is retried.
 - JSP:** The Java Server Page for the Java program executed by the Web server. Specify the name of the java server page if you want to use something other than the provided JSP. The default page is used if nothing is specified.
 - ♦ **Require Password:** Specifies whether to require a JSP password.
- 8 Click *OK*.

7.3 Configuring Authentication Methods

Authentication methods let you associate authentication classes with user stores. You use a particular authentication class to obtain credentials about an entity, and then validate those credentials against a list of user stores. After credentials are obtained, each user store is checked in order to validate the entity.

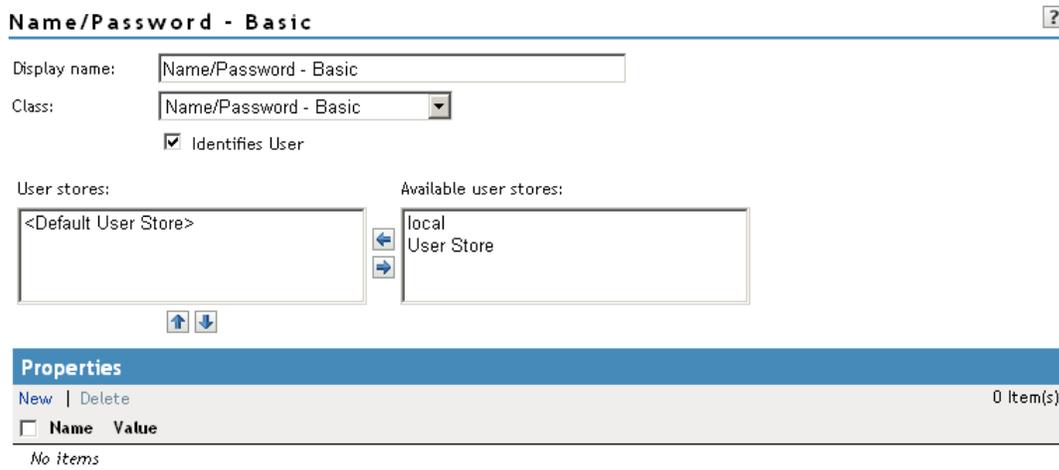
After the system locates the entity in a particular user store, no further checking occurs, even if the credentials fail to validate the entity. Typically, the entity being authenticated is a user, and the definition of an authentication method specifies whether this is the case. You can alter the behavior of an authentication class by specifying properties (name/value pairs) that override those of the authentication class.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Local > Methods*.



<input type="checkbox"/> Name	Identifies User	Class
<input type="checkbox"/> Name/Password - Basic	✓	Name/Password - Basic
<input type="checkbox"/> Name/Password - Form	✓	Name/Password - Form
<input type="checkbox"/> Secure Name/Password - Basic	✓	Secure Name/Password - Basic
<input type="checkbox"/> Secure Name/Password - Form	✓	Secure Name/Password - Form

- 2 Click one of the predefined authentication methods, or click *New* to create one.



Name/Password - Basic ?

Display name:

Class:

Identifies User

User stores:

Available user stores:

Properties

<input type="checkbox"/> Name	Value
No items	

- 3 Fill in the following fields:

Display Name: The name to be used to refer to the new method.

Class: The authentication class to use for this method. See [Section 7.2, “Creating Authentication Classes,”](#) on page 75.

Identifies User: Resolves to a user in the directory when credentials are provided. If this is not enabled, only the authentication is validated, such as the authentication of a computer. If multiple methods identify the user during the user session, all methods that identify the user must identify the same user in order for authentication to succeed.

4 Add user stores to search.

You can select from the list of all the user stores you have setup. If you have several user stores, the system searches through them based on the order specified here.

<**Default User Store**>: The default user store in your system. See [Section 7.5, “Specifying Authentication Defaults,” on page 84](#).

5 Under Properties, click *New*, then fill in the following fields:

Property Name: The name of the class property to be set. This value is case sensitive.

Examples of property names include:

- ♦ **Query:** Functions similar to form-based authentication. In the defined query, the system looks for a value between two %% and substitutes the value from the basic authentication name field on a login page.

For example, if you set the Query property to: (&(objectclass=person)(email=%EMail Value%)), the %Email Value% is replaced with the name entered in the basic authentication dialog box.

- ♦ **RADIUS_LOOKUP_ATTR:** Defines an LDAP attribute whose value is read and used as the ID is passed to the radius server. If not specified, the user name entered is used.
- ♦ **NAS_IP_ADDRESS:** Specifies an IP address used as a RADIUS attribute. You might use this property for situations in which service providers are using a cluster of small network access servers (NASs). The value you enter is sent to the RADIUS server.

Property Value: The values associated with the Property Name field.

You can use these authentication method property names and values to override the property settings specified on the authentication class (see [Section 7.2, “Creating Authentication Classes,” on page 75](#)) and to create custom login pages (see [Section 7.6, “Creating Custom Login Pages,” on page 85](#)).

6 Click *OK*.

7.4 Configuring Authentication Contracts

Authentication contracts define how authentication occurs. An Identity Server configuration might have several authentication contracts available, such as name/password or X.509. Resources at an Access Gateway or agent are protected by authentication contracts.

- ♦ [Section 7.4.1, “Creating a Local Contract,” on page 82](#)
- ♦ [Section 7.4.2, “Creating an External Contract,” on page 83](#)

NOTE: You cannot delete a contract if it is in use by an Access Gateway or J2EE agent.

7.4.1 Creating a Local Contract

Local contracts are executed by the identity provider when authenticating a user. A URI uniquely identifies each contract, and you can assign authentication methods to each contract. A single contract can be specified for local logins.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Local > Contracts*

2 Click *New > Local Contract*.

Create Authentication Contract ?

Configuration

Display name:

URI:

Password expiration servlet:

If you add more than one X509 method, only the first one will be used and it will automatically be moved to the top of the list.

Methods:

Available methods:

↑ ↓

3 Fill in the following fields:

Display name: The name of the local authentication contract.

URI: A value that uniquely identifies the contract from all other contracts. For example, as an identity provider, you might want to publish the details of a contract. In this case, you can use a URL so that the link resolves to a page. No spaces can exist in the URI field.

Password expiration servlet: A URL to a page where the user can change his or her password. This applies only to eDirectory™ servers and when the password is expired or within the grace login period. You must use eDirectory to change the number of grace logins.

Methods and Available Methods: The authentication method to use for the local contract. You can specify the order that the methods are executed for login; however, this is not a graded list, so all the methods you specify are required. *Available methods* are the authentication methods you have set up.

If you add more than one X.509 method, only the first one is used and is automatically moved to the top of the list.

When choosing a secure method, such as Secure Name/Password, ensure that you have enabled security for the Identity Server configuration by setting the protocol to HTTPS. See [Section 5.7, “Configuring Secure Communication on the Identity Server,” on page 63](#).

4 Click Finish.

7.4.2 Creating an External Contract

You use external contracts when your server can be authenticated by an external identity server. If your server and an external server are capable of performing an authentication, they can have the same local and external URI.

Additionally, the URI of the external authentication contract on the service provider must match the URI of the local authentication contract on the identity provider.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Local > Contracts*

- 2 Click *New > External Contract*.

- 3 Fill in the following fields:

Display name: The name for the external authentication contract.

URI: The unique identifier that references the contract at an external Identity Server. For example, as an identity provider, you might want to reference the details of a contract. In this case, you can use a URL so that the link resolves to a page. No spaces can exist in the URI field.

- 4 Click *Finish*.

7.5 Specifying Authentication Defaults

You can specify default values for how the system processes user stores and authentication contracts. The default contract is executed when users access the system without a specified contract, and when the Access Gateway is configured to use any authentication.

Additional default contracts can be specified for each authentication type that might be required by a service provider. These contracts are executed when a request for a specific authentication type comes from a service provider.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Configurations > [Configuration] > Local > Defaults*

Authentication Type	Default Contract
Name Password:	<None>
Secure Name Password:	<None>
X509:	<None>
Smart Card:	<None>
Smart Card PKI:	<None>
Token:	<None>

- 2 Configure the following fields as necessary:

User Store: The default user store for local authentication. If you selected *<Default Userstore>* when configuring an authentication method, the system uses the user store you specify here.

Authentication Contract: The default authentication contract to be used for local authentication. If you create a new contract and specify it as the default one, ensure that you

update the Access Gateway configuration if it is configured to accept the default (*Any*) contract. See [Section 12.4, “Configuring Protected Resources,” on page 150](#).

Authentication Type: The default authentication contracts to be used for each authentication type. The identity provider uses the default authentication contract specified here, when the identity provider receives an authentication request from a service provider for a specific authentication type.

You must create the authentication contracts prior to assigning them as defaults. (See [“Configuring Authentication Contracts” on page 82](#).)

3 Click *OK*.

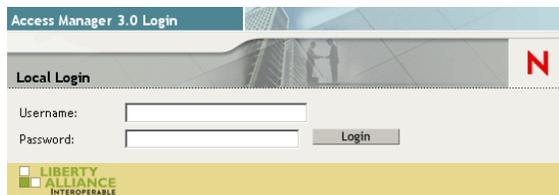
7.6 Creating Custom Login Pages

You can create custom login pages that refer to the Identity Server. You might want to rebrand the User Portal, authenticate users with non-default attributes (cn), or authenticate users based on multiple LDAP attributes. You also might be fronting several protected resources with an Access Gateway, and you need to create a unique login for each page.

- ◆ [Section 7.6.1, “Modifying the User Portal Page,” on page 85](#)
- ◆ [Section 7.6.2, “Creating Your Own Login or Error Page,” on page 86](#)

7.6.1 Modifying the User Portal Page

The following page is the default login page to the User Portal provided by the Access Manager. This page has been designed for the form-based authentication class.



Access Manager uses a JSP file as the default login page. You must be familiar with customizing `.jsp` files when creating custom login pages. The location of the `.jsp` file is:

```
/var/opt/novell/tomcat4/webapps/nidp/jsp.
```

You use the property name and values in the authentication classes and methods to customize the login. (See [Section 7.2, “Creating Authentication Classes,” on page 75](#).)

The Radius and Protected classes also support a JSP property. You can use other classes, but if you want to create a custom login page, you must select a class that supports the JSP property. You can add this property to either the class or to the method derived from the class.

Property Names and Values

The default Property Name is `Ecom_User_ID` with a value of `cn`. You could, for example, change this to `Ecom_User_ID` and `mail` if you want to authenticate using the user’s email address. If you want to authenticate with the current username and password credentials, as well as the user’s email address, you could modify the login page with an additional field in the form. For example:

```
input type="text" class="smalltext" name="Ecom_User_eMail" size="30"
```

You could then add an `Ecom_User_eMail` property name with “mail” as the property value. This is an example of an AND-based authentication request where you add the username AND user e-mail AND user password. You can OR fields if you add a Query property with a value similar to

```
(&(objectclass=person)(|(cn=%Ecom_User_ID%)(mail=%Ecom_User_Email%))).
```

This entry allows you to add a field to the login page and allow the user to log in with a username or an e-mail address.

7.6.2 Creating Your Own Login or Error Page

The easiest way to create a new login page is to copy the default JSP page, rename it, and then modify it to match your requirements.

Login requirements:

- ◆ **Post Action:** `https://IdentityServerDNS:8443/nidp/app/login`
- ◆ **User name input type = “text”:** `name=“Ecom_User_ID”`
- ◆ **User password input type = “password”:** `name=“Ecom_Password”`
- ◆ **Optional input type = “hidden”:** `name=“target”` with a value of a destination URL.

The default authentication contract is used if the post comes from an external page.

Logout links:

- ◆ **Identity Server:** `https://IdentityServerDNS:8443/nidp/app/logout`
- ◆ **Access Gateway:** `https://AGAuthDomain/nesp/app/plogout`

The location of the log out page for the Access Gateway:

- ◆ For the NetWare™ Access Gateway: `sys:tomcat\4\webapps\nesp\jsp\`
- ◆ For the Linux Access Gateway: `/var/opt/novell/tomcat4/webapps/nesp/jsp/`

To create a custom error page, you must modify the `err.jsp`. The location of this file is:

```
/var/opt/novell/tomcat4/webapps/nidp/jsp.
```

See Also

[Section 40.1.2, “Prompting Users to Refresh a Session before Expiration,” on page 469](#) in [Troubleshooting](#).

Configuring Trusted Providers

8

This section discusses configuring trust and how to reference internal and external trusted identity providers, service providers, and embedded service providers (ESPs). Steps for configuring trusted provider types are similar, and are also similar between the Liberty and SAML protocols. The interface pages in this section show the configuration of a Liberty trusted service provider.

- ◆ [Section 8.1, “Understanding the Trust Model,” on page 87](#)
- ◆ [Section 8.2, “Creating a Trusted Provider Reference,” on page 89](#)
- ◆ [Section 8.3, “Reimporting a Trusted Provider’s Metadata,” on page 91](#)
- ◆ [Section 8.4, “Editing a SAML 1.1 Trusted Identity Provider’s Metadata,” on page 92](#)
- ◆ [Section 8.5, “Editing a SAML 1.1 Trusted Service Provider’s Metadata,” on page 93](#)
- ◆ [Section 8.6, “Configuring Common Access Settings for a Trusted Provider,” on page 94](#)
- ◆ [Section 8.7, “Selecting Attributes for a Trusted Provider,” on page 97](#)

About SAML and Liberty

If you are unfamiliar with SAML 1.1, see [SAML Overview \(http://www.novell.com/documentation/saml/saml/data/ag8qdk7.html\)](http://www.novell.com/documentation/saml/saml/data/ag8qdk7.html) on the [Documentation \(http://www.novell.com/documentation/a-z.html\)](http://www.novell.com/documentation/a-z.html) Web site.

For conceptual information about Liberty, and to learn about what is new for SAML 2.0, see [Appendix A, “About Liberty and SAML 2.0,” on page 519](#).

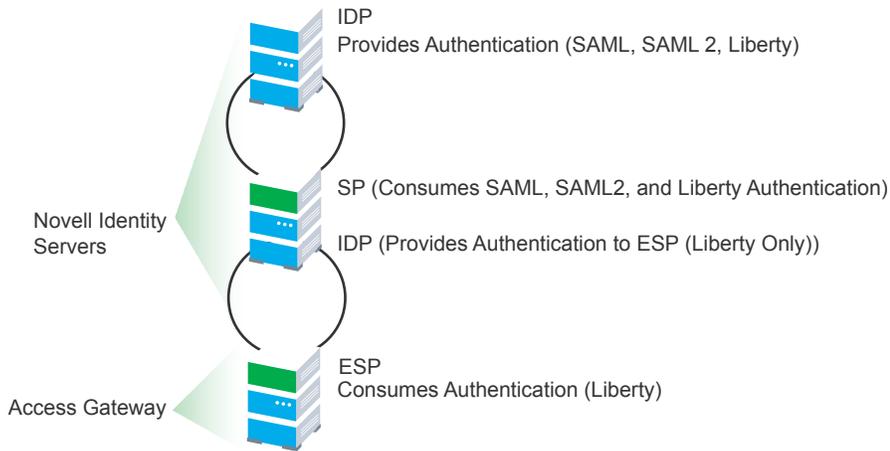
8.1 Understanding the Trust Model

Setting up trust involves system administrators agreeing on how to establish a secure method for providing and consuming authentication assertions between their Identity Servers. An Identity Server is always installed as an identity provider, which is used to provide authentication to trusted service providers and **embedded service providers** (ESPs).

An Identity Server also can be configured as an identity consumer (service provider), which enables the Identity Server to consume authentication assertions from trusted identity providers. [Figure 8-1](#)

depicts how two Identity Servers can be configured in a trust model using the SAML and Liberty protocols, in order to authenticate to an Access Gateway ESP.

Figure 8-1 Identity Server Trust

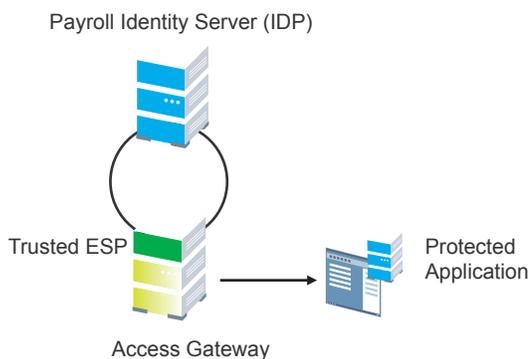


As an administrator, you determine whether your server is to be used as the identity provider or service provider in the trust relationship. You and the trusted partner agree to exchange Identity Server metadata, and then you create references to the trusted partner's Identity Server in your Identity Server configuration. You can obtain metadata via URL or XML document, then enter it in the system when you create the reference.

Embedded Service Providers

In addition to setting up trust with internal or external service providers, you can reference embedded service providers (ESPs) in your enterprise. An ESP uses the Liberty protocol and does not require metadata entry. The ESP comes with Access Manager and is embedded in the Access Gateway or application server agent (such as a J2EE agent). The ESP facilitates authentication between the Identity Server and the resource protected by the Access Gateway or agent, as shown in [Figure 8-2](#).

Figure 8-2 Embedded Service Provider



The components in this example reside in the same trust store and represent a typical Access Manager configuration used within an enterprise.

High-Level Steps

The following high-level steps describe what both administrators would perform to set up the trust model between an identity provider and a service provider. These steps assume that both providers are using the Novell[®] Identity Server provided with Access Manager.

1. Administrators at each company install and configure the Identity Server.

The Identity Server that consumes authentications must be enabled to run as a service provider. See [Section 5.1, “Creating an Identity Server Configuration \(Advanced Options\),” on page 53](#). (It is recommended that you are already familiar with the *Novell Access Manager 3.0 Installation Guide*.)

2. Administrators must exchange Identity Server metadata with the trusted partner.

Metadata is generated by the Identity Server and can be obtained via URL or XML document, then entered in the system when you create the reference. This step is not applicable if you are referencing an ESP. When you reference an ESP, the system lists the installed ESPs for you to choose, and no metadata entry is required.

3. Create the reference to the trusted identity provider or service provider.

This procedure associates the metadata with the new provider. See [Section 8.2, “Creating a Trusted Provider Reference,” on page 89](#).

4. Configure user authentication.

This procedure defines how your Identity Server interacts with the trusted provider during user authentication. Access Manager comes with default basic authentication settings already enabled. See [Chapter 9, “Configuring User Authentication and Federation,” on page 99](#).

Additional important steps for enabling authentication between trusted providers include:

- ♦ Setting up the necessary authentication contracts. See [Section 7.4, “Configuring Authentication Contracts,” on page 82](#).
- ♦ Enabling the profiles that you are using. See [Section 11.2, “Enabling Web Services and Profiles,” on page 116](#).
- ♦ Enabling the *Always Allow Interaction* option on the Web Service Consumer page. See [Section 11.8, “Configuring the Web Service Consumer,” on page 126](#).

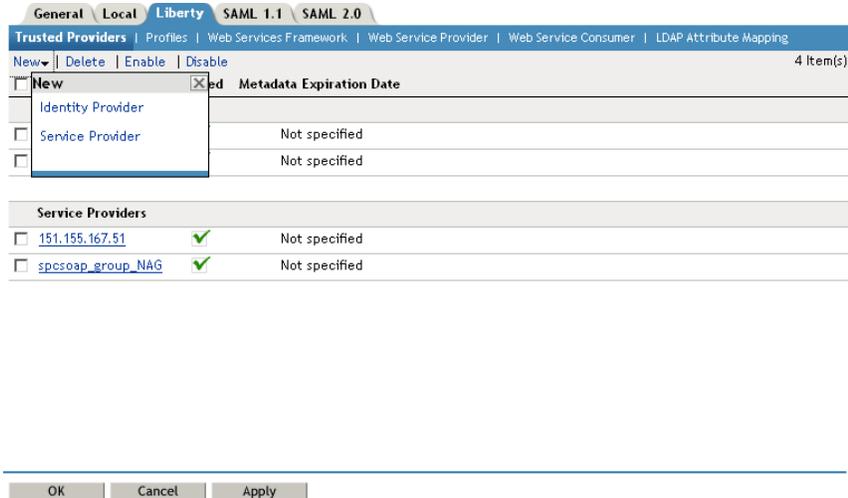
8.2 Creating a Trusted Provider Reference

The procedure for establishing trust between providers begins with obtaining metadata for the trusted provider. If you are using the Novell Identity Server, protocol-specific metadata is available via URL. Examples of metadata URLs for server 10.1.1.1 would be:

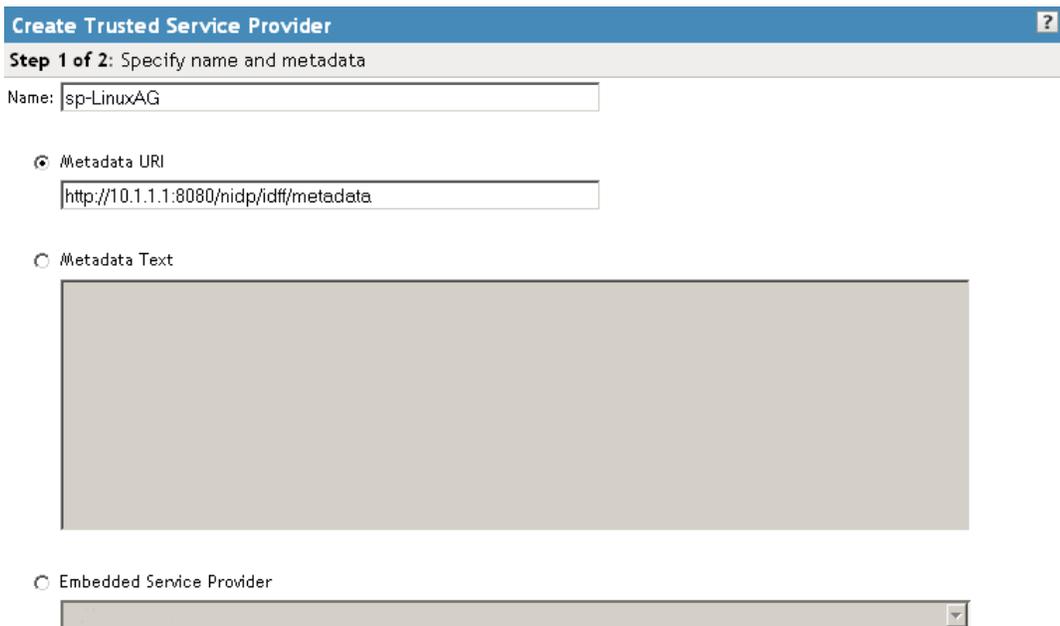
- ♦ **Liberty:** <http://10.1.1.1:8080/nidp/idff/metadata>
- ♦ **SAML 1.1:** <http://10.1.1.1:8080/nidp/saml/metadata>
- ♦ **SAML 2.0:** <http://10.1.1.1:8080/nidp/saml2/metadata>

The default values nidp and 8080 are established during product installation. Nidp is the Tomcat application name.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > [Protocol]*.



- 2 Click *New*, then click *Identity Provider* or *Service Provider*.



- 3 Fill in the following fields:

Name: The name by which you want to refer to the provider.

Metadata URL: The metadata URL for a trusted provider. The system retrieves protocol metadata using this URL.

Other methods for entering metadata are provided, depending on the type of provider you are creating, and the protocol used:

Metadata Text: An editable field in which you can paste copied metadata text from an XML document, assuming you obtained the metadata via e-mail or disk and are not using a URL. If you copy metadata text from a Web browser, you must copy the text from the page source.

Manual Entry: (SAML 1.1 only) Allows you to enter metadata values manually. When you select this option, the system displays the Enter Metadata Values page. See [Section 8.4, “Editing a SAML 1.1 Trusted Identity Provider’s Metadata,”](#) on page 92.

- 4 If you are creating a service provider for an Access Gateway or agent, click the following option:

Embedded Service Provider: Access Gateway and application server agents (J2EE or Windows) include an embedded service provider (ESP) that can be trusted by identity providers. ESPs run in the same enterprise as the identity provider, and are therefore created and configured in the same directory. The ESP enables all of the single-sign on functionality for Access Gateway or agent. Installed ESPs are displayed in a drop-down list for you to select as a trusted entity. You do not need to enter metadata for an ESP; it is automatically generated.

- 5 Click *Next*.
- 6 Review the metadata certificates, then click *Finish*.
- 7 The system displays the trusted provider on the Liberty page.

The screenshot shows the Liberty SAML 1.1 configuration page. At the top, there are tabs for 'General', 'Local', 'Liberty', 'SAML 1.1', and 'SAML 2.0'. Below the tabs, there is a navigation bar with links for 'Trusted Providers', 'Profiles', 'Web Services Framework', 'Web Service Provider', 'Web Service Consumer', and 'LDAP Attribute Mapping'. A 'New' dropdown menu is visible, along with 'Delete', 'Enable', and 'Disable' buttons. The main content area displays a table of trusted providers, categorized into 'Identity Providers' and 'Service Providers'. Each provider has a checkbox, a name, an 'Enabled' status (indicated by a green checkmark), and a 'Metadata Expiration Date' (all listed as 'Not specified').

<input type="checkbox"/>	Name	Enabled	Metadata Expiration Date
Identity Providers			
<input type="checkbox"/>	idp	✓	Not specified
<input type="checkbox"/>	idp-corporate	✓	Not specified
Service Providers			
<input type="checkbox"/>	10.10.157.30	✓	Not specified
<input type="checkbox"/>	sp-LinuxAG	✓	Not specified
<input type="checkbox"/>	spsoap_group_NAG	✓	Not specified

8.3 Reimporting a Trusted Provider’s Metadata

You might need to reimport a trusted provider’s metadata if you learn that it has changed, or if you have changed the base URL of the Identity Server configuration. The steps to do this are similar for Liberty and SAML protocols.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Configuration > [Liberty or SAML] > [Trusted Provider] > Metadata*.
- 2 Click the trusted provider, then click the *Metadata* tab.
- 3 Click *Reimport*.
Follow the prompts to import the metadata.
- 4 Specify the new metadata information as described in [Section 8.2, “Creating a Trusted Provider Reference,”](#) on page 89.
- 5 Confirm metadata certificates, then click *Finish*.

8.4 Editing a SAML 1.1 Trusted Identity Provider's Metadata

Access Manager allows you to obtain metadata for SAML 1.1 providers. However, metadata for SAML 1.1 might not be available for some trusted providers. Therefore, you can enter metadata manually. The page for this is available if you clicked the *Manual Entry* option when you **created the trusted provider**.

1 In the Administration Console, click *Access Manager > Identity Servers > Configuration > SAML 1.1 > Provider Metadata*.

2 To reimport the metadata from a URL or text:

2a Click *View*, then click *Reimport*.

The system displays the Create Trusted Identity Provider Wizard that lets you obtain the metadata. Follow the on-screen instructions to complete the steps in the wizard.

3 To edit the metadata manually, click *Edit*.

The screenshot shows a web form for editing SAML 1.1 provider metadata. It includes several text input fields and two 'Browse...' buttons. The fields are: 'Provider ID: *' (required), 'Source ID:', 'Metadata expiration:' (with a calendar icon), 'SAML attribute query URL:', and 'Artifact resolution URL:'. Below these is a section titled 'Signing Certificates' containing 'Attribute authority:' and 'Identity provider: *' (required), each with a 'Browse...' button.

4 Fill in the following fields as necessary:

Provider ID: (Required) The SAML 1.1 metadata unique identifier for the provider. For example, *https://dns.name:port/nidp/saml/metadata*.

Source ID: The SAML Source ID for the trusted provider. The Source ID is a 20-byte value that is used as part of the Browser/Artifact profile. It allows the receiving site to determine the source of received SAML Artifacts. If none is specified, the Source ID is auto-generated using a SHA-1 hash of the site provider ID.

Metadata expiration: The date upon which the metadata is no longer valid.

SAML attribute query URL: The URL location where an attribute query is to be sent to the partner. The attribute query requests a set of attributes associated with a specific object. A successful response contains assertions that contain attribute statements about the subject. A SAML 1.1 provider might use the base URL, followed by */saml/soap*. For example, *https://[dns:port]/nidp/saml/soap*.

Artifact resolution URL: The URL location where artifact resolution queries are sent. A SAML Artifact is included in the URL query string. The target URL on the destination site the user wants to access is also included on the query string. A SAML 1.1 provider might use the base URL, followed by */saml/soap*. For example, *https://[dns:port]/nidp/saml/soap*.

5 To specify signing certificate settings, fill in the following fields:

Attribute authority: The signing certificate of the partner SAML 1.1 attribute authority. The attribute authority relies on the identity provider to provide it with authentication information so that it can retrieve attributes for the appropriate entity or user. The attribute authority must know that the entity requesting the attribute has been authenticated to the system.

Identity provider: (Required) Appears if you are editing identity provider metadata. This field specifies the signing certificate of the partner SAML 1.1 identity provider. It is the certificate the partner uses to sign authentication assertions.

8.5 Editing a SAML 1.1 Trusted Service Provider's Metadata

Access Manager allows you to obtain metadata for SAML 1.1 providers. However, metadata for SAML 1.1 might not be available for some trusted providers. Therefore, Access Manager allows you to enter metadata manually. The page for this is available if you clicked the *Manual Entry* option when you **created the trusted provider**.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Configuration > SAML 1.1 > [Service Provider] > Metadata*.
- 2 If you want to reimport the metadata, click *View*, then click *Reimport*. Follow the on-screen instructions to complete the steps in the wizard.
- 3 Click *Edit*.

The screenshot shows a web form for editing SAML 1.1 metadata. It includes the following fields and options:

- Provider ID:** A required text input field with an asterisk.
- Metadata expiration:** A text input field with a calendar icon to its right.
- Want Assertion to be signed**
- Artifact consumer URL:** A text input field.
- Post Consumer URL:** A text input field.
- Signing Certificate** section with a sub-label **Service provider:** and a required text input field with an asterisk, followed by a **Browse...** button.

4 Fill in the following fields:

Provider ID: (Required) The SAML 1.1 metadata unique identifier for the provider. For example, *https://dns.name:port/nidp/saml/metadata*.

Metadata expiration: The date upon which the metadata is no longer valid.

Want assertion to be signed: Specifies that authentication assertions from the trusted provider must be signed.

Artifact consumer URL: Where the partner receives incoming SAML artifacts. For example, *https://[dns:port]/nidp/saml/spassertion_consumer*.

Post consumer URL: Where the partner receives incoming SAML POST data. For example, *https://[dns:port]/nidp/saml/spassertion_consumer*.

Service Provider: The public key certificate used to sign SAML data. You can browse to locate the service provider certificate.

- 5 Click *Finish*.

8.6 Configuring Common Access Settings for a Trusted Provider

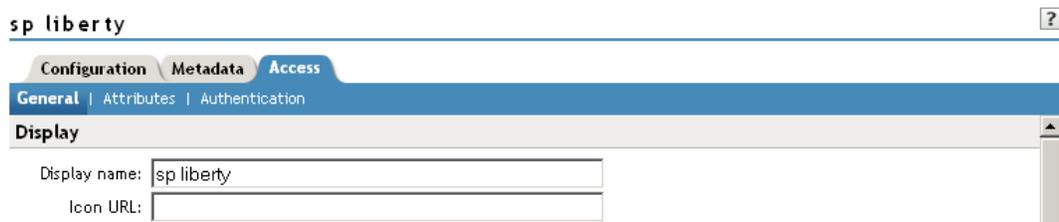
Common access settings for a trusted provider include specifying how to display the provider on the Identity Server's login page, and specifying security settings for communication.

- ♦ [Section 8.6.1, "Configuring Display and Access Settings," on page 94](#)
- ♦ [Section 8.6.2, "Configuring Communication Security Settings," on page 95](#)
- ♦ [Section 8.6.3, "Specifying the Intersite Transfer Service URL," on page 96](#)

8.6.1 Configuring Display and Access Settings

You can configure how you want to display the provider on the Identity Server's login page. The fields that are displayed on this page vary depending on the protocol and provider type you selected for configuration.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > [Protocol] > [Provider Name]*.
- 2 Click *Access > General*.



The screenshot shows the Administration Console interface for configuring a trusted provider. The browser address bar shows 'sp liberty'. The page has three main tabs: 'Configuration', 'Metadata', and 'Access'. Under the 'Access' tab, there are three sub-tabs: 'General', 'Attributes', and 'Authentication'. The 'Display' section is currently selected, showing two input fields: 'Display name' with the value 'sp liberty' and 'Icon URL' which is empty.

- 3 Fill in the following fields:

Display name: The display name seen by the end user for this trusted provider. The default name is the name you entered when creating the trusted provider.

Icon URL: The URL of the icon to display for this trusted provider. If you add an icon, the system displays the icon as the link, rather than the text in the *Display name* field.

Login URL: (Displayed for a SAML 1.1 trusted identity provider) The URL required by the identity provider to authenticate the user from the service provider.

For Liberty and SAML 2.0, the URL for this **Intersite Transfer service** is automatically generated and can be displayed on the service provider's login page as a link to the identity provider. For SAML 1.1, you manually enter the URL in the *Login URL* field.

Destination URL: (Displayed for a SAML 1.1 trusted service provider) Specifies the target URL used in the SAML 1.1 identity provider's login URL (Intersite Transfer service).

Advertise (Display) on Login Dialog: Displays the identity provider's link on the Login page in the User Portal.

- 4 Click *OK*.

- 5 Click *OK* on the Trusted Providers page.
- 6 Click *Update Servers* on the Setup page.

8.6.2 Configuring Communication Security Settings

You can configure the security settings to control direct communication between the Identity Server and a trusted provider across the SOAP back channel. These methods apply to the trusted identity provider and are similar between Liberty and SAML.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > [Protocol] > [Provider Name]*.
- 2 Click *Access > General*.

- 3 Fill in the following fields:

Encrypt name identifiers: (SAML 2.0 identity provider security). Adds a level of security by encrypting name identifiers so that they cannot be meaningfully interpreted at an intermediate entity.

The following settings specify how to validate messages received from trusted providers over the SOAP back channel.

Message Signing: Specifies no security but rather relies upon message signing using a digital signature.

Mutual SSL: Specifies that this trusted provider provides a digital certificate (mutual SSL) when it sends a SOAP message.

SSL communication requires only the client to trust the server. For mutual SSL, trust must exist between the client and the server. The client provides a certificate containing its identity and private key in its keystore. The client also creates a version of this certificate, which the server

stores in its trust store. In turn, the client needs to trust the server, which is accomplished by importing the server's CA certificate into the client trust store.

Basic Authentication: Standard header-based authentication. This method assumes that a name and password for authentication are sent and received over the SOAP back channel.

Send: The name and password to be sent for authentication to the trusted partner. The partner expects this password for all SOAP back-channel requests, which means that the name and password must be agreed upon.

Verify: The name and password used to verify data that the trusted provider sends.

- 4 Click *OK*.
- 5 Click *OK* on the Trusted Providers page.
- 6 Click *Update Servers* on the Setup page.

8.6.3 Specifying the Intersite Transfer Service URL

The Novell Identity Server provides access to an Intersite Transfer Service for each of the three supported protocols (Liberty, **SAML 1.1** and SAML 2.0). The Intersite Transfer Service directs the Identity Server to authenticate a user at a service provider without being directly requested to do so by the service provider.

You can place the Intersite Transfer Service URL on any Web page that accesses services provided by a service provider that can be authenticated by the identity provider.

The parameters required include:

- ♦ **SAML 1.1:** *[SAML 1.1 IDP Base URL]/saml/idpsend?PID=[The SAML 1.1 SP Provider ID]&TARGET=[final destination URL]*
- ♦ **SAML 2.0:** *[SAML 2.0 IDP Base URL]/saml2/idpsend?PID=[The SAML 2.0 SP Provider ID]&TARGET=[final destination URL]*
- ♦ **Liberty:** *[Liberty IDP Base URL]/idff/idpsend?PID=[The Liberty SP Provider ID]&TARGET=[final destination URL]*

The values in the URL are described below.

PID: The ID of the service provider that consumes the assertion to authenticate a user.

TARGET: The URL the service provider will redirect to after receiving an assertion from the Identity Provider. When Access Gateway is being used, this can be the URL of a protected resource.

8.7 Selecting Attributes for a Trusted Provider

You can select attributes that an identity provider sends and a service provider receives in an authentication. You can also create attribute sets or select attribute sets that you created globally in [Section 6.1, “Configuring Attribute Sets,” on page 65](#).

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Liberty [or SAML] > [Provider] > Access > Attributes*.

Identity Servers > sp-401k > Corporate IDP

Configuration Metadata Access

General | Attributes | Authentication

Attribute set: attributeset1

Obtain at authentication:

Available:

- Common First Name
- Common Last Name
- Every Day Name

OK Cancel Apply

- 2 To create an attribute set, select *New Attribute Set* from the *Attribute Set* drop-down menu. An attribute set is a group of attributes that can be exchanged with the trusted provider. For example, you can specify that the local attribute of any attribute in the Liberty profile (such as Informal Name) matches the remote attribute specified at the service provider.
- 3 Specify a set name, then click *Next*.
- 4 On the Define Attributes page, click *New*.
- 5 Select a local attribute.
- 6 Optionally, you can provide the name of the remote attribute.
- 7 Click *OK*, then click *Finish*.

After you select attributes, the system displays them on the Attributes page.

IDP

Configuration Metadata Access

General | **Attributes** | Authentication

Attribute set:

Obtain at authentication:

- Common First Name
- Common Last Name

Available:

- Every Day Name

You can select attributes from the *Available Attributes* field, and move them to the left side of the page. If you are an identity provider setting up a service provider, the left side of the page is used for attributes to be sent in an assertion to a service provider.

If you are a service provider setting up an identity provider, the attributes that you move to the left side of the page are those you want to be obtained by the service provider during authentication.

Configuring User Authentication and Federation

9

Configuring authentication involves determining how the trusted service provider interacts with the trusted identity provider during user authentication and federation. Examples include when authentication occurs and which authentication contracts to use. You can also configure the identification methods a service provider uses for provisioning unknown users.

- ◆ [Section 9.1, “Configuring Authentication for a Trusted Identity Provider,” on page 99](#)
- ◆ [Section 9.2, “Configuring User Identification Methods,” on page 101](#)
- ◆ [Section 9.3, “Configuring Authentication for a Trusted Service Provider,” on page 108](#)
- ◆ [Section 9.4, “Configuring User Identification Methods for SAML 1.1 Trusted Identity Providers,” on page 109](#)
- ◆ [Section 9.5, “Specifying a SAML Audience URI,” on page 111](#)

9.1 Configuring Authentication for a Trusted Identity Provider

When users authenticate to a service provider, they can be given the option to federate their account identities with their preferred identity provider. This process creates an account association between the identity provider and service provider that enables single sign-on and single logout.

- 1 In the Administration Console, click *Access Manager Identity Servers > Setup > [Configuration] > Liberty [or SAML 2.0] > [Identity Provider] > Access > Authentication*.

2 Click *Authentication*.

The screenshot shows the 'IDP' configuration window with the 'Authentication' tab selected. The 'Federation' section has the 'Allow users to federate' checkbox checked. Below it, 'Allow after authentication' is checked, and 'Allow before authentication' is unchecked. The 'Authentication Context' section has 'Do not specify' selected. The 'Authentication types' section shows 'Secure Name Password' in the 'Types' list and 'Name Password', 'X509', 'Smart Card', and 'Smart Card PKI' in the 'Available types' list. The 'Authentication contracts' section is empty.

Configuration | Metadata | Access

General | Attributes | **Authentication**

Federation Allow users to federate

Allow after authentication

Allow before authentication (Select to view methods)

Authentication Context Use Types Use Contracts Do not specify

Authentication types

Types: Secure Name Password

Available types: Name Password, X509, Smart Card, Smart Card PKI

Authentication contracts

Contracts:

Available contracts:

OK Cancel Apply

3 Enable the following option:

Allow users to federate: Enables account federation. By enabling this option you assume that a user account exists at the service provider and that the account can be associated with a user's account at the identity provider. If you do not use this feature, authentication is permitted but is not associated with a particular user account.

4 Specify when the federation request occurs:

Allow after authentication: Sends the federation request after the user has authenticated (logged in) to the service provider. When you set this option, users can federate from the Federations page in the Access Manager User Portal.

Allow before authentication: Specifies whether federation can occur when the user clicks the login link to the identity provider. Allowing federation in this method means that a user must be identified at a later time during the federation process. For this reason, when you click this option, the system displays additional options on the Authentication page, under *User Identification Methods*.

These options are discussed in [Section 9.2, "Configuring User Identification Methods," on page 101](#).

5 Under *Authentication Context*, configure the following fields:

Use Types: Specifies whether to use authentication types. Select the types from the *Available types* field to specify which type to use for authentication between trusted service providers and identity providers. Standard types include Name/Password, X.509, Token, and so on.

Use Contracts: Specifies whether to use authentication contracts. Select the contract from the *Available contracts* list. An Identity Server configuration might have several authentication contract types available, such as Name/Password or X.509.

Do not specify: Specifies that the identity provider can send any type of authentication to satisfy a service provider's request, and instructs a service provider to not send a request for a specific authentication type or contract.

6 Under the *Options* heading, configure the following fields, as necessary:

Response Protocol Binding: Select *Artifact* or *Post* or *None*. Artifact and Post are the two methods for transmitting assertions between the authenticating system and the target system.

If you select *None*, you are letting the identity provider determine the protocol.

Identity provider proxy redirects: Specifies whether or not the trusted identity provider can proxy the authentication request to another identity provider. A value of zero specifies that the trusted identity provider cannot redirect an authentication request. Values 1-5 determine the number of times the request can be proxied. Select *Configured on IDP* to let the trusted identity provider decide how many times the request can be proxied.

Force authentication at the IDP: Specifies that the trusted identity provider must prompt the user for authentication, even if they are already logged in.

Use automatic introduction: Automatically attempts single sign-on to this trusted identity provider.

7 Click *OK*.

8 On the Trusted Providers page, click *OK*.

9 Update the Identity Server configuration on the *Setup* tab.

9.2 Configuring User Identification Methods

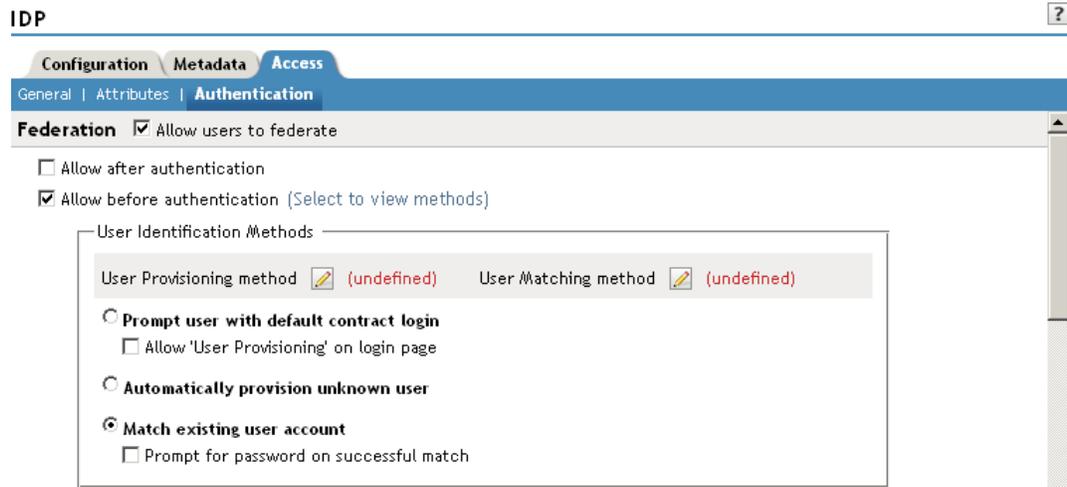
Three methods exist for you to identify users from a trusted identity provider. You can authenticate users by using the default authentication contract, match existing user accounts, or create new account with user provisioning.

- ♦ [Section 9.2.1, “Selecting a User Identification Method,” on page 101](#)
- ♦ [Section 9.2.2, “Configuring the User Matching Method,” on page 102](#)
- ♦ [Section 9.2.3, “Defining the User Provisioning Method,” on page 103](#)
- ♦ [Section 9.2.4, “User Provisioning Error Messages,” on page 107](#)

9.2.1 Selecting a User Identification Method

1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Liberty [or SAML 2.0] > [Identity Provider] > Access > Authentication*.

- 2 Enable *Allow before authentication*, then configure user provisioning or account matching as necessary as described below.



The system displays the following options on the Authentication page under *User Identification Methods*. These options are used to further configure how the service provider can authenticate an unrecognized user.

Authenticate user with default contract: Executes the default authentication contract.

- ♦ **Allow User Provisioning on login page:** Provides a button that the user can click to create an account.

Automatically provision unknown users: Enables a service provider to trust unknown users that have authenticated to the trusted identity provider. User provisioning is used when no user account for federation exists at the service provider.

You must click *User Provisioning Method* to define user provisioning. See [Section 9.2.3, “Defining the User Provisioning Method,”](#) on page 103.

Match existing user accounts: Enables account matching. The service provider can uniquely identify a user in its directory by obtaining specific user attributes sent by the trusted identity provider.

You must click *User Matching Method* to define the match method. See [Section 9.2.2, “Configuring the User Matching Method,”](#) on page 102.

- ♦ **Prompt for password on successful match:** (Optional) Specifies whether to prompt the user for a password when the user’s name is matched to an account, to ensure that the account matches.

- 3 Click *OK*.
- 4 Click *OK* on the Trusted Providers page.
- 5 Click *Update Servers* on the Setup page.

9.2.2 Configuring the User Matching Method

If you enabled the *Match existing user account* option when [selecting an identification method](#), you must configure the matching method.

Before you begin, enable the Liberty Personal Profile. See [Section 11.2, “Enabling Web Services and Profiles,”](#) on page 116.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Liberty [or SAML 2.0] > [Identity Provider Name] > Access > Authentication.*
- 2 Click *Allow before authentication.*
- 3 Click *Match existing user account.*
- 4 Click *User Matching Method.*

Identity Servers ▶ sp-401k ▶ Corporate IDP ▶

User Matching Method ?

Select User Stores to search

User stores:	Available user stores:
Installed User Store	

↑ ↓ ← →

User Matching Expression: <Select User Matching Expression>

If match not found: Do nothing

- 5 Select and arrange the user stores you want to use.
- 6 Set the matching expression as the default, or click *New* to create a look-up expression. See [Section 6.3, “Configuring User Matching Expressions,”](#) on page 67.
- 7 Specify what action to take if no match is found.

You perform account matching before user provisioning, in order to prevent the creation of multiple accounts for one user. If no match is found, you can specify whether to:

- ♦ Do nothing
- ♦ Prompt the user for authentication
- ♦ Automatically provision the user account

- 8 Click *Finish.*
- 9 On the Authentication page, click *OK.*
- 10 On the Trusted Providers page, click *OK.*
- 11 On the Setup page, click *Update Servers* to update the Identity Server configuration.

9.2.3 Defining the User Provisioning Method

If you enabled *Automatically provision unknown users* when [selecting an identification method](#), you must define the user provisioning method. This procedure involves selecting required and optional attributes that the service provider requests from the identity provider during provisioning.

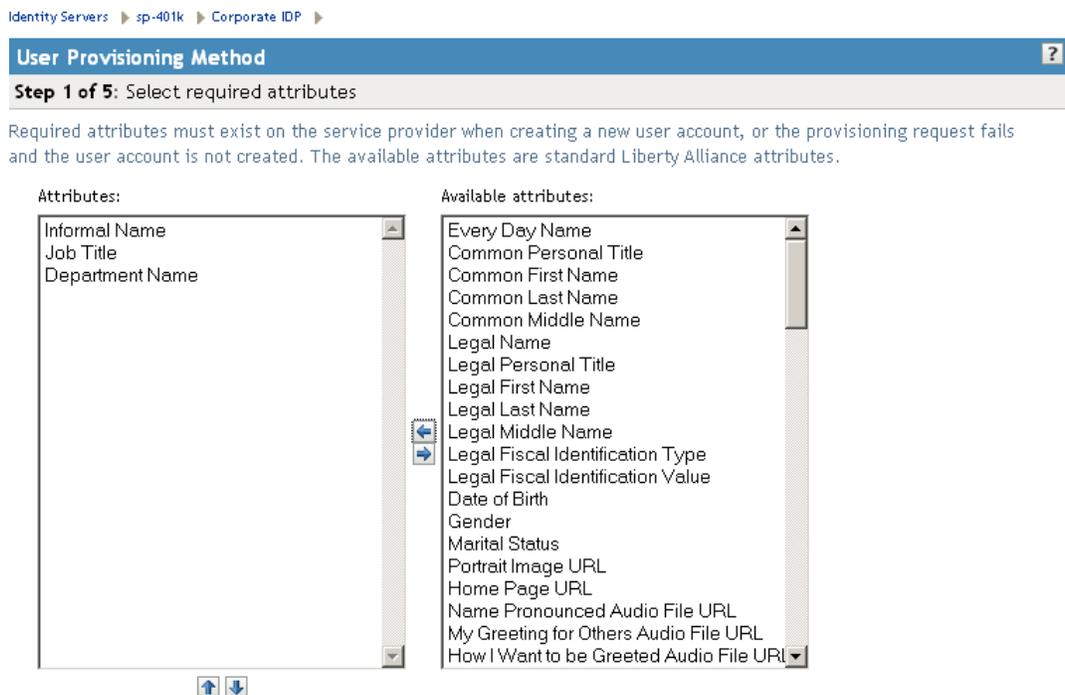
Attribute Considerations

When a user object is created in the directory, some attributes are initially created with the value of NAM Generated. Afterwards, an attempt is made to write the required and optional attributes to the

new user object. Because required and optional attributes are profile attributes, the system checks the write policy for the profile's Data Location Settings (specified in *Liberty > Web Service Provider*) and writes the attribute in either LDAP or the configuration store. In order for the LDAP write to succeed, each attribute must be properly mapped as an LDAP Attribute. Additionally, you must enable the read/write permissions for each attribute in the Liberty/LDAP attribute maps. See [Section 11.9, "Mapping LDAP and Liberty Attributes," on page 127](#).

To configure user provisioning:

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Liberty [or SAML 2.0] > [Identity Provider] > Access > Authentication*.
- 2 Click *Allow before authentication*, then click *User Provisioning method*.



- 3 Select the required attributes from the *Available Attributes* list and move them to the *Attributes* list.

Required attributes are those used in the creation of a user name, or that are required when creating the account.

- 4 Click *Next*.
- 5 Select optional attributes from the *Available Attributes* list and move them to the *Attributes* list.

This step is similar to selecting required attributes. However, the user provisioning request creates the user account whether or not optional attributes exist on the service provider.

- 6 Click *Next*.

7 Define how to create the user name.

Identity Servers > sp-401k > Corporate IDP >

User Provisioning Method ?

Step 3 of 5: Define user name creation

Selecting an attribute for the user name segments from the required attributes list will improve the chances the new user name will be created.

Maximum length: character(s)

Prompt for user name

Automatically create user name

Segment 1: Length: character(s)

Junction:

Segment 2: Length: character(s)

Ensure name is unique

You can specify whether users are prompted to create their own usernames or whether the system automatically creates usernames. Selecting an attribute for the username segments from the required attributes list improves the chances that a new username is successfully created.

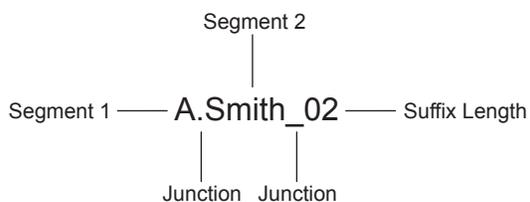
Maximum length: The maximum length of the user name. This value must be between 1 and 50.

Prompt for user name: Enables users to create their own usernames.

Automatically create user name: Specifies that the system creates usernames. You can configure the segments for the system to use when creating usernames and configure how the names are displayed.

For example, if you are using the required attributes of Common First Name and Common Last Name, a username for Adam Smith might be generated as A.Smith_02, as shown in [Figure 9-1](#):

Figure 9-1 Attribute Segments



Use the following settings to specify how this is accomplished:

- ◆ **Segment 1:** The required attribute to use as the first segment for the user name. The values displayed in this drop-down menu correspond to the required attributes you selected. For example, you might select Common First Name to use for *Segment 1*.
- ◆ **Length:** The length of the first attribute segment. For example, if you selected Common First Name for the *Segment 1* value, setting the length to 1 specifies that the system uses the first letter of the Common First Name attribute. Therefore, Adam Smith would be ASmith.
- ◆ **Junction:** The type of junction to use between the attributes of the user name, such as no space, or a hyphen, or a period. Adam Smith would display as A.Smith.

- ♦ **Segment 2:** The required attribute to use as the second segment for the user name. The values displayed in this drop-down menu correspond to the required attributes you selected. For example, you might select Common Last Name to use for *Segment 2*.
- ♦ **Length:** The length of the second attribute segment. For example, if you selected Common Last Name for the *Segment 2* value, you might set the length to *All*, so that the full last name is displayed. However, the system does not allow more than 20 characters for the length of segment 2.
- ♦ **Ensure name is unique:** Applies a suffix to the colliding name until a unique name is found, if using attributes causes a collision with an existing name. If no attributes are provided, or the lengths for them are 0, and this option is selected, the system creates a unique name.

8 Click *Next*.

9 Specify password settings.

Identity Servers > sp-401k > Corporate IDP >

User Provisioning Method ?

Step 4 of 5: Define new user password creation

The new user account will not be valid after the initial use if the user is not given the generated password.

Min. password length:

Max. password length:

Prompt for password

Automatically create password

Use this page to specify whether to prompt the user for a password or to create a password automatically.

Min. password length: The minimum length of the password.

Max. password length: The maximum length of the password.

Prompt for password: Prompts the user for a password.

Automatically create password: Specifies whether to automatically create passwords.

10 Click *Next*.

11 Specify the user store and context in which to create the account.

Identity Servers > sp-401k > Corporate IDP >

User Provisioning Method ?

Step 5 of 5: Select User Store where new user account is created

The selected User Store will be the target directory. Specify the directory context where the new user accounts will be created.

User Store:

Context: (ex. ou=users,o=novell)

Delete user provisioning accounts if federation is terminated

User Store: The user store in which to create the new user account.

Context: The context in the user store you want accounts created.

The system creates the user within a specific context; however, uniqueness is not guaranteed across the directory.

Delete user provisioning accounts if federation is terminated: Specifies whether to automatically delete the provisioned user account at the service provider if the user terminates his or her federation between the identity provider and service provider.

- 12 Click *Finish*.
- 13 On the Authentication page, click *OK*.
- 14 On the Trusted Providers page, click *OK*.
- 15 On the Setup page, click *Update Servers* to update the Identity Server configuration.

9.2.4 User Provisioning Error Messages

The following error messages are displayed for the end user if there are problems during provisioning.

Table 9-1 *Provisioning Error Messages*

Error Message	Cause
Username length cannot exceed (?) characters.	The user entered more characters for a user name than is allowed, as specified by the administrator.
Username is not available.	The user entered a name that already exists in the directory.
Passwords don't match.	The user provided two password values that do not match.
Passwords must be between (x) and (y) characters in length.	The user provided password values that are either too short or too long.
Username unavailable.	The provisioned user account was deleted without first defederating the user. Remove orphaned identity objects from the LDAP user store.

IMPORTANT: Only experienced LDAP users should removed orphaned identity objects from the LDAP user store. You must ensure that the objects you are removing are orphaned. Otherwise, you will create orphaned objects by mistake.

9.3 Configuring Authentication for a Trusted Service Provider

After you create a trusted service provider, you can configure how your Identity Server responds to authentication requests from the service provider.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Liberty > [Service Provider] > Access > Authentication*.

The screenshot shows the Administration Console interface for configuring authentication. At the top, there are three tabs: Configuration, Metadata, and Access. Below the tabs is a breadcrumb trail: General | Attributes | Authentication. The main content area shows the 'Authentication response binding' dropdown menu set to 'Artifact'. Below this is a table for 'Supported identity formats' with columns for 'Use' and 'Default'. The table has two rows: 'Persistent identifier format' and 'Transient identifier format'. Both have a checked 'Use' checkbox. The 'Default' column has a radio button selected for 'Persistent identifier format' and an unselected one for 'Transient identifier format'. Below the table are two checked checkboxes: 'Use proxied requests' and 'Provide Discovery Services'.

Supported identity formats	Use	Default
Persistent identifier format:	<input checked="" type="checkbox"/>	<input checked="" type="radio"/>
Transient identifier format:	<input checked="" type="checkbox"/>	<input type="radio"/>

Use proxied requests
 Provide Discovery Services

- 2 Fill in the following fields as required:

Authentication Response Binding: Specifies whether to use *Artifact* or *Post* if the request from the trusted service provider does not specify a response binding. Select *Artifact* to provided an increased level of security by using a back-channel means of communication between the two servers. Select *Post* to use HTTP redirection to accomplish communication between servers.

- ◆ **Persistent Identifier Format:** Specifies whether to use this format and make it the default identifier format. A persistent identifier is written to the directory and remains intact between sessions.
- ◆ **Transient Identifier Format:** Specifies whether to use this format and make it the default identifier format. A transient identifier expires between sessions.

Use Proxied Requests: Enables proxying for the service provider. If disabled, no proxying is allowed.

For example, the service provider can authenticate a user to IDP B through IDP A, when no trust relationship exists between the service provider and IDP B. This feature is allowed by default. However, you can disable the service provider's ability to use proxied requests. In order to use this, you must specify Silent Login on IDP A.

Proxying can also be used to achieve single sign-on when the trust authentication types and contracts differ between identity providers, or when identity providers are using multiple protocols, such as when one identity provider communicates via SAML 2.0, and another uses Liberty.

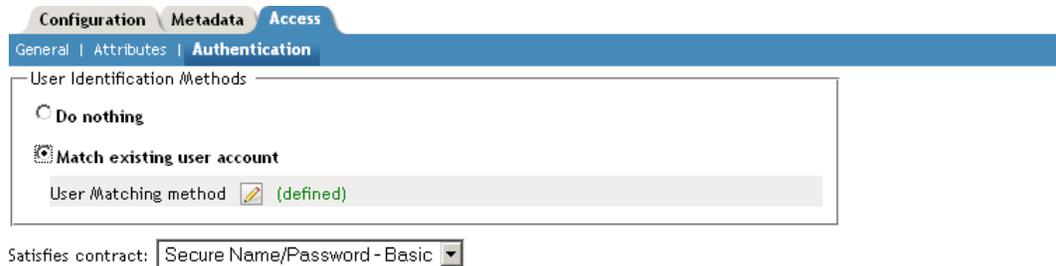
Provide Discovery Services: Advertises to the service provider the Web services available at the Identity Server. This option is required if the identity provider is to provide services to the service provider.

- 3 Click *OK*.

9.4 Configuring User Identification Methods for SAML 1.1 Trusted Identity Providers

Two methods exist for identifying users from a trusted identity provider. You can specify that no account matching needs to occur, or you can configure a match method. You configure a match method when you want to use attributes from this trusted identity provider to uniquely identify a user.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > SAML 1.1 > [Identity Provider] > Access > Authentication*.



The screenshot shows the Administration Console interface for configuring SAML 1.1 Trusted Identity Providers. The 'Authentication' tab is selected, and the 'User Identification Methods' section is visible. The 'Match existing user account' option is selected, and the 'User Matching method' is set to '(defined)'. The 'Satisfies contract' dropdown is set to 'Secure Name/Password - Basic'.

- 2 Configure the following options as necessary:

Do nothing: Specifies that the service provider does not match user accounts. This option allows you to authenticate the session without identifying a user account.

Match existing user accounts: Authenticates a user by matching a user account. This option requires that you set up the match method. (See [Step 3](#).)

Satisfies contract: The contract that is satisfied by the assertion received from the identity provider. Because SAML 1.1 does not use contracts and because the Identity Server is contract-based, this setting permits an association to be made between a contract and a SAML 1.1 assertion.

Use caution when assigning the contract to associate with the assertion, because it is possible to imply that authentication has occurred, when it has not. For example, if a contract is assigned to the assertion, and the contract has two authentication methods (such as one for name/password and another for X.509), the server sending the assertion might use only name/password, but the service provider might assume that X.509 authentication took place and then incorrectly assert it to another server.

3 To configure the match method, click *User Matching Method*.

User Matching Method



Select User Stores to search

User stores: Installed User Store	Available user stores:
--------------------------------------	------------------------

User Matching Expression: Dept_Users

OK Cancel Apply

4 To configure user matching, fill in the following fields:

Select User Stores to search: Select and order the user stores you want to use in the search.

User Matching Expression: Set the matching expression as the default, or click *New* to create a look-up expression.

Create User Matching Expression



Specify name and attributes

A user matching expression is a set of logic groups with attributes that uniquely identify a user. The "Type" designation (AND or OR) applies only between groups. Attributes within a group are always "AND" comparisons.

Name: Dept_Users

User Matching Expression

New Logic Group | Delete

3 Item(s)

Groups Type AND (all groups)

Logic Group 1

Legal Name

AND

Logic Group 2

Department Name

<< Back Finish Cancel

A user matching expression is a set of logic groups with attributes that uniquely identify a user. User matching expressions enable you to map the Liberty attributes to the correct LDAP

attributes during searches. You must know the LDAP attributes that are used to name the users in the user store and create the user's distinguished name.

In order to use user matching, you must enable the Personal Profile on the identity provider and the service provider. See [Section 11.2, “Enabling Web Services and Profiles,” on page 116](#).

- 5 Click *Finish*.
- 6 Select the new expression on the User Method Matching page, then click *OK*.
- 7 Click *OK* on the Authentication page, then click *OK* on the Trusted Providers page.
- 8 Update the Identity Server configuration on the Setup page, as prompted.

9.5 Specifying a SAML Audience URI

When an identity provider sends an assertion to a service provider, the assertion can be restricted to an intended audience. The intended audience is defined to be any abstract URI in SAML 1.1. The URI reference can also identify a document that describes the terms and conditions of audience membership.

In the Liberty specification, which uses SAML assertions, the audience is the provider ID. When you first set up a SAML partnership, adding audience restrictions conditions can add unnecessary complexity.

- 1 In the Administration Console, click *Access Manager > Identity Servers > [Configuration Assignment] > SAML 1.1 > [Service Provider] > Access Audiences*.
- 2 Click *New*.
- 3 Specify the *SAML Audience URI* value, then click *OK*.

Configuring Communication Profiles

10

You can configure the methods of communication that are available at the server for requests and responses sent between providers. These settings affect the metadata for the server and should be determined prior to publishing to other sites.

In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > [Protocol] > Profiles*.

Artifact Resolution: (SAML 2.0 only) The assertion consumer service at the service provider performs a back-channel exchange with the artifact resolution service at the identity provider. Artifacts are small data objects pointing to larger SAML protocol messages. They are designed to be embedded in URLs and conveyed in HTTP messages.

Login: Specifies whether to support Artifact or Post binding for login. The Artifact binding provides an increased level of security by using a back-channel means of communication between the two servers during authentication. The Post method uses HTTP redirection to accomplish communication between servers.

Single Logout: Enables the identity provider or service provider to accept HTTP and SOAP requests. Typically, you select both of these options. SOAP is used if both options are selected, or if the service provider has not specified a preference.

HTTP Redirect: A browser-based method that uses HTTP 302 redirects or HTTP GET requests to communicate requests from this identity site to the service provider. SAML messages are transmitted within URL parameters.

Federation Termination: (Liberty only) Specifies whether to use HTTP or SOAP profiles. Typically, you select both of these options, which enables the identity provider or service provider to accept both HTTP and SOAP requests. SOAP is the default setting if the service provider has not specified a preference.

Register Name: (Liberty only) Specifies whether to use HTTP or SOAP profiles. Typically, you select both of these options, which enables the identity provider or service provider to accept both HTTP and SOAP requests. SOAP is the default setting if the service provider has not specified a preference.

Name Management: (SAML 2.0 only) Specifies the binding protocol for the SAML Name Identifier Management profile. Name management is how the system manages the sharing of common identifiers for a principal between identity and service providers. When an identity provider has exchanged a persistent identifier for the principal with a service provider, the providers share the common identifier for a length of time. When either the identity or service provider changes the format or value to identify the principal, the system can ensure that the new format or value is properly transmitted.

Configuring Liberty Web Services

11

A Web service uses Internet protocols to provide a service. It is an XML-based protocol transported over SOAP, or a service whose instances and data objects are addressable via URIs.

Access Manager consists of several elements that comprise Web services:

- ♦ **Web Service Framework:** Manages all Web services. The framework defines SOAP header blocks and processing rules that enable identity services to be invoked via SOAP requests and responses.
- ♦ **Web Service Provider:** An entity that provides data via a Web service. In Access Manager, Web Service Providers host Web service profiles, such as the Employee Profile, Credential Profile, Personal Profile, and so on.
- ♦ **Web Service Consumer:** An entity that uses a Web service to access data. Web Service Consumers discover resources at the Web Service Provider, and then retrieve or update information about a user, or on behalf of a user. Resource discovery among trusted partners is necessary because a user might have many kinds of identities (employee, spouse, parent, member of a group), as well as several identity providers (employers or other commercial Web sites).
- ♦ **Discovery Service:** The service assigned to an identity provider that enables a Web Service Consumer to determine which Web Service Provider provides the required resource.
- ♦ **LDAP Attribute Mapping:** Access Manager's solution for mapping Liberty attributes with established LDAP attributes.

This section describes the following topics:

- ♦ [Section 11.1, “Configuring the Web Services Framework,” on page 115](#)
- ♦ [Section 11.2, “Enabling Web Services and Profiles,” on page 116](#)
- ♦ [Section 11.3, “Editing Web Service Descriptions,” on page 117](#)
- ♦ [Section 11.4, “Configuring Credential Profile Security and Display Settings,” on page 118](#)
- ♦ [Section 11.5, “Configuring Service and Profile Details,” on page 120](#)
- ♦ [Section 11.6, “Customizing Attribute Names,” on page 122](#)
- ♦ [Section 11.7, “Editing Web Service Policies,” on page 123](#)
- ♦ [Section 11.8, “Configuring the Web Service Consumer,” on page 126](#)
- ♦ [Section 11.9, “Mapping LDAP and Liberty Attributes,” on page 127](#)

For additional resources about the Liberty Alliance specifications, visit the [Liberty Alliance Specification \(http://www.projectliberty.org/resources/specifications.php\)](http://www.projectliberty.org/resources/specifications.php) page.

11.1 Configuring the Web Services Framework

The Web Services Framework page lets you edit and manage all the details that pertain to all Web services. This includes the framework for building interoperable identity services, permission-based attribute sharing, identity service description and discovery, and the associated security mechanisms.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Liberty > Web Service Framework*.

- 2 Fill in the following fields:

Enable Framework: Enables Web Services Framework.

Axis SOAP Engine Settings: Axis is the SOAP engine that handles all Web service requests and responses. Web services are deployed using XML-based files known as Web service deployment descriptors (WSDD). On startup, Access Manager automatically creates the server-side and client-side configuration for Axis to handle all enabled Web services. If you need to override this default configuration, use the *Axis Server Configuration WSDD XML* field and the *Axis Client Configuration WSDD XML* field to enter valid WSDD XML. If either or both of these controls contain valid XML, then Access Manager does not automatically create the configuration (server or client) on startup.

- 3 Click *OK*.

11.2 Enabling Web Services and Profiles

After a service has been discovered and authorization data has been received from a trusted identity provider, the Web service consumer can invoke the service at the Web service provider. A Web service provider is the hosting or relying entity on the server side that can make access control decisions based on this authorization data and upon its business practices and preferences.

- 1 In the Administration Console click *Identity Servers > Setup > [Configuration] > Liberty > Web Service Providers*.

- 2 Select one of the following services:

Credential Profile: Allows users to define information to keep secret. It uses encryption to store the data in the directory the user profile resides in.

Custom Profile: Used to create custom attributes for general use.

Discovery: Allows requesters to discover where the resources they need are located. Entities can place resource offerings in a discovery resource, allowing other entities to discover them. Resources might be a user's credit card information, a personal profile, calendar, travel preferences, and so on. Providers can place these resource offerings in a discovery resource, so that these by other entities.

Employee Profile: Allows you to manage employment-related information and how the information is shared with others. A company address book that provides names, phones, office locations, and so on, is an example of an employee profile.

LDAP Profile: Allows you to use LDAP attributes for authorization and general use.

Personal Profile: Allows you to manage personal information and to determine how to share that information with others. Shopping portals that manage the user's account number is an example of a personal profile.

Authentication Profile: Allows the system to access the roles and authentication contracts in use by current authentications. In normal circumstances, this is used only by the system. Do not delete this profile.

User Interaction: Allows you to set up a trusted user interaction service, used for identity services that must interact with the resource owner to get information or permission to share data with another Web service consumer. This profile enables a Web service consumer and Web service provider to cooperate in redirecting the resource owner to the Web service provider and back to the Web service consumer.

- 3 Click *Enable*, then click *OK*.
- 4 On the Setup page, click *Update Servers* to update the Identity Server configuration.

11.3 Editing Web Service Descriptions

All of the Description pages on each profile are identical. You can define how a service provider gains access to portions of the user's identity information that can be distributed across multiple providers. The service provider uses the Discovery Service to ascertain the location of a specific identity service for a user. The Discovery Service enables various entities to dynamically and securely discover a user's identity service, and it responds, on a permission basis, with a service description of the desired identity service.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Liberty > Web Service Provider*.
- 2 Click profile or service.
- 3 Click *Descriptions*.
- 4 Click the description name, or click *New*.
- 5 Fill in the following fields:

Name: The Web Service Description name.

Security Mechanism: (Required) Liberty uses channel security (TLS 1.0) and message security in conjunction with the security mechanism. Channel security addresses how communication between identity providers, service providers, and user agents is protected. For authentication, service providers are required to authenticate identity providers using identity provider server-side certificates. Identity providers have the option to require authentication of service providers using service provider client-side certificates.

Message security addresses security mechanisms applied to the discrete Liberty protocol messages passed between identity providers, service providers, and user agents.

Select the mechanism for message security. Message authentication mechanisms indicate which profile is used to ensure the authenticity of a message.

- ♦ **X.509:** Used for message exchanges that generally rely upon message authentication as the principle factor in making authorization decisions.
 - ♦ **SAML:** Used for message exchanges that generally rely upon message authentication as well as the conveyance and attestation of authorization information.
 - ♦ **Bearer:** Based on the presence of the security header of a message. In this case, the bearer token is verified for authenticity rather than proving the authenticity of the message.
- 6 Under Select Service Access Method, click either *Brief Service Access Method* or *WSDL Service Access Method*.

Brief Service Access Method: Provides the information necessary to invoke basic SOAP-over-HTTP-based service instances without using WSDL.

- ◆ **EndPoint URL:** This is the SOAP endpoint location at the service provider to which Liberty SOAP messages are sent. An example of this for the Employee Profile is [BASEURL]/services/IDSISEmployeeProfile. If the service instance exposes an endpoint that is different from the logically generated concrete WSDL, you must use the WSDL URI instead.

A WSF service description endpoint cannot contain double-byte characters.

- ◆ **SOAP Action:** The SOAP action HTTP header required on HTTP-bound SOAP messages. This header can be used to indicate the intent of a SOAP message to the recipient.

WSDL Service Access Method: Specify the method used to access the WSDL service. WSDL (Web Service Description Language) describes the interface of a Web service.

- ◆ **Service Name Reference:** A reference name for the service.
- ◆ **WSDL URI:** Provides a URI to an external concrete WSDL resource containing the service description. URIs need to be constant across all implementations of a service to enable interoperability.

7 Click *OK*.

8 Update the Identity Server configuration.

11.4 Configuring Credential Profile Security and Display Settings

On the Credential Profile Details page, you can specify whether this profile is displayed for end users, and determine how you control and store encrypted secrets. You can store and access secrets locally or on remote eDirectory servers that are running Novell[®] Secret Store[®]. For general information about this product, see the [Novell SecretStore Administration Guide \(http://www.novell.com/documentation/secretstore33/pdfdoc/nssadm/nssadm.pdf\)](http://www.novell.com/documentation/secretstore33/pdfdoc/nssadm/nssadm.pdf).

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Liberty > Web Service Providers*.

2 Click *Credential Profile*.

Identity Servers ► IDSConfig ►

Credential Profile

Edit the details about the web service.

Details Descriptions Custom Attribute Names

General Settings

Display name:

Have Discovery Encrypt This Service's Resource Ids

Credential Profile Settings

Allow End Users to See Credential Profile

Local Storage of Secrets

Access Manager controls the storage and encryption of secrets.

Encryption Password Hash Key:

Preferred Encryption Method:

Extended Schema User Store References

New 0 Item(s)

<input type="checkbox"/> User Store

OK Cancel Apply

3 On the Credential Profile Details page, fill in the following fields as necessary:

Display name: The name you want to display for the Web service.

Have Discovery Encrypt This Service's Resource Ids: Specifies whether the Discovery Service encrypts resource IDs. A resource ID is an identifier used by Web services to identify a user. The Discovery Service returns a list of resource IDs when a trusted service provider queries for the services owned by a given user. The Discovery Service has the option of encrypting the resource ID or sending it unencrypted. Encrypting resource IDs is disabled by default.

4 Under *Credential Profile Settings*, enable the following option if necessary:

Allow End Users to See Credential Profile: Specifies whether to display or hide the Credential Profile in the Access Manager User Portal. Profiles are viewed on the My Profile page, where the user can modify his or her profile.

5 Specify how you want to control and store secrets:

5a To locally control and store secrets, configure the following fields:

Encryption Password Hash Key: (Required) Encrypts and decrypts secret values in order to keep data confidential. You must change this key from the default value.

Preferred Encryption Method: Specify the preferred encryption method:

- ◆ **Password Based Encryption With MD5 and DES:** MD5 is an algorithm that is used to verify data integrity.

- ♦ **DES:** Data Encryption Standard (DES) is a widely used method of data encryption using a private key. Like other private key cryptographic methods, both the sender and the receiver must know and use the same private key.
- ♦ **Triple DES:** A variant of DES in which data is encrypted three times with standard DES using two different keys.

If you click *OK*, the secret data is stored in the user profile object in the configuration data store.

- 5b** To specify where to store secret data, click *New* under *Extended Schema User Store References*.

User Store: Allows you to select a user store where secret data is stored.

Attribute Name: Specifies the LDAP attribute to store in the User object. When a user authenticates using the user store specified here, the secret data is stored in this attribute in the User object. The attribute must be a valid extended attribute of the directory schema user object.

- 5c** Click *OK*.

- 6** To use Novell SecretStore to remotely store secrets:

- 6a** Click *New* under *Novell Secret Store User Store References*.

This adds a reference to a user store where SecretStore has been installed. Click the user store that you configured for SecretStore.

Secure LDAP must be enabled in order to add this user store reference.

- 7** Click *OK*.

11.5 Configuring Service and Profile Details

The settings on the Details page are identical for the Employee, Custom, and Personal Profiles. This page allows you to specify the display name, resource ID encryption, and how the system reads and writes data.

- 1** In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Liberty > Web Service Provider*.
- 2** Click either Custom Profile, Employee Profile, or Personal Profile, depending on which profile you want to edit.
- 3** Click the Details tab (it is displayed by default).

Edit the details about the web service.

The screenshot shows a web interface with three tabs: 'Details', 'Descriptions', and 'Custom Attribute Names'. The 'Details' tab is active. Below the tabs is a section titled 'General Settings'. It contains a text input field labeled 'Display name:' with the value 'Custom test Profile'. Below this is a checkbox labeled 'Have Discovery Encrypt This Service's Resource Ids' which is checked.

- 4** Specify the general settings, as necessary:

Display Name: The Web service name. This specifies how the profile is displayed in the Administration Console.

Have Discovery Encrypt This Service's Resource Ids: Specifies whether the Discovery Service encrypts resource IDs. A resource ID is an identifier used by Web services to identify a

user. The Discovery Service returns a list of resource IDs when a trusted service provider queries for the services owned by a given user. The Discovery Service has the option of encrypting the resource ID or sending it unencrypted.

5 Specify data location settings:

Identity Servers ▶ TradingCo-ids ▶

Custom Profile

Edit the details about the web service.

Details | Descriptions | Custom Attribute Names

General Settings

Display name:

Have Discovery Encrypt This Service's Resource Ids

Data Locations Settings

Selected Read Locations: Configuration Datastore, LDAP Data Mappings, Remote Attributes

Available Read Locations:

Selected Write Locations: LDAP Data Mappings, Configuration Datastore

Available Write Locations:

Data Model Extensions

Extend the service data model by defining new data types.

OK Cancel Apply

The following settings apply only to the Custom, Employee, and Personal Profiles.

Selected Read Locations: The list of selected locations from which the system reads attributes containing profile data. If you add multiple entries to this list, the system searches attributes in each location in the order you specify. Use the Up/Down and Left/Right arrows to control which locations are selected and the order in which to read them. Read locations can include:

- ♦ **Configuration Datastore:** The Identity Server configuration user store. If you place the Configuration Datastore first in this list, all writes go to this data store. Subsequent write locations are ignored.
- ♦ **LDAP Data Mappings:** The LDAP attribute maps.

Remote Attributes: Attributes pushed to the identity provider during authentication. For example, if a service provider requests a user's first name from an identity provider's Personal Profile service, and the identity provider does not have the first name attribute, the identity provider checks whether the user has authentications with other remote identity providers. If there are remote services that match the requested service, the identity provider requests the data from the remote identity provider to see if that service can fill the request.

Available Read Locations: The list of available locations from which the system reads attributes containing profile data. The Configuration Datastore is the directory containing all of

the configuration information for Access Manager. The LDAP Data Mappings reference an external directory containing attribute information to be displayed by the profile.

Selected Write Locations: The list of selected locations to write attribute data to. If you add multiple entries to this list, the system searches attributes in each location in the order you specify. Use the Up/Down and Left/Right arrows to control which locations are selected and the order in which to read them.

Available Write Locations: The list of available locations to write attributes containing profile data, similar to the available read locations.

- 6 (Optional) Specify data model extensions.

Data Model Extension XML: The data model for some Web services is extensible. You can enter XML definitions of data model extensions in this field. Data model extensions hook into the existing Web service data model at predefined locations.

All schema model extensions reside inside of a schema model extension group. The group exists to bind model data items together under a single localized group name and description. Schema model extension groups can reside inside of a schema model extension root or inside of a schema model extension. There can only be one group per root or extension. Each root is hooked into the existing Web service data model. Multiple roots can be hooked into the same location in the existing Web service data model. This conceptual model applies to the structure of the XML that is required to define data model extensions.

See [Appendix C, “Data Model Extension XML,”](#) on page 523 for more information.

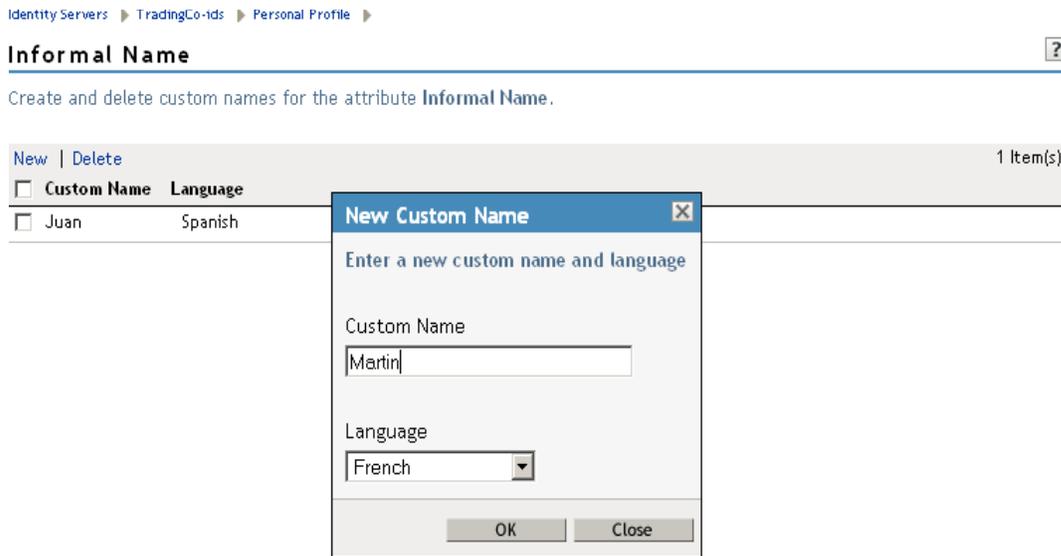
- 7 Click *OK*, then click *OK* on the Web Service Provider page.
- 8 Update the Identity Server configuration on the Setup page.

11.6 Customizing Attribute Names

You can change the display name of an attribute names for the Credential, Custom, Employee, and Personal profiles. The customized names are displayed on the My Profile page in the User Portal. The users see the custom names applicable to their language. Custom Attributes are displayed on the My Profile page in the User Portal in place of the corresponding English attribute name when the language in the drop-down list is the accepted language of the browser.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Liberty > Web Service Provider > [Profile] > Custom Attribute Names*.

- 2 Click the data item name to view the customized attribute names.



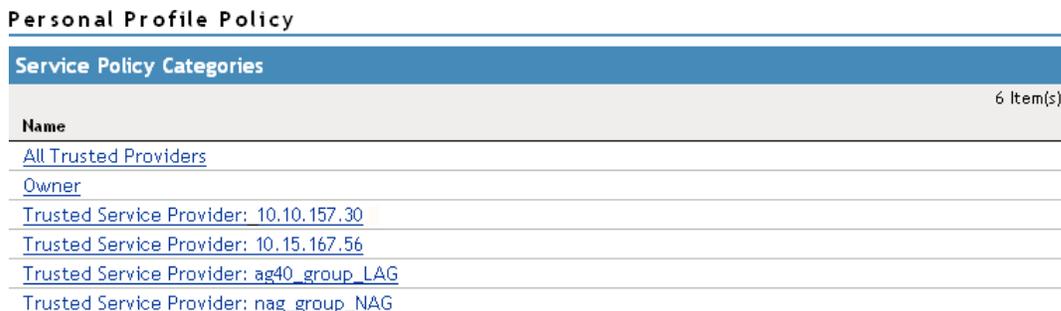
- 3 Click *New* to create a new custom name.
- 4 Type the name and select a language.
- 5 Click *OK*.
- 6 On the Custom Attribute Names page, click *OK*.
- 7 On the Web Service Provider page, click *OK*.
- 8 Update the Identity Server configuration on the Setup page.

11.7 Editing Web Service Policies

Web Service Policies are permission policies (query and modify) that govern how identity providers share end-user data with service providers. Administrators and policy owners (users) can control whether private information is always allowed to be given, never allowed, or must be requested.

As an administrator, you can configure this information for the policy owner, for specific service providers, or globally for all service providers. You can also specify what policies are displayed for the end user in the User Portal, and whether users are allowed to edit them.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Liberty > Web Service Provider*.
- 2 Click the *Policy* link next to the service name.



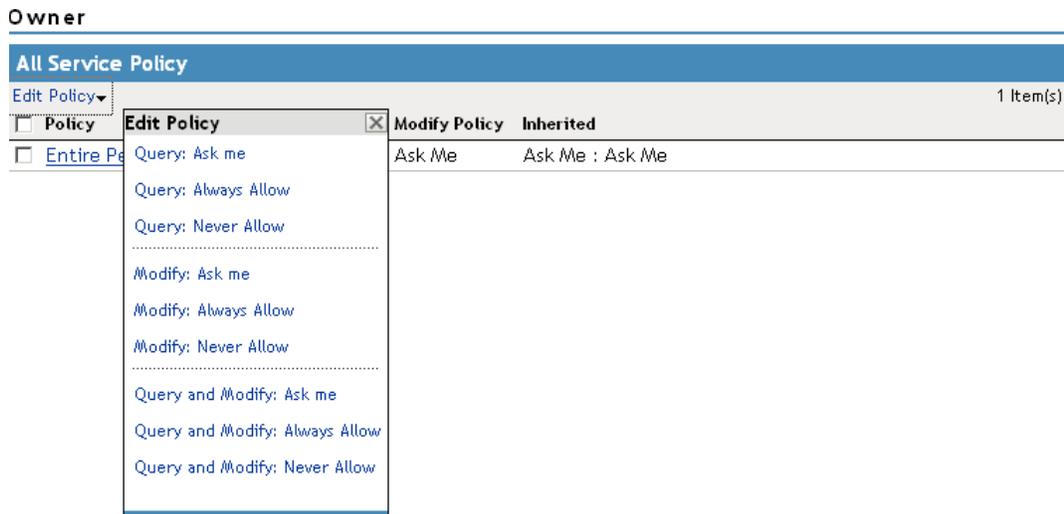
- 3 Click the category you want to edit.

All Trusted Providers: Policies that are defined by the service provider’s ability to query and modify the particular Liberty attributes or groups of attributes for the Web service. When All Trusted Providers permissions are established, and a service provider needs data, the system first looks here to determine whether user data is allowed, never allowed, or must be asked for. If no solution is found in All Trusted Providers, the system examines the permissions established within the specific service provider.

Owners: Policies that limit the end user’s ability to modify or query data from his own profile. The settings you specify in the *Owner* group are reflected on the My Profile page in the User Portal. Portal users have the authority to modify the data items in their profiles. The data items include Liberty and LDAP attributes for personal identity, employment, and any customized attributes defined in the Identity Server configuration. Any settings you specify in the Administration Console override what is displayed in the User Portal. Overrides are displayed in the *Inherited* column.

If you want the user to have Write permission for a given data item, and that data item is used in an LDAP Attribute Map, then you must configure the LDAP Attribute Map with Write permission.

- 4 On the All Service Policy page, select the policy’s check box, then click *Edit Policy*.



This lets you modify the parent service policy attribute. Any selections you specify on this page are inherited by child policies.

Query Policy: Allows the service provider to query for the data on a particular attribute. This is similar to read access to a particular piece of data.

Modify Policy: Allows the service provider to modify a particular attribute. This is similar to write access to a particular piece of data.

Query and Modify: Allows you to set both options at once.

- 5 To edit child attributes of the parent, click the policy.

In the following example, child attributes are inheriting Ask Me permission from the parent *Entire Personal Identity* attribute. The *Postal Address* attribute, however, is modified to never allow permission for sharing.

Entire Personal Identity

Personal Identity			
Edit Policy▼	12 Item(s)		
<input type="checkbox"/> Policy	Query Policy	Modify Policy	Inherited
<input type="checkbox"/> Informal Name	Ask Me	Ask Me	Ask Me : Ask Me
<input type="checkbox"/> Localized Informal Name	Ask Me	Ask Me	Ask Me : Ask Me
<input type="checkbox"/> Entire Common Name	Ask Me	Ask Me	Ask Me : Ask Me
<input type="checkbox"/> Entire Legal Identity	Ask Me	Ask Me	Ask Me : Ask Me
<input type="checkbox"/> Employment Identity	Ask Me	Ask Me	Ask Me : Ask Me
<input type="checkbox"/> Postal Addresses	Never Allow	Never Allow	Ask Me : Ask Me
<input type="checkbox"/> Contact Profiles	Ask Me	Ask Me	Ask Me : Ask Me
<input type="checkbox"/> Internet Identity	Ask Me	Ask Me	Ask Me : Ask Me

If you click the *Postal Address* attribute, all of its child attributes have inherited the *Never Allow* setting. You can specify different permission attributes for *Address Type* (for example), but the inherited policy still overrides changes made at the child level, as shown below.

Postal Addresses

Postal Addresses			
Edit Policy▼	6 Item(s)		
<input type="checkbox"/> Policy	Query Policy	Modify Policy	Inherited
<input type="checkbox"/> Address Type	Always Allow	Always Allow	Never Allow : Never Allow
<input type="checkbox"/> NickName	Ask Me	Ask Me	Never Allow : Never Allow
<input type="checkbox"/> Localized NickNames	Ask Me	Ask Me	Never Allow : Never Allow
<input type="checkbox"/> Comment	Ask Me	Ask Me	Never Allow : Never Allow
<input type="checkbox"/> Postal Address	Ask Me	Ask Me	Never Allow : Never Allow
<input type="checkbox"/> Postal Addresses Extensions	Ask Me	Ask Me	Never Allow : Never Allow

The interface allows these changes in order to simplify switching between configurations if, for example, you want to remove an inherited policy.

Inherited: Specifies the settings inherited from the parent attribute policy, when you view a child attribute. In the User Portal, settings displayed under *Inherited* are not modifiable by the user. At the top-level policy in the User Portal, the values are inherited from the settings in the Administration Console. Thereafter, inheritance can come from the service policy or the parent data item's policy.

Ask Me: Specifies that the service provider requests from the user what action to take.

Always Allow: Specifies that the identity provider always allows the attribute data to be sent to the service provider.

Never Allow: Specifies that the identity provider never allows the attribute data to be sent to the service provider.

When a request for data is received, the Identity Server examines policies to determine what action to take. For example, if a service provider like DigitalAirlines.com requires a postal address for the user, the Identity Server performs the following actions:

- ♦ Checks the settings specified in *All Service Providers*.
- ♦ If no solution is found, checks for the policy settings configured for the service provider.

- 6 Click *OK* until the Web Service Provider page is displayed.
- 7 Click *OK*, then update the Identity Server as prompted.

11.8 Configuring the Web Service Consumer

The Web service consumer is the component within the identity provider that request attributes from Web service providers. The identity provider and Web services consumer cooperate to redirect the user or resource owner to the identity provider, allowing interaction. You can configure an interaction service, which allows the identity provider to pose simple questions to a user. This service can be offered by trusted Web services consumers, or by a dedicated interaction service provider that has a reliable means of communication with the users.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Liberty > Web Service Consumers*

The following general settings configure time limits and processing speed:

Protocol Timeout (seconds): Limits the time the transport protocol allows.

Provider Timeout (seconds): Limits the request processing at the Web service provider. This value must always be equal to or longer than the *Provider Timeout* value.

Attribute Cache Enabled: A sub-system of the Web service consumer that caches attribute data that the Web service consumer requests. For example, if the Web service consumer has already requested a first name attribute from a Web service provider, the Web service consumer does not need to request the attribute again. This setting improves performance when enabled. However, you can disable this option to increase system memory.

- 2 Specify how and when the identity provider interacts with the user:

Always Allow Interaction: Allows interaction to take place between users and service providers.

Never Allow Interaction: Never allows interaction between users and service providers.

Always Allow Interaction for Permissions, Never for Data: Allows interaction for permissions, never for data.

Maximum Allowed Interaction Time: Specifies the allowed time (in seconds).

- 3 To specify the allowable methods that a Web Service Provider can use for user interaction, click one of the following options:

Redirect to a User Interaction Service: Allows the Web service consumer to redirect the user agent to the Web service provider to ask questions. After the Web service provider has obtained the information it needs, it can redirect the user back to the Web service consumer.

Call a Trusted User Interaction Service: Allows the Web service provider to trust the Web service consumer to act as proxy for the resource owner.

- 4 Under *Security Settings*, fill in the following fields:

WSS Security Token Type: Instructs the Web service consumer/requestor how to place the token in the security header as outlined in the Liberty ID-WSF Security Mechanisms.

Signature Algorithm: The signature algorithm to use for signing the payload.

- 5 Click *OK*, then update the Identity Server configuration as prompted.

11.9 Mapping LDAP and Liberty Attributes

You can create an LDAP attribute map or edit an existing one. Attribute mapping involves specifying how single-value and multi-value data items map to single-value and multi-value LDAP attributes. A single-value attribute can contain no more than one value, and a multi-value attribute can contain more than one. An example of a single-value attribute might be a person's gender, while an example of a multi-value attribute might be a person's various e-mail addresses, phone numbers, or titles.

The following fields are common among all attribute maps and are defined here:

Type: Specifies the map type. Access Manager comes with a predefined "one-to-one" mapping type for the Liberty profiles of Personal, Employee, and General. However, the following sections describe how to create additional map types:

- ◆ [Section 11.9.1, "Configuring One-to-One Attribute Maps," on page 127](#)
- ◆ [Section 11.9.2, "Configuring Employee Type Attribute Maps," on page 131](#)
- ◆ [Section 11.9.3, "Configuring Employee Status Attribute Maps," on page 132](#)
- ◆ [Section 11.9.4, "Configuring Postal Address Attribute Maps," on page 133](#)
- ◆ [Section 11.9.5, "Configuring Contact Method Attribute Maps," on page 134](#)
- ◆ [Section 11.9.6, "Configuring Gender Attribute Maps," on page 136](#)
- ◆ [Section 11.9.7, "Configuring Marital Status Attribute Maps," on page 136](#)

Name: The name you want to give the map.

Description: A description of the map.

Access Rights: A drop-down menu that provide the broadest control for the page. If you set this to *Read/Write*, you can specify rights for individual data items.

In order for user provisioning to succeed, you must select *Read/Write* from the *Access Rights* drop-down menu for any maps that use an attribute during user provisioning.

User Stores: The user store that a map applies to. If a user logs into a user store that is not in the map's user store list, that map is not used to read or write attributes for that user.

LDAP Attribute Name: The LDAP attribute name that you want to map to the Liberty attribute.

LDAP Attribute Value: The predefined LDAP attribute values that you want to map to the Liberty values. These LDAP values are those you want to store in the LDAP attribute for each given Liberty attribute value. The LDAP attribute map then maps the actual Liberty URI value, back and forth, to this supplied value. Values must match the attribute as it appears in the directory exactly. For example, "givenName" must be entered as "givenName" in the text field or the mapping does not work.

11.9.1 Configuring One-to-One Attribute Maps

A one-to-one map enables you to map single-value and multiple-value LDAP attribute names to standard Liberty attributes. A default one-to-one attribute map is provided with Access Manager, but you can also define your own.

An example of a one-to-one attribute map might be the single-valued Liberty attribute Common Name (CommonName) used by the Personal Profile that is mapped to the LDAP attribute givenName. The attribute value CN might be mapped to the LDAP fullName. You can further configure the various Liberty values to map to any LDAP attribute names that you use.

- 1 In the Administration Console, click *Access Manager > Identity Server > Setup > [Configuration] > Liberty > LDAP Attribute Mapping > New > One to One.*
- 2 Configure the map using the following guidelines:
 - ◆ [Mapping Personal Profile Single-Value Data Items to LDAP Attributes](#)
 - ◆ [Mapping Personal Profile Multiple-Value Data Items to LDAP Attributes](#)
 - ◆ [Mapping Employee Profile Single-Value Data Items to LDAP Attributes](#)
 - ◆ [Mapping Employee Profile Multiple-Value Data Items to LDAP Attributes](#)
 - ◆ [Mapping Custom Profile Single-Value Data Items to LDAP Attributes](#)
 - ◆ [Mapping Custom Profile Multiple-Value Data Items to LDAP Attributes](#)
- 3 After you create the mapping, click *Finish.*
- 4 On the LDAP Attribute Mapping page, click *OK.*
- 5 Update the Identity Server configuration on the Setup page as prompted.

Mapping Personal Profile Single-Value Data Items to LDAP Attributes

The data items displayed are single-value Liberty Personal Profile attributes that you can map to the single-valued LDAP attributes that you have defined for your directory.

Default One-To-One Ldap Attribute Mapping ?

Personal Profile Single Valued Data Items to LDAP Attributes

Data Item Name:	Ldap Attribute Name:	Access Rights:
Informal Name	<input type="text"/>	Read Only ▾
Every Day Name	fullName	Read Only ▾
Common Personal Title	title	Read Only ▾
Common First Name	givenName	Read Only ▾
Common Last Name	sn	Read Only ▾
Common Middle Name	<input type="text"/>	Read Only ▾
Legal Name	<input type="text"/>	Read Only ▾
Legal Personal Title	<input type="text"/>	Read Only ▾
Legal First Name	<input type="text"/>	Read Only ▾
Legal Last Name	<input type="text"/>	Read Only ▾
Legal Middle Name	<input type="text"/>	Read Only ▾
Legal Fiscal Identification Type	<input type="text"/>	Read Only ▾
Legal Fiscal Identification Value	<input type="text"/>	Read Only ▾

OK Cancel

Mapping Personal Profile Multiple-Value Data Items to LDAP Attributes

Use the fields on this page to map multiple-value attributes from the Liberty Personal Profile to the multiple-value LDAP attributes you have defined for your directory. For example, you can map the

Liberty attribute Alternate Every Day Name (AltCN) to the LDAP attribute you have defined for this purpose in your directory.

Default One-To-One Ldap Attribute Mapping ?

Personal Profile Multiple Valued Data Items to LDAP Attributes

Data Item Name:	Ldap Attribute Name:	Access Rights:
Alternate Every Day Name	<input type="text"/>	Read Only ▾
Alternate Department Names	<input type="text"/>	Read Only ▾
Spoken or Understood Languages	<input type="text"/>	Read Only ▾

Employee Profile Single Valued Data Items to LDAP Attributes

Data Item Name:	Ldap Attribute Name:	Access Rights:
Id	<input type="text"/>	Read Only ▾
Date of Hire	<input type="text"/>	Read Only ▾
Job Start Date	<input type="text"/>	Read Only ▾
Status	<input type="text"/>	Read Only ▾
Type	<input type="text"/>	Read Only ▾
Internal Job Title	<input type="text"/>	Read Only ▾
Department	ou	Read Only ▾

Mapping Employee Profile Single-Value Data Items to LDAP Attributes

Map the Liberty Employee Profile single-value attributes to the LDAP attributes you have defined in your directory for entries such as ID, Date of Hire, Job Start Date, Department, and so on.

Mapping Employee Profile Multiple-Value Data Items to LDAP Attributes

Map the Liberty Employee Profile multiple-value attributes to the LDAP attributes you have defined in your directory.

Mapping Custom Profile Single-Value Data Items to LDAP Attributes

Map custom Liberty profile single-value attributes to LDAP attributes you have defined in your directory. These attributes are customizable strings associated with the Custom Profile.

Default One-To-One Ldap Attribute Mapping

Custom Profile Single Valued Data Items to LDAP Attributes		
Data Item Name:	Ldap Attribute Name:	Access Rights:
Customizable String One	<input type="text"/>	Read Only
Customizable String Two	<input type="text"/>	Read Only
Customizable String Three	<input type="text"/>	Read Only
Customizable String Four	<input type="text"/>	Read Only
Customizable String Five	<input type="text"/>	Read Only
Customizable String Six	<input type="text"/>	Read Only
Customizable String Seven	<input type="text"/>	Read Only
Customizable String Eight	<input type="text"/>	Read Only
Customizable String Nine	<input type="text"/>	Read Only
Customizable String Ten	<input type="text"/>	Read Only

Custom Profile Multiple Valued Data Items to LDAP Attributes		
Data Item Name:	Ldap Attribute Name:	Access Rights:
Customizable Multi-Valued Strings One	<input type="text"/>	Read Only
Customizable Multi-Valued Strings Two	<input type="text"/>	Read Only

Customizable String (1 - 10): The Custom Profile allows custom single-value and multiple-value attributes to be defined without using the [Data Model Extension XML](#) to extend a service's schema. To use a customizable attribute, navigate to the Custom Attribute Names tab on the Custom Profile Details page (see [Section 11.6, "Customizing Attribute Names," on page 122](#)). There you can customize the name of any of the predefined single-value or multiple-value customizable attributes in the Custom Profile. After you customize a name, you can use that attribute in the same way you would use any other profile attribute.

Mapping Custom Profile Multiple-Value Data Items to LDAP Attributes

Customizable Multi-Valued Strings (1 - 5): Similar to customizable strings for single-value attributes, except these attributes can have multiple values. Use this list of fields to map directory attributes that can have multiple values (like SN) to multiple-value strings from the Custom Profile.

11.9.2 Configuring Employee Type Attribute Maps

You can map the LDAP attribute name and values to the Liberty profile values for Employee Type. This is an Employee Profile attribute. Examples of Liberty values appended to this attribute include Contractor Part Time, Contractor Full Time, Full Time Regular, and so on.

- 1 In the Administration Console, click *Access Manager > Identity Server > Setup > [Configuration] > Liberty > LDAP Attribute Mapping > New > Employee Type*.

New Employee Type LDAP Attribute Mapping ?

Specify name, description, user stores and mapping data.

Name:

Description:

Access Rights:

User stores:

Available user stores:

Employee Type to LDAP Attribute

LDAP Attribute Name:

Liberty Profile Values to LDAP Attribute Values

Employee Type Value: LDAP Attribute Value:

Contractor Part Time:

Contractor Full Time:

<< Back Finish Cancel

- 2 Specify a name and description for the map.
- 3 Choose the type of access rights you want.
Select *Read/Write* for any attributes used in user provisioning.
- 4 In the *LDAP Attribute Name* field, type the LDAP attribute name that you want to map to the Liberty Employee Type attribute.
- 5 In the *LDAP Attribute Value* fields, type your predefined LDAP attribute values that you want to map to the Liberty Employee Type values.
These are the values that you want to store in the LDAP attribute for each given Liberty attribute value. The LDAP attribute map then maps the actual Liberty URI value, back and forth, to this supplied value.
- 6 Click *Finish*.
- 7 On the LDAP Attribute Mapping page, click *OK*.
- 8 Update the Identity Server configuration on the Setup page as prompted.

11.9.3 Configuring Employee Status Attribute Maps

You can map the LDAP attribute name and values to the Liberty profile values for Employee Status. This is an Employee Profile attribute. Examples of the values appended to this Liberty attribute include Active, Trial, Retired, Terminated, and so on.

- 1 In the Administration Console, click *Access Manager > Identity Server > Setup > [Configuration] > Liberty > LDAP Attribute Mapping > New > Employee Status*.

New Employee Status LDAP Attribute Mapping ?

Specify name, description, user stores and mapping data.

Name:

Description:

Access Rights:

User stores: Available user stores:

Employee Status to LDAP Attribute

LDAP Attribute Name:

Liberty Profile Values to LDAP Attribute Values

Employee Status Value:	LDAP Attribute Value:
Active:	<input type="text" value="Active"/>
Trial:	<input type="text" value="Trial"/>
Laid Off:	<input type="text" value="Laid Off"/>
Retired:	<input type="text" value="Retired"/>
Stop Pay:	<input type="text" value="Stop Pay"/>

- 2 Specify a name and description for the map.
- 3 Choose the type of access rights you want.
Select *Read/Write* for any attributes used in user provisioning.
- 4 In the *LDAP Attribute Name* field, type the LDAP attribute name that you want to map to the Liberty Employee Status element.
- 5 In the *LDAP Attribute Value* fields, type the predefined LDAP attribute values that you want to map to the *Liberty Employee Status* values.
These are the values that you want to store in the LDAP attribute for each given Liberty attribute value. The LDAP attribute map then maps the actual Liberty URI value, back and forth, to this supplied value.
- 6 Click *Finish*.
- 7 On the LDAP Attribute Mapping page, click *OK*.
- 8 Update the Identity Server configuration on the Setup page as prompted.

11.9.4 Configuring Postal Address Attribute Maps

You can map the LDAP attribute name and values to the Liberty profile values for Postal Address. The PostalAddress element refers to the local address, including street or block with a house number, and so on. This is a Personal Profile attribute.

- 1 In the Administration Console, click *Access Manager > Identity Server > Setup > [Configuration] > Liberty > LDAP Attribute Mapping > New > Postal Address*.

The screenshot shows the 'New Postal Address LDAP Attribute Mapping' configuration page. The page title is 'New Postal Address LDAP Attribute Mapping' with a help icon. Below the title is a subtitle 'Specify name, description, user stores and mapping data.' The form includes the following fields and controls:

- Name:** A text input field.
- Description:** A text input field.
- Access Rights:** A dropdown menu with 'Read' selected.
- User stores:** A list box containing '<Default User Store>'.
- Available user stores:** A list box containing 'Installed User Store'.
- Mode of Operation:** A dropdown menu with 'Multiple Ldap Attributes' selected.
- Postal Address to LDAP Attribute(s):** A group of five text input fields for mapping:
 - Postal Address Ldap Attribute:
 - Postal Code Attribute:
 - City Ldap Attribute:
 - State Ldap Attribute:
 - Country Ldap Attribute:

- 2 Specify a name and description for the map.
- 3 Choose the type of access rights you want.
Select *Read/Write* for any attributes used in user provisioning.
- 4 In the *Mode* drop-down menu, select either *Multiple LDAP Attributes* or *Single Delimited LDAP Attributes*.

Multiple LDAP Attributes: Allows you to map multiple LDAP attributes to multiple Liberty Postal Address elements. When you select this option, the following Liberty Postal Address elements are displayed under the *Postal Address to LDAP Attributes* group. Type the LDAP attributes that you want to map to the Liberty elements.

- ◆ Postal Address
- ◆ Postal Code
- ◆ City
- ◆ State
- ◆ Country

Single Delimited LDAP Attributes: Allows you to specify one LDAP attribute that is used to hold multiple elements of a Liberty Postal Address in a single delimited value. When you select this option, the page displays the following fields:

- ♦ **Delimited LDAP Attribute Name:** The delimited LDAP attribute name you have defined for the LDAP postal address that you want to map to the Liberty Postal Address attribute.
- ♦ **Delimiter:** The character to use to delimit single-value entries. A \$ sign is the default delimiter.

5 Under *Postal Address Template Data*, fill in the following options:

Nickname: (Required) A Liberty element name used to identify the Postal Address object.

Contact Method Type: Select the contact method type, such as *Domicile, Work, Emergency*, and so on.

6 Click *Finish*.

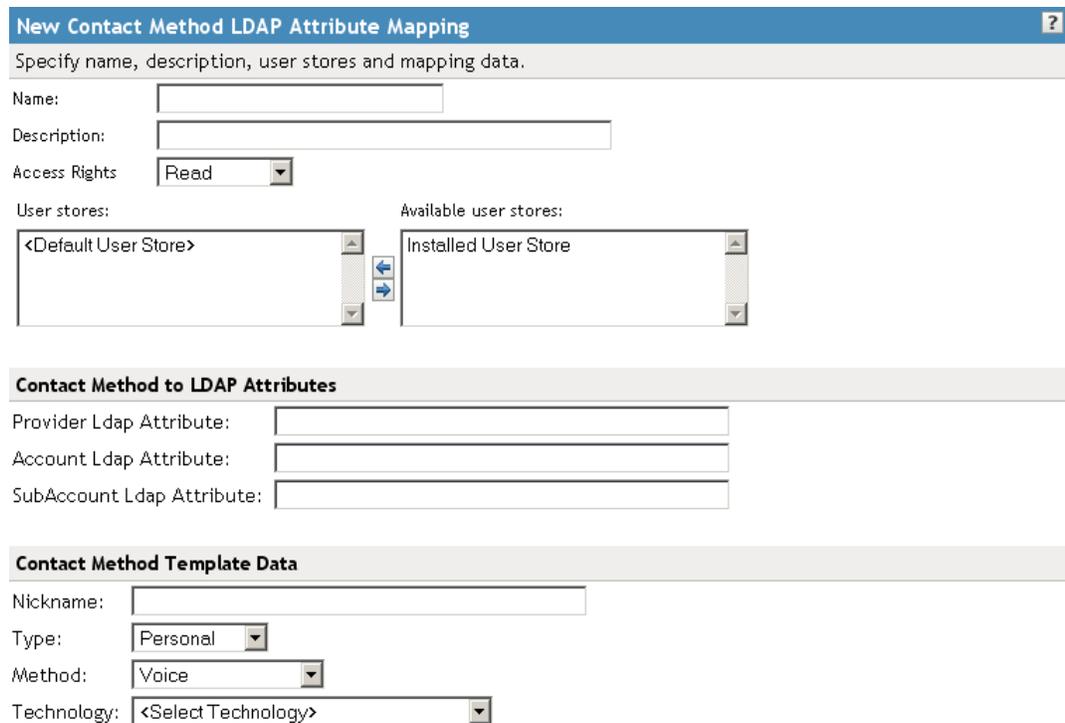
7 On the LDAP Attribute Mapping page, click *OK*.

8 Update the Identity Server configuration on the Setup page as prompted.

11.9.5 Configuring Contact Method Attribute Maps

You can map the LDAP attribute you have defined for contact methods to the Liberty attribute Contact Method (MsgContact).

1 In the Administration Console, click *Access Manager > Identity Server > Setup > [Configuration] > Liberty > LDAP Attribute Mapping > New > Contact Method*.



New Contact Method LDAP Attribute Mapping ?

Specify name, description, user stores and mapping data.

Name:

Description:

Access Rights:

User stores: Available user stores:

Contact Method to LDAP Attributes

Provider Ldap Attribute:

Account Ldap Attribute:

SubAccount Ldap Attribute:

Contact Method Template Data

Nickname:

Type:

Method:

Technology:

2 Specify a name and description for the map.

- 3 Choose the type of access rights you want.
Select *Read/Write* for any attributes used in user provisioning.
- 4 Under *Contact Method to LDAP Attributes*, fill in the following fields to map to the Liberty Contact Method attribute:
 - ♦ **Provider LDAP Attribute:** Maps to the Liberty attribute *MsgProvider*, which is the service provider or domain that provides the messaging service.
 - ♦ **Account LDAP Attribute:** Maps to the Liberty attribute *MsgAccount*, which is the account or address information within the messaging provider.
 - ♦ **SubAccount LDAP Attribute:** Maps to the Liberty *MsgSubaccount*, which is the subaccount within a messaging account, such as voice mail box associated with a phone number.
- 5 Under *Contact Method Template Data*, specify the settings for the Liberty attribute values of:
 - ♦ **Nickname:** Maps to the Liberty attribute *Nick*, which is an informal name for the contact.
 - ♦ **Type:** Maps to the Liberty attribute *MsgType* (such as *Mobile*, *Personal*, or *Work*).
 - ♦ **Method:** Maps to the Liberty *MsgMethod* (such as *Voice*, *Fax*, or *E-mail*).
 - ♦ **Technology:** Maps to the Liberty attribute *MsgTechnology* (such as *Pager*, *VOIP*, and so on).
- 6 Click *Finish*.
- 7 On the LDAP Attribute Mapping page, click *OK*.
- 8 Update the Identity Server configuration on the Setup page as prompted.

11.9.6 Configuring Gender Attribute Maps

You can map the LDAP attribute name and values to the Liberty profile values for the Gender attribute. You can use gender to differentiate between people with the same name, especially in countries where national ID numbers cannot be collected. This is a Personal Profile attribute.

- 1 In the Administration Console, click *Access Manager > Identity Server > Setup > [Configuration] > Liberty > LDAP Attribute Mapping > New > Gender*.

New Gender LDAP Attribute Mapping ?

Specify name, description, user stores and mapping data.

Name:

Description:

Access Rights:

User stores: Available user stores:

Gender to LDAP Attribute

LDAP Attribute Name:

Liberty Profile Values to LDAP Attribute Values

Gender Value: LDAP Attribute Value:

Male:

Female:

- 2 Specify a name and description for the map.
- 3 Choose the type of access rights you want.
Select *Read/Write* for any attributes used in user provisioning.
- 4 In the *LDAP Attribute Name* field, type the LDAP attribute name that you want to map to the Liberty element Gender.
- 5 In the *LDAP Attribute Value* fields, type your predefined LDAP attribute values that you want to map to the Gender values.
These are the values that you want to store in the LDAP attribute for each given Liberty attribute value. The LDAP attribute map then maps the actual Liberty URI value, back and forth, to this supplied value.
- 6 Click *Finish*.
- 7 On the LDAP Attribute Mapping page, click *OK*.
- 8 Update the Identity Server configuration on the Setup page as prompted.

11.9.7 Configuring Marital Status Attribute Maps

You can map the LDAP marital status attribute to the Liberty attribute. The Liberty Marital Status (MaritalStatus) element includes appended values such as single, married, divorced, and so on. For

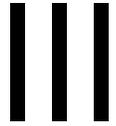
example, `urn:liberty:id-sis-pp:maritalstatus:single`. This is a Personal Profile attribute.

- 1 In the Administration Console, click *Access Manager > Identity Server > Setup > [Configuration] > Liberty > LDAP Attribute Mapping > New > Marital Status*.

Marital Status Value:	LDAP Attribute Value:
Single:	<input type="text" value="Single"/>
Married:	<input type="text" value="Married"/>
Common Law Marriage:	<input type="text" value="Common Law Marriage"/>
Separated:	<input type="text" value="Separated"/>
Divorced:	<input type="text" value="Divorced"/>

- 2 Specify a name and description for the map.
- 3 Choose the type of access rights you want.
Select *Read/Write* for any attributes used in user provisioning.
- 4 In the *LDAP Attribute Name* field, type the LDAP attribute name that you want to map to the Liberty element Marital Status (`MaritalStatus`).
- 5 In the *LDAP Attribute Value* fields, type your predefined LDAP attribute values that you want to map to the MaritalStatus values.
These are the values that you want to store in the LDAP attribute for each given Liberty attribute value. The LDAP attribute map then maps the actual Liberty URI value, back and forth, to this supplied value.
- 6 Click *Finish*.
- 7 On the LDAP Attribute Mapping page, click *OK*.
- 8 Update the Identity Server configuration on the Setup page as prompted.

Access Gateway Configuration



This section describes how you configure and manage the Novell® Access Gateway. The procedures in this section assume that you have already done the following:

- ◆ Installed the Access Gateway. (See *Novell Access Manager 3.0 Installation Guide*).
- ◆ Logged in to the Administration Console as the admin user and that you have expanded the Access Manager task. (See “Logging in to the Administration Console” in the *Novell Access Manager 3.0 Installation Guide*.)
- ◆ Created an Identity Server configuration. (See Chapter 5, “Configuring an Identity Server,” on page 53.)

You should be familiar with the steps documented in the “Setting Up a Basic Access Manager Configuration” in the *Novell Access Manager 3.0 Setup Guide*. These guides are designed to work together.

When you click *Access Gateways* in the Administration Console, the following screen appears.

Access Gateways

Servers		Groups			
<input type="checkbox"/> Server	Server Status	Alerts	Command Status - Group	Statistics	Configuration
<input type="checkbox"/> 10.10.159.190		0	Succeeded - Succeeded	View	Edit: spc
<input type="checkbox"/> 10.10.159.72		0	Succeeded	View	Edit
<input type="checkbox"/> 10.15.167.45		4	Succeeded - Succeeded	View	Edit: aq45

The tabs and the links on this page allow you to manage the Access Gateways on your network. The following sections describe these tasks.

- ◆ Chapter 12, “Configuring the Access Gateway to Protect Resources,” on page 141
- ◆ Chapter 13, “Configuring the Access Gateway for SSL,” on page 171
- ◆ Chapter 14, “Additional Configuration Settings,” on page 181
- ◆ Chapter 15, “Configuring the Cache Settings,” on page 201
- ◆ Chapter 16, “Protecting Multiple Resources,” on page 215
- ◆ Chapter 17, “Configuring Access Manager for Citrix Clients,” on page 233
- ◆ Chapter 18, “Virtualization on the Linux Access Gateway,” on page 237

For auditing and logging, reviewing statistics, command status, and alerts, see Part VII, “Monitoring Access Manager Components,” on page 399.

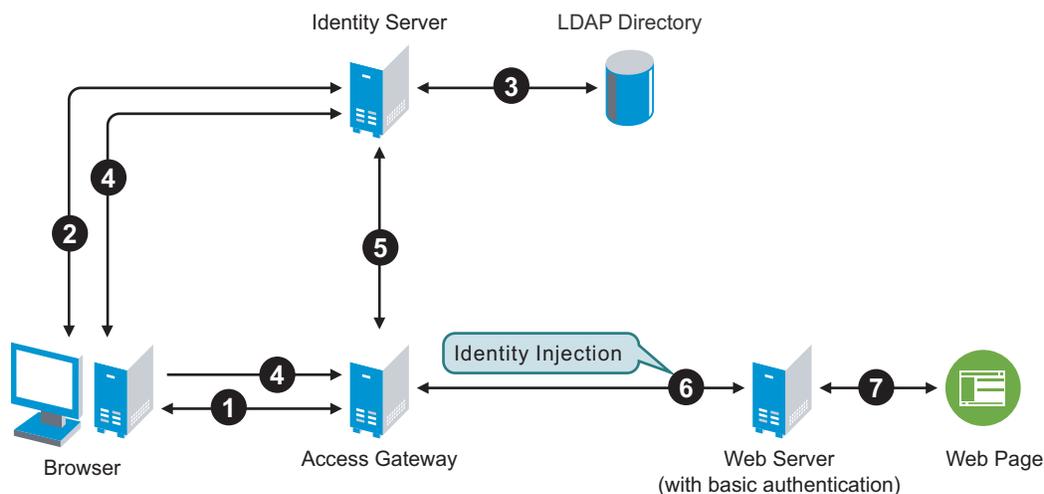
Configuring the Access Gateway to Protect Resources

12

The Novell® Access Gateway is a reverse proxy server (protected site server) that restricts access to Web-based content, portals, and Web applications that employ authentication and access control policies. It also provides single sign-on to multiple Web servers and Web applications by securely providing the credential information of authenticated users to the protected servers and applications. The Access Gateway lets you simplify, secure, and accelerate your Internet business initiatives.

A typical Access Manager configuration includes an Identity Server with LDAP directories and an Access Gateway with a protected Web server. **Figure 12-1** illustrates the process flow that allows an authorized user to access the protected resource.

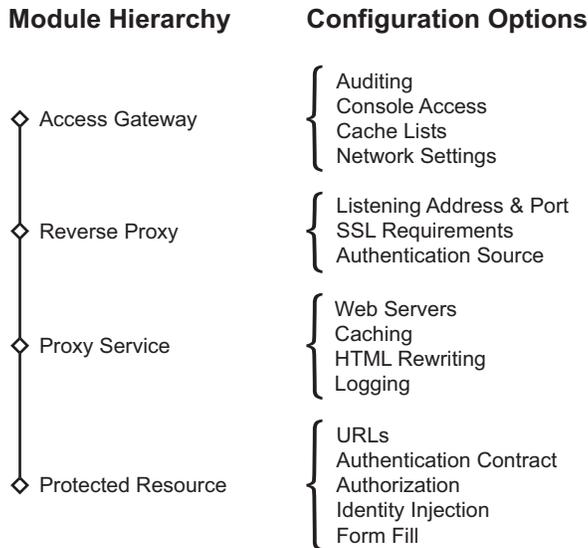
Figure 12-1 Accessing a Web Resource



1. The user requests access to a resource protected by the Access Gateway.
2. The Access Gateway redirects the user to the Identity Server, which prompts the user for a username and password.
3. The Identity Server verifies the username and password against an LDAP directory (eDirectory™, Active Directory, or Sun ONE).
4. The Identity Server returns an authentication success to the browser and the browser forwards the resource request to the Access Gateway.
5. The Access Gateway verifies that the user is authenticated and retrieves the user's credentials from the Identity Server.
6. The Access Gateway uses an Identity Injection policy to insert the basic authentication credentials in the HTTP header of the request and sends it to the Web server.
7. The Web server grants access and sends the requested page to the user.

When you are setting up the Access Gateway to protect Web resources, you create and configure reverse proxies, proxy services, and protected resources. The following figure illustrates the hierarchy of these modules and the major configuration tasks you perform on each module.

Figure 12-2 Access Gateway modules and their configuration options



This hierarchy allows you to have precise control over what is required to access a particular resource, while at the same time allowing you to provide a single sign-on solution for all the resources protected by the Access Gateway. The authentication contract and the Authorization, Identity Injection, and Form Fill policies are configured at the resource level so that you can enable exactly what the resource requires. This allows you to decide whether you want the Access Gateway to control access to the resources, whether you want the Web server configured for access control and the Access Gateway configured to supply the required information, or whether you want to use the first method for some resources and the second method for other resources.

This section describes the following tasks:

- ◆ [Section 12.1, “Creating a Reverse Proxy and Proxy Service,” on page 142](#)
- ◆ [Section 12.2, “Configuring a Proxy Service,” on page 146](#)
- ◆ [Section 12.3, “Configuring the Web Servers of a Proxy Service,” on page 148](#)
- ◆ [Section 12.4, “Configuring Protected Resources,” on page 150](#)
- ◆ [Section 12.5, “Configuring HTML Rewriting,” on page 155](#)
- ◆ [Section 12.6, “Configuring Connection Limits,” on page 166](#)

12.1 Creating a Reverse Proxy and Proxy Service

A reverse proxy acts as the front end to your Web servers on your Internet or intranet and off-loads frequent requests, thereby freeing up bandwidth. The proxy also increases security because the IP addresses of your Web servers are hidden from the Internet.

To create a reverse proxy, you must create at least one proxy service with a protected resource. You must supply a name for each of these components. Reverse proxy names and proxy service names must be unique to the Access Gateway because they are configured for global services such as IP

addresses and TCP ports. For example, if you have a reverse proxy named `products` and another reverse proxy named `library`, only one of these reverse proxies can have a proxy service named `corporate`.

Protected resource names need to be unique to the proxy service, but they don't need to be unique to the Access Gateway because they are always accessed through their proxy service. For example, if you have a proxy service named `account` and a proxy service named `sales`, they both can have a protected resource named `public`.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Reverse Proxy / Authentication*.

The *Edit* link is either *Edit* for a single Access Gateway or *Edit: <Name of the Group>* for a group of Access Gateways.

The screenshot shows a configuration panel with two main sections. The top section is titled "Authentication Settings" and contains a label "Identity Server Configuration:" followed by a dropdown menu currently set to "[None]". The bottom section is titled "Reverse Proxy List" and has a blue header. Below the header are links for "New...", "Delete", "Enable", and "Disable". There is a checkbox and a table header with columns: "Name", "Enabled", "Listening Address", and "Port". Below the table, it says "No items". At the bottom of the panel, there is a note: "Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel." and two buttons: "OK" and "Cancel".

- 2 Select an *Identity Server Configuration*.

Identity Server Configuration: Specifies the Identity Server you want the Access Gateway to trust for authentication. Select the configuration you have assigned to the Identity Server.

Whenever the Identity Server is assigned a new trust relationship, it needs to be updated. This process is explained following the step that saves this configuration setting.

- 3 In the *Reverse Proxy List*, click *New*, specify a display name for the reverse proxy, then click *OK*.

Group Member: 

Listening Address(es): 10.10.159.206

[TCP Listen Options](#)

Enable SSL with Embedded Service Provider

Enable SSL between Browser and Access Gateway

Redirect Requests from Non-Secure Port to Secure Port

Server Certificate: 

Non-Secure Port: * (Used for Trusted IDS Communication, HTTP Listening)

Secure Port: (Unused)

- 4 Enable a listening address. Fill in the following fields:

Group Member: (Available only if the Access Gateway is a member of a group.) Select the server you want to configure from the list of servers. The *Listening Address(es)* and *TCP Listen Options* modifications apply to the selected server. Modifications made to any other options on the page apply to all servers in the group.

Listening Address(es): Displays a list of available IP addresses. If the server has only one IP address, only one is displayed and it is automatically selected. If the server has multiple addresses, you can select one or more IP addresses to enable. You must enable at least one address by selecting its check box.

If the Access Gateway is in a group, you must select a listening address for each group member.

TCP Listen Options: Provides options for configuring how requests are handled between the reverse proxy and the client browsers. You cannot set up the listening options until you create and configure a proxy service. For information about these options, see [Section 12.6.1, “Configuring TCP Listen Options for Clients,” on page 167](#).

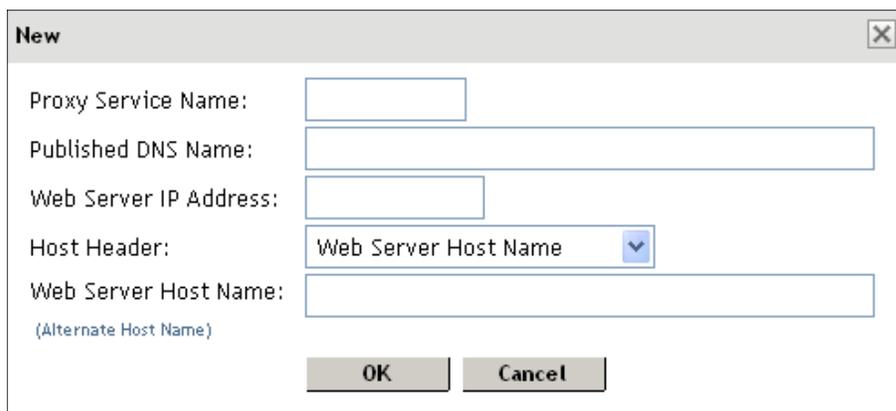
- 5 Configure the listening ports:

Non-Secure Port: Specifies the port on which to listen for HTTP requests; the default port for HTTP is 80. Depending upon your configuration, this port might also handle other tasks. These tasks are listed to the right of the text box.

Secure Port: Specifies the port on which to listen for HTTPS requests; the default port for HTTPS is 443.

For information about the SSL options, see [Chapter 13, “Configuring the Access Gateway for SSL,” on page 171](#).

6 In the *Proxy Service List* section, click *New*.



The first proxy service of a reverse proxy is considered the master (or parent) proxy. Subsequent proxy services can use domain-based, path-based, or virtual multi-homing, relative to the published DNS name of the master proxy service. If you are creating a second proxy service for a reverse proxy, see [Section 16.2, “Using Multi-Homing to Access Multiple Resources,”](#) on page 216.

7 Fill in the fields:

Proxy Service Name: Specify a display name for the proxy service, which the Administration Console uses for its interfaces.

Published DNS Name: Specify the DNS name you want the public to use to access your site. This DNS name must resolve to the IP address you set up as the listening address.

Web Server IP Address: Specify the IP address of the Web server you want this proxy service to manage. You can specify additional Web server IP addresses by clicking the Web Server Addresses link once you have finished creating the proxy service.

Host Header: Specify whether the HTTP header should contain the name of the back-end Web server (*Web Server Host Name* option) or whether the HTTP header should contain the published DNS name (the *Forward Received Host Name* option).

Web Server Host Name: Specify the DNS name of the Web server that the Access Gateway should forward to the Web server. If you have set up a DNS name for the Web server and it requires its DNS name in the HTTP header, specify that name in this field. If the Web server has absolute links referencing its DNS name, include this name in this field. If you selected *Forward Received Host Name*, this option is not available.

NOTE: For iChain[®] administrators, the *Web Server Host Name* is the alternate host name when configuring a Web Server Accelerator.

8 Click *OK*.

9 Continue with [Section 12.2, “Configuring a Proxy Service,”](#) on page 146 or select one of the following tasks:

- For instructions on creating multiple reverse proxies, see [Section 16.3, “Managing Multiple Reverse Proxies,”](#) on page 224.
- For instructions on creating multiple proxy services for a reverse proxy, see [Section 16.2, “Using Multi-Homing to Access Multiple Resources,”](#) on page 216.

12.2 Configuring a Proxy Service

A reverse proxy can have multiple proxy services, and each proxy service can protect multiple resources. You can modify the following features of the proxy service:

- ♦ Web servers
- ♦ HTML rewriting
- ♦ Logging
- ♦ Protected resources
- ♦ Caching

- 1 To configure a proxy service, click *Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service]*.

Proxy Service Web Servers HTML Rewriting Protected Resources Logging

Published DNS Name: sales.mycompany.com

Description:

Cookie Domain: mycompany.com

[HTTP Options](#)

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

OK Cancel

- 2 Fill in the following fields:

Published DNS Name: Displays the value that users are currently using to access this proxy service. This DNS name must resolve to the IP address you set up as a listening address on the Access Gateway. You should modify this field only if you have modified the DNS name you want users to use to access this resource.

This name determines the possible values of the *Cookie Domain*.

Description: (Optional). Provides a field where you can describe the purpose of this proxy service or specify any other pertinent information.

Cookie Domain: Specifies the domain for which the cookie is valid. The Web server the user is accessing must be configured to be part of this domain.

The *Cookie Domain* field can be used to allow single authentication across multiple reverse proxies where the cookie domain by default would be different for each reverse proxy.

For example, if one reverse proxy has a DNS name of `www.support.novell.com` and the second reverse proxy has a DNS name of `www.developernet.novell.com`, the cookie domains would be `support.novell.com` for the first reverse proxy and `developernet.novell.com` for the second reverse proxy. The cookie domains are different for each reverse proxy, so users are prompted for authentication when accessing the second reverse proxy, even if they have already

authenticated to the first reverse proxy. If the cookie domain is changed to novell.com on each reverse proxy, then users do not need to authenticate again when accessing the second reverse proxy if they have already authenticated through the first reverse proxy.

If you have multiple proxy services, the cookie domain is automatically set to the subdomain common to the group. If you select a more specific domain for the cookie, the users are forced to log in when accessing the resources of this proxy service.

A cookie is specific to a single authentication contract, so even though the cookie domains are the same, the cookie set by the first reverse proxy is not valid for the second reverse proxy if you use different authentication contracts.

HTTP Options: Allows you to set up global caching and custom caching options for this proxy service. See the following:

- ♦ [Section 15.2, “Controlling Browser Caching,” on page 204](#)
- ♦ [Section 15.3, “Configuring Custom Cache Control Headers,” on page 205](#)
- ♦ [Section 15.1, “Configuring Global Caching Options,” on page 202](#)

3 At the bottom of the page, click *Configuration Panel*, then click *OK*.

4 On the Server Configuration page, click *Apply Changes*, then click *Okay*.

Until this step, nothing has been saved. The *Apply Changes* button pushes the configuration to the server. When the configuration process has completed, the server returns the status of the changes.

5 Update the Identity Server to accept the new trusted relationship. Click *Identity Servers > Setup > Update Servers*.

6 Continue with one of the following.

- ♦ If the Web server that contains the resources you want to protect does not use the standard HTML port (port 80), you need to configure the Web server. See [Section 12.3, “Configuring the Web Servers of a Proxy Service,” on page 148](#).
- ♦ Until you configure a protected resource, the proxy service blocks access to all services on the Web server. To configure a protected resource, see [Section 12.4, “Configuring Protected Resources,” on page 150](#).

12.3 Configuring the Web Servers of a Proxy Service

The Web server configuration determines how the Access Gateway handles connections and packets between itself and the Web servers.

- 1 Click *Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > Web Servers* tab.

Proxy Service Web Servers HTML Rewriting Protected Resources Logging

Host Header: Forward Received Host Name

Web Server Host Name:
(Alternate Host Name)

Error on DNS Mismatch

Enable Force HTTP 1.0 to Origin

Enable Forwarding of Encoding Header

Connect Using SSL

Web Server Trusted Root: Any in Reverse Proxy Trust Store

SSL Mutual Certificate:

Connect Port: * 443

[TCP Connect Options](#)

- 2 Specify the host name that is placed in the HTTP header of the packets being sent to the Web servers. In the *Host Header* field, select one of the following:
 - ♦ **Forward Received Host Name:** Indicates that you want the HTTP header to contain the published DNS name that the user sent in the request.
 - ♦ **Web Server Host Name:** Indicates that you want the published DNS name that the user sent in the request to be replaced by the DNS of the Web server. Use the *Web Server Host Name* field to specify this name.
- 3 To report errors when the name in the HTTP header does not match what the Web server expects, select *Error on DNS Mismatch*. This is automatically enabled when you select to send the *Web Server Host Name* in the HTTP header.
- 4 If your browsers are capable of sending HTTP 1.1 requests, configure the following fields to match your Web servers.

Enable Force HTTP 1.0 to Origin: Indicates whether HTTP 1.1 requests from browsers are translated to HTTP 1.0 requests before sending them to the Web server. If your browsers are sending HTTP 1.1 requests and your Web server can only handle HTTP 1.0 requests, you should enable this option.

When the option is enabled, the Access Gateway translates an HTTP 1.1 request to an HTTP 1.0 request.

Enable Forwarding of Encoding Header: Determines whether the HTTP 1.1 header is sent to the Web server:

- ♦ If you enable this option, the entire HTTP 1.1 header is sent to the Web server. If your browsers are sending HTTP 1.1 requests and your Web servers are HTTP 1.1 compliant, this is the configuration you should use.
 - ♦ If you enable this option and you have also enabled the *Enable Force HTTP 1.0 to Origin* option, a few select fields of the HTTP 1.1 header (such as the content encoding header for compression) are sent to the Web server. If your Web server is not HTTP 1.1-compliant, but it can handle a few HTTP 1.1 fields, you should enable this option.
 - ♦ If your Web server can handle only HTTP 1.0 headers, you should not enable this option.
- 5** To enable SSL connections between the proxy service and its Web servers, select *Connect Using SSL*. For configuration information for this option, *Web Server Trusted Root*, and *SSL Mutual Certificate*, see [Section 13.4, “Configuring SSL between the Proxy Service and the Web Servers,”](#) on page 176.
- 6** In the *Connect Port* field, specify the port that the Access Gateway should use to communicate with the Web servers. The following table lists some default port values for common types of Web servers.

Server Type	Non-Secure Port	Secure Port
Web server with HTML content	80	443
SSL VPN	8080	8443
WebSphere	9080	9443
JBoss	8080	8443

- 7** To control how idle and unresponsive Web server connections are handled and to optimize these processes for your network, select *TCP Connect Options*. For more information, see [Section 12.6.2, “Configuring TCP Connect Options for Web Servers,”](#) on page 168.
- 8** To add a Web server, click *New* in the *Web Server List* and specify the IP address of the Web server.
- The Web servers added to this list must contain identical Web content. Configuring your system with multiple servers with the same content adds fault tolerance and increases the speed for processing requests. For more information about this process, see [Section 16.1, “Setting Up a Group of Web Servers,”](#) on page 216.
- 9** To delete a Web server, select the Web server, then click *Delete*.
- This deletes the Web server from the list so that the Access Gateway no longer sends requests to the deleted Web server. At least one Web server must remain in the list. You must delete the proxy service to remove the last server in the list.
- 10** To save your changes and push them to the Access Gateway, click *Configuration Panel*, then *OK*.
- 11** Click *Apply Changes*.

12.4 Configuring Protected Resources

A protected resource configuration specifies the directory (or directories) on the Web server that you want to protect. The protected resource configuration specifies the authorization contract and the policies that should be used to enforce protection. The authorization contract and policies (Authorization, Identity Injection, and Form Fill) enable the single sign-on environment for the user. The type of protections a resource requires depends upon the resource, the Web server, and the conditions you define for the resource.

You can select from the following types of protection:

Authentication Contract: Specifies the type of credentials the user must use to log in (such as name and password or secure name and password). You can select *None* for the contract, which allows the resource to be a public resource, with no login required.

Authorization Policy: Specifies the conditions a user must meet to be allowed access to a protected resource. You define the conditions, and the Access Gateway enforces the Authorization policies. For example, you can assign roles to your users, and use these roles to grant and deny access to resources.

Identity Injection Policy: Specifies the information that must be injected into the HTTP header. If the Web application has been configured to look for certain fields in the header and the information cannot be found, the Web application determines whether the user is denied access or redirected. The Web application defines the requirements for Identity Injection. The Identity Injection policies allow you to inject the required information into the header.

Form Fill Policy: Allows you to manage forms that Web servers return in response to client requests. Form fill allows you to pre-populate fields in a form on first login and then securely save the information in the completed form to a secret store for subsequent logins. The user is prompted to re-enter the information only when something changes, such as a password.

These policies allow you to design a custom policy for each protected resource:

- ◆ Resources that share the same protection requirements can be configured as a group. You set up the policies, and then add the URLs of each resource that requires these policies.
- ◆ A resource that has specialized protection requirements can be set up as a single protected resource. For example, a page that uses Form Fill is usually set up as a single protected resource.

This section describes the following tasks:

- ◆ [Section 12.4.1, “Setting Up a Protected Resource,” on page 150](#)
- ◆ [Section 12.4.2, “Assigning an Authorization Policy to a Protected Resource,” on page 152](#)
- ◆ [Section 12.4.3, “Assigning an Identity Injection Policy to a Protected Resource,” on page 153](#)
- ◆ [Section 12.4.4, “Assigning a Form Fill Policy to a Protected Resource,” on page 154](#)

12.4.1 Setting Up a Protected Resource

To configure a protected resource:

- 1 Click *Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > Protected Resources* tab.

- 2 Either click the name of an existing resource or click *New*, then specify a display name for the resource.

Overview Authorization Identity Injection Form Fill

Protected Resource: basic

Description:

Contract:

URL Path List	
New... Delete	1 item(s)
<input type="checkbox"/> URL Path	
<input type="checkbox"/> /*	

- 3 (Optional) Specify a description for the protected resource. You can use it to briefly describe the purpose for protecting this resource.
- 4 Select the type of contract, which determines the information a user must supply for authentication. By default, the Administration Console allows you to select from the following contracts and options when specifying whether a resource requires an authentication contract:
- ♦ **None:** If you want to allow public access to the resource and not require an authentication contract, select *None*.
 - ♦ **Any Contract:** If the user has authenticated, allows any contract defined for the Identity Server to be valid, or if the user has not authenticated, prompts the user to authenticate using the default contract assigned to the Identity Server configuration.
 - ♦ **Name/Password - Basic:** Specifies basic authentication over HTTP using a standard login pop-up screen provided by the Web browser.
 - ♦ **Name/Password - Form:** Specifies a form-based authentication over HTTP using the Access Manager login form.
 - ♦ **Secure Name/Password - Basic:** Specifies basic authentication over HTTPS using a standard login pop-up screen provided by the Web browser.
 - ♦ **Secure Name/Password - Form:** Specifies a form-based authentication over HTTPS using the Access Manager login form.

You can configure other types of contract types. See [Section 7.4, “Configuring Authentication Contracts,”](#) on page 82.

If these default contracts are not available, you have not configured a relationship between the Access Gateway and the Identity Server. See [Section 12.1, “Creating a Reverse Proxy and Proxy Service,”](#) on page 142.

- 5 Configure the *URL Path*.

The default path is */**, which indicates everything on the Web server. Modify this if you need to restrict access to a specific directory on your Web server. If you have multiple directories on

your Web server that require the same authentication contract and access control, add each directory as a URL path.

- ♦ **New:** To add a path, click *New*, specify the path, then click *OK*. For example, to allow access to all the pages in the public directory on the Web server, specify the following path:

/public/*

To use this protected resource to protect a single page, specify the path and the filename. For example, to protect the login.html page in the /login directory, specify the following:

/login/login.html

This is the type of URL path you want to specify when you create a Form Fill policy for a protected resource. The *URL Path List* normally contains only this one entry. If you have multiple pages that the Form Fill policy applies to, list each one separately in the list. For optimum speed, you want the Access Gateway to be able to quickly identify the page and not search other pages to see if the policy applies to them.

- ♦ **Modify:** To modify a path, click the path link, then modify the *URL Path*.
- ♦ **Delete:** To delete a path, select the path, then click *Delete*.

6 Click *OK*.

7 Select the protected resource you created and click *Enable*.

8 To add policies for protecting this resource, continue with one of the following:

- ♦ “Assigning an Authorization Policy to a Protected Resource” on page 152
- ♦ “Assigning an Identity Injection Policy to a Protected Resource” on page 153
- ♦ “Assigning a Form Fill Policy to a Protected Resource” on page 154

12.4.2 Assigning an Authorization Policy to a Protected Resource

An Authorization policy specifies conditions that a user must meet in order to access a resource. The Access Gateway enforces these conditions. The policy can specify the criteria a user must meet either to allow access or to deny access.

1 Click *Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > Protected Resources > [Name of Protected Resource] > Authorization*.

Name	Enabled	Policy Container	Description
deny_but_manager	<input type="checkbox"/>	Master_Container	

The Authorization Policy List contains all the Access Gateway Authorization policies that have been created on this Administration Console.

2 Select one of the following:

- ♦ To enable an existing policy, select the policy, then click *Enable*. Continue with [Step 4](#).
- ♦ To disable an existing policy, select the policy, then click *Disable*. Continue with [Step 4](#).
- ♦ To edit an existing policy, click the name of the policy. Remember that policies can be assigned to multiple protected resources. If you modify the policy, you are also affecting how this policy protects those resources. Continue with [Step 4](#).
- ♦ To create a new policy, click *Manage Policies*. On the Policies page, click *New*, specify a display name, select *Access Gateway: Authorization* as the type, then click *OK*. For configuration information, see [Section 30.2, “Creating Access Gateway Authorization Policies,”](#) on page 344.

When you have created your policy, continue with [Step 3](#).

3 To enable the policy you just created, select the policy, then click *Enable*.

Only the policies that are enabled are applied to this resource. All available Authorization policies are listed. If you use the same policy for multiple protected resources, use the policy description field to indicate this.

4 To save your changes and push them to the Access Gateway, click *Configuration Panel*, then click *Apply Changes*.

12.4.3 Assigning an Identity Injection Policy to a Protected Resource

The Web application defines the requirements for Identity Injection. If a Web application has been configured to look for certain fields in the header and the information cannot be found, the Web application determines whether the user is denied access, granted access, or redirected. You configure an Identity Injection policy to inject into the HTTP header the information that the Web application requires.

1 Click *Access Gateways > Edit > [Reverse Proxy Name] > [Name of Proxy Service] > Protected Resources > [Name of Protected Resource] > Identity Injection*.

<input type="checkbox"/>	Name	Enabled	Policy Container	Description
<input type="checkbox"/>	cred_ii	✓	Master_Container	
<input type="checkbox"/>	custom_ii	✓	Master_Container	
<input type="checkbox"/>	SSLVPN_Default		Master_Container	

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

The *Identity Injection Policy List* contains all the Identity Injection policies that have been created on this Administration Console.

2 Select one of the following:

- ♦ To enable an existing policy, select the policy, then click *Enable*. Only the policies that are enabled are applied to this resource. Continue with [Step 4](#).

- ♦ To disable an existing policy, select the policy, then click *Disable*. Continue with [Step 4](#).
- ♦ To edit an existing policy, click the name of the policy. Remember that policies can be assigned to multiple protected resources. If you modify the policy, you are also affecting how this policy protects those resources. Continue with [Step 4](#).
- ♦ To create a new policy, click *Manage Policies*. On the Policies page, click *New*, specify a display name, select *Access Gateway: Identity Injection* as the type, then click *OK*. For configuration information, see [Chapter 31, “Creating Identity Injection Policies,” on page 373](#).

When you have created your policy, continue with [Step 3](#).

- 3** To enable the policy you just created, select the policy, then click *Enable*.

Only the policies that are enabled are applied to this resource. All available Identity Injection policies are listed. If you use the same policy for multiple protected resources, use the policy description field to indicate this.

- 4** To save your changes and push them to the Access Gateway, click *Configuration Panel*, then click *Apply Changes*.

12.4.4 Assigning a Form Fill Policy to a Protected Resource

Some client requests cause the Web server to return a form. Sometimes this form contains a request to log in. If you create a Form Fill policy, you can have the Access Gateway fill in the form. When a user first logs in, the Access Gateway prepopulates some fields and prompt the users for the others. The Access Gateway securely saves the information, so that on subsequent logins, the Access Gateway can fill in the form. The user is only prompted to fill in the form when something changes, such as a password expiring.

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 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)

You must create the policy before you can assign it to a resource (see [Chapter 32, “Creating Form Fill Policies,” on page 385](#)). To assign a Form Fill policy to a protected resource:

- 1** In the Administration Console, click *Access Manager > Access Gateways > Edit > [Reverse Proxy Name] > [Name of Proxy Service] > Protected Resources > [Name of Protected Resource]*.
- 2** Examine the entries in the *URL Path List*.

Ideally, the URL to which you are assigning a Form Fill policy should be a single HTML page or a few HTML pages. If at all possible, it should not be a URL that ends in a wildcard (for example, an asterisk) and therefore matches many pages. When the URL ends in a wildcard, the Access Gateway has to search each page that matches the URL and check to see if it contains the form. This adds extra processing overhead for all the pages that match the URL, but do not contain the form.

3 > Form Fill.

Form Fill Policy List			
Manage Policies Enable Disable			
<input type="checkbox"/>	Name	Enabled	Policy Container Description
<input type="checkbox"/>	simple_ff		Master_Container

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

The Form Fill Policy List contains all the Form Fill policies that have been created on this Administration Console.

4 Select one of the following:

- ♦ To enable an existing policy, select the policy, then click *Enable*. Only the policies that are enabled are applied to this resource. Continue with [Step 4](#).
- ♦ To disable an existing policy, select the policy, then click *Disable*. Continue with [Step 4](#).
- ♦ To edit an existing policy, click the name of the policy. Remember that policies can be assigned to multiple protected resources. If you modify the policy, you are also affecting how this policy protects those resources. Continue with [Step 4](#).
- ♦ To create a new policy, click *Manage Policies*. On the Policies page, click *New*, specify a display name, select *Access Gateway: Form Fill* as the type, then click *OK*. For configuration information, see, see [Chapter 32, “Creating Form Fill Policies,” on page 385](#).

When you have created your new policy, continue with [Step 3](#).

5 To enable the policy you just created, select the policy, then click *Enable*.

Only the policies that are enabled are applied to this resource. All available Form Fill policies are listed. If you use the same policy for multiple protected resources, use the policy description field to indicate this.

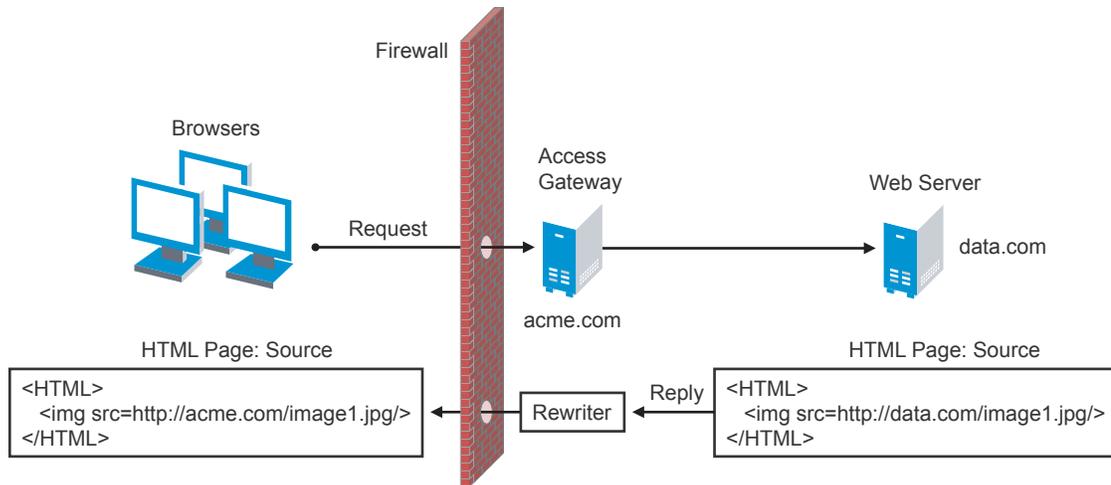
6 To save your changes and push them to the Access Gateway, click *Configuration Panel*, then click *Apply Changes*.

12.5 Configuring HTML Rewriting

Access Gateway configurations require HTML rewriting because the Web servers are not aware that the Access Gateway machine is obfuscating their DNS names. Any reference URLs contained in their pages must be checked so that these references contain the DNS names that the client browser expects. On the other end, the client browsers are not aware that the Access Gateway is obfuscating the DNS names of the resources they are accessing. The URL requests coming from the client browsers that use published DNS names must be rewritten to the DNS names that the Web servers

expect. [Figure 12-3](#) illustrates how the reply message is rewritten to use the published DNS name of the Web server.

Figure 12-3 *Rewriting a Reply Page*



The following sections describe this process:

- ◆ [Section 12.5.1, “Understanding the Rewriting Process,”](#) on page 156
- ◆ [Section 12.5.2, “Configuring the HTML Rewriter,”](#) on page 162

12.5.1 Understanding the Rewriting Process

The Access Gateway needs to rewrite URL references under the following conditions:

- ◆ To ensure URL references contain the proper scheme (HTTP or HTTPS).

If your Web servers and Access Gateway machines are behind a secure firewall, you might not require SSL connections between them, and only require SSL between the client browser and the Access Gateway. For example, an HTML file being accessed through the Access Gateway for the Web site mynovell.com might have a URL reference to `http://mynovell.com/file1.html`. If the reverse proxy for mynovell.com is using SSL sessions between the browser and Access Gateway, 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 ensure URL references that contain private IP addresses or private DNS names are changed to the published DNS name of the Access Gateway or hosts.

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 the Access Gateway using a published 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 HTML rewriter can resolve this issue. The DNS name field in the Access Gateway configuration is set to mynovell.com, which users can resolve through a public DNS server to the Access Gateway. The rewriter parses Web content retrieved through the Access Gateway, and any URL references matching the private DNS name or private IP address listed in the Web

server address field of the Access Gateway configuration are changed (rewritten) with the published DNS name `mynovell.com` and the port number of the Access Gateway.

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 sensitive information.

- ◆ To ensure that the Host header in incoming HTTP packets contains 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 published DNS name `mynovell.com`, the Host field of the packets in those requests received by the Access Gateway is set to `mynovell.com`. The Access Gateway can be configured to rewrite this public name to the private name expected by the Web server by setting the Web Server Host Name option to `internal.web.site.com`. Before the Access Gateway forwards packets to the Web server, the Host field is changed (rewritten) from `mynovell.com` to `internal.web.site.com`. For information about configuring this option, see [“Configuring the Web Servers of a Proxy Service” on page 148](#).

The sections below describe additional features of the rewriter:

- ◆ [“Evaluating Criteria for URLs” on page 157](#)
- ◆ [“Determining Whether You Need to Specify Additional DNS Names” on page 158](#)
- ◆ [“Determining Whether You Need to Exclude DNS Names from Being Rewritten” on page 159](#)
- ◆ [“Types of Rewriter Profiles” on page 160](#)
- ◆ [“String Replacement Rules for Character Profiles” on page 161](#)
- ◆ [“String Replacement Rules for Word Profiles” on page 162](#)

Evaluating Criteria for URLs

The rewriter parses and searches the Web content that passes through the Access Gateway for URL references that qualify to be rewritten. URL references are rewritten only if they first meet the following conditions:

- ◆ URL references containing DNS names or IP addresses matching those in the Web server address list are rewritten with the *Published DNS Name*.
- ◆ URL references matching the *Web Server Host Name* are rewritten with the *Published DNS Name*.
- ◆ URL references matching entries in the *Additional DNS Name List* of the host are rewritten with the *Published DNS Name*. The Web Server Host Name does not need to be included in this list because the rewriter is always configured to look for this name.
- ◆ The DNS names in the *Exclude DNS Name List* specify the names that the rewriter should skip and not rewrite.

In addition to conditions listed above, the URLs with the following HTML types must meet additional criteria:

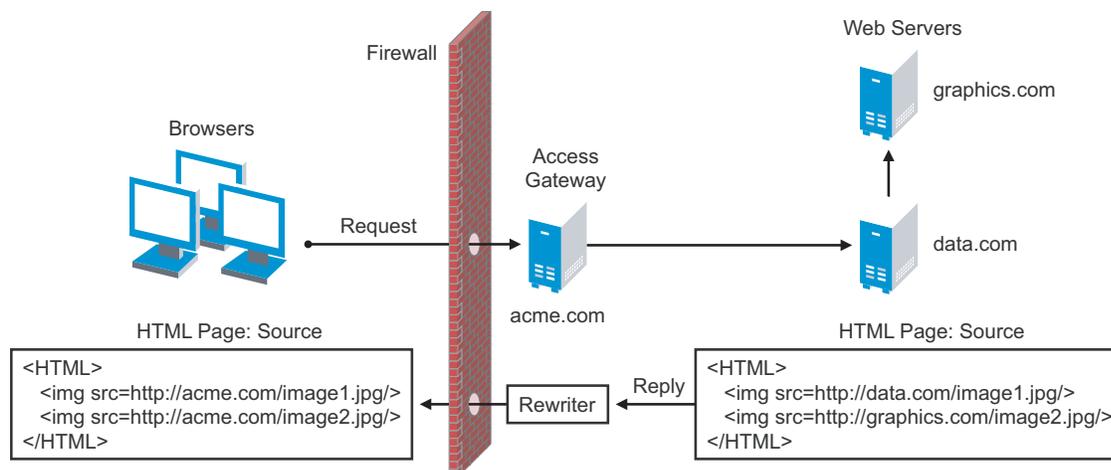
HTML Type	Criteria																		
Query Strings	URL references contained within query strings are not rewritten. Only the hostname portion of the reference is evaluated for rewriting.																		
HTTP Headers	Qualified URL references occurring within certain types of HTTP response headers such as Location and Content-Location are rewritten. 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*, absolute references are always evaluated for rewriting. Relative references (such as <code>index.html</code>) are not attempted. Relative paths (such as <code>/docs/file.html</code>) are evaluated if the file is read from a path-based multi-homing Web server and the reference follows an HTML tag. For example, the string <code>href='/docs/file.html'</code> is rewritten if <code>/docs</code> is a multi-homing path.																		
HTML Tags	<p>URL references occurring within the following HTML tag attributes are evaluated for rewriting:</p> <table border="0"> <tr> <td><code>action</code></td> <td><code>archive</code></td> <td><code>background</code></td> </tr> <tr> <td><code>base</code></td> <td><code>borderimage</code></td> <td><code>cite</code></td> </tr> <tr> <td><code>code</code></td> <td><code>codebase</code></td> <td><code>data</code></td> </tr> <tr> <td><code>dynscr</code></td> <td><code>href</code></td> <td><code>longdesc</code></td> </tr> <tr> <td><code>lowsrc</code></td> <td><code>onclick</code></td> <td><code>pluginspage</code></td> </tr> <tr> <td><code>src</code></td> <td><code>usemap</code></td> <td></td> </tr> </table> <p>The value attribute is not evaluated.</p>	<code>action</code>	<code>archive</code>	<code>background</code>	<code>base</code>	<code>borderimage</code>	<code>cite</code>	<code>code</code>	<code>codebase</code>	<code>data</code>	<code>dynscr</code>	<code>href</code>	<code>longdesc</code>	<code>lowsrc</code>	<code>onclick</code>	<code>pluginspage</code>	<code>src</code>	<code>usemap</code>	
<code>action</code>	<code>archive</code>	<code>background</code>																	
<code>base</code>	<code>borderimage</code>	<code>cite</code>																	
<code>code</code>	<code>codebase</code>	<code>data</code>																	
<code>dynscr</code>	<code>href</code>	<code>longdesc</code>																	
<code>lowsrc</code>	<code>onclick</code>	<code>pluginspage</code>																	
<code>src</code>	<code>usemap</code>																		
Mime Types	<p>The following Mime Content-Types are parsed regardless of the file extension:</p> <table border="0"> <tr> <td><code>text/html</code></td> <td><code>text/javascript</code></td> </tr> <tr> <td><code>text/xml</code></td> <td><code>application/javascript</code></td> </tr> <tr> <td><code>text/css</code></td> <td><code>application/x-javascript</code></td> </tr> </table> <p>If an HTTP or HTTPS response has a Mime Content-Type set to any of the above types, the page is parsed for possible rewriting. It is also parsed if the file extension is <code>html</code>, <code>htm</code>, <code>shtml</code>, <code>jhtml</code>, <code>asp</code>, or <code>jsp</code>.</p>	<code>text/html</code>	<code>text/javascript</code>	<code>text/xml</code>	<code>application/javascript</code>	<code>text/css</code>	<code>application/x-javascript</code>												
<code>text/html</code>	<code>text/javascript</code>																		
<code>text/xml</code>	<code>application/javascript</code>																		
<code>text/css</code>	<code>application/x-javascript</code>																		
Absolute and Relative References	<p>An absolute reference is a reference that has all the information needed to locate a resource, including the hostname, such as <code>http://internal.web.site.com/index.html</code>. The rewriter always attempts to rewrite absolute references.</p> <p>The rewriter attempts to rewrite a relative path only when it is defined in a path-based multi-homing host. For example, <code>/docs/file1.html</code> is rewritten if <code>/docs</code> is a multi-homing path.</p>																		
Path-based Multi-Homing	When a host is configured for path-based multi-homing, absolute references and absolute paths are evaluated for rewriting. Relative references are not attempted.																		

Determining Whether You Need to Specify Additional DNS Names

Sometimes Web pages contain URL references to a host name that does not meet the default criteria for being rewritten. That is, the URL reference does not match the *Web Server Host Name* or any

value (IP address) in the *Web Server List*. If these names are sent back to the client, they are not resolvable. **Figure 12-4** illustrates a scenario that requires an entry in the *Additional DNS Name List*.

Figure 12-4 Rewriting a URLs for Web Servers



The page on the data.com Web server contains two links, one to an image on the data.com server and one to an image on the graphics.com server. The link to the data.com server is automatically rewritten to acme.com, when rewriting is enabled. The link to the image on graphics.com is not rewritten, until you add this URL to the *Additional DNS Name List*. When the link is rewritten, the browser knows how to request it, and the Access Gateway knows how to resolve it.

You need to include names in this list if your Web servers have the following configurations:

- ◆ If you have a cluster of Web servers that are not sharing the same DNS name, you need to add their DNS names to this list.
- ◆ If your Web server obtains content from another Web server, the DNS name to this additional Web server needs to be rewritten.
- ◆ If the Web server listens on one port (for example, 80), and redirects the request to a secure port (for example, 443). The response to the user comes back on https://<DNS_name>:443. This does not match the request which was sent on http://<DNS_name>:80. If you add the DNS name to the list, the response can be sent in the format that the user expects.
- ◆ Sometimes an application is written to use a private DNS name. For example, assume that an application URL reference contains the host name of home (http://home/index.html). This host name would need to be added to the *Additional DNS Name List*.

When you enter a URL in the list, it can use any of the following formats:

```
DNS_name
DNS_name:port
scheme://DNS_name
scheme://DNS_name:port
```

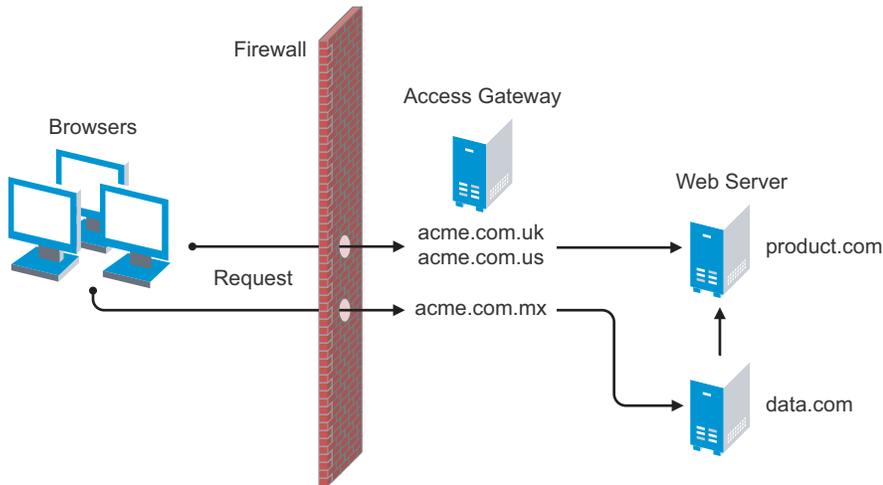
These entries are not case sensitive.

Determining Whether You Need to Exclude DNS Names from Being Rewritten

If you have two reverse proxies protecting the same Web server, the rewriter correctly rewrites the references to the Web server so that browser always uses the same reverse proxy. In other words, if

the browser requests a resource using `acme.com.uk`, the response is returned with references to `acme.com.uk` and not `acme.com.usa`. If you have a third reverse proxy protecting a Web server, the rewriting rules can become ambiguous. For example, consider the configuration illustrated in [Figure 12-5](#).

Figure 12-5 *Excluding URLs*



A user accesses `data.com` through the published DNS name of `acme.com.mx`. The `data.com` server has references to `product.com`. The `acme.com.mx` proxy has two ways to get to the `product.com` server because this Web server has two published DNS names (`acme.com.uk` and `acme.com.usa`). The rewriter could use either of these names to rewrite references to `product.com`.

- ◆ If you want all users coming through `acme.com.mx` to use the `acme.com.usa` proxy, you can add `acme.com.uk` to the Exclude DNS Names List of the resource protecting `data.com`, and the rewriter will always replace references to `product.com` with `acme.com.usa`.
- ◆ If you do not care which proxy is returned in the reference, you do not need to add anything to the Exclude DNS Names List.

Types of Rewriter Profiles

In the HTML profile, you specify the following:

- ◆ The pages you want searched for URLs to rewrite
- ◆ The data items on the page you want searched for URLs to rewrite
- ◆ The text you want rewritten

The Access Gateway uses two types of rewriter profiles:

- ◆ **Word:** A Word profile searches for matches on words. For example, “to” matches only the word “to” and not the “to” in “total”.

The Access Gateway comes with a default Word profile, which is designed to be applied to all pages protected by the Access Gateway. It is not specific to a reverse proxy or its proxy services. When you modify its behavior, remember its scope.

It comes preconfigured to rewrite the *Web Server Host Name* and any other names listed in the *Additional URL List*. It searches for these names in the following document content-type headers:

text/html	text/javascript
text/xml	application/javascript
text/css	application/x-javascript

If this default behavior does not match your requirements for a particular protected resource, create your own Word profile and position it first in the list of profiles. If the page matches your profile, the default profile is not executed. The default profile is executed only for the pages that do not match your profile. Only one Word profile is executed per page. If you create multiple Word profiles, order is important. The first Word profile that matches the page is executed. Profiles lower in the list are ignored.

- ◆ **Character:** A Character profile searches for matches on a specified set of characters. For example, “to” matches both the word “to” and the “to” in “total”. If you want to add functionality to the default profile, create a Character profile. It has all the functionality of a Word profile, except searching for attribute names and Java variables and methods. If you create multiple Character profiles, order is important. The first Character profile that matches the page is executed. Profiles lower in the list are ignored.

If you enable HTML rewriting, but do not enable a profile with a search boundary of Word, URLs are rewritten in the following Content-Type headers:

text/html
text/xml
text/css
text/javascript
application/javascript
application/x-javascript

String Replacement Rules for Character Profiles

When you configure multiple strings for replacement, the rewriter uses the following rules for determining how characters are replaced in strings:

- ◆ String replacement is done as a single pass.
- ◆ String replacement is not performed recursively. Suppose you have listed the following search and replacement strings:

DOG to be replaced with CAT
A to be replaced with O

All occurrences of the string DOG are replaced with CAT, regardless of whether it is the word DOG or the word DOGMA. This one replacement occurs. The rewritten CAT is not replaced with COT.

- ◆ Because string replacement is done in one pass, the string that matches first takes precedence. Suppose you have listed the following search and replacement strings:

ABC to be replaced with XYZ
BCDEF to be replaced with PQRSTUVWXYZ

If the original string is ABCDEFGH, the replaced string is XYZDEFGH.

- ◆ If two specified search strings match the data portion, the search string of longer length is used for the replacement except for the case detailed above. Suppose you have listed the following search and replacement strings:

ABC	to be replaced with	XYZ
ABCDEF	to be replaced with	PQRSTUVWXYZ

If the original string is ABCDEFGH, the replaced string is PQRSTUVWXYZGH.

String Replacement Rules for Word Profiles

If you are configuring a Word profile, you can use the following special tokens:

- ◆ [w] to indicate white space
- ◆ [ow] to indicate optional white space

You use the [w] and the [ow] to specify where white space might occur in the string. For example:
[ow]my[w]string[w]to[w]replace[ow]

If you don't know, or don't care, whether the string has zero or more white characters at the beginning and at the end, use [ow] to specify this. The [w] specifies exactly one white character.

12.5.2 Configuring the HTML Rewriter

You configure the HTML rewriter for a proxy service, and these values are applied to all Web servers that are protected by this proxy service.

To configure the HTML rewriter:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > HTML Rewriting*.

Proxy Service **Web Servers** **HTML Rewriting** **Protected Resources** **Logging**

Enable HTML Rewriting

Additional DNS Name List
New... | Delete 1 item(s)
 DNS Name
 [rewriteMe](#)

Exclude DNS Name List
New... | Delete 0 item(s)
 DNS Name
No items

HTML Rewriter Profile List
New... | Delete | Enable | Disable ⓘ 1 item(s)
 Name **Enabled** **Search Boundary**
 [default](#) Word

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

OK **Cancel**

The HTML Rewriting page specifies what is to be rewritten. The HTML Rewriter Profile specifies which pages to search for DNS names that need to be rewritten.

- 2 Select *Enable HTML Rewriting*.

This option is enabled by default. When it is disabled, no rewriting occurs. When enabled, this option activates the internal HTML rewriter. This rewriter replaces the name of the Web server with the published DNS name when sending data to the browsers. It replaces the published DNS name with the Web Server Host Name when sending data to the Web server. It also makes sure the proper scheme (HTTP or HTTPS) is included in the URL. This is needed because you can configure the Access Gateway to use HTTPS between itself and client browsers and to use HTTP between itself and the Web servers.

- 3 In the *Additional DNS Name List* section, click *New*, specify a DNS that appears on the Web pages of your server (for example a DNS name other than the Web server's DNS name), then click *OK*.

For more information, see “[Determining Whether You Need to Specify Additional DNS Names](#)” on page 158.

- 4 In the *Exclude DNS Name List* section, click *New*, specify a DNS name that appears on the Web pages of your server that you do not want rewritten, then click *OK*.

For more information, see “[Determining Whether You Need to Exclude DNS Names from Being Rewritten](#)” on page 159.

5 Use the *HTML Rewriter Profile List* to configure a profile. Select one of the following actions:

- ♦ **New:** To create a profile, click *New*. Specify a display name for the profile and select either a *Word* or *Character* for the *Search Boundary*. Continue with **Step 6**.
 - ♦ **Word:** A Word profile searches for matches on words. For example, “to” matches only the word “to” and not the “to” in “total”.
 The Access Gateway comes with a default Word profile, which is designed to be applied to all pages protected by the Access Gateway. It is not specific to a reverse proxy or its proxy services. When you modify its behavior, remember its scope.
 - ♦ **Character:** A Character profile searches for matches on a specified set of characters. For example, “to” matches both the word “to” and the “to” in “total”. If you want to add functionality to the default profile, create a Character profile. It has all the functionality of a Word profile, except searching for attribute names and Java variables and methods.
- ♦ **Delete:** To delete a profile, select the profile, then click *Delete*. Continue with **Step 13**.
- ♦ **Enable:** To enable a profile, select the profile, then click *Enable*. Continue with **Step 13**.
- ♦ **Disable:** To disable a profile, select the profile, then click *Disable*. Continue with **Step 13**.
- ♦ **Modify:** To view or modify the current configuration for a profile, click the name of the profile. Continue with **Step 6**.

The default profile is designed to be applied to all pages protected by the Access Gateway. It is not specific to a reverse proxy or its proxy services. If you modify its behavior, remember its scope. Rather than modify the default profile, you should create your own customized Word profile and enable it

6 Use the *Requested URLs to Search* section to set up a policy for finding the pages that need rewriting.

Requested URLs to Search	
If Requested URL Is	
New... Delete	0 item(s)
<input type="checkbox"/> Include URL	
Att	
And Requested URL Is Not	
New... Delete	0 item(s)
<input type="checkbox"/> Exclude URL	
No items	
And Document Content-Type Header Is	
New... Delete Restore Defaults	6 item(s)
<input type="checkbox"/> Content-Type Header	
<input type="checkbox"/> text/html [default]	
<input type="checkbox"/> text/xml [default]	
<input type="checkbox"/> text/css [default]	
<input type="checkbox"/> text/javascript [default]	
<input type="checkbox"/> application/javascript [default]	
<input type="checkbox"/> application/x-javascript [default]	

The URLs in the search policy should use the following formats:

Sample URL	Description
<code>http://www.a.com/content</code>	Matches only if the URL does not contain a trailing slash.
<code>http://www.a.com/content/</code>	Matches only if the URL does contain a trailing slash.
<code>http://www.a.com/content/index.html</code>	Matches only this specific file.
<code>http://www.a.com/content/*</code>	Matches the URL whether or not it has a trailing slash and matches all files in the directory.
<code>http://www.a.com/*</code>	Matches the proxy service and everything it is protecting.

Fill in the following fields:

If Requested URL Is: Specify the URLs of the pages you want to rewrite data. You can use the asterisk wildcard to include a few pages you really don't want to rewrite data, and use the *And Requested URL Is Not* section to exclude them.

And Requested URL Is Not: Specify the URLs of pages where you do not want any data rewritten. If a page matches both the *If Requested URL Is* and *And Requested URL Is Not*, the page is excluded from rewriting.

And Document Content-Type Is: Select the data types you want searched for URLs to rewrite.

- 7 Use the *Enable Rewriter Actions* option to specify the action the rewriter should perform if the page matches the criteria in the Requested URLs to Search section:
 - ♦ Select it to have the rewriter use the profile to rewrite references and data on the page.
 - ♦ Leave it unselected to disable rewriting on the page. This allows you to create a profile for the pages you do not want rewritten.
- 8 (Not available for Character profiles) If your pages contain JavaScript, use the *Additional Names to Search for URL Strings to Rewrite with Host Name* section to specify JavaScript variables or methods. You can also add HTML attribute names. (For the list of attribute names that are automatically searched, see [“HTML Tags” on page 158.](#))

Additional Names to Search for URL Strings to Rewrite with Host Name

Then Variable or Attribute Name to Search for Is
i

New... | Delete
0 item(s)

Variable or Attribute Name

No items

And JavaScript Method to Search for Is
i

New... | Delete
0 item(s)

JavaScript Method

No items

Fill in the following fields:

Then Variable or Attribute Name to Search for Is: Lists the name of an HTML attribute or JavaScript variable to search to see if its value contains a URL string. Click *New* to add a name to the list.

And JavaScript Method to Search for Is: Lists the name of Java methods to search to see if its parameters contain a URL string. Click *New* to add a method to the list.

- 9 Use the *Additional Strings to Replace* section to specify a string to search for and specify the text it should be replaced with. The search boundary (word or character) that you specified when creating the profile is used when searching for the string.

Search	Replace With
<input type="checkbox"/> No items	

To add a string, click *New*, then fill in the following:

Search: Specify the string you want to search for. The profile type controls the matching and replacement rules. For more information, see one of the following:

- ♦ [“String Replacement Rules for Character Profiles” on page 161](#)
- ♦ [“String Replacement Rules for Word Profiles” on page 162](#)

Replace With: Specify the string you want to use in place of the search string.

- 10 Click *OK*.
- 11 If you have more than one profile in the *HTML Rewriter Profile List*, use the arrows to order the profiles.

If you create more than one profile, order becomes important. For example if you want to rewrite all pages with a general rewriter profile (with a URL such as */**) and one specific set of pages with another rewriter profile (with a URL such as */doc/100506/**), you need to have the specific rewriter profile listed before the general rewriter profile. Only one Word profile and one Character profile is executed per page. The first one in the list that matches a page is executed, and the others are ignored.

- 12 Enable the profiles you want to use for this protected resource. Select the profile, then click *Enable*.

The default profile cannot be disabled.

- 13 To save your changes and push them to the Access Gateway, click *Configuration Panel*, then click *Apply Changes*.

12.6 Configuring Connection Limits

The Access Gateway establishes connections with clients and with Web servers. For most networks, the default values for the connection limits provide adequate performance, but these options allow you to fine tune the performance for your network and its requirements.

- ♦ [Section 12.6.1, “Configuring TCP Listen Options for Clients,” on page 167](#)
- ♦ [Section 12.6.2, “Configuring TCP Connect Options for Web Servers,” on page 168](#)
- ♦ [Section 12.6.3, “Configuring Connection and Session Persistence,” on page 170](#)

12.6.1 Configuring TCP Listen Options for Clients

The TCP listen options allow you to control how idle and unresponsive browser connections are handled and to optimize these processes for your network. For most networks, the default values provide adequate performance. But if your network is congested and slow, you might want to increase some of the limits.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy] > TCP Listen Options*.

Enable Persistent Connections

TCP Listen Options

Connection Handshake Timeout: Second(s) (1-120)

Keep Alive Interval: Second(s) (0-1440)

Data Read Timeout: Second(s) (1-3600)

Idle Timeout: Second(s) (1-1800)

Retransmit Limit: (1-50)

Enable Nagle's Algorithm (For Coalescing Packets)

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

- 2 Select *Enable Persistent Connections* to allow the Access Gateway to establish a persistent HTTP connection between the Access Gateway and the browser. Usually, HTTP connections service only one request and response sequence. A persistent connection allows multiple requests to be serviced before the connection is closed.

This option is enabled by default.

- 3 Specify values for the following fields:

Connection Handshake Timeout: Sets a timeout limit for a connecting device that stops responding after having initiated the TCP handshake process. If an expected handshake response is not received from the connecting device in this amount of time, an error occurs. Setting the value lower might help defend against SYN attacks. The timeout can be set from 1 to 120 seconds. The default is 30 seconds.

Keep Alive Interval: Sets the length of time between packets being sent to a connected device to determine if the connection is still alive. If a response is not received within the Data Read Timeout value, the connection is closed. On an idle connection, sending these ping packets continues until the Idle Timeout value is reached. Setting the value to zero prevents the sending of keep-alive packets. The value can be set from 0 to 1440 seconds (24 minutes). The default is 300 seconds (5 minutes).

Data Read Timeout: Determines when an unresponsive connection is closed. When exchanging data, if an expected response from the connected device is not received within this amount of time, the connection is closed. This value might need to be increased for slow or

congested network links. The value can be set from 1 to 3600 seconds (1 hour). The default is 120 seconds (2 minutes).

Idle Timeout: Determines when an idle connection is closed. If no application data is exchanged over a connection for this amount of time, the connection is closed. This value limits how long an idle persistent connection is kept open. This setting is a compromise between freeing resources to allow additional inbound connections, and keeping connections established so that new connections from the same device do not need to be re-established. The value can be set from 1 to 1800 seconds (30 minutes). The default is 180 seconds (3 minutes).

Retransmit Limit: Determines how many times data is resent. When exchanging data, if the expected acknowledgement (ACK) response is not received, this is the number of times the device attempts to resend the data before closing the connection. You can set the value from 1 - 50. The default is 8.

Enable Nagle's Algorithm: Determines where small buffer messages can be concatenated into one large message. When this option is enabled, small buffer messages are automatically concatenated. This process increases the efficiency of a network application system by decreasing the number of packets that must be sent. Enabling this feature delays data transmission until a full TCP packet can be sent.

- 4 On a Linux Access Gateway, you can also configure the encryption key. (For the NetWare[®] Access Gateway, the encryption key is set globally for all reverse proxies. See [Section 13.6, "Configuring the Encryption Key,"](#) on page 179.) Select one or more of the following:

Enable 128-Bit Encryption between Browser and Access Gateway: When this option is selected, the Access Gateway requires all its server connections with client browsers to use 128-bit encryption. If the encryption key is less than 128, regardless of the cipher suite, the connection is denied.

Enable 128-Bit Encryption between Access Gateway and Web Server: When this option is selected, the Access Gateway requires all its client connections to Web servers to use 128-bit encryption. If the encryption key is less than 128, regardless of the cipher suite, the connection is denied.

- 5 To save your changes and push them to the Access Gateway, click *Configuration Panel*, then click *Apply Changes*.

12.6.2 Configuring TCP Connect Options for Web Servers

Connect options are specific to the group of Web servers configured for a proxy service. They allow you to control how idle and unresponsive Web server connections are handled and to optimize these

processes for your network. For most networks, the default values provide adequate performance. But if your network is congested and slow, you might want to increase some of the limits.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > Web Servers > TCP Connect Options*.

Enable Persistent Connections

TCP Connect Options

Connection Handshake Timeout: Second(s) (1-120)

Keep Alive Interval: Second(s) (0-1440)

Data Read Timeout: Second(s) (1-3600)

Idle Timeout: Second(s) (1-1800)

Retransmit Limit: (1-50)

Enable Nagle's Algorithm

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

- 2 Select *Enable Persistent Connections* to allow the Access Gateway to establish a persistent HTTP connection between the Access Gateway and the Web server. Usually, HTTP connections service only one request and response sequence. A persistent connection allows multiple requests to be serviced before the connection is closed.

This option is enabled by default.

- 3 To modify the connection timeouts between the Access Gateway and the Web servers, configure the following fields:

Connection Handshake Timeout: Sets a timeout limit for a connecting device that stops responding after having initiated the TCP handshake process. If an expected handshake response is not received from the connecting device in this amount of time, an error occurs. Setting the value lower might help defend against SYN attacks. The timeout can be set from 1 to 120 seconds. The default is 30 seconds.

Keep Alive Interval: Sets the length of time between packets being sent to a connected device to determine if the connection is still alive. If a response is not received within the Data Read Timeout value, the connection is closed. On an idle connection, sending these ping packets continues until the Idle Timeout value is reached. Setting the value to zero prevents the sending of keep-alive packets. The value can be set from 0 to 1440 seconds (24 minutes). The default is 300 seconds (5 minutes).

Data Read Timeout: Determines when an unresponsive connection is closed. When exchanging data, if an expected response from the connected device is not received within this amount of time, the connection is closed. This value might need to be increased for slow or congested network links. The value can be set from 1 to 3600 seconds (1 hour). The default is 120 seconds (2 minutes).

Idle Timeout: Determines when an idle connection is closed. If no application data is exchanged over a connection for this amount of time, the connection is closed. This value limits how long an idle persistent connection is kept open. This setting is a compromise between freeing resources to allow additional inbound connections, and keeping connections established so that new connections from the same device do not need to be re-established. The value can be set from 1 to 1800 seconds (30 minutes). The default is 180 seconds (3 minutes).

Retransmit Limit: Determines how many times data is resent. When exchanging data, if the expected acknowledgement (ACK) response is not received, this is the number of times the device attempts to resend the data before closing the connection. You can set the value from 1 - 50. The default is 8.

Enable Nagle's Algorithm: Determines where small buffer messages can be concatenated into one large message. When this option is enabled, small buffer messages are automatically concatenated. This process increases the efficiency of a network application system by decreasing the number of packets that must be sent. Enabling this feature delays data transmission until a full TCP packet can be sent.

- 4 To save your changes and push them to the Access Gateway, click *Configuration Panel*, then click *Apply Changes*.

12.6.3 Configuring Connection and Session Persistence

The Access Gateway establishes three types of connections:

- ♦ Access Gateway to browser
- ♦ Access Gateway to Web server
- ♦ Browser to Web server

The Access Gateway to the browser connections and the Access Gateway to the Web server connections involve setting up a TCP connection for an HTTP request. HTTP connections usually service only one request and response sequence, and the TCP connection is opened and closed during the sequence. A persistent connection allows multiple requests to be serviced before the connection is closed and saves a significant amount of processing time. To configure this type of persistence, see the following:

- ♦ **Access Gateway to Browser:** Click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy] > TCP Listen Options* and configure the *Enable Persistent Connections* option.
- ♦ **Access Gateway to Web Server:** Click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > Web Servers > TCP Connect Options* and configure the *Enable Persistent Connections* option.

The persistence of the browser to Web server connection is always enabled and is not configurable. This feature allows a browser to use the same Web server after an initial connection has been established. Most Web applications are designed to expect this type of behavior.

Configuring the Access Gateway for SSL

13

SSL provides the following security features:

- ◆ 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.

To ensure the validity of X.509 certificates, Access Manager supports both Certificate Revocation Lists (CRLs) and Online Certificate Status Protocol (OCSP) methods of verification.

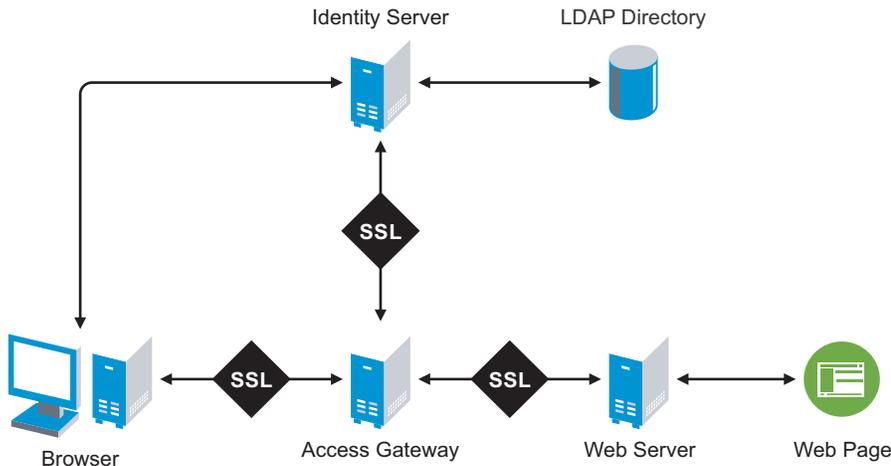
This section describes how the Access Gateway can use SSL in its interactions with other Access Manager components and how you can enable SSL on an Access Gateway to these components:

- ◆ [Section 13.1, “Using SSL on the Access Gateway Communication Channels,” on page 172](#)
- ◆ [Section 13.2, “Prerequisites for SSL,” on page 173](#)
- ◆ [Section 13.3, “Configuring SSL Communication with the Browsers and the Identity Server,” on page 174](#)
- ◆ [Section 13.4, “Configuring SSL between the Proxy Service and the Web Servers,” on page 176](#)
- ◆ [Section 13.5, “Managing the Certificates of the Embedded Service Provider,” on page 178](#)
- ◆ [Section 13.6, “Configuring the Encryption Key,” on page 179](#)

13.1 Using SSL on the Access Gateway Communication Channels

You can configure the Access Gateway to use SSL in its connections to the Identity Server, to the browsers, and to its Web servers. [Figure 13-1](#) illustrates these communication channels.

Figure 13-1 Setting up SSL for the Access Gateway communication channels



This section only describes how to set up SSL for the Access Gateway communication channels. The Identity Server needs to be configured for SSL before the Access Gateway can be configured for SSL. See “[Configuring Secure Communication on the Identity Server](#)” in the *Novell Access Manager 3.0 Setup Guide*.

When the user logs in to the Identity Server, the Identity Server verifies the user’s credentials, usually with the credentials stored in an LDAP directory but other methods are available. If the login is successful, the Identity Server sends an artifact to the browser, and the browser forwards it to the Access Gateway. The Access Gateway uses the artifact to retrieve the user’s name and password from the Identity Server. The Access Gateway and Identity Server channel is probably the first communication channel you should enable for SSL. The Access Gateway uses an embedded service provider to communicate with the Identity Server. When you enable SSL between the two, the Access Manager distributes the necessary certificates to set up SSL. However, if you have configured the Identity Server to use certificates from an external CA, you need to import the public certificate of this CA into the trust store of the Access Gateway. If you have set up the Access Gateway to use a certificate from an external CA, you need to import the public certificate of this CA into the trust store of the Identity Server.

SSL must be enabled between the Access Gateway and the browsers before you can enable SSL between the Access Gateway and its Web servers. If you enable SSL between the Access Gateway and the browsers, SSL is automatically enabled for the Access Gateway embedded service provider that communicates with the Identity Server. After you have enabled SSL between the Access Gateway and the browsers, you can select whether to enable SSL between the Access Gateway and the Web servers. By not enabling SSL to the Web servers, you can save processing overhead if the data on the Web servers is not sensitive or if it is already sufficiently protected.

Whether you need the added security of SSL or mutual SSL between the Access Gateway and its Web servers depends upon how you have set up your Web servers. If you have configured the Web servers so that they can only accept connections with the Access Gateway, mutual SSL is probably

not needed. If the Access Gateway is injecting authentication credentials into HTTP headers, you should enable SSL.

13.2 Prerequisites for SSL

The following SSL configuration instructions assume that you have already created or imported the certificate that you are going to use for SSL. This certificate must have a subject name (cn) that matches the published DNS name of the proxy service that you are going to use for authentication. You can obtain this certificate one of two ways:

- You can use the Access Manager CA to create this certificate. See [Section 26.1.1, “Creating a Locally Signed Certificate,” on page 290.](#)
- You can create a certificate signing request (CSR), send it to an external CA, then import the returned certificates into Access Manager. See [Section 26.1.2, “Generating a Certificate Signing Request,” on page 294](#) and [Section 26.5, “Importing Public Key Certificates \(Trusted Roots\),” on page 297.](#)

13.2.1 Prerequisites for SSL Communication between the Access Gateway and the Web Servers

If you are going to set up SSL between the Access Gateway and the Web servers, you need to configure your Web servers for SSL. Your Web servers must supply a certificate that clients (in this case, the Access Gateway) can import. See your Web server documentation for information on how to configure the Web server for SSL.

For mutual SSL, the proxy service must supply a certificate that the Web server can trust. This certificate can be the same one you use for SSL between the browsers and the reverse proxy.

13.2.2 Prerequisite for SSL Communication between the Identity Server and the Access Gateway

If you are going to set up SSL communication between the Identity Server and the Access Gateway for authentication and you have configured the Identity Server to use certificates created by an external CA, you need to import the public certificate of this CA into the trusted root keystore of the Access Gateway.

- 1 If you haven't already imported the public certificate of this CA into the trusted root store of the Identity Server, do so now. For instructions, see [Section 26.5, “Importing Public Key Certificates \(Trusted Roots\),” on page 297.](#)
- 2 In the Administration Console, click *Access Manager > Access Gateways > Edit > Service Provider Certificates > Trusted Roots.*
- 3 In the *Trusted Roots* section, click *Add.*
- 4 Click the *Select trusted root(s)* icon, select the public certificate of the CA that signed the Identity Server certificates, then click *OK.*
- 5 Specify an alias, then click *OK* twice.
- 6 To save the changes, click *Apply Changes.*

If you schedule the changes, these changes must take effect before enabling SSL communication between the Identity Server and the Access Gateway.

13.3 Configuring SSL Communication with the Browsers and the Identity Server

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy]*.

Listening Address(es): 10.10.167.50
 10.10.167.51

[TCP Listen Options](#)

Enable SSL with Embedded Service Provider

Enable SSL between Browser and Access Gateway

Redirect Requests from Non-Secure Port to Secure Port

Server Certificate: 

[Auto-generate Key](#)

[Auto-Import Embedded Service Provider Trusted Root](#)

Non-Secure Port: * (Redirected to Secure Port)

Secure Port: * (Used for Trusted IDS Encryption, HTTPS Listening)

- 2 Configure the reverse proxy for SSL. Fill in the following fields:

Enable SSL with Embedded Service Provider: Select this option to encrypt the data exchanged for authentication (the communication channel between the Identity Server and the Access Gateway). This option is only available for the reverse proxy which has been assigned to perform authentication.

If you enable SSL between the browsers and the Access Gateway, this option is automatically selected for you. You can enable SSL with the Embedded Service Provider without enabling SSL between the Access Gateway and the browsers. This allows the authentication and identity information that the Access Gateway and the Identity Server exchange to use a secure channel, but allows the data that the Access Gateways retrieves from the back-end Web servers and sends to users to use a non-secure channel. This saves processing overhead if the data on the Web servers is not sensitive.

Enable SSL between Browser and Access Gateway: Select to require SSL connections between your clients and the Access Gateway. SSL must be configured between the browsers and the Access Gateway before you can configure SSL between the Access Gateway and the Web servers.

Redirect Requests from Non-Secure Port to Secure Port: Determines whether browsers are redirected to the Secure Port and allowed to establish an SSL connection. If this option is not selected, browsers that connect to the non-secure port are denied service.

This option is only available if you have selected *Enable SSL with Embedded Service Provider*.

3 Select the certificate to use for SSL between the Access Gateway and the browsers. Select one of the following methods:

- ◆ To auto-generate a certificate key using the Access Manager CA, click *Auto-generate Key*, then click *OK* twice. The generated certificate appears in the *Server Certificate* text box.

The generated certificate uses the published DNS name of the first proxy service for the Subject name of the certificate. If there is more than one proxy service, the CA generates a wildcard certificate (*.Cookie Domain).

If you have not created a proxy service for this reverse proxy, wait until you have created a proxy service before generating the key. This allows the CN in the Subject field of the certificate to match the published DNS name of the proxy service.

- ◆ To select a certificate, click the *Select Certificate* icon, select the certificate you have created for the DNS name of your proxy service, then click *OK*. The certificate appears in the *Server Certificate* text box. For SSL to work, the CN in the Subject field of the certificate must match the published DNS name of the proxy service.

4 (Conditional) If you have selected a certificate in **Step 3** that was created by an external CA, click *Auto-Import Embedded Service Provider Trusted Root*, click *OK*, specify an alias name, click *OK*, then click *Close*.

This option imports the public key from the embedded service provider into the trust store of the Identity Servers in the selected Identity Server Configuration. This sets up a trusted SSL relationship between the Identity Server and the embedded service provider.

5 Configure the ports for SSL:

Non-Secure Port: Specifies the port on which to listen for HTTP requests. The default port for HTTP is 80. .

- ◆ If you have selected the *Redirect Requests from Non-Secure Port to Secure Port* option, requests sent to this port are redirected to the secure port. If the browser can establish an SSL connection, the session continues on the secure port. If the browser cannot establish an SSL connection, the session is terminated.
- ◆ If you do not select the *Redirect Requests from Non-Secure Port to Secure Port* option, this port is not used when SSL is enabled.

Secure Port: Specifies the port on which to listen for HTTPS requests (which is usually 443). This port needs to match the configuration for SSL. If SSL is enabled, this port is used for all communication with the browsers. The listening address and port combination must not match any combination you have configured for another reverse proxy or tunnel.

6 Click *OK*.

When you enable the SSL options, you need to save the configuration, then restart the embedded service provider and update the Identity Server. See the last two steps in this procedure.

7 On the *Configuration* page, click *Reverse Proxy / Authentication*.

8 In the *Embedded Service Provider* section, click *Auto-Import Identity Server Trusted Root*, click *OK*, specify an alias, click *OK* twice, then click *Close*.

This option imports the public key of the Identity Server into the trust store of the embedded service provider. This sets up a trusted SSL relationship between the embedded service provider and the Identity Server.

9 Click *Configuration Panel*, then *OK*.

10 On the *Configuration* page, click *Apply Changes*.

- 11 When the changes have been applied, restart the embedded service provider. Click *Access Gateways > [Name of Access Gateway] > Actions > Restart Service Provider >OK*.
If the Access Gateway is part of a group, click *Access Gateways > Groups > [Name of Group] > Actions > Restart Service Provider >OK*.
- 12 Update the Identity Server so that it uses the new SSL configuration. Click *Identity Servers > Setup > Update Servers*.

13.4 Configuring SSL between the Proxy Service and the Web Servers

SSL must be enabled between the Access Gateway and the browsers before you can enable it between the Access Gateway and its Web servers.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > Web Servers*.

- 2 To configure SSL, select *Connect Using SSL*.
This option is not available if you have not set up SSL between the browsers and the Access Gateway. See [Section 13.3, “Configuring SSL Communication with the Browsers and the Identity Server,” on page 174](#) and select the *Enable SSL between Browser and Access Gateway* field.
- 3 In the *Connect Port* field, specify the port that your Web server uses for SSL communication. The following table lists some common servers and their default ports.

Server Type	Non-Secure Port	Secure Port
Web server with HTML content	80	443
SSL VPN	8080	8443

Server Type	Non-Secure Port	Secure Port
WebSphere	9080	9443
JBoss	8080	8443

4 Configure how you want the certificate verified. The Access Gateway platforms support different options:

4a (Conditional) If you are configuring a Linux Access Gateway, select one of the following options:

- ♦ To not verify this certificate, select *Do not verify* for the *Web Server Trusted Root*. Continue with [Step 9](#).
- ♦ To allow the certificate to match any certificate in the trust store, select *Any in Reverse Proxy Trust Store* for the *Web Server Trusted Root*. Continue with [Step 9](#).
- ♦ To add a certificate to the trust store for the Web server, click the *Manage Reverse Proxy Trust Store* icon. Continue with [Step 4c](#).

4b (Conditional) If you are configuring a NetWare[®] Access Gateway, all the certificates in the certificate chain of the Web server must be in its trust store. To add these certificates to the trust store, click *Any in Reverse Proxy Trust Store*. Continue with [Step 4c](#).

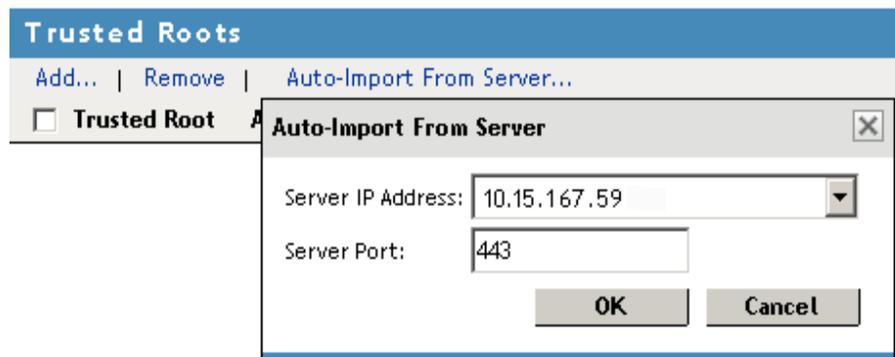
4c The auto import screen appears.

Trust Store: Proxy Trust Store

Trust store name: Proxy Trust Store

Trust store type: DER

Device: 10.10.159.206



5 Ensure that the IP address of the Web server and the port match your Web server configuration. If these values are wrong, you have entered them incorrectly on the Web server page. If this is true, click *Cancel* and reconfigure them before continuing.

6 Click *OK*.

The server certificate, the root CA certificate, and any CA certificates from a chain are listed.

If the whole chain is not displayed, import what is displayed. You then need to manually import the missing parents in the chain. A parent is missing if the chain does not include a certificate where the Subject and the Issuer have the same CN.

- 7 Specify an alias, then click *OK*.
All the certificates displayed are added to the trust store.
- 8 Click *Close*.
- 9 (Optional) For mutual authentication, the Access Gateway platforms support different options:
 - 9a (Conditional) If you are configuring a Linux Access Gateway, you need to select the certificate. Click the *Select Certificate* icon, select the certificate you created for the reverse proxy, then click *OK*.

This is only part of the process. You need to import the trusted root certificate of the CA that signed the proxy service's certificate to the Web servers assigned to this proxy service.
 - 9b (Conditional) If you are configuring a NetWare Access Gateway, the text box displays the certificate that is sent to the Web server if the Web server requires it. If the Web server is not set up for mutual SSL, the certificate is not sent.

To set up the Web server for mutual SSL, you need to import the trusted root certificate of the CA that signed the certificate displayed in the text box.
- 10 Click *Configuration Panel*, then *OK*.
- 11 On the *Configuration* page, click *Apply Changes*.

13.5 Managing the Certificates of the Embedded Service Provider

The Access Gateway uses an embedded service provider module to communicate with the Identity Server. The Service Provider Certificates page allows you to view the private keys, CA certificates, and certificate containers associated with this module. These keystores do not contain the certificates that the Access Gateway uses for SSL connections to browsers or to back-end Web servers.

To view or modify these certificates:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Service Provider Certificates*.
- 2 Configure the following:
 - Signing:** The signing certificate keystore. Click this link to access the keystore and replace the signing certificate as necessary. The signing certificate is used to sign the assertion or specific parts of the assertion.
 - Trusted Roots:** The trusted root certificate container for the CA certificates associated with the Access Gateway. Click this link to access the trust store, where you can change the password or add trusted roots to the container.

The embedded service provider must trust the certificate of the Identity Server that the Access Gateway has been configured to trust. The public certificate of the CA that generated the Identity Server certificate must be in this trust store. If you configured the Identity Server to use a certificate generated by a CA other than the Access Manager CA, you must add the public certificate of this CA to the Trusted Roots store.
- 3 To save your modifications, click *OK*, then on the Configuration page, click *Apply Changes*.

13.6 Configuring the Encryption Key

You can specify the size of the encryption key that the Access Gateway requires when it establishes connections. For the Linux Access Gateway, these options are set per reverse proxy (see [Section 12.6.1, “Configuring TCP Listen Options for Clients,” on page 167](#)). For the NetWare Access Gateway, these options are set globally for all reverse proxies.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Security Options*.

- 2 Configure the following fields:

Enable 128-Bit Encryption between Browser and Access Gateway: When this option is selected, the Access Gateway requires all its server connections with client browsers to use 128-bit encryption. If the encryption key is less than 128, regardless of the cipher suite, the connection is denied.

Enable 128-Bit Encryption between Access Gateway and Web Server: When this option is selected, the Access Gateway requires all its client connections to Web servers to use 128-bit encryption. If the encryption key is less than 128, regardless of the cipher suite, the connection is denied.

- 3 To save your modifications, click *OK*, then on the Configuration page, click *Apply Changes*.

Additional Configuration Settings

14

This section describes the configuration settings that affect the Access Gateway as a server, such as changing its name or setting the time.

- ◆ Section 14.1, “Changing the Name of an Access Gateway and Modifying Other Descriptive Details,” on page 181
- ◆ Section 14.2, “Modifying the Base URL of the Identity Server,” on page 182
- ◆ Section 14.3, “Setting Date and Time,” on page 182
- ◆ Section 14.4, “Setting Up a Tunnel,” on page 184
- ◆ Section 14.5, “Configuring Error Page Presentation,” on page 185
- ◆ Section 14.6, “Configuring Console Access,” on page 186
- ◆ Section 14.7, “Configuring Network Settings,” on page 188
- ◆ Section 14.8, “Customizing Log Out,” on page 198
- ◆ Section 14.9, “Configuring X-Forwarded-For Headers,” on page 198
- ◆ Section 14.10, “Upgrading the Access Gateway Software,” on page 199

14.1 Changing the Name of an Access Gateway and Modifying Other Descriptive Details

The default name of an Access Gateway is its IP address. You can change this to a more descriptive name as well as adding other details that can help you identify one Access Gateway from another.

- 1 In the Administration Console, click *Access Manager > Access Gateways > [Name of Access Gateway] > Edit*.

Name:	<input type="text" value="NetWareAG"/>		
Management IP Address:	<input type="text" value="10.15.167.59"/> <input type="button" value="v"/>	Port:	<input type="text" value="1443"/>
Location:	<input type="text"/>		
Description:	<input type="text"/>		

- 2 Modify the values in the following fields:

Name: Specifies the Administration Console display name for the Access Gateway. The default name is the IP address of the Access Gateway. This is required.

The name must use alphanumeric characters and can include spaces, hyphens, and underscores.

Location: Specifies the location of the Access Gateway server. This is optional, but useful if your network has multiple Access Gateway servers.

Description: Describes the purpose of this Access Gateway. This is optional, but useful if your network has multiple Access Gateways.

For information on changing the Management IP Address, see [Section 3.3, “Changing the IP Address of the Access Gateway,”](#) on page 42.

- 3 Click *OK* twice, then click *Close*.

14.2 Modifying the Base URL of the Identity Server

When you change the base URL of the Identity Server, you destroy two trusted relationships:

- ♦ The trusted relationship that the Identity Server has established with each device that has been configured to use the Identity Server for authentication
- ♦ The trusted relationship that each device has established with the Identity Server when the Identity Server configuration was selected.

Your Web site is down until you re-establish these trust relationships. The sessions of any logged in users are destroyed and no user can log in and access resources until the trust relationships are re-established.

For information on reconfiguring the Identity Server, see [Section 5.4, “Modifying the Base URL,”](#) on page 59. After you have completed these steps, you need to re-establish the trusted relationship of the Access Gateway with the Identity Server.

For each Access Gateway or Access Gateway group:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Reverse Proxy / Authentication*.
- 2 Set the *Identity Server Configuration* to *[None]*.
This clears the current trusted relationship when the change is applied.
- 3 Click *OK*, then *Apply Changes*.
For a group of Access Gateways, select to apply the change to all servers in the group.
- 4 Click *Access Gateways > Edit > Reverse Proxy / Authentication*.
- 5 From the list, select the *Identity Server Configuration*.
- 6 Click *OK*, then *Apply Changes*.
For a group of Access Gateways, select to apply the change to all servers in the group.
- 7 Update the Identity Server. Click *Identity Servers > Setup > Update Servers*.

14.3 Setting Date and Time

The *Date & Time* option lets you set the system time for the Access Gateway. The time between the Identity Server and the Access Gateway must be either synchronized or set to be within 1 minute of each other for trusted authentication to work.

To configure the date and time options:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Date & Time*

Group Member:

Server Date and Time

June 13, 2006 11:20:25 [Set Date & Time Manually](#)

Network Time Protocol

Use Network Time Protocol [Set Up NTP](#)

Time Zone

Name:

- (GMT-12:00) Eniwetok, Kwajalein
- (GMT-11:00) Midway Island, Samoa
- (GMT-10:00) Hawaii
- (GMT-09:00) Alaska
- (GMT-08:00) Pacific Time (US & Canada)
- (GMT-07:00) Mountain Time (US & Canada)
- (GMT-07:00) Arizona
- (GMT-06:00) Central Time (US & Canada)
- (GMT-06:00) Saskatchewan
- (GMT-06:00) Mexico City, Tegucigalpa

- 2 (Conditional) If the Access Gateway belongs to a group of Access Gateways, select the Access Gateway from the list displayed in the *Group Member* field. If the Access Gateway does not belong to a group, this option is not available.
- 3 Fill in the following fields:

Server Date and Time: Displays the current time and allows you to set the current time. Click *Set Date & Time Manually*, then select the current year, month, day, hour, and minute.

WARNING: If the date is set to a time before the Access Gateway certificates are valid, communication to the Access Gateway is lost. This error cannot be corrected from the Administration Console. You need to correct it at the console of the Access Gateway machine.

- ♦ For a NetWare® Access Gateway server, see [Section 42.3.2, “Setting the Date and Time at the Console,” on page 508.](#)
- ♦ For the Linux Access Gateway, use `yast` and select *System > Date and Time*.

Network Time Protocol: Allows you to configure the Access Gateway to use a Network Time Protocol Server. Configure the following options:

- ♦ **Use Network Time Protocol:** If this option is selected, the Access Gateway uses the NTP server to keep its time accurate.
- ♦ **Set Up NTP:** Click this option to specify the DNS name or IP address of an NTP server. The installation program enters the name of `pool.ntp.org`, the DNS name of a public NTP server.

Time Zone: Select your time zone, then click *OK*. Regardless of the method you used to set the time, you must select a time zone.

- 4 (NetWare only) Configure daylight saving time. Configure the following options:

Use Daylight Saving: Select this option to enable daylight saving.

Offset: Select the hours and minutes that daylight saving varies from standard time.

Start: Select the month, day, hour, and day of month when daylight savings starts.

End: Select the month, day, hour, and day of month when daylight saving ends.

- 5 To save your changes and push them to the Access Gateway, click *Configuration Panel*, then click *Apply Changes*.

14.4 Setting Up a Tunnel

The tunnel option lets you create one or more services for the specific purpose of tunneling non-HTTP traffic through the Access Gateway to the Web server. To do this, the non-HTTP traffic must use a different IP address and port combination than the HTTP traffic.

An Access Gateway usually processes HTTP requests in order to fill them. However, it is not unusual that some of the traffic coming through the gateway is not HTTP-based. Web servers sometimes handle Telnet, FTP, chat, or other kinds of traffic without attempting to process it. If your Web servers are handling this type of traffic, you should set up a tunnel for it.

Reverse proxies and tunnels cannot share the same IP address and port combination. You can either configure a reverse proxy for an IP address and port or a tunnel for that IP address and port.

To set up a tunnel:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Tunneling*.
- 2 Click *New*, enter a display name for the tunnel, then click *OK*.

Enable Tunnel

Tunnel SSL Traffic Only

Group Member: 

Listening Address(es): 10.10.159.206

[TCP Listen Options](#)

Listening Port: *

Connect Port: *

[TCP Connect Options](#)

Web Server List

[New...](#) | [Delete](#)

Web Server

No items

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

OK

Cancel

3 Fill in the following fields:

Enable Tunnel: Specifies that the Access Gateway should set up a tunnel for all incoming traffic. This option must be enabled to configure a tunnel.

Tunnel SSL Traffic Only: Allows you to configure the Access Gateway to tunnel only SSL traffic. If this option is selected, the Access Gateway verifies that the address and port being accessed are actually an SSL Web site. If verification fails, the service tears down the connection. The SSL port number for the SSL tunnel is specified via the *Listening Port* and the *Connect Port*.

Published DNS Name: Specify the DNS name you want the public to use to access your tunnel. This DNS name must resolve to the IP address you set up as the listening address for the tunnel.

4 Configure the communication options between the browsers and the tunnel by configuring the following fields:

Group Member: (Available only if the Access Gateway is a member of a group.) Select the server you want to configure from the list of servers. The *Listening Address(es)* modifications apply to the selected server. Any other modifications apply to all servers in the group.

Listening Address(es): Displays a list of available IP addresses. If the Access Gateway has only one IP address, only one is displayed. If it has multiple addresses, you can select one or more addresses to enable. You must enable at least one address by selecting its box.

TCP Listen Options: Provides additional options for configuring how requests are handled. See [Section 12.6.1, “Configuring TCP Listen Options for Clients,” on page 167](#). At least one Web server must be configured before you can modify these options.

Listening Port: Specifies the port on which to listen for requests from browsers. The listening address and port combination must not match any combination you have configured for a reverse proxy.

5 Configure the communication options between the tunnel and the Web servers by configuring the following fields:

Connect Port: Specifies the port that the Access Gateway uses to communicate with the Web server.

TCP Connect Options: Allows you to control how idle and unresponsive Web server connections are handled and to optimize these processes for your network. See [Section 12.6.2, “Configuring TCP Connect Options for Web Servers,” on page 168](#).

6 Specify a Web server to receive the traffic. In the Web Server List section, click *New*, specify an IP address, then click *OK*.

At least one Web server must be specified in the list to create a tunnel.

7 To save your changes and push them to the Access Gateway, click *Configuration Panel*, then click *Apply Changes*.

14.5 Configuring Error Page Presentation

The Error Page option allows you to specify how the error pages generated by the Access Gateway are published to the browsers.

- 1** In the Administration Console, click *Access Manager > Access Gateways > Edit > Error Page*.
- 2** In the *Error Page Language* field, select the language in which the error page is published.

- 3 To save your changes and push them to the Access Gateway, click *Configuration Panel*, then click *Apply Changes*.

For information on how to customize the error messages for a specific language on a NetWare Access Gateway, see “[Customizing Error Pages](#)” in the *Novell Access Manager 3.0 Installation Guide*.

14.6 Configuring Console Access

The following options control access to the NetWare Access Gateway console:

- [Section 14.6.1, “Setting Up an FTP Listening Address,” on page 186](#)
- [Section 14.6.2, “Enabling Console Access with SSH and Telnet Sessions,” on page 187](#)
- [Section 14.6.3, “Setting the Password for the admin and config Users,” on page 188](#)

14.6.1 Setting Up an FTP Listening Address

(NetWare only) The Mini FTP option allows you to configure an FTP listening address for management. If this option is enabled, you can upload and download files using FTP.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Mini FTP*.
- 2 Fill in the following fields:

Group Member: (Available only if the Access Gateway is a member of a group.) Select the server you want to configure from the list of servers. All changes made to this page apply to the selected server.

Listening Addresses: To enable this feature, select an IP address for FTP listening.

If the Access Gateway server has only one IP address, only one is displayed for selection. If the server has multiple IP addresses, you can select one or more.

- 3 To save your changes and push them to the Access Gateway, click *Configuration Panel*, then click *Apply Changes*.

When logging in to an FTP session, the username must be `config`, and the password is empty unless you have configured a password. If you enable FTP, we strongly recommend that you set up a password for the config user. See [Section 14.6.3, “Setting the Password for the admin and config Users,” on page 188](#).

14.6.2 Enabling Console Access with SSH and Telnet Sessions

(NetWare only) The Console Access option allows you to control whether administrators can set up SSH or Telnet sessions with the NetWare Access Manager and use command line options to configure it.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Console Access*.

Group Member:

Enable SSH on Server

Enable Telnet on Server

 Warning: Enabling SSH will also open an LDAP listener on port 636 on the server. When disabling, a restart of the server is required to fully close the LDAP listener.

Change Password

Console User:

Old Password:

New Password:

Confirm New Password:

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

- 2 Fill in the following fields:

Group Member: (Available only if the Access Gateway is a member of a group.) Select the server you want to configure from the list of servers. All changes made to this page apply to the selected server.

Enable SSH on Server: If this option is selected, SSH is enabled. SSH sets up a secure, encrypted connection between the Access Gateway and the client. Enabling this option opens an LDAP listener on the Access Gateway for port 636. Disabling this option does not fully close the listener. You must restart the Access Gateway to fully close the LDAP listener.

Enable Telnet on Server: If this option is selected, Telnet is enabled.

IMPORTANT: Telnet is inherently insecure. All information is sent in clear text, including passwords.

- 3 To save your modifications, click *OK*, then on the Configuration page, click *Apply Changes*.

You can use SSH client software or a terminal window to set up a session. When prompted, log in as the admin user for the NetWare Access Gateway console.

If you enable Telnet, use the client software on your workstation to set up a session. When prompted, you can log in as either the config or the admin user for the NetWare Access Gateway console.

14.6.3 Setting the Password for the admin and config Users

(NetWare only) Access Manager sets up an admin user when you install the Administration Console, and you are prompted to supply a name for this user. During installation, the NetWare Access Gateway sets up an `admin` and a `config` user, for managing the NetWare Access Gateway console. These names are not configurable.

The `admin` user is the NetWare Access Gateway console user that has been created for accessing the console over SSH. It is assigned a default password of `novell`.

The `config` user is the NetWare Access Gateway console user that has been created for accessing the console over FTP and Telnet. If you enable FTP or Telnet, you should set up a password for the `config` user. When an Access Gateway is installed, the `config` user is not assigned a password.

To set or modify the password for the `config` or `admin` user:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Console Access*.
- 2 (Conditional) If the Access Gateway is a member of a group, select the server you want to configure from the list of servers. All changes made to this page apply to the selected server.
- 3 In the *Change Password* section, select the *Console User*, either *config* or *admin*.
- 4 Fill in the following fields:

Old Password: Specifies the current password for the console user. When used in conjunction with the *New Password* and *Confirm New Password* fields, this field allows you to change the console password. When the `admin` user was created, it was assigned a default password of `novell`. When you install the NetWare Access Gateway, no password is assigned to the `config` user. To create a password the first time for the `config` user, leave this field blank.

New Password: Specifies a new password. The password must be at least six characters long.

Confirm New Password: Specifies a new password that must match the value in the *New Password* field.

- 5 To apply the changes, click *OK*, then click *Apply Changes*.

14.7 Configuring Network Settings

After initial setup, you seldom need to change the network settings unless something in your network changes, such as you add a new gateway or DNS server. This section describes the following tasks:

- ◆ [Section 14.7.1, “Viewing and Modifying Adapter Settings,” on page 189](#)
- ◆ [Section 14.7.2, “Viewing and Modifying Gateway Settings,” on page 191](#)
- ◆ [Section 14.7.3, “Viewing and Modifying DNS Settings,” on page 194](#)
- ◆ [Section 14.7.4, “Configuring Hosts,” on page 196](#)
- ◆ [Section 14.7.5, “Configuring IPQoS,” on page 197](#)

14.7.1 Viewing and Modifying Adapter Settings

The adapter settings allow you to view the current configuration for the network adapters installed in the Access Gateway machine and manage the IP addresses that are assigned to them. If you want to configure an adapter to use more than one IP address, you can use this option to add them.

To view or modify your current adapter settings:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Adapter List*.

Group Member:

Adapter eth0

[New](#) | [Delete](#)

<input type="checkbox"/> Subnet	Subnet Mask	Addresses
10.0.164.0	255.255.252.0	10.15.167.59
10.0.0.0	255.0.0.0	10.0.0.1

Adapter List Options

Speed: Duplex: NAT:

Custom load parameters:

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

- 2 If the Access Gateway is a member of a group, select the server you want to configure from the list of servers in the *Group Member* field. All changes made to this page apply to the selected server.
- 3 Select the adapter you want to modify, then select one of the following actions:
 - ♦ To add a new subnet to an existing adapter, click *New*.
 - ♦ To delete a subnet, select a subnet, then click *Delete*. More than one must be configured for you to delete a subnet.
 - ♦ To modify an existing subnet, click the IP address of the subnet.

- 4 To configure a new subnet or a new IP address for a subnet, configure the following fields:

Adapter eth0

Subnet: 10.10.167.50

Subnet Mask: *

IP Address List *

New... | Delete | Change IP Address...

<input type="checkbox"/>	IP Addresses
<input type="checkbox"/>	10.15.167.50

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

Subnet: Displays the address of the subnet that you are modifying. This is empty if you are creating a new one.

Subnet Mask: (Required) Specifies the subnet mask address for this subnet.

IP Addresses: Allows you to manage the IP addresses assigned to the subnet.

- ◆ To add an address, click *New*, specify the address, then click *OK*.
- ◆ To delete an address, select the address, then click *Delete*.
- ◆ To change the IP address, see [Section 3.3, “Changing the IP Address of the Access Gateway,” on page 42](#).

- 5 Click *OK*.

- 6 Configure the *Adapter List Options*.

These options let you change settings for the network adapters on the Access Gateway to ensure compatibility with an existing LAN. Modify the default settings only if your LAN requires specialized adapter card changes.

- ◆ **Speed:** Select *Default*, *10 MB*, or *100 MB*.
- ◆ **Duplex:** Select *Default*, *Half*, or *Full*.

IMPORTANT: Some network adapter drivers do not correctly detect duplex settings. This is a general industry problem with Fast Ethernet technology.

If your Access Gateway 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 Access Gateway and your Ethernet switch or hub.

- ◆ **NAT:** Select *Dynamic* or *Disabled*.

If the Access Gateway is serving as a router, and your network employs non-unique private IP addresses, you can configure the Access Gateway 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 Access Gateway, provided that the *Dynamic* option is selected in the NAT drop-down list for the eth1 adapter.

The Access Gateway 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 reverse proxy on an IP address assigned to an adapter that has the *Dynamic* option set for NAT. NAT and a reverse proxy cannot coexist on the same adapter.

Custom load parameters: (NetWare only) Allows you to specify non-standard load parameters for a custom driver. If you used the custom driver option during installation and the documentation for this driver specified some custom load parameters, enter these parameters in the text box.

- 7 To save your changes and push them to the Access Gateway, click *Configuration Panel*, then click *Apply Changes*.

14.7.2 Viewing and Modifying Gateway Settings

The gateway settings display the current gateway configuration that the Access Gateway is using to route packets. From this page, you can also to configure additional gateways. During installation, you could specify only a default gateway. You must have at least one gateway defined for the Access Gateway to function.

The Access Gateway routes requests to specific destinations through these gateways. If a request could be routed through multiple gateways, the Access Gateway 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. The Access Gateway uses additional gateways only when the *Act As Router* option is selected.

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.

To modify your current gateway configuration:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Gateways*.

Group Member:

Enable RIP

Act as Router

Enable Gateway Statistics Monitoring

Default Gateway

[New...](#) | [Delete](#)

<input type="checkbox"/> Next Hop	Metric	Type
<input type="checkbox"/> 10.10.167.50	1	Passive

Host Gateway

[New...](#) | [Delete](#)

<input type="checkbox"/> Next Hop	Host	Metric	Type
No items in list			

Network Gateway

[New...](#) | [Delete](#)

<input type="checkbox"/> Next Hop	Network Address	Mask	Metric	Type
No items in list				

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

- 2 If the Access Gateway is a member of a group, select the server you want to configure from the list of servers in the *Group Member* field. All changes made to this page apply to the selected server.
- 3 Fill in the following fields:

Enable RIP: Allows you to turn on the Routing Information Protocol 1. Through this protocol, the Access Gateway is able to learn routes.

Act as Router: Select this option if the Access Gateway functions as the default gateway for clients on the network. If you select this option, you can specify additional gateways.

Enable Gateway Statistics Monitoring: Select this option if you want to gather statistics and monitor the traffic on the gateways.
- 4 Configure your default gateway, which specifies the gateway to use when no other routes apply. When you select *New* from the *Default Gateway* list, you are asked for the following information:
 - ♦ **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. Specifying a number higher than 1 makes this resource more expensive and alters the gateway logic used. Valid numbers include 1 through 16.
 - ♦ **Type:** Gateways are active if they publish their presence, or passive if they do not.

- 5 Configure your host gateways, which specify the gateways to be used for packets being sent to specific hosts. When you select *New* from the *Host Gateway* list, you are asked for the following information:
 - ♦ **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 relative number indicating the bias you can add to the normal flow of gateway logic. Specifying a number higher than 1 makes this resource more expensive and alters the gateway logic used. Valid numbers include 1 through 16.
 - ♦ **Type:** Gateways are active if they publish their presence, or passive if they do not.
- 6 Configure your network gateways, which specify the gateways to be used for packets being sent to specific subnets. When you select *New* from the *Network Gateway* list, you are asked for the following information
 - ♦ **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 Access Gateway calculates 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 relative number indicating the bias you can add to the normal flow of gateway logic. Specifying a number higher than 1 makes this resource more expensive and alters the gateway logic used. Valid numbers include 1 through 16.
 - ♦ **Type:** Gateways are active if they publish their presence, or passive if they do not.
- 7 To save your modifications, click *OK*, then on the Configuration page, click *Apply Changes*.

14.7.3 Viewing and Modifying DNS Settings

The DNS page displays the current configuration for domain name services and allows you to modify it. To view or modify your DNS configuration:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > DNS*.

Group Member:
Server Hostname:
Domain:

DNS Server IP Addresses	
New... Delete	3 item(s)
<input type="checkbox"/> IP Address	
<input type="checkbox"/> 10.10.1.1	
<input type="checkbox"/> 10.10.1.3	
<input type="checkbox"/> 10.10.1.4	

DNS Cache Settings		
Negative Lookup: *	<input type="text" value="120"/>	(0 - 3600 Second(s))
Minimum Time to Live per Entry: *	<input type="text" value="120"/>	(0 - 3600 Second(s))
Maximum Time to Live per Entry: *	<input type="text" value="168"/>	(0 - 744 Hour(s))
Maximum Entries: *	<input type="text" value="5000"/>	(2000 - 100000)
DNS Transport Protocol:	<input type="text" value="UDP"/> <input type="button" value="v"/>	
<input checked="" type="checkbox"/> Monitor DNS Server		

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

- 2 (Conditional) If the Access Gateway is a member of a group, select the server you want to configure from the list of servers in the *Group Member* field. All changes made to this page apply to the selected server.
- 3 Fill in the following fields:

Server Hostname: Displays the unique host or computer name that you have assigned to the Access Gateway machine. If you modify this name, you need to modify the entry for the Access Gateway in your DNS server to resolve this new name.

Domain: Specifies the domain name for your network. Your DNS server must be configured to resolve the combination of the server hostname and the domain name to the Access Gateway machine. This field assumes you are using dotted names for your machines, such as sales.mytest.com, where sales is the *Server Hostname* and mytest.com is the *Domain*.

DNS Server IP Addresses: Displays the IP addresses of the servers on your network that resolve DNS names to IP addresses. You can have up to three servers in the list. If you specified

any addresses during installation, they appear in this list. To manage the servers in this list, select one of the following options:

- ♦ **New:** To add a server to the list, click this option and specify the IP address of a DNS server.
- ♦ **Delete:** To delete a server from the list, select the address of a server, then click this option.

- 4 Configure the DNS Cache Settings. These options allow you to control the refresh of DNS information. These are all standard DNS options.

Negative Lookup: Specifies how long a failed DNS lookup domain name remains in cache. If the Access Gateway cannot resolve a domain name, it stores that information in its cache for the specified amount of time. If the Access Gateway 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. The default is 120 seconds.

Minimum Time To Live per Entry: Specifies the minimum amount of time that DNS entries remain in cache before they expire. This is the minimum value the Access Gateway uses regardless of the value the DNS server returns. Valid field values include 0–3600 seconds. The default is 120 seconds.

Maximum Time To Live per Entry: Specifies the maximum amount of time that DNS entries remain in cache before they expire. This is the maximum value the Access Gateway uses regardless of the value the DNS server returns. Valid field values include 0–744 hours. The default is 168 hours.

Maximum Entries: Specifies the maximum number of DNS cache entries. When this number is reached, the Access Gateway deletes old entries to make room for newer ones. Valid field values include 2000–100000. The default is 5000.

DNS Transport Protocol: Specifies the transport protocol that DNS uses on the network where the Access Gateway is installed. Valid values are UDP and TCP. The default is UDP.

Monitor DNS Server: If selected, allows the Access Gateway to monitor DNS server availability by pinging the configured servers every minute. This ensures timely handling of DNS requests. You should deselect this item if the Access Gateway accesses DNS through a connection that is not kept continually open, such as a dial-up phone line or ISDN connection. Keep in mind, however, that deselecting this option causes the DNS configuration on the Health tab to fail.

- 5 To save your modifications, click *OK*, then on the Configuration page, click *Apply Changes*.

14.7.4 Configuring Hosts

(Linux only) You can configure the Linux Access Gateway to have multiple host names:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Hosts*.

Group Member:

Host IP Address List	
New...	Delete
<input type="checkbox"/> Host IP Address	Host Name
<input type="checkbox"/> 10.10.0.1	localhost
<input type="checkbox"/> 10.10.167.50	lag45, lag45.provo.novell.com
<input type="checkbox"/> 10.10.167.52	spa-unstable.provo.novell.com

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

This page displays a list of host IP addresses.

- 2 (Conditional) If the Access Gateway is a member of a group, select the server you want to configure from the list of servers in the Group Member field. All changes made to this page apply to the selected server.
- 3 To add a new host name to an existing IP address, click the name of a *Host IP Address*.

Host IP Address: 10.10.159.206

Host Name(s): *

(Place each Host Name on a separate line.)

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

- 4 In the *Host Name(s)* text box, specify a name for the host. Place each host name on a separate line. Then click *OK*.
- 5 To add a new IP address and host name, click *New*, then specify the IP address. In the *Host Name(s)* text box, specify a host name, then click *OK*.
- 6 To delete a host, select the check box next to the host you want to delete, then click *Delete*.
- 7 To save changes, click *OK*, then on the Configuration page, click *Apply Changes*.

14.7.5 Configuring IPQoS

(Linux only) The IP quality-of-service (IPQoS) feature enables you to prioritize, control, and gather accounting statistics. Using IPQoS, you can provide consistent levels of service to users of your network, and manage traffic to avoid network congestion. This configuration can be performed only for the Linux Access Gateway.

- 1 Click *Access Gateways > Edit > IPQoS*.

Mode:

Requests to Server: (0x0-0x3f/6bit)

Replies to Clients: (0x0-0x3f/6bit)

System Errors: (0x0-0x3f/6bit)

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

- 2 Fill in the following fields:

Mode: Specify whether the mode is 6-bit Differentiated Service Bits or 8-Bit TOS Bits. The 8-Bit TOS Bits mode is not widely used.

Requests to Server: Specify the number of QoS bits to be applied to requests sent to servers. If you have selected the *Pass-through* option, it specifies that any existing QoS bits are sent without changing them.

Replies to Clients: Specify the number of QoS bits to be applied to replies sent to clients. If you have selected the *Pass-through* option, it specifies that any existing QoS bits are sent without changing them.

System Errors: Specify the number of QoS bits to be applied error to responses sent to clients.

- 3 To save changes, click *OK*, then on the Configuration page, click *Apply Changes*.

14.8 Customizing Log Out

If any of your protected resources have a logout page or button, you need to redirect the user's logout request to the Access Gateway logout page. The Access Gateway can then clear the user's session and log the user out of any other resources that have been enabled for single sign-on. If you do not redirect the user's logout request, the user is logged out of one resource, but the user's session remains active until inactivity closes the session. If the user accesses the resource again before the session is closed, single sign-on re-authenticates the user to the resource, and it appears that the logout did nothing.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Reverse Proxy / Authentication*.

- 2 In the *Embedded Service Provider* section, view the path in the *Logout URL* option.

The Logout URL displays the URL that you need to use for logging users out of protected resources. This option is not displayed until you have created at least one reverse proxy with a proxy service. If you create two or more reverse proxies, you can select which one is used for authentication, and the logout URL changes to match the assigned reverse proxy. For more information on changing the authentication proxy, see [Section 16.3.2, "Changing the Authentication Proxy Service,"](#) on page 226.

- 3 Use this path to redirect application logout requests to this page.

- 4 Click *OK*.

14.9 Configuring X-Forwarded-For Headers

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.

Deciding whether to enable x-forwarded-for headers 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 sites. This option is disabled by default. To enable it:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > HTTP Options > Header Options*.

Header Options | **Global Cache Options**

Allow Pages to be Cached by the Browser

Enable X-Forwarded-For

Enable Custom Cache Control Header

When Objects Reach the Custom Cache Control Expiration Time:

Revalidate the object with a "Get-If-Modified"

Always obtain a fresh copy of the object

Cache Control Header List

New... | Delete

Cache Control Header

MYCACHE

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

OK Cancel

- 2 Select the *Enable X-Forwarded-For* option.

With this option selected, the proxy service either adds information to an existing X-Forwarded-For or Forwarded-For header, or creates a header if one doesn't already exist. Leaving the option deselected causes the proxy service to remove X-Forwarded-For headers from any Web requests passing through the proxy service.

- 3 Save the changes. Click *Configuration Panel*, then click *Apply Changes*.

14.10 Upgrading the Access Gateway Software

You can upgrade the software currently running on Access Gateway to a newer version without losing configuration information and with down time limited to the time it takes the Access Gateway to restart. See “[Upgrading the Linux Access Gateway](#)” and “[Upgrading the NetWare Access Gateway](#)” in the *Novell Access Manager 3.0 Installation Guide*.

Configuring the Cache Settings

15

One of the major benefits of using an Access Gateway to protect Web resources is that it can cache the requested information and send it directly to the client browser rather than contacting the origin Web resource and waiting for the requested information to be sent. This can significantly accelerate access to the information.

The object cache on an Access Gateway is quite different from a browser's cache, which all users access when they click the Back button and which can serve stale content that doesn't accurately reflect the fresh content on the origin Web server.

The Access Gateway caching system uses a number of methods to ensure cache freshness. Most time-sensitive Web content is flagged by Web masters in such a way that it cannot become stale unless a caching system ignores the Web master's settings. The Access Gateway honors all flags that affect cache freshness, including Time to Expire, Don't Cache, and Must Revalidate directives.

In addition, the Access Gateway can be fine-tuned for cache freshness in the following ways:

- ◆ Accelerated checking of objects that have longer than desirable Time to Expire headers
- ◆ Delayed checking of objects that have shorter than desirable Time to Expire headers
- ◆ Checking for freshness of objects that do not include Time to Expire headers

The following sections describe the features available to fine-tune this process for your network:

- ◆ [Section 15.1, "Configuring Global Caching Options," on page 202](#)
- ◆ [Section 15.2, "Controlling Browser Caching," on page 204](#)
- ◆ [Section 15.3, "Configuring Custom Cache Control Headers," on page 205](#)
- ◆ [Section 15.4, "Configuring a Pin List," on page 207](#)
- ◆ [Section 15.5, "Configuring a Purge List," on page 211](#)
- ◆ [Section 15.6, "Purging Cached Content," on page 212](#)
- ◆ [Section 15.7, "Preventing a Web Site from Being Cached," on page 212](#)

15.1 Configuring Global Caching Options

Caching is configured at the proxy service level. This gives you a great deal of control in specifying what you want cached.

- 1 Click *Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > HTTP Options > Global Cache Options*.

Header Options **Global Cache Options**

Cache Management

- Enable Caching of Objects with a Question Mark
- Enable Caching of Objects with CGI in The Path

Cache Tuning

Refresh Requests from Browser:

- Enable Filter Cookies
- 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:

Cache Freshness

HTTP Maximum:

HTTP Default:

HTTP Minimum:

Continue Fill Time:

HTTP Retries:

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

- 2 Configure the *Cache Management* options:

Enable Caching of Objects with a Question Mark: If this option is selected, a cacheable object is cached if it has a question mark in the URL.

Enable Caching of Objects with CGI in The Path: If this option is selected, a cacheable object is cached if it has `/cgi` in its URL.

Objects that meet these criteria are only cached if they are also cacheable objects. Web server administrators can mark objects as non-cacheable. When so marked, these objects are not cached, even when the above options are selected.

If you disable both of these options, it does not mean that objects with question marks or `cgi` in their paths cannot be cached. These objects can match some other criteria and be cached.

- 3 Configure the Cache Tuning options.

These options restrict or enable functionality that affects all the resources protected by a proxy service.

Refresh Requests From Browser: When a user clicks *Refresh* or *Reload* in the browser, this action sends a new request to the Web server. Select one of the following options to control how the proxy service handles the request:

- ◆ **Refill:** Causes the proxy service to send the request to the Web server
- ◆ **Revalidate:** Causes the proxy service to check whether the current information is valid. If it is, the currently cached information is returned. If it isn't valid, the request is forwarded to the Web server.
- ◆ **Ignore:** Causes the proxy service to ignore the request and send the data from cache without checking to see if the cached data is valid.

Enable Filter Cookies: If this option is selected, removes all cookie HTTP request headers from requests forwarded to Web servers. It also removes all set-cookie HTTP reply headers from replies coming from Web servers.

Enable Initial Splash Screen: If this option is enabled, browsers receive first-time and periodic notification that their requests are being processed by the Access Gateway. You can customize the splash screen so that ISPs, for example, can advertise the fact that they are providing accelerated Web services. The splash screen is disabled (turned off) by default.

Act as Single User (private) Cache: If this option is enabled, the proxy service caches objects that have been flagged for private caches only.

Enable Read-Ahead Images Embedded in the Page: If this option is selected, the proxy service retrieves and caches objects that have been flagged Read-Ahead. You specify the maximum number of read-ahead objects the proxy service retrieves in the *Maximum Number of Concurrent Read-Ahead Requests* field.

Maximum Number of Concurrent Read-Ahead Requests: Sets a limit on the number of read-ahead images that can be cached.

- 4 (Optional) Modify the Cache Freshness settings. Use the *Reset* button to return these settings to their default values.

These options govern when the proxy service revalidates requested cached objects against those on their respective origin Web servers. If the objects have changed, the proxy service re-caches them.

HTTP Maximum: Specifies the maximum time the proxy service 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 directive specified by the Webmaster if he or she specified a longer time.

You use this value to reduce the maximum time the proxy service waits before checking whether requested objects need to be refreshed. The default is 6 minutes.

HTTP Default: Specifies the maximum time the proxy service serves HTTP data for which Web masters have not specified a freshness or Time to Expire directive. The default is 2 minutes.

HTTP Minimum: Specifies the minimum time the proxy service 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 directive specified by the Web master if he or she specified a shorter time.

You can use this value to increase the minimum time the proxy service 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.

The default value is 0, which allows the proxy service to honor the Time To Expire directive of each object (unless it is longer than the *HTTPM Maximum* option). If the *HTTP Minimum* option is set to a value other than 0, the value overrides any object's Time to Expire directive that is shorter than the value set. The default is 0.

Continue Fill Time: Specifies the how long the proxy service ignores browser request cancellations and continues downloading objects from the target Web server until the download is complete. The default is 1 second.

HTTP Retries: Specifies the number of retry requests to issue to a Web server. The default is 4.

- 5 Save the changes. Click *Configuration Panel*, then click *Apply Changes*.

15.2 Controlling Browser Caching

Web masters control how browsers cache information by adding the following cache-control directives to the HTTP headers:

```
Cache-Control: no-store  
Cache-Control: no-cache  
Cache-Control: private  
Cache-Control: public  
Pragma: no-cache
```

You can configure a proxy service to overwrite these directives in the HTTP header.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > HTTP Options > Header Options*.

Header Options | **Global Cache Options**

Allow Pages to be Cached by the Browser

Enable X-Forwarded-For

Enable Custom Cache Control Header

When Objects Reach the Custom Cache Control Expiration Time:

Revalidate the object with a "Get-If-Modified"

Always obtain a fresh copy of the object

Cache Control Header List

New... | Delete

Cache Control Header

MYCACHE

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

OK Cancel

- 2 To mark all pages coming through this host as cacheable on the browser, select *Allow Pages to be Cached by the Browser*.

If this option is not selected, all pages are marked as non-cacheable on the browser. This forces the browser to request a resend of the data from the Access Gateway when a user returns to a previously viewed page. This is the default setting.

- 3 To configure custom caching instructions, see [Section 15.3, “Configuring Custom Cache Control Headers,” on page 205](#).
- 4 Save the changes. Click *Configuration Panel*, then click *Apply Changes*.

15.3 Configuring Custom Cache Control Headers

In addition to fine-tuning cache freshness using the global HTTP timers, as explained in [Section 15.1, “Configuring Global Caching Options,” on page 202](#), you can configure each proxy service to recognize custom headers in HTTP packets. Your Web server can then use these headers for transmitting caching instructions that only the Access Gateway can recognize and follow.

- ♦ [Section 15.3.1, “Understanding How Custom Cache Control Headers Work,” on page 205](#)
- ♦ [Section 15.3.2, “Enabling Custom Cache Control Headers,” on page 206](#)

15.3.1 Understanding How Custom Cache Control Headers Work

Only the proxy service containing the custom header definition follows the cache policies specified in the custom headers.

All other proxy services, requesting browsers, and external proxy caches (transparent caches, client accelerators, etc.), do not recognize the custom headers. They follow only the cache policies specified by the standard cache control headers.

This means that you have the following options for configuring your Web server:

- ♦ You can specify that browsers and/or external caches cannot cache the objects, but the proxy service can.

This lets you offload request-processing from the origin Web server while still requiring that users return to the site each time they request an object.
- ♦ You can also specify separate cache times for browsers, external caches, and the proxy service.

To implement custom cache control headers, you must do the following:

- ♦ Configure a proxy service to use custom cache control headers by enabling the feature and specifying a header string such as MYCACHE (see [Section 15.3.2, “Enabling Custom Cache Control Headers,” on page 206](#)).
- ♦ Configure the Web servers of the proxy service to send an HTTP header containing the defined string and the time in seconds that the object should be retained in cache (for example, MYCACHE: 60).

If the number is non-zero, the Access Gateway treats the reply as if it has the following headers:

```
Cache-Control: public  
Cache-Control: max-age=number
```

If the number is zero (0), the Access Gateway treats the reply as if it has the following header:

```
Cache-Control: no-cache
```

- ◆ Ensure that the Web server continues to send standard HTTP cache-control headers so that browsers and external caches follow the caching policies you intend them to.

For example, you can configure the following:

- ◆ Use an Expires or Cache-Control: Max-Age header to specify that browsers should cache an object for two minutes.
- ◆ Use a Cache-Control: Private header to prevent external caches from caching the object at all.
- ◆ Use a custom cache control header, such as MYCACHE: 1800, to indicate that the accelerator should cache the object for 30 minutes.

Custom Cache Control Headers override the following standard HTTP cache-control headers on the Access Gateway, but they do not affect how browsers and external caches respond to them:

```
Cache-Control: no-store  
Cache-Control: no-cache  
Cache-Control: max-age=number  
Cache-Control: private  
Cache-Control: public  
Pragma: no-cache  
Expires: date
```

15.3.2 Enabling Custom Cache Control Headers

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > HTTP Options > Header Options*.

Header Options **Global Cache Options**

Allow Pages to be Cached by the Browser

Enable X-Forwarded-For

Enable Custom Cache Control Header

When Objects Reach the Custom Cache Control Expiration Time:

Revalidate the object with a "Get-If-Modified"

Always obtain a fresh copy of the object

Cache Control Header List	
New... Delete	
<input type="checkbox"/> Cache Control Header	
<input type="checkbox"/> MYCACHE	

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

OK Cancel

- 2 To enable the use of custom headers, select *Enable Custom Cache Control Header*.

With this option selected, the proxy service searches HTTP packets for custom cache control headers, and caches the objects according to its policies. The policy contains a timer, which specifies how long the object can be cached before checking with the Web server for updates.

- 3 Select one of the following options to specify what occurs when the custom cache control expiration time expires.
 - ♦ **Revalidate the object with a “Get-If-Modified”:** Causes the proxy service to update the object in cache only if the object has been modified.
 - ♦ **Always obtain a fresh copy of the object:** Causes the proxy service to update the object in cache, even if the object has not been modified.
- 4 In the *Cache Control Header List*, select *New* and specify a name for the header, for example MYCACHE.
- 5 Save the changes. Click *Configuration Panel*, then *Apply Changes*.
- 6 Modify the pages on the Web server that you want to set custom caching intervals for the Access Gateway. To the HTTP header, add a string similar to the following:
MYCACHE: 600

The numeric value indicates the number of seconds the Access Gateway can retain the object in cache. A value of zero prevents the Access Gateway from caching the object. This cache interval can be different than the value set for browsers (see [Section 15.3.1, “Understanding How Custom Cache Control Headers Work,”](#) on page 205).
- 7 Ensure that the Web server continues to send the following standard HTTP cache-control headers:
 - ♦ Cache-Control: Max-Age headers that cause browsers to cache object for no longer than two minutes.
 - ♦ Cache-Control: Private headers that cause external caches to not cache the objects.

When your Web server sends an object with the MYCACHE header in response to a request made through the Access Gateway, the proxy service recognizes the custom header and caches the object for 10 minutes. Requesting browsers cache the object for only two minutes, and external caches do not cache the object.

Thus, the Access Gateway off-loads a processing burden from the Web server by caching the frequently requested objects for 10 minutes (the value you specified in [Step 6](#)). Browsers, on the other hand, must always access the Access Gateway to get the objects if their previous requests are older than two minutes. And the objects in the cache of the Access Gateway are kept fresh due to their relatively brief time-to-live value.

15.4 Configuring a Pin List

A pin list contains URL patterns for identifying objects on the Web. The Access Gateway uses the list to prepopulate the cache, before any requests have come in for the content. This accelerates user access to the content because it is retrieved from a local cache rather than from an exchange with the Web server, which would read it from disk.

You can use the pin list to specify the following:

- ♦ Which objects you want always to remain in cache
- ♦ Which objects you never want cached
- ♦ Whether you want the Access Gateway to follow links on the cached pages and cache these linked objects
- ♦ How often you want the Access Gateway to check for modified content (new and deleted objects)

The pin list is global to the Access Gateway and affects all protected resources. The pinned objects remain in cache indefinitely unless the cache fills up. This ensures that the objects are available from cache and are not bumped out by more recently requested objects. You configure each pinned object with a URL pattern and specific handling instructions.

To configure a pin list:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Pin List*.

Enable Pin List

Default Refresh Frequency/Time:

Once Immediately

Each Sunday / 12 MID : 00

Every 01 Hour(s)

Pin List				
New... Delete				
<input type="checkbox"/> URL Mask	Pin Type	Follow Links	Follow to Other Hosts	Refresh Frequency/Time
No items				

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

OK Cancel

- 2 Fill in the following fields:

Enable Pin List: Select this option to enable the use of pinned objects. If this option is not selected, the pinned objects in the Pin List are not used.

Default Refresh Frequency/Time: Sets a default refresh interval for checking the URL patterns and seeing if any new objects need to be cached (or deleted objects removed from cache). This default refresh interval can be overwritten by selecting a different refresh interval for a specific pinned object. Select one of the following for the default value:

- ♦ **Once Immediately:** Select this option to refresh the list as soon as the changes to this page are pushed to the server.
- ♦ **Day and Hour:** Select a day and a time for the refresh.
- ♦ **Hourly Interval:** Select an interval, specified in hours, for refreshing the pin list.

- 3 In the *Pin List* section, click *New*.

- 4 Fill in the following fields.

- ♦ **URL Mask:** Specifies the URL pattern to match. For more information, see [Section 15.4.1, “URL Mask,” on page 209](#).
- ♦ **Pin Type:** Specifies how the URL is to be used to cache objects. Select from *Normal*, *Cache*, *Memory*, and *Bypass*. For more information, see [Section 15.4.2, “Pin Type,” on page 210](#).
- ♦ **Follow Links:** Indicates whether the Access Gateway can follow links and limits nested links to the value specified. A value of zero indicates that links should not be followed. For more information, see [Section 15.4.3, “Follow Links,” on page 211](#).

- ♦ **Other Hosts:** Indicates whether the Access Gateway can follow links to other hosts and cache pages from these hosts. This is only available if the *Follow Links* field is set between 1 and 4.
- ♦ **Refresh Frequent/Time:** Sets a default refresh interval for checking the URL patterns and seeing if any objects have been modified. You can select *Use Default* to use the refresh interval set for all URL patterns or you can specify one for this object, whose value overrides the default setting.

5 Save the changes. Click *Configuration Panel*, then click *Apply Changes*.

15.4.1 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.

The Access Gateway 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.

The Access Gateways recognizes four levels of specificity using the following format:

Level	Examples
hostname	<pre>http://www.foo.gov/documents/picture.gif http://www.foo.gov/documents/ http://www.foo.gov foo.gov/documents/ foo.gov/ *.foo.gov/</pre> <p>All of these are classified as hostnames, and they are ordered by specificity. The first item in the list is considered the most specific and would be processed first. The last item is the most general and would be processed last.</p>
path	<pre>/documents/picture.gif /documents/pictures.gif/ /documents/</pre> <p>Path entries are processed after hostnames. A leading forward slash must always be used when specifying a path, and the entry that follows must always reference the root directory of the Web server. In these examples, <code>documents</code> is the root directory.</p> <p>The trailing forward slash indicates that the entry is a directory. Its absence indicates that the entry is a file. In these examples, <code>picture.gif</code> is a file and <code>pictures.gif/</code> and <code>documents/</code> are directories.</p> <p>These path entry examples are ordered by specificity. The objects in the <code>/documents/picture.gif/</code> directory are processed before the objects in the <code>/documents/</code> directory.</p>

Level	Examples
filename	<pre>/picture.gif /widget.js</pre> <p>Filenames are processed after paths. A leading forward slash must always be used when specifying a filename. If a path is included with a filename, the path must start with the root directory of the Web server (and the entry is processed as a path entry, not as a filename entry).</p>
file extension	<pre>/*.gif /*.js /*.htm</pre> <p>File extensions are processed last. They consist of a leading forward slash, an asterisk, a period, and a file extension.</p>

Specific rules have precedence over less specific rules. Thus, objects matched by a more specific rule are always processed according to its conditions. If a less specific rule also matches the object, the less specific rule is ignored for the object. For example, assume the following two entries in the pin list:

URL Mask	Pin Type	Pin Links
<code>http://www.foo.gov/documents/</code>	cache	1
<code>www.foo*</code>	bypass	n/a

The first entry, because it is most specific, caches the pages in the documents directory and follows any links on those pages and caches the linked pages. The second entry does not affect what the first entry caches, but it prevents any other domain extensions (.com, .net, .org, etc.) whose DNS names begin with www.foo from being cached.

15.4.2 Pin Type

The pin type specifies how the Access Gateway caches objects that match the URL mask.

- ◆ **Normal:** The Access Gateway handles objects matching the mask in the same way it handles any other requested objects. In other words, the objects are cached but not pinned.

Administrators often use this pin type in combination with a broad URL mask that has a bypass pin type. This allows them to insulate specific objects from the effects of the bypass rule.

For example, you could specify a URL mask of `/*.jpg` with a pin type of bypass and a second URL mask of `www.foo.gov/graphics/*.jpg` with a pin type of normal. This causes the 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:** The Access Gateway keeps the pinned objects in cache as long as possible, although they might be written to the hard disk.

- ◆ **Memory:** The Access Gateway 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:** The Access Gateway does not cache the objects. In other words, you can use this option to prevent objects from being cached. URL masks that are set with a type of bypass operate differently than URL masks set with a type of cache. For example, `http://www.digitalair.com/*` with `bypass` as its type bypasses only the root directory. To bypass the entire URL, you need to add as many directory levels as are found on the remote Web server. For example, `www.digitalair.com/*`, `www.digitalair.com/*/*`, `www.digitalair.com/*/*/*`.

You can use two formats to bypass the root directory: the `http://www.digitalair.com/*` format as described above or the DNS name without a trailing forward slash, the `http://www.digitalair.com` format.

15.4.3 Follow Links

The *Follow Links* field specifies the number of links the Access Gateway can follow as it caches objects that match the URL pattern. For example, if the requested object is an HTML page and you have specified a *Follow Links* level of 1, the HTML page is downloaded and cached along with all the items linked from the page. These cached objects are also refreshed at the frequency and time specified. If there are links on the linked pages, these links are not followed and those pages are not cached. To add these objects, you would need to specify 2 for the *Follow Links* option.

To use a level other than 0, you must specify an absolute address, including the scheme, host, and path for the URL mask, for example:

```
http://www.foo.gov/documents/
```

15.5 Configuring a Purge List

The purge list is global to the Access Gateway and affects all protected resources. This option allows you to specify URL patterns or masks for the pages and sites whose objects you want to purge from cache.

When defining the masks, keep in mind that the Access Gateway 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. For example:

URL Mask	Effects
<code>/*.pdf</code>	Causes all PDF files to be purged from cache.
<code>www.foo.gov/contracts/*</code>	Causes all objects in the <code>contracts</code> directory and beyond to be purged from cache.

This option also allows you to purge cached 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` purges all objects with the text “=SPORTS” or any other combination of uppercase and lowercase letters for “=SPORTS” following the question mark in the URL.

IMPORTANT: If you also configure a pin list, carefully select the objects that you add to the pin and purge lists. You can configure the Access Gateway to use the pin list to add objects to the cache and to use the purge list to remove the same objects.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Purge List*.

Purge List
New... | Delete
 URL Mask
No items

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

OK Cancel

- 2 Click *New* and enter a URL pattern.
- 3 (Optional) Repeat Step 2 to add additional URL patterns.
- 4 Save the changes. Click *Configuration Panel*, then click *Apply Changes*.

15.6 Purging Cached Content

You can select to purge the content of the purge list or all content cached on the server.

- 1 In the Administration Console, click *Access Manager > Access Gateways > [Name of Server] > Actions*.
- 2 Select one of the following actions:
 - Purge List Now** Click this action to cause all objects in the current purge list to be purged from the cache.
 - Purge All Cache** Click this action to purge the server cache. All cached content, including items cached by the pin list, is purged.
- 3 Click either *OK* or *Cancel*

15.7 Preventing a Web Site from Being Cached

The Access Gateway is designed to cache Web pages. However, sometimes you need to use the Access Gateway to protect a Web site and provide single sign-on, but you do not want the content of the Web server cached.

To prevent the caching of a Web site, you need to add the site to the Pin List with a pin type of Bypass.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Pin List*.
- 2 Make sure the *Enable Pin List* option is selected.
- 3 In the Pin List section, click *New* and fill in the following fields:
 - ♦ **URL Mask:** The URL pattern to match. Specify the published DNS name of the Web server that should not have its content cached. For example:
`http://myserver.mycompany.com`

This type of entry prevents the caching of pages on the Web site when accessed over HTTP. To block both HTTP and HTTPS, you can add a second entry for HTTPS or remove the scheme from the URL pattern.

- ♦ **Pin Type:** The caching action. To prevent caching, select *Bypass*.

Accept the default values for the other fields, or configure them to fit your needs. For more information, see [Section 15.4, “Configuring a Pin List,” on page 207](#).

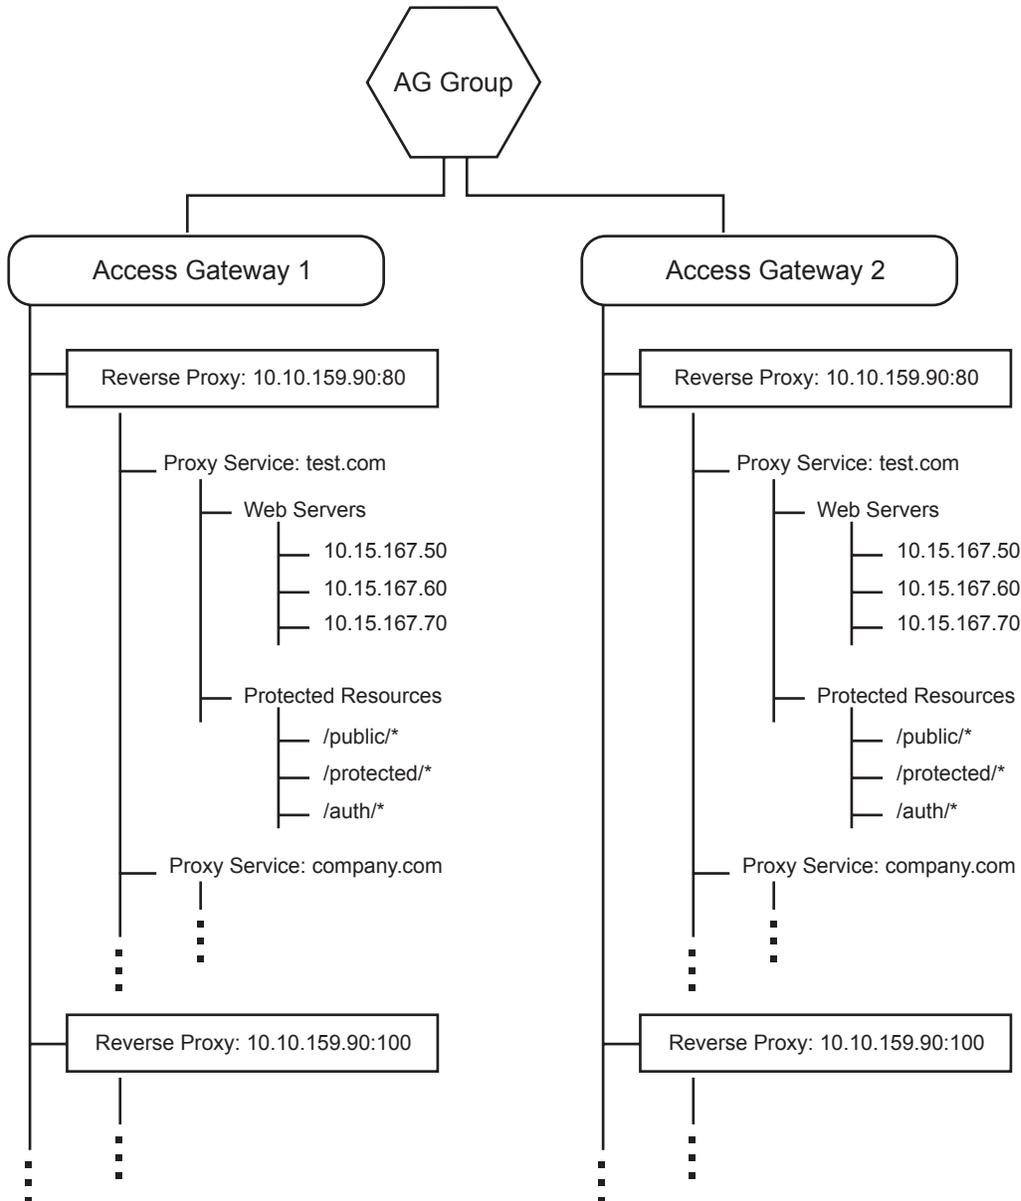
- 4 Click *OK*.
- 5 Save the changes. Click *Configuration Panel*, then click *Apply Changes*.
- 6 To purge any pages that might have been cached while you were configuring the Pin List, purge the existing cache. See [Section 15.6, “Purging Cached Content,” on page 212](#).

Protecting Multiple Resources

16

This section describes how to create multiple resources for the various Access Gateway components, including setting up a group of Access Gateways for fault tolerance. [Figure 16-1](#) illustrates the relationships that Access Gateways, reverse proxies, proxy services, Web servers, and protected resources have with each other.

Figure 16-1 Hierarchical View of the Access Gateway Configured Objects



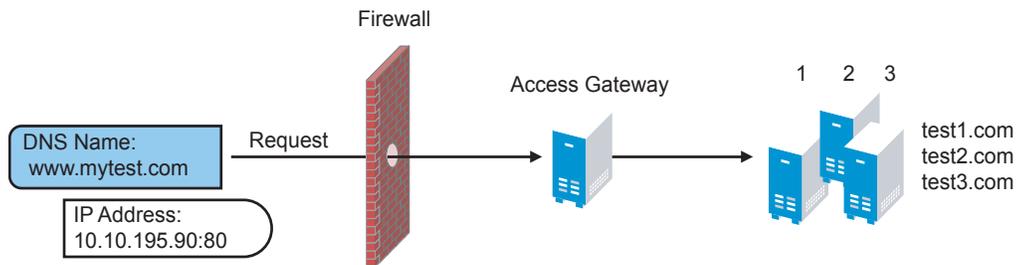
In [Figure 16-1](#), Access Gateway 1 and Access Gateway 2 have the same configuration because they are members of the AG Group. This section explains how to create a group of Web servers and how to add multiple proxy services and reverse proxies to an Access Gateway.

- ◆ [Section 16.1, “Setting Up a Group of Web Servers,”](#) on page 216
- ◆ [Section 16.2, “Using Multi-Homing to Access Multiple Resources,”](#) on page 216
- ◆ [Section 16.3, “Managing Multiple Reverse Proxies,”](#) on page 224
- ◆ [Section 16.4, “Managing a Group of Access Gateways,”](#) on page 226

16.1 Setting Up a Group of Web Servers

You can configure a proxy service to service a “virtual” group of Web servers, which adds load balancing and redundancy. Each Web server in the group must contain the same material. When you create the proxy service, you set up the first server by specifying the URLs you want users to access and the rights the users need for each URL. When you add additional Web servers to the proxy service, these servers automatically inherit everything you have configured for the first Web server.

Figure 16-2 Adding Redundant Web servers



For this configuration, you use a single reverse proxy and proxy service. To add multiple Web servers to a host:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > Web Servers*.
- 2 In the *Web Server List* section, click *New*.
- 3 Specify the IP address of another Web server for the “virtual” group, then click *OK*.
- 4 Repeat Steps 2 and 3 to add additional Web servers to the group.
- 5 Click *OK* twice.
- 6 On the Configuration page, click *Apply Changes*, then click *OK*.

Session persistence, which allows a browser to persistently use the same Web server after an initial connection has been established, is enabled whenever a second Web server is added to the list. For more information on persistent connections, see [Section 12.6.3, “Configuring Connection and Session Persistence,”](#) on page 170.

16.2 Using Multi-Homing to Access Multiple Resources

You can configure the Access Gateway to use one public IP address to protect multiple types of Web resources. This is one of the major benefits of Access Gateway; it conserves valuable resources such

as IP addresses. This feature also makes the Access Gateway a multi-homing device because it becomes a single endpoint supporting multiple back-end resources.

You can select to use only one multi-homing method, or you can use multiple methods. Select the methods that meet the needs of your network and the resources you are protecting. The first proxy service configured for a reverse proxy is always configured to use the DNS name of the resource. Subsequent proxy services can be configured to use one of the following methods:

- ◆ [Section 16.2.1, “Domain-Based Multi-Homing,” on page 217](#)
- ◆ [Section 16.2.2, “Path-Based Multi-Homing,” on page 218](#)
- ◆ [Section 16.2.3, “Virtual Multi-Homing,” on page 221](#)

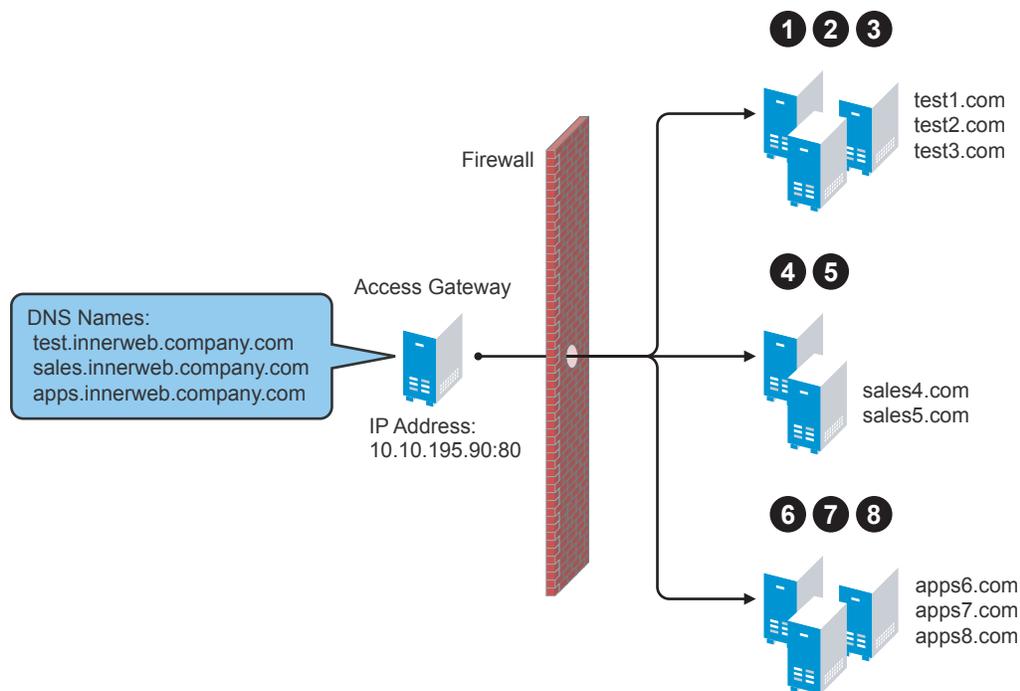
This section describes these multi-homing methods, then explains the following:

- ◆ [Section 16.2.4, “Creating a Second Proxy Service,” on page 221](#)
- ◆ [Section 16.2.5, “Configuring a Path-Based Multi-Homing Proxy Service,” on page 223](#)

16.2.1 Domain-Based Multi-Homing

Domain-based multi-homing is based on the cookie domain. For example, if you have a cookie domain of `company.com`, you can prepend host names to cookie domain name. For a test resource, you can prepend `test` to `company.com` and have `test.company.com` resolve to the IP address of the Access Gateway. The Access Gateway configuration for the `test.company.com` proxy service contains the information for accessing its Web servers (`test1.com`). [Figure 16-3](#) illustrates this type of configuration for three proxy services.

Figure 16-3 Using a Base Domain Name with Host Names



Domain-based multi-homing has the following characteristics:

- ◆ If you are using SSL, the back-end servers can all listen on the same SSL port (default for HTTPS is 443).
- ◆ If you are using SSL, the back-end servers can share the same SSL certificate. Instead of using a specific host name in the SSL certificate, the certificate can use a wildcard name such as *.company.com, which matches all the servers.

Before configuring the Access Gateway, you need to complete the following:

- ◆ Create the published DNS names with a common domain name for public access to the back-end resources. For example, the table below lists three DNS that use company.com as a common domain name and then lists the IP address that these DNS names resolve to and the Web servers they are going to protect.

Published DNS Name	Access Gateway IP Address	Web Server Host Name	Web Server IP Address
test.company.com	10.10.195.90:80	test.internal.com	10.15.0.10
sales.company.com	10.10.195.90:80	sales.internal.com	10.15.0.20
apps.company.com	10.10.195.90:80	apps.internal.com	10.15.0.30

- ◆ Configure your DNS server to resolve the published DNS names to the IP address of the Access Gateway.
- ◆ Set up the back-end Web servers.

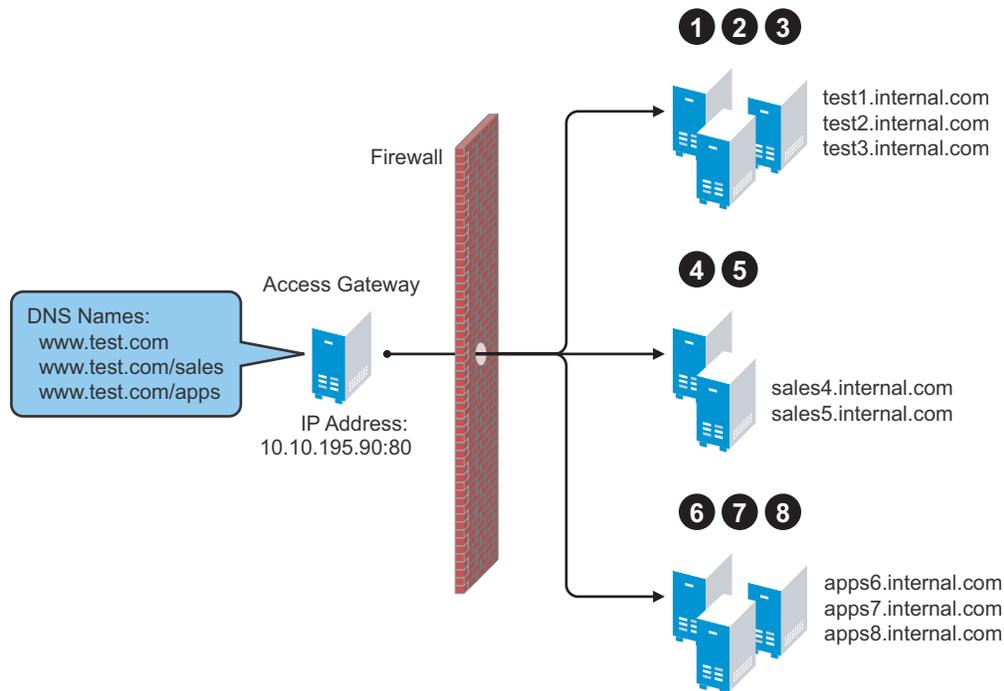
To create a domain-based multi-homing proxy service, see [Section 16.2.4, “Creating a Second Proxy Service,” on page 221](#), and select domain-based for the multi-homing type.

16.2.2 Path-Based Multi-Homing

Path-based multi-homing uses the same DNS name for all resources, but each resource, or resource group, must have a unique path appended to the DNS name. For example, if the DNS name is test.com, you would append /sales to test.com. When the user enters the URL of www.test.com/

sales, the Access Gateway resolves the URL to the sales resource group. **Figure 16-4** illustrates this type of configuration.

Figure 16-4 Using a Domain Name with Path Elements



Path-based multi-homing has the following characteristics:

- ◆ It is considered to be more secure than domain-base multi-homing, because some security experts consider wildcard certificates less secure than a certificate with a specific hostname.
- ◆ Each resource or group of resources must have a unique starting path.
- ◆ JavaScript applications might not work as designed if they obscure the URL path. The Access Gateway needs access to the URL path, and if it is obscured, the path cannot be resolved to the correct back-end resource.
- ◆ The protected resources for each path-based child come from the parent proxy service.
- ◆ The proxy service can be configured to remove the path or use the path when accessing the back-end resource. See the next sections.

Removing the Path from the URL

The path that is part of the published DNS name (/sales or /apps) can be used to identify a resource and if this is the case, the path is not part of directory configuration on the Web server. For example, suppose you use the following configuration:

Browser URL Using the Published DNS Name	Web Server URL
http://www.test.com/sales	http://sales4.internal.com/

In this case, the path needs to be removed from the URL that the Access Gateway sends to the Web server. This configuration assumes that the first page on Web server provides navigation for the site. The Access Gateway does not allow you to set up multiple paths to the Web server.

Using the Path in the URL

The path in the published DNS name can be passed to the Web server as part of the URL. For example, suppose you use the following configuration:

Browser URL Using the Published DNS Name	Web Server URL
http://www.test.com/sales	http://sales4.internal.com/sales

If the path component is a directory on the Web server where the content begins, you need to select to include the path. The Access Gateway then includes the path as part of the URL it sends to the Web server. This configuration allows you to set up multiple paths to the Web server, such as

- ◆ sales/payroll
- ◆ sales/reports
- ◆ sales/products

Configuring Path-Based Multi-Homing

Before configuring the Access Gateway, you need to complete the following:

- ◆ Create the published DNS names with paths for public access to the back-end resources. For example, the table below uses test.com as the domain name. It lists three published DNS names (two with paths), the IP address these names resolve to, and the Web servers that they are going to protect:

Published DNS Name	Access Gateway IP Address	Web Server Host Name	Web Server IP Address
test.com	10.10.195.90:80	test.internal.com	10.15.0.10
test.com/sales	10.10.195.90:80	sales.internal.com	10.15.0.20
test.com/apps	10.10.195.90:80	apps.internal.com	10.15.0.30

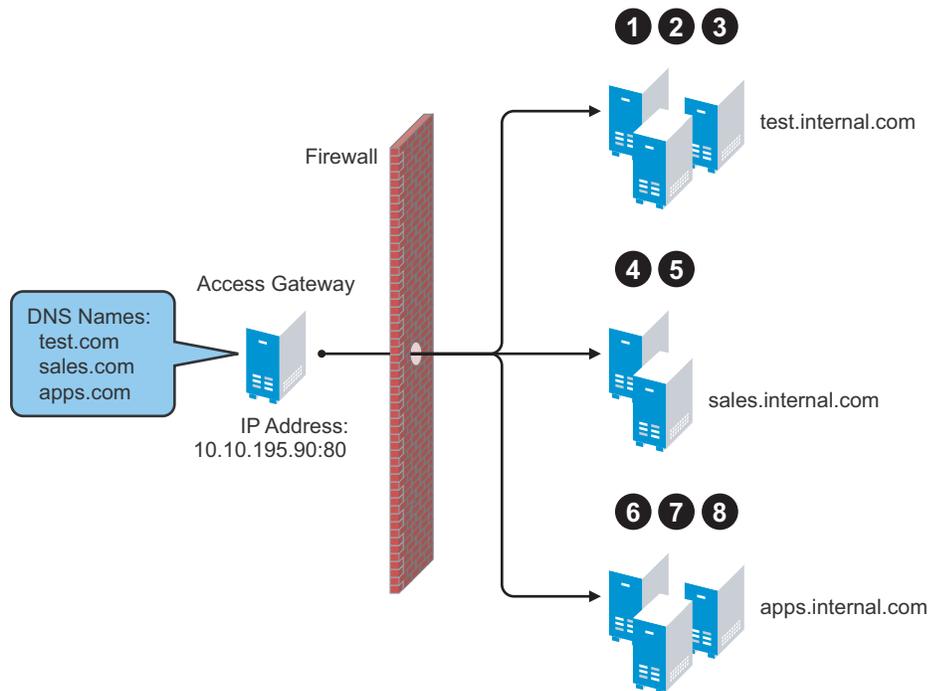
- ◆ Configure your DNS server to resolve the published DNS names to the IP address of the Access Gateway.
- ◆ Set up the back-end Web servers.

To create a path-based multi-homing proxy service, see [Section 16.2.4, “Creating a Second Proxy Service,”](#) on page 221, and select path-based for the multi-homing type.

16.2.3 Virtual Multi-Homing

Virtual multi-homing allows you to use DNS names from different domains (for example `test.com` and `sales.com`). Each of these domain names must resolve to the Access Gateway host. [Figure 16-5](#) illustrates this type of configuration.

Figure 16-5 Using Multiple DNS Names



Virtual multi-homing cannot be used with SSL. You should use this configuration with resources that need to be protected, but the information exchanged should be public information that does not need to be secure. For example, you could use this configuration to protect your Web servers that contain the catalog of your shipping products. It isn't until the user selects to order a product that you need to switch the user to a secure site.

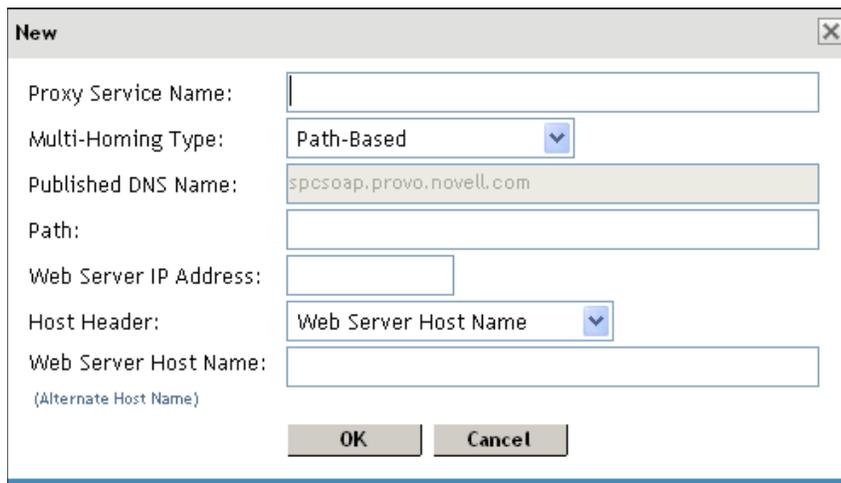
Whether a client can use one DNS name or multiple DNS names to access the Access Gateway depends upon the set up of your DNS server. Once you have configured your DNS server to allow this, you are ready to configure the Access Gateway.

To create a Virtual multi-homing proxy service, see [Section 16.2.4, “Creating a Second Proxy Service,”](#) on page 221, and select *Virtual* for the multi-homing type.

16.2.4 Creating a Second Proxy Service

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy]*.

2 In the *Proxy Service List*, select *New*.



3 Fill in the fields.

Proxy Service Name. Specify a display name for the proxy service. For the sales group, you might use sales. For the group of application servers, you might use apps.

Multi-Homing Type: Specify the multi-homing method that the Access Gateway should use to identify this proxy service. Select one of the following:

- ♦ **Domain-Based:** Uses the published DNS name (www.test.com) with a host name (www.newsite.test.com). For more information, see [Section 16.2.1, “Domain-Based Multi-Homing,”](#) on page 217.
- ♦ **Path-Based:** Uses the published DNS name (www.test.com) with a path (www.test.com/path). For more information, see [Section 16.2.2, “Path-Based Multi-Homing,”](#) on page 218.
- ♦ **Virtual:** Uses a unique DNS name (www.newsite.newcompany.com). Virtual multi-homing cannot be used with SSL. For more information, see [Section 16.2.3, “Virtual Multi-Homing,”](#) on page 221. If you need a unique DNS name and SSL, you need to create a reverse proxy rather than a proxy service. For information on creating a second reverse proxy, see [Section 16.3, “Managing Multiple Reverse Proxies,”](#) on page 224.

Published DNS Name: Specify the DNS name you want the public to use to access your site. This DNS name must resolve to the IP address you set up as the listening address. This option is not available when path-based multi-homing is selected.

Path: Specify the path to use for this proxy service. This option is available only when path-based multi-homing is selected.

Web Server IP Address: Specify the IP address of the Web server you want this proxy service to manage.

Host Header: Specify whether the HTTP header should contain the name of the back-end Web server (*Web Server Host Name* option) or whether the HTTP header should contain the published DNS name (the *Forward Received Host Name* option).

Web Server Host Name: Specify the DNS name of the Web server that the Access Gateway should forward to the Web server. If you have set up a DNS name for the Web server and the Web server requires its DNS name in the HTTP header, specify that name in this field. If you selected *Forward Received Host Name*, this option is not available.

NOTE: For iChain[®] administrators, the *Web Server Host Name* is the alternate host name when configuring a Web Server Accelerator.

4 Click *OK*.

5 To continue, select one of the following:

- ♦ To configure a virtual or domain-based proxy service, see [Section 12.2, “Configuring a Proxy Service,”](#) on page 146.
- ♦ To configure a path-based proxy service, see [Section 16.2.5, “Configuring a Path-Based Multi-Homing Proxy Service,”](#) on page 223.

16.2.5 Configuring a Path-Based Multi-Homing Proxy Service

To configure a path-based proxy service:

1 In the Administration Console, click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Path-Based Multi-Homing Proxy Service]*.

The screenshot shows the configuration interface for a Path-Based Multi-Homing proxy service. At the top, there are tabs for 'Path-Based Multi-Homing', 'Web Servers', 'HTML Rewriting', and 'Logging'. The 'Path-Based Multi-Homing' tab is active. Below the tabs, the following fields are visible:

- Published DNS Name: spcsoap.provo.novell.com/ ... (1) path(s)
- Description: [Empty text box]
- Cookie Domain: provo.novell.com
- [HTTP Options](#)
- Remove Path on Fill
 - Reinsert Path in "set-cookie" Header

Below these options is a 'Path List' table with a blue header. The table contains one entry:

<input type="checkbox"/> Path	Protected Resource
<input type="checkbox"/> /apps	base

At the bottom of the panel, there is a note: 'Changes made on this panel must be applied or scheduled from the [Configuration](#)' and two buttons: 'OK' and 'Cancel'.

The following fields display information that must be configured on the parent proxy service (the first proxy service created for this reverse proxy).

- ♦ **Published DNS Name:** Displays the value that users are currently using to access this proxy service. This DNS name must resolve to the IP address you set up as a listening address on the Access Gateway.
- ♦ **Cookie Domain:** Displays the domain for which the cookie is valid. The Web server that the user is accessing must be configured to be part of this domain.

2 Configure the following options:

Description: (Optional) Provides a field where you can describe the purpose of this proxy service or specify any other pertinent information.

HTTP Options: Determines how the proxy service handles HTTP headers and caching. For more information, see [Section 15.3, “Configuring Custom Cache Control Headers,”](#) on

page 205, Section 15.2, “Controlling Browser Caching,” on page 204, and Section 15.1, “Configuring Global Caching Options,” on page 202.

3 Configure the path options:

Remove Path on Fill: Determines whether the multi-homing path is removed from the URL before forwarding it to the Web server. If the path is not a directory at the root of the Web server, the path must be removed. If this option is selected, the path is stripped from the request before the request is sent to the Web server.

If you enable this option, this proxy service can protect only one path. If you have configured multiple paths in the *Path List*, you cannot enable this option until you have deleted all but one path.

Reinsert Path in “set-cookie” Header: Determines whether the path is inserted into the “set cookie” header. This option is only available if you enable the *Remove Path on Fill* option.

4 Determine whether you need to create a protected resource for your path.

In the *Path List*, the path you specified is listed along with the protected resource that best matches its path.

4a In the *Path List* section, click the *Protected Resource* link.

4b Examine the contract, Authorization, Identity Injection, and Form Fill policies assigned to this protected resource.

4c To return to the Path-Based Multi-Homing page, click the *Overview* tab, then click *OK*.

- ♦ If the protected resource meets your needs, continue with **Step 5**
- ♦ If it does not meet your needs, you must create a protected resource for the path-based proxy service. Continue with **Step 4d**.

4d Click *OK*, the name of the parent proxy service, then *Protected Resources*.

4e In the *Protected Resource List*, click *New*, specify a name, then click *OK*.

4f Assign a contract.

4g In the *URL Path List*, specify the path you used when creating the path-based proxy service followed by a */**. For example, if your path was */apps*, specify */apps/** in the *URL Path List*.

4h (Optional) Enable the policies the path-based proxy service requires. Click *Authorization*, *Identity Injection*, or *Form Fill* and enable the appropriate policies.

4i Click *OK*.

5 To save the changes, click *Configuration Panel*, then click *Apply Changes*.

16.3 Managing Multiple Reverse Proxies

Each reverse proxy must have a unique IP address and port combination. If your Access Gateway has only one IP address, you must select unique port numbers for each additional reverse proxy that you create. You can configure the Access Gateway to use multiple IP addresses. These addresses can be configured to use the same network interface card, or if you have installed multiple network cards, you can assign the IP addresses to different cards. To configure IP addresses and network interface cards, see **Section 14.7.1, “Viewing and Modifying Adapter Settings,” on page 189**.

If you are creating more than one reverse proxy, you must select one to be used for authentication. By default, the first reverse proxy you create is assigned this task. Depending upon your Access Gateway configuration, you might want to set up one reverse proxy specifically for handling

authentication. The authentication reverse proxy is also used for logout. If you have Web applications that contain logout options, these options need to be redirected to the Logout URL of the authentication proxy.

- ♦ [Section 16.3.1, “Managing Entries in the Reverse Proxy List,” on page 225](#)
- ♦ [Section 16.3.2, “Changing the Authentication Proxy Service,” on page 226](#)

16.3.1 Managing Entries in the Reverse Proxy List

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Reverse Proxy / Authentication*

The screenshot displays the 'Authentication Settings' and 'Reverse Proxy List' sections of the Administration Console. The 'Authentication Settings' section includes a dropdown for 'Identity Server Configuration' set to 'IDS-BF-Provo'. The 'Embedded Service Provider' section shows 'Reverse Proxy' set to 'myCompany', with 'Metadata URL', 'Health-Check URL', and 'Logout URL' all pointing to 'https://jw.provo.novell.com:443/nesp/app/heartbeat', 'https://jw.provo.novell.com:443/nesp/app/heartbeat', and 'https://jw.provo.novell.com:443/nesp/app/plogout' respectively. Below this is a link for 'Auto-Import Identity Server Configuration Trusted Root'. The 'Reverse Proxy List' section features a table with columns for 'Name', 'Enabled', 'Listening Address', and 'Port'. Two entries are listed: 'myCompany' (Enabled, 10.10.159.206, 443) and 'mySales' (Enabled, 10.10.159.206, 81). At the bottom, there are 'OK' and 'Cancel' buttons.

Authentication Settings

Identity Server Configuration:

Embedded Service Provider

Reverse Proxy:

Metadata URL:

Health-Check URL:

Logout URL:

[Auto-Import Identity Server Configuration Trusted Root](#)

Reverse Proxy List

[New...](#) | [Delete](#) | [Enable](#) | [Disable](#)

<input type="checkbox"/>	Name	Enabled	Listening Address	Port
<input type="checkbox"/>	myCompany	✓	10.10.159.206	443
<input type="checkbox"/>	mySales	✓	10.10.159.206	81

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

- 2 In the *Reverse Proxy List*, select one of the following actions:
 - ♦ **New:** To create a new reverse proxy, click *New*. You are prompted to enter a display name for the proxy. For configuration information, see [Section 12.1, “Creating a Reverse Proxy and Proxy Service,” on page 142](#).
Reverse proxy names and proxy service names must be unique to the Access Gateway. Protected resource names need to be unique to the proxy service, but they don’t need to be unique to the Access Gateway.
 - ♦ **Delete:** To delete a reverse proxy, select the check box by a specific reverse proxy, then click *Delete*. To delete all reverse proxies, select the check box by the *Name* column, then click *Delete*.
 - ♦ **Enable:** To enable a reverse proxy, select the check box by a specific reverse proxy, then click *Enable*. To enable all reverse proxies, select the check box by the *Name* column, then click *Enable*.

- ♦ **Disable:** To disable a reverse proxy, select the check box by a specific reverse proxy, then click *Disable*. To enable all reverse proxies, select the check box by the *Name* column, then click *Disable*.
- 3 To save your changes and push them to the Access Gateway, click *Configuration Panel*, then click *Apply Changes*.

16.3.2 Changing the Authentication Proxy Service

If you have multiple reverse proxies, you can select the reverse proxy that users are redirected to for login and logout.

IMPORTANT: Changing the reverse proxy that is used for authentication is not a trivial task. For example, if you have customized the logout options on your Web servers to redirect the logout request to the Logout URL of the current authentication reverse proxy, you need to modify these options to point to a new Logout URL.

If you have set up SSL connections, you need to change your certificate configurations.

To select the reverse proxy to use for authentication:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Reverse Proxy / Authentication*.
- 2 In the *Embedded Service Provider* section, select the reverse proxy that you want to use for authentication.
The screen is refreshed and the *Logout URL*, *Metadata URL*, and *Health-Check URL* are rewritten to use the new reverse proxy.
- 3 In the *Reverse Proxy List*, click the name of the reverse proxy that you have selected for authentication.
- 4 If you have enabled SSL between the embedded service provider and the Identity Server, you need to import the trusted root of the Identity Server into the trusted root keystore of embedded service provider. Click *Auto-Import Identity Server Configuration Trusted Root*, click *OK*, specify an alias, click *OK*, then click *Close*.
- 5 Click *OK*, then click *Apply Changes*.
If you are applying the changes to a group, a group scheduled command page appears. To apply the changes, you must select each server in the *Group Member Status* list, then click *Apply to server(s)*. You can select to apply to changes to one or more servers at the same time.
- 6 (Conditional) If you have customized Web logout pages, update them to use the new Logout URL.

16.4 Managing a Group of Access Gateways

Most of the configuration tasks are the same for a single Access Gateway and a group of Access Gateways. (For information on how to create a group of Access Gateways, see “[Configuring Access Gateways for Fault Tolerance](#)” in the *Novell Access Manager 3.0 Setup Guide*.) This section describes the tasks that are specific to managing a group:

- ♦ [Section 16.4.1, “Applying Configuration Changes to a Group,” on page 227](#)
- ♦ [Section 16.4.2, “Issuing Commands to a Group of Access Gateways,” on page 228](#)

- ◆ Section 16.4.3, “Managing the Servers in the Group,” on page 229
- ◆ Section 16.4.4, “Adding Servers to an Existing Group,” on page 229
- ◆ Section 16.4.5, “Changing the Primary Cluster Server,” on page 230
- ◆ Section 16.4.6, “Deleting Members from a Group,” on page 230
- ◆ Section 16.4.7, “Viewing the Health of the Group,” on page 230
- ◆ Section 16.4.8, “Viewing Group Statistics,” on page 231
- ◆ Section 16.4.9, “Viewing Group Command Status,” on page 231
- ◆ Section 16.4.10, “Viewing the Status of Group Alerts,” on page 232

16.4.1 Applying Configuration Changes to a Group

Most configuration changes require the same steps, regardless of whether the Access Gateway is a single instance or whether it belongs to a group. However when you are configuring a group, you must select an Access Gateway for a few configuration options because the changes, such as modifying the time, apply only to one machine. The biggest difference between configuring a single machine and a group occurs when you click *Apply Changes* on the main configuration page.

When you apply changes to a group, you are taken to the Group Command Information page.

Command Information

[Refresh](#) | [Disable](#) | [Delete](#)

Name: ag60 Configuration
 Type: Group Configuration
 Admin: cn=admin,o=novell
 Group: ag60
 Description: ag60 Configuration
 Status: PENDING
 Last Executed On: Oct 12, 2006 2:07 PM

Group Members

[Apply to server\(s\)](#)

<input type="checkbox"/> Server	Apply Status	Server Command Status	Date & Time	Health	Details
<input type="checkbox"/> 10.15.167.59	[None]	[None]	Oct 12, 2006 2:07 PM		View
<input type="checkbox"/> 10.15.167.63	[None]	[None]	Oct 12, 2006 2:07 PM		View

From this page you must decide how you want the changes applied. The Group Members section lists all the servers in the group. You can select to apply the command to all servers in the group by selecting the servers, then clicking *Apply to server(s)*. Or you can select to manually stagger the application of the command by applying the command to one server, and waiting until the command succeeds before applying the command to the next server. Staggering the application of the command allows you to verify that the command does what you want. If the command causes the Health icon of the first server to turn red or white, your site remains up while you fix the problem before applying the command to the other servers in the group.

Even if the configuration change applies to only one server in the group, the command must be applied to all servers in the group for the status of the command to change from pending to success.

If you disable a command or fail to apply it to all members of the group, you cannot make any additional changes to the group configuration until the command has been either deleted or applied to all servers in the group.

16.4.2 Issuing Commands to a Group of Access Gateways

- 1 In the Administration Console, click *Access Manager > Access Gateways > Groups* tab > *[Name of Group] > Actions*.
- 2 Select one of the following actions, then click *OK* or *Cancel*.
 - ◆ **Shutdown Now**
 - ◆ **Schedule Shutdown**
 - ◆ **Restart Now**
 - ◆ **Schedule Restart**
 - ◆ **Upgrade** (NetWare only)
 - ◆ **Purge All Cache**
 - ◆ **Start Service Provider**
 - ◆ **Stop Service Provider**
 - ◆ **Restart Service Provider**

Shutdown Now

Shuts down all the Access Gateways in the group. You need to have physical access to each Access Gateway server to restart them.

Restart Now

Click this action to stop and start all the Access Gateways in the group.

Purge All Cache

Click this action to purge the server cache on all servers in the group. All cached content is lost on all servers.

Start Service Provider

Click this option to start the service providers associated with the Access Gateways in this group. The service provider is the module within the Access Gateway that communicates with the Identity Server.

Stop Service Provider

Click this option to stop the service providers associated with the Access Gateways in this group. The service provider is the module within the Access Gateway that communicates with the Identity Server.

Restart Service Provider

Click this option to restart the service providers associated with the Access Gateways in this group. This command stops the service provider and then starts it. The service provider is the module within the Access Gateway that communicates with the Identity Server.

The service provider should be restarted whenever you enable or modify logging on the Identity Server.

When an Access Gateway group is not functioning correctly, you should always try restarting the service providers before stopping and starting the Access Gateways.

16.4.3 Managing the Servers in the Group

To view the servers that are currently members of the group.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Groups* tab > *[Group Name] > Servers*
- 2 To add a server to the group, click *Add*. If other servers are available, a list of available servers appears. Select the server or servers, then click *Add Selected Servers*. If you decide not to add any servers, click *Cancel*.
- 3 To remove a server from the group, select the server, then click *Delete*.

Usually when you delete a member from a group, you have discovered that traffic is lighter than anticipated and that it can be handled with fewer machines while another group is experiencing higher traffic and can benefit from having another member. When the member is deleted, its configuration object maintains all the configuration settings from the group. When it is added to a new group, its configuration object is updated with the configuration settings of the new group. If your groups are behind an L4 switch, you need to reconfigure the switch so that the server is assigned to the correct group.

When a member is deleted from a group, its embedded service provider is stopped. When you have reconfigured the member so that it is protecting resources other than the ones it did in the group, you need to restart the embedded service provider. See [“Restarting the Access Gateway Service Provider” on page 28](#).

- 4 To modify which server is the primary cluster server, see [Section 16.4.5, “Changing the Primary Cluster Server,” on page 230](#).
- 5 To view detailed information about a server in the group, click the name of the server.
- 6 To view detailed health information about a server, click the health icon of the server. For more information, see [Section 36.3, “Monitoring the Health of an Access Gateway,” on page 446](#).
- 7 Click *Close*.

16.4.4 Adding Servers to an Existing Group

- 1 In the Administration Console, click *Access Manager > Access Gateways > Groups* tab > *[Name of Group]*.
- 2 Click the *Servers* tab, then click *Add*.
If any servers are available to add to the group, they are listed.
- 3 Select the server, then click *Add Selected Server(s)*.
The server is added to the group.

- 4 Configure your L4 switch to add this server to its virtual group.

16.4.5 Changing the Primary Cluster Server

If the current primary cluster server is down and will be down for an extended period of time, you should select another server to be the primary cluster server

- 1 In the Administration Console, click *Access Manager > Access Gateways > Groups* tab > *[Group Name] > Servers*.
- 2 Click *Primary Cluster Server*, select the IP address of a server, then click *OK*.
- 3 Click *Access Gateways > Edit > Reverse Proxy / Authentication*.
- 4 Set the *Identity Server Configuration* to *None*.
- 5 Click *OK*.
- 6 Click *Apply Changes* and apply the changes to all the servers in the group.
- 7 Click *Identity Servers > [Configuration Assignment] > Liberty*.
- 8 In the list of *Service Providers*, select the Access Gateway group name, then click *Delete*.
- 9 Click *Identity Servers > Setup > Update Servers*.
- 10 Click *Access Gateways > Edit > Reverse Proxy / Authentication*.
- 11 Set the *Identity Server Configuration* to the configuration name of the Identity Server.
- 12 Click *OK*.
- 13 Click *Apply Changes* and apply the changes to all the servers in the group.
- 14 Click *Identity Servers > Setup > Update Servers*.

16.4.6 Deleting Members from a Group

Each Access Gateway leaves the group with an identical reverse proxy, proxy service, and protected resource configuration. You need to assign them to a new group, reconfigure them so that they have unique resources they are protecting, or delete them from the Administration Console. If you placed the group behind an L4 switch, you need to reconfigure the switch so that the switch matches your reconfiguration.

16.4.7 Viewing the Health of the Group

The Group Health icon displays the status of the least healthy member of the group. Four health states are possible:

- ◆ A green status indicates that all the Access Gateways in the group are functioning has not detected any problems.
- ◆ A yellow status indicates that one or more of the Access Gateways might be functioning sub-optimally because of configuration discrepancies.
- ◆ A red status indicates that at least one Access Gateway server has an incomplete or wrong configuration.
- ◆ A white status indicates that one or more of the Access Gateway servers are not communicating with the Administration Console.

To view details about the status of the group:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Groups* tab > *[Group Name] > Health*.
- 2 To ensure that the information is current, click *Refresh*.
- 3 To view specific information about the status of an Access Gateway, click the health icon. For more information, see [Section 36.3, “Monitoring the Health of an Access Gateway,” on page 446](#).
- 4 Click *Close*.

16.4.8 Viewing Group Statistics

To view general performance statistics of the servers assigned to the selected group:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Groups* tab > *[Group Name] > Statistics*.
- 2 To determine performance, analyze the following statistics:

Column	Description
Server Name	Lists the name of the Access Gateways that belong to the group. To view additional statistical information about a specific Access Gateway, click the name of an Access Gateway.
CPU %	Displays the current CPU utilization rate. Use this statistic for capacity planning.
Cache Hit Rate %	Displays the current cache hit rate. A high cache hit rate indicates that the caching system is off-loading significant request processing from the Web server whose objects have been cached. If the percentage is low, you might want to configure a pin list. For this and other caching options, see Chapter 15, “Configuring the Cache Settings,” on page 201 .
Bytes per second to/from Server	Displays the rate at which the Access Gateway is requesting Web objects from the Web servers it is protecting.
Bytes per second to/from Browser	Displays the rate at which browser clients are requesting Web objects.
Current Connections	Displays the total number of TCP connections that are active, idle, or closing.
Statistics	Allows you to view all the statistics for a selected server. Click <i>View</i> to see these additional statistics. For more information, see Section 35.2, “Monitoring Access Gateway Statistics,” on page 429 .

- 3 Click *Close*.

16.4.9 Viewing Group Command Status

To view the command status of the Access Gateway group as a whole:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Groups* tab > *[Group Name] > Command Status*.

- 2 Analyze the information displayed about each command.

Column	Description
Name	Contains the display name of the server. Select this link to view the commands that have been recently issued to this server.
Status	Specifies the status of the command, and includes such states as Pending, Incomplete, Executing, Succeeded, and Failed.
Type	Specifies the type of command.
Admin	Specifies who issued the command, either the system or a user. If a user issued the command, the field contains the DN of the user.
Date & Time	Specifies when the command was issued. The date and time are displayed in local time.

- 3 To delete a command, select the check box for the command, then click *Delete*. The selected command is cleared.
- 4 To update the current cache of commands, click *Refresh*.
- 5 Click *Close*.

16.4.10 Viewing the Status of Group Alerts

To view information about current alerts:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Groups* tab > *[Group Name] > Alerts*.
- 2 To view general information about all alerts, analyze the data displayed in the table.

Column	Description
Server Name	Lists the name of the Access Gateway that sent the alert. To view additional information about the alerts for a specific Access Gateway, click the name of an Access Gateway.
Severe	Lists the number of critical alerts that have been sent and not acknowledged.
Warning	Lists the number of warning alerts that have been sent and not acknowledged.
Information	Lists the number of informational alerts that have been sent and not acknowledged.

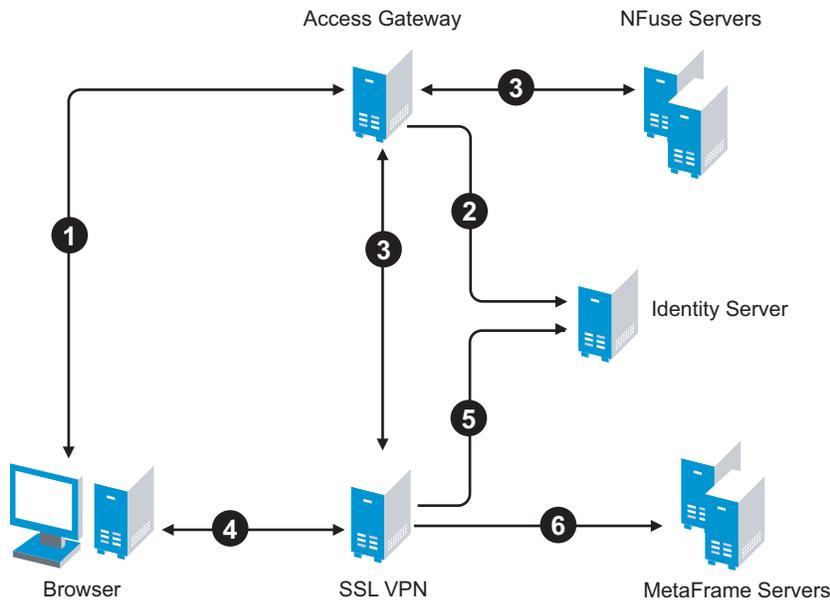
- 3 To view information about a particular alert, click the server name. For information about a specific alert, see [Section 38.2.1, “Reviewing Java Alerts,” on page 455](#).
- 4 To acknowledge an alert, select the check box for the alert, then click *Acknowledge Alert(s)*. When you acknowledge an alert, you clear the alert from the list.
- 5 Click *Close*.

Configuring Access Manager for Citrix Clients

17

The Access Manager can be configured to provide single sign-on for Citrix* clients. [Figure 17-1](#) illustrates this process for the Citrix Web client.

Figure 17-1 Citrix Client Configuration



1. The client specifies the URL of the Citrix login page that is hosted on the NFuse* server.
2. The Access Gateway, which has been configured to protect the NFuse server, requests login credentials from the Identity Server.
3. The user enters his or her login credentials. Access Gateway has been configured with a Form Fill policy so that it can auto fill the login form of the NFuse server. The user is authenticated to both the NFuse Server and to the SSL VPN server.
4. The NFuse server displays the application screen so that the user can request a MetaFrame* application.
5. The SSL VPN server verifies the user's credentials.
6. The SSL VPN contacts the MetaFrame server and delivers the application to the user.

17.1 Prerequisites

- NFuse server
- MetaFrame server
- Identity Server

- ❑ Access Gateway
- ❑ SSL VPN configured to use the same Identity Server as the Access Gateway.
- ❑ SSL VPN and the MetaFrame server need to be on the same network. The SSL VPN needs to use its private network interface adapter to communicate with the private address of the MetaFrame server.

17.2 Configuring the Access Gateway for Citrix Clients

- 1 Create a protected resource for the Citrix login page.
 - 1a In the Administration Console, click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy]*.
The reverse proxy can be set up to require SSL or not.
 - 1b Click *Name of Proxy Service > Protected Resources > New*.
 - 1c When configuring the protected resource, set up the following:
 - ♦ Select a contract that requires authentication. Usually this is a Name/Password contract, but it can be a certificate contract if your NFuse server is configured to use certificates.
 - ♦ For the URL Path List, specify the URL to the Citrix login page. This URL should include the file name of this login page.

For more information, see [Section 12.4, “Configuring Protected Resources,” on page 150](#).
- 2 Create a Form Fill policy and assign it to the protected resource for the Citrix login page.
 - 2a Click *Form Fill > Manage Policies > New*.
 - 2b Name the Citrix policy, select *Access Gateway: Form Fill* as the type, then click *OK*.
 - 2c In the *Actions* section, click *New > Form Fill*.
 - 2d In the *Form Selection* section, identify the form on the Citrix login page.
 - 2e In the *Fill Options* section, create the following:
 - ♦ Username input field
 - ♦ Password input field
 - ♦ (Optional). If your login page requires a domain, add a domain input field.
 - 2f In the *Submit Options* section, configure the following:
 - ♦ Select *Auto Submit*.
 - ♦ Select *Enable JavaScript Handling*.
 - ♦ Click *Statements to Execute on Post* and specify the following string:

```

window.open ("http://<dns-name>/sslvpn/custom-login",
"sslvpnWindow", "scrollbars=yes, resizable=yes,
HEIGHT=800,WIDTH=800");

```

Replace *<dns_name>* with the public DNS name that your users specify to log in the SSL VPN server.
 - 2g Configure any other options to match your form and your network.
For more information, see [Section 32.3.1, “Creating a Form Fill Policy,” on page 391](#).

2h In the *Actions* section, click *New > Form Login Failure*.

Specify the procedures you want followed when login fails. For more information, see [Section 32.3.2, “Creating a Login Failure Policy,” on page 395](#).

Citrix displays login failures via the query string, so you’ll need to use CGI matching.

2i Click *OK*, then click *Apply Changes*.

3 Click *Close*.

You should return to the Form Fill page for the protected resource.

4 Select the policy you just created, then click *Enable*.

5 Click *Configuration Panel*, then click *OK*.

6 On the Configuration page, click *Apply Changes*.

Virtualization on the Linux Access Gateway

18

Application virtualization, a feature available with the Linux Access Gateway is the capacity to run multiple instances of a proxy process or Virtual Machines (VM). This section contains the following information on virtualization:

- ♦ [Section 18.1, “Virtualization Overview,” on page 237](#)
- ♦ [Section 18.2, “Configuring VMs,” on page 237](#)
- ♦ [Section 18.3, “Modifying a VM,” on page 239](#)
- ♦ [Section 18.4, “Deleting a VM,” on page 240](#)
- ♦ [Section 18.5, “Setting Maximum Memory Limit,” on page 240](#)
- ♦ [Section 18.6, “Scan,” on page 241](#)

18.1 Virtualization Overview

VMs lead to higher utilization of resources as the multiple instances of proxy VMs share the available memory, CPU, and the disk space based on the configuration. Each VM can run its own tasks in parallel and each VM can address its own 4 GB address space. Virtualization is required in one of the following situations:

- ♦ An Access Gateway machine with 4 GB RAM and Dual Core or SMP machines on which you have enabled SSL. For more information on configuring SSL, see [Chapter 13, “Configuring the Access Gateway for SSL,” on page 171](#).
- ♦ Systems with more than 4 GB RAM and more than one SCSI disk.

Only one VM instance is visible to the components external to the proxy. A default VM is created at the time of installation, which acts as the master. The master VM listens on a socket port, which distributes the incoming connections to the VM instances. The master VM cannot be deleted.

A browser can connect to any of the VMs depending on the IP address of the client. After authentication, the client is forced to use the same VM. The current implementation of VMs does not have the load balancing capabilities.

18.2 Configuring VMs

You can add, delete, choose partitions, CPU, and system memory percentage for VMs using the Command Line Interface (CLI). The configuration changes are stored in the `.current/virtual.xml` file. Each VM you create must be assigned at least one COS partition. The following sections describe the tasks associated with creating a VM:

- ♦ [Section 18.2.1, “Creating a Novell COS Partition,” on page 238](#)
- ♦ [Section 18.2.2, “Configuring a VM,” on page 238](#)

18.2.1 Creating a Novell COS Partition

- 1 In the *Choose* section of the Suggested Partitioning page, select *Create Custom Partition Setup*.
- 2 In the *Choose a Hard Disk* section, select *Custom Partitioning*, then click *Next* to display the Expert Partitioner page.
- 3 On the Expert Partitioner page, select *Create*.
- 4 At the prompt, select *Primary Partition* to display the *Create Primary Partition* section.
- 5 In the *Format* section, click *Do Not Format*.
- 6 From the drop-down list, select *0x68 Novell COS*.
- 7 Make sure you do not provide any value for *Mount Point*.
- 8 Click *OK*.
- 9 Repeat the above steps to create multiple COS partitions.

18.2.2 Configuring a VM

- 1 To enter nash shell, enter the following command from the command prompt:

```
nash
```

- 2 To enter the configuration mode, enter the following command:

```
configure .current
```

- 3 To configure a VM, enter the following command:

```
vm vmlist add <profile name>
```

Replace <profile name> with a name.

- 4 To configure memory for VM, enter the following command:

```
memory <percentage of memory to allocate>
```

Replace <percentage of memory to allocate> with a value greater than or equal to 10. The default value is 10%.

The sum of memory percentages of all the VMs cannot exceed the total memory limit set. For more information on setting total memory limit, see [Section 18.5, “Setting Maximum Memory Limit,” on page 240](#).

- 5 To configure CPU for VM, enter the following command:

```
cpu <number>
```

Replace <number> with a value. The default value is 0.

- 6 To set partition, enter the following command:

```
partition <path>
```

Replace <path> with path of COS partition.

A COS partition cannot be shared with another VM. To configure partition for a particular VM, enter a unique COS partition which has not been assigned to another VM.

- 7 To view configuration for a particular VM, enter the following command:

```
vm vmlist show <profile name>
```

- 8 To view the configuration for multiple VMs, enter the following command:

```
vm vmlist show
```

- 9 To exit the add subshell, enter the following command:

```
exit
```

- 10 To exit the vmlist subshell, enter the following command:

```
exit
```

- 11 To save the VM configuration, enter the following command:

```
save
```

The configuration is saved to the `virtual.xml` file in the `.current` directory

18.3 Modifying a VM

To modify the configuration of an existing VM:

- 1 To enter the nash shell, enter the following command from the command prompt:

```
nash
```

- 2 To enter the configuration mode, enter the following command:

```
configure .current
```

- 3 To edit an existing VM, enter the following command:

```
vm vmlist edit <profile name>
```

- 4 To view the VM configuration, enter the following command:

```
show
```

- 5 To edit CPU for VM, enter the following command:

```
cpu <number indicating cpu>
```

The entered value replaces the old value.

- 6 To edit the memory for the VM, enter the following command:

```
memory <percentage of memory to allocate>
```

- 7 To edit the partition for the VM, enter the following command:

```
partition <path of cos partition>
```

The entered COS partition appends to the list of COS partitions assigned to the VM only if the entered COS partition is not in use by another VM.

- 8 To remove partition for a particular VM, enter the following command:

```
no partition <path of cos partition>
```

- 9 To exit the add subshell, enter the following command:

```
exit
```

- 10 To exit the vmlist subshell, enter the following command:

```
exit
```

- 11 To save the VM configuration, enter the following command:

```
save
```

The configuration is saved to the `virtual.xml` file in the `.current` directory

18.4 Deleting a VM

If there is only one VM existing, you cannot delete it. It is the master VM.

To delete a secondary VM:

- 1 To enter the nash shell, enter the following command from the command prompt:

```
nash
```

- 2 To enter the configuration mode, enter the following command:

```
configure .current
```

The configure changes are stored in the `config.xml` file in the `.current` directory.

- 3 To delete a secondary VM, enter the following command:

```
vm vmlist delete <profile name>
```

- 4 To exit the add subshell, enter the following command:

```
exit
```

- 5 To exit the vmlist subshell, enter the following command:

```
exit
```

- 6 To save the VM configuration, enter the following command:

```
save
```

The configuration is saved to the `virtual.xml` file in the `.current` directory.

18.5 Setting Maximum Memory Limit

The `set maximum limit` parameter sets the maximum amount of memory that can be used by all the VMs. The memory limit for each VM is set using a percentage of the maximum memory limit specified. The default maximum memory limit is 70 per cent.

To set the maximum memory percentage limit for sum of memory percentage of all individual VMs enter the following command:

- 1 To enter nash shell, enter the following command from the command prompt:

```
nash
```

- 2 To enter the configuration mode, enter the following command:

```
configure .current
```

- 3 To set the maximum memory percentage limit for sum of memory percentage of all individual VMs, enter the following command:

```
vm set-vm-mem <percentage of memory>
```

NOTE: An error message is displayed when the sum of memory percentages of the VMs exceeds the set maximum limit.

- 4 To save the configuration, enter the following command:

```
save
```

The configuration is saved to the `virtual.xml` file in the `.current` directory.

18.6 Scan

You can delete all VMs and create a single VM with the following configuration:

- ♦ Memory percentage of 20
- ♦ Assigned to first COS partition in the partition table, if available
- ♦ CPU assigned to 0

1 To enter nash shell, enter the following command from the command prompt:

```
nash
```

2 To enter the configuration mode, enter the following command:

```
configure .current
```

3 To delete all VMs and create a single VM with the following configuration enter the following command:

```
vm scan
```

NOTE: This command is usually used during installation to populate `virtual.xml` with the appropriate values.

4 To save the configuration, enter the following command:

```
save
```

The configuration is saved to the `virtual.xml` file in the `.current` directory.

Configuring the SSL VPN Gateway

IV

This section describes how you can configure and manage the Novell® SSL VPN Gateway. These sections assume that you have already done the following:

- ◆ Installed the Access Gateway. For more information in installing Access Gateway, see *Novell Access Manager 3.0 Installation Guide*.
- ◆ You are logged in to the Administration Console as the admin user for the Access Manager and that you have expanded the Novell Access Manager task.

This section has the following information:

- ◆ [Chapter 19, “Overview of SSL VPN,” on page 245](#)
- ◆ [Chapter 20, “Managing SSL VPN Servers,” on page 247](#)
- ◆ [Chapter 21, “Configuring the SSL VPN Servers,” on page 251](#)
- ◆ [Chapter 22, “Load Balancing and Fault Tolerance,” on page 269](#)
- ◆ [Chapter 23, “Accessing the SSL VPN User Portal,” on page 273](#)
- ◆ [Chapter 24, “Understanding the SSL VPN User Interface,” on page 279](#)

SSL VPN is a server that uses encryption and other security mechanisms to ensure that only authorized users can access the network and the data cannot be intercepted. This functionality uses secure socket layers (SSL) as the underlying security protocol for network transmissions.

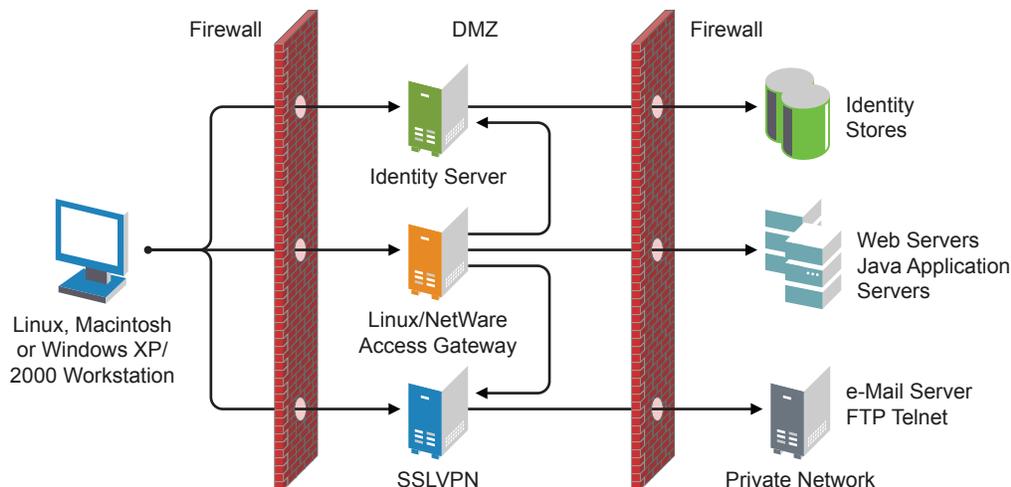
SSL VPN allows authorized access to applications and services that are behind a firewall. It also provides secure access to HTTP and non-HTTP based applications and performs single sign-on when authenticated to the Access Gateway.

- ◆ [Section 19.1, “Understanding How SSL VPN Works,” on page 245](#)
- ◆ [Section 19.2, “SSL VPN Servers,” on page 246](#)

19.1 Understanding How SSL VPN Works

SSL VPN server comprises of SSL VPN servlet and the SSL VPN Server module. The following figure explains how the SSL VPN works:

Figure 19-1 SSL VPN Functionality



- ◆ The user accesses the SSL VPN URL.
- ◆ Proxy accelerates the SSL VPN server.
- ◆ Proxy redirects users to NIDP. Users are prompted to authenticate.
- ◆ The NIDP server verifies the username and password with the directory service.
- ◆ After identity validation, the Access Gateway retrieves the user's common name and role.
- ◆ The identity information is carried forward to the SSL VPN server, allowing the user to access SSL VPN server if the proxy gets the role and passes it to SSL VPN server.
- ◆ After authentication, SSL VPN server retrieves policies for the user.
- ◆ The client components are carried forward to client desktop through Java* applet and ActiveX*, along with the policies and the required client components.

- ◆ Applet installs the components and establishes SSL VPN tunnel between SSL VPN client and SSL VPN server for the networks chosen, based on the policies for this particular role.
- ◆ Through the tunnel, access to TCP and UDP applications in the protected networks is made possible if the policies from the SSL VPN server match.
- ◆ When the SSL VPN session closes, all the client components are automatically uninstalled from the workstation.

19.2 SSL VPN Servers

When a user sends a connection request to the SSL VPN server through a browser, Access Gateway recognizes it as an SSL VPN request and redirects it to the Identity Server for authentication. Once the user is successfully authenticated to the Identity server, the information is sent to the Access Gateway. The Access Gateway forwards the request to the SSL VPN servlet, which handles the request.

While processing a new connection request, the servlet interfaces with the SSL VPN server module to get the client-side verification information and policies that are applicable for the user. The SSL VPN server adds the new connection information to its connection database. The servlet pushes the Java applet or ActiveX to the browser along with the policy and the verification information. The SSL VPN connection is then established.

The servlet communicates with the SSL VPN server over TCP port 2010. The SSL VPN server module can run either on the same server or on another server. The server module is installed by default on the server in which the SSL VPN servlet is running. The administrator can configure several SSL VPN servers on different servers to form a server failover group which would enable Load Balancing and Fault Tolerance. For more information on Load Balancing and Fault Tolerance, see [Section 22.2, "Load Balancing," on page 270](#).

Managing SSL VPN Servers

20

SSL VPN servers are auto-imported into the Administration Console during installation. You can use the SSL VPNs page in Administration Console to view information about the current status of all SSL VPN servers and to configure the SSL VPN servers.

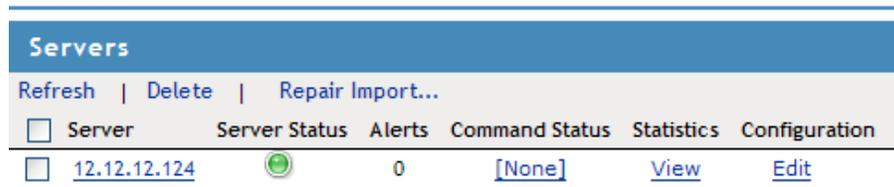
- ◆ [Section 20.1, “Viewing SSL VPN Servers,” on page 247](#)
- ◆ [Section 20.2, “Viewing SSL VPN Server Details,” on page 248](#)
- ◆ [Section 20.3, “Monitoring Alerts,” on page 250](#)

20.1 Viewing SSL VPN Servers

To view the SSL VPN page:

- 1 In Administration Console, click *Access Manager > SSL VPNs*.

SSL VPNs



Servers						
<input type="checkbox"/>	Server	Server Status	Alerts	Command Status	Statistics	Configuration
<input type="checkbox"/>	12.12.12.124		0	[None]	View	Edit

The following server information is displayed:

- ◆ **Server:** Displays a list of servers added to Administration Console. Click the link of a particular server to view or modify its configuration. For more information, see [Section 20.2, “Viewing SSL VPN Server Details,” on page 248](#)
- ◆ **Server Status:** Indicates whether the server is functional. Click the icon to view additional information about the operational status of the server. For more information, see [Section 36.4, “Monitoring the Health of an SSL VPN Server,” on page 448](#).
- ◆ **Alerts:** Indicates whether any alerts have been sent. This option is not available to you if the alert count is 0. For more information, see [Section 20.3, “Monitoring Alerts,” on page 250](#).
- ◆ **Command Status:** Indicates the status of commands issued to all servers. For more information, see [Section 37.3, “Viewing Command Status of the SSL VPN Server,” on page 453](#).
- ◆ **Statistics:** Indicates the number of active client connections and the time when the server was started. Click *View* to get the statistics information. [Section 35.3, “Viewing SSL VPN Statistics,” on page 438](#).
- ◆ **Configuration:** Click *Edit* in the *Configuration* column of SSL VPNs page to view and modify the configuration of the SSL VPN server. This page specifies the date and time when the last modification was made and lists the full distinguished name of the user who made the last modification. For more information, see [Chapter 21, “Configuring the SSL VPN Servers,” on page 251](#).

- 2 To refresh a server, select the check box next to the Gateway that you want to refresh, then click *Refresh*.
- 3 To delete a server, select the check box next to the Gateway that you want to delete, then click *Delete*.
- 4 To repair an import of a server you have recently installed that does not appear in the list. Give the system at least 10 minutes after an install, then click *Refresh*. If the server still does not appear, click *Repair Import*, specify the IP address of the SSL VPN gateway that is not appearing in the list, then click *OK*.
- 5 To shut down a server, click *Shutdown*. Click *OK* in the Confirmation dialog box.
- 6 To restart the server, click *Restart*. Click *OK* in the Confirmation dialog box.

20.2 Viewing SSL VPN Server Details

To edit the Gateway information:

- 1 In Administration Console, click *Access Manager > SSL VPNs > [Server Name]*.

The Server Details page is displayed.

[Servers](#) ▶ **General**

Server Details: 12.12.12.124

General | Health | Alerts | Command Status | Statistics

Shutdown | Restart | Edit

Name: [12.12.12.124](#)

Management IP Address: [12.12.12.124](#) Port: [1443](#)

Location:

Server Version: SSL VPN 3.0.1

Description:

Close

The *General* tab of Server Details page displays information such as name, Management IP address, Port, Location, and the server version of the selected server.

2 Click *Edit*. The Server Details Edit page is displayed.

Servers ▸ General ▸ Edit

Server Details Edit: 12.12.12.123

Name:

Management IP Address: Port:

Location:

Description:

You can edit the information in the following fields:

- ◆ **Name:** Specify the IP address of the server. This field is mandatory.
- ◆ **Management IP Address:** Specify the IP address used to manage the server. If the system on which the agent is installed has multiple IP addresses, you can select one from the drop-down list.
- ◆ **Port:** Specify the port used for management. This field is mandatory.
- ◆ **Location:** Specify the location of the SSL VPN server.
- ◆ **Description:** (Optional) You can provide a brief description of the purpose of this SSL VPN Gateway or any other relevant information.

3 Click *OK* to save changes or *Cancel* to discard the changes.

20.3 Monitoring Alerts

The *Alerts* page allows you to view information about current system alerts and to clear them. An alert is generated whenever the SSL VPN Gateway detects a condition that prevents it from performing normal system services.

- 1 In Administration Console, click *Access Manager > SSL VPNs > [Server Name] > Health*.

Servers ▶ Alerts

Server Alert Detail: 10.10.12.123

General	Health	Alerts	Command Status	Statistics
Acknowledge Alert(s)				
<input type="checkbox"/>	Severity	Date & Time	Message	
<input type="checkbox"/>	Information	Aug 16, 2006 3:09 PM	SSLVPN Servlet is registered	
<input type="checkbox"/>	Information	Aug 16, 2006 5:46 PM	VCC Started	
<input type="checkbox"/>	Information	Aug 16, 2006 5:47 PM	SSLVPN Servlet is registered	
<input type="checkbox"/>	Information	Aug 17, 2006 4:19 PM	VCC Started	
<input type="checkbox"/>	Information	Aug 17, 2006 4:20 PM	SSLVPN Servlet is registered	
<input type="checkbox"/>	Information	Aug 17, 2006 6:27 PM	VCC Started	
<input type="checkbox"/>	Information	Aug 17, 2006 6:28 PM	SSLVPN Servlet is registered	
<input type="checkbox"/>	Information	Aug 18, 2006 2:43 PM	SSLVPN Servlet is registered	
<input type="checkbox"/>	Information	Aug 21, 2006 4:44 PM	SSLVPN Servlet is registered	
<input type="checkbox"/>	Information	Aug 21, 2006 5:29 PM	SSLVPN Servlet is registered	
<input type="button" value="Close"/>				

The following information is displayed:

- ♦ **Severity:** Describes the type of alert. An alert can be informational, critical, or a warning.
 - ♦ **Date & Time:** Indicates the date and time when an alert was issued. The date and time are given in the local time.
 - ♦ **Message:** Displays the message that was sent with the alert. This information is optional.
- 2 To send an acknowledgement, select the check box next to the alert, then click *Acknowledge Alert(s)*. When you acknowledge an alert, the alert is cleared from the list.
 - 3 Click *Close* to close the Alerts page.

Configuring the SSL VPN Servers

21

The configuration page in the Administration Console allows you to configure the SSL VPN server. The page specifies the date and time the last modification was made and lists the full distinguished name of the user who made the last modification.

- 1 In Administration Console, click *Access Manager > SSL VPNs > [Server Name]*.
- 2 Click *Edit* in the Configuration column of Server.

The Server configuration page is displayed.

Basic Gateway Configuration	Last Changed	Change By
Gateway Configuration	Jun 2, 2006 8:52 AM	cn=admin,o=novell
DNS Servers List	Jun 5, 2006 11:35 PM	cn=admin,o=novell
Policies	Last Changed	Change By
Traffic Policies	Jun 2, 2006 8:52 AM	cn=admin,o=novell
Client Integrity Check Policies	Jun 2, 2006 8:32 AM	cn=admin,o=novell
Novell Audit and Alerts	Last Changed	Change By
Novell Audit Settings	Jun 2, 2006 8:52 AM	cn=admin,o=novell
Alerts Settings	Jun 2, 2006 8:52 AM	cn=admin,o=novell
Security Settings	Last Changed	Change By
SSL VPN Certificates		

You can configure the following fields:

Basic Gateway Configuration: Use this section to configure Gateway and DNS Server List. For more information, see [Section 21.1, “Viewing and Modifying Server Configuration,” on page 252](#) and [Section 21.2, “Configuring the DNS Server List,” on page 253](#).

Policies: Use this section to configure traffic policies and client integrity check policies. For more information, see [Section 21.3, “Viewing and Modifying Traffic Policies,” on page 254](#) and [Section 21.6, “Creating New Client Integrity Check Policies,” on page 260](#).

Novell Audit and Alerts: Use this section to configure Novell® Audit Settings and Alerts Settings. For more information see, [Section 33.4, “Enabling SSL VPN Audit Events,” on page 406](#) and [Section 21.7, “Configuring Alerts Settings,” on page 261](#).

Security Settings: This section contains certificate details for SSL VPN. For more information, see [Section 21.8, “Adding Certificates to the SSL VPN Keystore,” on page 262](#).

All configuration changes must be applied from the Configuration page. The links from this page allow you to OK or Cancel any changes, but the changes are not sent to the SSL VPN from those pages.

- 3 If you have applied, scheduled, or canceled changes, click *Refresh* to refresh the current display. This option is available when no changes are pending.
- 4 To apply changes to the SSL VPN server, click *Apply Changes*.
- 5 To select when the changes should be sent to SSL VPN, click *Schedule Changes*.
- 6 To cancel all configured changes click *Cancel*. If you have configured multiple changes and need to cancel only some of the changes, use the *Cancel Changes* link by the specific section to cancel individual changes.

The following sections describe SSL VPN configuration in detail:

- ♦ [Section 21.1, “Viewing and Modifying Server Configuration,” on page 252](#)
- ♦ [Section 21.2, “Configuring the DNS Server List,” on page 253](#)

- ◆ Section 21.3, “Viewing and Modifying Traffic Policies,” on page 254
- ◆ Section 21.4, “Configuring Full Tunneling,” on page 256
- ◆ Section 21.5, “Viewing Client Integrity Check Policies,” on page 257
- ◆ Section 21.6, “Creating New Client Integrity Check Policies,” on page 260
- ◆ Section 21.7, “Configuring Alerts Settings,” on page 261
- ◆ Section 21.8, “Adding Certificates to the SSL VPN Keystore,” on page 262
- ◆ Section 21.9, “Adding Trusted Roots for SSL VPN,” on page 263
- ◆ Section 21.10, “Configuring SSL VPN to Download Applet on Internet Explorer,” on page 263
- ◆ Section 21.11, “Accelerating SSL VPN,” on page 264
- ◆ Section 21.12, “Moving SSL VPN Server to a Different Administration Console,” on page 267
- ◆ Section 21.13, “Accelerating SSL VPN with a New Access Gateway,” on page 268

21.1 Viewing and Modifying Server Configuration

The Gateway Configuration page displays the current configuration of the SSL VPN server, such as the external IP address if the SSL VPN server is behind NAT, listening IP address, TCP encryption port, connection manager port and the type of encryption used. You can modify the server configuration as follows:

- 1 In Administration Console, click *Access Manager > SSL VPNs > Edit*.

The Server configuration page is displayed.

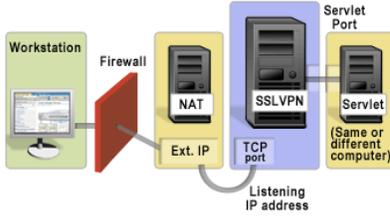
- 2 Select *Gateway Configuration* from the *Basic Gateway Configuration* section.

The SSL VPN Gateway Basic Configuration page is displayed.

Servers > Configuration > Basic Configuration

SSL VPN Gateway Basic Configuration: 12.12.12.123

<input checked="" type="checkbox"/> Behind NAT	
External IP Address:	<input type="text" value="11.11.11.123"/>
Listening IP Address:	<input type="text" value="12.12.12.123"/>
Private Address(es):	<input type="text" value="10.10.10.1"/>
TCP Encryption Port:	<input type="text" value="7777"/>
Connection Manager Port:	<input type="text" value="2010"/>
Inactivity Timeout (Minutes):	<input type="text" value="30"/>
Encryption:	<input type="text" value="AES256"/>
Debug Level:	<input type="text" value="On"/>



Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

Fill in the following fields:

- ◆ **Behind NAT:** Specify whether the SSL VPN Gateway is behind NAT (Network Address Translation).
- ◆ **External IP Address:** When the *Behind NAT* check box is selected, you are prompted to enter the external address. Specify the IP address by which the external user on the Internet must be able to access the SSL VPN server.

- ♦ **Listening IP address:** Specify the IP address that SSL VPN listens on.
- ♦ **Private Address(es):** Specifies the IP address of the private interface of the network card. If you have multiple private networks, specify the private IP addresses of the servers separated by a comma.
- ♦ **TCP Encryption Port:** The port to encrypt the TCP traffic through STunnel. The default encryption port is 7777.
- ♦ **Connection Manager Port:** The port on which the connection manager listens to. The default port number is 2010.

If you change the connection manager port from 2010 to any other port, do the following:

- a. Open `config.txt`, located at `/var/opt/novell/tomcat4/webapps/sslvpn/WEB-INF` and change the port number found in the first line of the file to the new number.
- b. Restart the SSL VPN server manually by entering the following commands:


```
/etc/init.d/novell-sslvpn stop
/etc/init.d/novell-sslvpn start
```
- c. Restart Tomcat manually by entering the following commands:


```
/etc/init.d/novell-tomcat4 stop
/etc/init.d/novell-tomcat4 start
```

- ♦ **Inactivity Timeout (Minutes):** Configure the time in minutes after which an idle connection should be closed. If no data exchange takes place during the stipulated time, the connection is closed. An inactive connection is closed after a stipulated time so that the resources are freed to allow additional incoming connections. The inactivity timeout period can be one minute to 120 minutes. The default inactive timeout period is 30 minutes.
 - ♦ **Encryption:** Select the type of encryption. It can be either AES 128 or AES 256.
 - ♦ **Debug Level:** Specifies whether you want to receive debug logs or not. The default option is *Off*. If you choose to receive the debug logs, you can set it to *On*.
- 3** To save your modifications, click *OK*, then click *Apply Changes* or *Schedule Changes* on the Configuration page.

21.2 Configuring the DNS Server List

This section contains a list of configured DNS servers. The servers listed here are pushed to the client from the SSL VPN server during the connection.

- 1** In the Administration Console, click *Access Manager > SSL VPNs > Edit*.
The Server configuration page is displayed.
- 2** Select *DNS Server List* from the *Basic Gateway Configuration* section.

The DNS server list page is displayed.

DNS Servers
New... | Delete
 DNS Servers
 10.1.1.10

Domains
New... | Delete
 Search Domains
 example.com

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

OK Cancel

- 3 Click *New* in the *DNS server* section to add new DNS servers.
- 4 Click *New* in the *Domains* section to add new Domain names.
- 5 To delete a DNS server or a domain, select the check box next to the field and click *Delete* in the section.
- 6 To save your modifications, click *OK*, then click *Apply Changes* or *Schedules Changes* on the Configuration page.

21.3 Viewing and Modifying Traffic Policies

Traffic policies are a set of rules and regulations, administered to regulate user access to the protected network resources. Novell SSL VPN traffic policies are role-based policies. The access to the protected network is restricted based on the role to which the user belongs.

NOTE: You can configure a maximum of 32 traffic rules per role. If you have configured multiple traffic policies, the policies are prioritized based on the order of their creation.

You can configure a different set of traffic policies for different roles as follows:

- 1 In the Administration Console, click *Access Manager > SSL VPNs > Edit*.
- 2 Select *Traffic Policies* from the *Policies* section. The SSL VPN Traffic Policies page is displayed.

Servers > Configuration > Traffic Policies

SSL VPN Traffic Policies: 12.12.12.123 ?

List of Traffic Policies							
Policy Name	Enabled	Role	Dst. Network	Protocol	Application	Port	Action
<input type="checkbox"/> Any Role TCP Modify Network	✓	Any	0.0.0.0/255.0.0.0	TCP	AnyTCP	0	Encrypt
<input type="checkbox"/> Any Role UDP Modify Network	✓	Any	0.0.0.0/255.0.0.0	UDP	AnyUDP	0	Encrypt

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

OK Cancel

The traffic policies page contains the following columns:

Column	Description
Policy Name	Specifies the name of the traffic rule.
Enabled	Specifies the status of the traffic rule.
Role	Specifies the role to which the traffic rule applies.
Dst. Network	Specifies the host or network IP address of the destination.
Protocol	Specifies if the protocol is TCP or UDP or Any.
Application	Specifies the name of the application.
Port	Specifies the port number on which the service you select listens.
Action	Specifies if a service can be encrypted or denied.

- 3 To create a new traffic policy, click *New*. The New dialog box is displayed.
- 4 Specify the traffic policy name in the *Traffic Policy Name* field, then click *OK*.
- 5 Click the newly added traffic policy. The Edit Traffic Policy page is displayed.

[Servers](#) ▶ [Configuration](#) ▶ [Traffic Policies](#) ▶ [Edit](#)

Edit Traffic Policy: 12.12.12.123

Traffic Policy : "my_policy"

Policy Name

Scope of Policy

Role	<input type="text" value="Any"/> 
Destination Network	<input type="text" value="10.0.0.0"/>
Network Mask	<input type="text" value="255.0.0.0"/>
Predefined Applications	<input type="text" value="Any"/>
Name	<input type="text" value="Any"/>
Protocol	<input type="text" value="ANY"/>
Port	<input type="text" value="0"/>

Action

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

Fill in the following fields:

- ♦ **Policy Name:** Specifies the name for the traffic policy.
- ♦ **Role:** The role to which the traffic rule applies. Select the role from the drop-down list. If the role is not listed, click the role icon to add new roles.
- ♦ **Destination Network:** Specify the host or network IP address of the destination.

- ♦ **Network Mask:** The network mask is displayed by default when you specify the destination address. However, you can edit the mask.
- ♦ **Predefined Application:** Select a predefined application from the drop-down list.
- ♦ **Name:** Specify a name for the application. This information is optional.
- ♦ **Protocol:** Select a protocol from the drop-down list. You can select the protocol to be TCP, UDP, or Any.
- ♦ **Port:** Specify the port number on which the service you select must listen.

NOTE: Specify 0 to allow all ports depending on the protocol.

- ♦ **Action:** Specify if a service can be allowed or denied. Select *Encrypt* to allow the service in encrypted form. Select *Deny* if you do not want to allow the service.

- 6 To delete a traffic rule, select the rule that, then click *Delete*.
- 7 To enable a traffic rule, select the rule that, then click *Enable*.
- 8 To disable a traffic rule, select the rule that, then click *Disable*.
- 9 To save your modifications, click *OK*, then click *Apply Changes* or *Schedule Changes* on the Configuration page.

21.4 Configuring Full Tunneling

Novell SSL VPN is configured for split tunneling by default. When SSL VPN is configured for split tunneling, only that traffic which is destined for the protected network goes through the VPN tunnel. But, if you want all traffic in the client machine to go through the tunnel (full tunneling), after connecting to SSL VPN, do the following:

- 1 In the Administration Console, click *Access Manager > SSL VPNs > Edit*.
- 2 Create a new traffic policy. For more information on adding new traffic policy, see [Section 21.3, “Viewing and Modifying Traffic Policies,” on page 254](#).
- 3 Click the newly added traffic policy. The Edit Traffic Policy page is displayed.

Configure the following fields:

- ♦ **Destination Network:** Specify 0.0.0.0 as the destination network IP address.
- ♦ **Protocol:** Select *Any* as the protocol.
- ♦ **Port:** Specify the port number as 0.
- ♦ **Action:** Select *Encrypt* to allow the service in encrypted form.

Leave the default values in the other fields unchanged.

- 4 Click *OK* to save changes.
- 5 In the Edit page, select *Gateway Configuration* from the *Basic Gateway Configuration* section. The SSL VPN Gateway Basic Configuration page is displayed.
- 6 In the *Private IP Address(es)* field, specify all the IP address using which the SSL VPN server can access the public resources.
- 7 To save your modifications, click *OK*, then click *Apply Changes* or *Schedule Changes* on the Configuration page.

21.5 Viewing Client Integrity Check Policies

Novell SSL VPN runs client integrity check on the client workstations before establishing a tunnel to the SSL VPN gateway. The check ensures that the users have specified software installed and running in their systems.

You can configure the Client Integrity Check policy to check for application categories such as Firewall, Antivirus and Mail clients depending on your requirements.

To configure the Client Integrity Check policy do the following:

- 1. Configure Category:** A category is a group of similar software. For example, a firewall category, can contain a list of firewall such as Windows Firewall, and Zone alarm firewall. You can configure multiple software categories in the client integrity check policy. For more information on configuring category, see [“Configuring Category” on page 257](#). The client workstation is checked to see if the software specified under these category are installed in the workstation, before the SSL VPN connection is established.
- 2. Configure Applications Names for a Category:** After you have created a category, you must add application names to that category. Application name is the name of the software configured under a particular category. You can add more than one software under a category. A client workstation is checked for the presence of any one of the software in the category. If none of the software specified in the category is present, then the Client Integrity Check fails and the tunnel to the SSL VPN gateway is not established. For more information on configuring applications names for a category, see [Section 21.5.2, “Configuring Application Names,” on page 258](#).
- 3. Configure Application Details:** After you have added an application to a category, you have to configure the attributes of that particular application. For more information on application attributes and to know how to configure them, see [Section 21.5.3, “Configuring Application Attributes and Details,” on page 259](#).

21.5.1 Configuring Category

- 1** In Administration Console, click *Access Manager > SSL VPNs > Edit*.
- 2** Select *Client Integrity Check Policies* from the *Policies* section. The Client Integrity Check Policies page is displayed.

Servers ▶ Configuration ▶ Client Integrity Check

Client Integrity Check Policies: 12.12.12.123

For Operating System: Windows	
New... Delete Enable Disable	
<input type="checkbox"/> Category Name	Enabled
<input type="checkbox"/> Firewall_Windows	✓
<input type="checkbox"/> Antivirus_Windows	✓

OK Cancel

3 You can perform the following action:

- ♦ **For Operating System:** Select an operating system from the drop-down list.
- ♦ **New:** To enter a new software category, click *New*, enter a *Category Name* and an *Application Name* and click *OK*.
- ♦ **Delete:** To delete a category, select the category name that you want to delete, then click *Delete*.
- ♦ **Enable:** To enable a software category name, select the check box next to category name, then click *Enable*.
- ♦ **Disable:** To disable a software category name, select the check box next to category name, then click *Disable*.

21.5.2 Configuring Application Names

Application name is the name of the software configured under a particular category. You can add more than one software under a category. A client workstation is checked for the presence of any one of the software in the category. If none of the software specified in the category is present, then the Client Integrity Check fails and the tunnel to the SSL VPN gateway is not established.

To configure applications listed under each category:

- 1 In Administration Console, click *Access Manager > SSL VPNs > Edit*.
- 2 Select *Client Integrity Check Policies* from the *Policies* section. The Client Integrity Check Policies page is displayed.
- 3 Click a category to add applications to it. The Client Integrity Check - Category page is displayed.

Servers ▶ Configuration ▶ Client Integrity Check ▶ Application Category

Client Integrity Check - Category: 12.12.12.123

Operating System	: Windows
Category	<input type="text" value="Firewall_Windows"/>
Applications under this category	
New... Delete Enable Disable	
<input type="checkbox"/>	Application Name Enabled
<input type="checkbox"/>	Zone Alarm Personal Firewall 6.0.631.003 ✓
Changes made on this panel must be applied or scheduled from the Configuration Panel.	
<input type="button" value="OK"/>	<input type="button" value="Cancel"/>

You can perform the following actions in this page:

- ♦ **New:** To add a new application to the category, click *New*, add an application name, then click *OK*.
- ♦ **Delete:** To delete an application, select the application that you want to delete, then click *Delete*.
- ♦ **Enable:** To enable an application, select the check box next to application, then click *Enable*.

- ♦ **Disable:** To disable an application, select the check box next to application, then click *Disable*.

21.5.3 Configuring Application Attributes and Details

- 1 In Administration Console, click *Access Manager > SSL VPNs > Edit*.
- 2 Select *Client Integrity Check Policies* from the *Policies* section. The Client Integrity Check Policies page is displayed.
- 3 Click a category to add applications to it. The Client Integrity Check - Category page is displayed.
- 4 Click an application to add application details and attributes to it. The Application Details and Attributes page is displayed.

Servers ▶ Configuration ▶ Client Integrity Check ▶ Application Category ▶ Application Details

Application Attributes and Details: 12.12.12.123

Operating System	: Windows
Category	: Firewall_Windows
Application	Zone Alarm Personal Firewall

Definition of the Application		
New... Delete		
<input type="checkbox"/> Attribute Type	Attribute Name	Attribute Value
	Name	zclient234.exe
	Version	6.0.631.003
<input type="checkbox"/> Process	RegistryKey	HKEY_LOCAL_MACHINE\SOFTWARE\Zone Labs\ZoneAlarm
	RegistryKeyValue	InstallDirectory

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

OK Cancel

This page specifies the operating system, application category and name and details of the application. The following application details are listed in this page:

- ♦ **Attribute Type:** Specifies whether the attribute is a Process, Package, AbsoluteFile, Registry Key, or an RPM, based on the type of operating system you select.
 - ♦ **Attribute Name:** Specifies attribute names for different attribute types. For more information, see [“Application Attribute Name” on page 259](#).
 - ♦ **Attribute Value:** Specifies the value of each attribute name.
- 5 You can perform the following actions in this page:
 - ♦ **New:** To add a new application to the category, click *New*, add an application name, then click *OK*.
 - ♦ **Delete:** To delete an application, select the application that you want to delete, then click *Delete*.

Application Attribute Name

The following table lists the attributes for applications on different Operating System:

Operating System	Attribute Type	Attribute Name
Linux	RPM	Name: Specifies the name of the RPM. Version: Specifies the RPM version.
	Process	Name: Specifies the name of the process. Owner: Specifies the owner of the process.
	Absolute File	Name: Specifies the Name and absolute path of the file.
Windows	Process	Name: Specifies the name of the executable file if the application is a process, Version: Specifies the software version. RegistryKey: Specifies the registry key path. RegistryKeyValue: Specifies the registry key value. The value data found in this key value should be the absolute path of the folder where the process file is present.
	RegistryKey	Name: Specifies the name of the RegistryKey.
	Absolute File	Name: Specifies the name of the absolute path of the file name. Version: Specifies the owner of the process.
	Package	Name: Specifies the name of the software package. Version Specifies the version of the software package
Macintosh	Process	Name: Specifies the name of the process Owner: Specifies the owner of the process.
	Absolute File	Name: Specifies the name of the executable file if the application is a process,

21.6 Creating New Client Integrity Check Policies

As an administrative user, you can configure Client Integrity Check policies as follows:

- 1 In Administration Console, click *Access Manager > SSL VPNs > Edit*.
- 2 Select *Client Integrity Check Policies* from the *Policies* section. The Client Integrity Check Policies page is displayed.
- 3 Select the operating system from the *Operating System* drop-down list.
- 4 Click *New* to enter a new software category. New dialog box is displayed.
- 5 Enter a *Category Name* and an *Application Name* and click *OK*.
- 6 Click the newly added category to add applications to it. The Client Integrity Check - Category page is displayed.
- 7 Click *New* to add a new application to the category. The new dialog box is displayed.
- 8 Add an application name, then click *OK*.

- 9 Click an application to add application details and attributes to it. The Application Details and Attributes page is displayed.
- 10 Click *New* to add an attribute to the application, add an attribute name, then click *OK*.
- 11 Click *OK*, to save your modifications. Click *Apply Changes* or *Schedule Changes* on the Configuration page.

21.7 Configuring Alerts Settings

To configure the server to receive alerts when a particular event occurs:

- 1 In the Administration Console, click *Access Manager > SSL VPNs > Edit*.
- 2 Select *Alert Settings* from the *Novell Audit and Alerts* section. The Alert Settings for SSL VPN page is displayed.

Servers > Configuration > Alert Settings

Alert Settings for SSL VPN: 12.12.12.123

Alerts	
<input type="checkbox"/> Select All	
<input type="checkbox"/> SSL VPN Gateway UP	<input type="checkbox"/> SSL VPN Gateway DOWN
<input type="checkbox"/> Concurrent Connections Reached 200	<input type="checkbox"/> Concurrent Connections Reached Maximum Limit (249)
<input type="checkbox"/> Invalid Configuration	<input type="checkbox"/> Invalid Certificate
<input type="checkbox"/> Webserver Servlet Down	<input type="checkbox"/> SSL Encryptor Down
<input type="checkbox"/> Socks Protocol Daemon Down	

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

- 3 Check the *Select All* option to send alerts for all the events. Otherwise, select one or more of the following:

Alerts	Description
Webserver servlet down	Sends an alert whenever a Webserver servlet is down.
SSL encryptor down	Sends an alert whenever the SSL encryptor is down.
Socks Protocol Daemon down	Sends an alert whenever the socket protocol daemon is down.
200 concurrent connections reached	Sends an alert when the number of concurrent connection reaches 200. The maximum limit is 249.
Concurrent connections reached maximum limit (249)	Sends an alert when number of concurrent reaches 249.
Invalid configuration	Sends an alert when the configuration is not valid.
Invalid certificate	Sends an alert when the SSL VPN certificate used for encryption and communication is invalid.

Alerts	Description
SSL VPN Gateway up	Sends an alert when the SSL VPN Gateway is up and running.
SSL VPN Gateway down	Sends an alert when the SSL VPN Gateway is down and is not functional.

- 4 To save your modifications, click *OK*, then click *Apply Changes* or *Schedule Changes* on the Configuration page.

21.8 Adding Certificates to the SSL VPN Keystore

Access Manager components and agents can access the keystore to retrieve certificates, keys, and trusted roots as needed.

To add certificates to SSL VPN keystore:

- 1 In Administration Console, select *Access Manager > SSL VPN > Edit*.
- 2 Select *SSL VPN Certificates* from the *Security settings* section. The Certificates for SSL VPN page is displayed.

[Servers](#) ▶ [Configuration](#) ▶ [Certificates](#)

Certificates for SSL VPN: 12.12.12.124

[Stunnel](#)

[Trusted Root](#)

Close

- 3 Click *STunnel*. The Keystore: SSL VPN Secure Tunnel page is displayed.

[Servers](#) ▶ [Configuration](#) ▶ [Keystore Details](#)

Keystore: SSLVPN Secure Tunnel ?

Keystore name: SSLVPN Secure Tunnel

Keystore type: PKCS12

Device: 12.12.12.124

Certificates

Replace...

Certificate Alias

1 item(s)

[bhav](#) stunnel

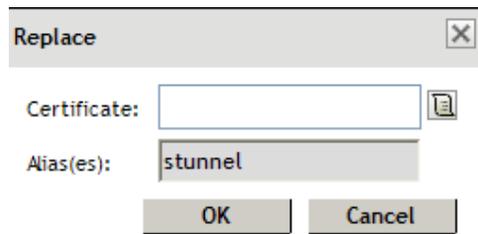
Close

Certificates in the SSL VPN STunnel are used by SSL VPN services for encryption. This page contains the following information:

- ♦ **Keystore name:** Specifies the name of the keystore to which the certificate belongs.
- ♦ **Keystore type:** Specifies the type of keystore. It can be Java, PEM, or PKCS12
- ♦ **Device:** Specifies the IP address of the SSL VPN device.

NOTE: Every imported SSL VPN device has a default certificate.

- 4 To replace the default certificate, click *Replace*. The Replace dialog box is displayed.



Fill in the following fields:

- ♦ **Certificates:** Click the *Select Certificate* icon to browse and select the certificate that you want to associate with SSL VPN.
 - ♦ **Alias(es):** You can suggest an alternate name for the certificate you are importing.
- 5 Click *OK* to save changes.
 - 6 To save your modifications, click *OK*. Click *Apply Changes* or *Schedules Changes* on the Configuration page.

21.9 Adding Trusted Roots for SSL VPN

A trust store contains certificates from a certificate authority (CA). These certificates are self-signed and are recognized as representing a CA that is trusted. When creating a trust store, you can assign trust stores to devices and add trusted root certificates to the new trust stores.

NOTE: Trusted roots need not be configured for SSL VPN.

21.10 Configuring SSL VPN to Download Applet on Internet Explorer

The SSL VPN client components are carried forward to client desktop through Java applet or ActiveX, along with the policies and the required client components.

Some Windows clients do not allow ActiveX controls to run in the Internet Explorer. In such scenarios, the user can force the Windows client to load a Java-based applet instead of the ActiveX controls. In order to force load the applet, enter the following URL to launch the SSL VPN user interface:

`https:<DNS-Name>/sslvpn/login?forcejre`

If your company's policy does not allow ActiveX controls to be downloaded from Internet Explorer, you can change the SSL VPN configuration to always download the applet based client. You can change the value within the `<param-value>` tags in the `web.xml` file from 'true' from 'false' as follows:

- 1 Login as root.
- 2 Open the `web.xml` file found in the following location:
`/var/opt/novell/tomcat4/webapps/sslvpn/WEB-INF/`

3 In the `<context-param>` section, change the `<param-value>` to 'true' as follows:

```
<context-param>
<param-name>forcejre</param-name>
<param-value>>true</param-value>
<description>My organization does not allow activex ? enter true
if so</description>
</context-param>
```

Save the `web.xml` file.

4 Restart the Tomcat server by entering the following command:

```
/etc/init.d/novell-tomcat restart
```

21.11 Accelerating SSL VPN

The SSL VPN server requires a user credential profile consisting of the following elements:

- ◆ Username and password information
- ◆ A proxy session cookie
- ◆ The roles assigned to the current user for authentication information

Each element added to the custom header requires a name with an "X-" prefix. The name you enter is specific to the application using the custom header, and might be case sensitive. You need to obtain this information from the application before creating the custom header.

The SSL VPN server requires the following three headers:

- ◆ Authentication header containing the username and password credential profile
- ◆ Custom header containing a proxy session cookie element named X-SSLVPN-PROXY-SESSION-COOKIE
- ◆ Custom header containing roles for current user element, named X-SSLVPN-ROLE

The policy engine allows you to add these and other elements to a custom header and the reverse proxy injects these headers into the SSL VPN server.

21.11.1 Injecting the SSL VPN Header

The example in this section explains how to accelerate SSL VPN server in a Path-Based Multi-homing configuration.

Before you begin, make sure you have already created a proxy service and an authentication procedure. For more information on creating proxy service and authentication procedure, see [Chapter 12, "Configuring the Access Gateway to Protect Resources,"](#) on page 141.

1 In Administration Console, click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy]*.

2 In the *Proxy Service List* section, click *New*.

New [X]

Proxy Service Name: sslvpn

Multi-Homing Type: Path-Based

Published DNS Name: www.proxy140.com

Path: /sslvpn/

Web Server IP Address: 10.1.1.10

Host Header: Web Server Host Name

Web Server Host Name: www.proxy140.com
(Alternate Host Name)

OK Cancel

3 Fill in the following fields.

Proxy Service Name: Specify a name for proxy service.

Multi-Homing Type: Specify the method for finding a second resource on the reverse proxy. For this example configuration, Path-based has been selected.

Published DNS Name: This field is populated by default with the published DNS name.

Path: Specify the URL to the SSL VPN resource as `/sslvpn/`

Web Server IP Address: This is the IP address of the SSL VPN server.

Host Header: Specify a host header name. This name is forwarded to the Web server in the host header.

Web Server Host Name: Specify the alternate host name.

4 Click *OK*.

- To configure the default Identity Injection policy and protected resources, click the newly added proxy service.

[Path-Based Multi-Homing](#) | [Web Servers](#) | [HTML Rewriting](#) | [Logging](#)

Published DNS Name: jwilson.provo.novell.com/ ... (1) path(s)

Description:

Cookie Domain: provo.novell.com

[HTTP Options](#)

Remove Path on Fill
 Reinsert Path in "set-cookie" Header

Path List	
New...	Delete Enable SSL VPN...
1 item(s)	
<input type="checkbox"/> Path	Protected Resource
<input type="checkbox"/> /sslvpn	public

Changes made on this panel must be applied or scheduled from the [Configuration](#)

- In the *Path List* section, make sure the *Path* is */sslvpn*.
- In the *Path List* section, select the */sslvpn* check box, then click *Enable SSL VPN*. The *Enable SSL VPN* pop-up is displayed.

Enable SSL VPN

Identity Injection Policy (for SSL VPN)

Policy Container:

Policy:

Protected Resource (for SSL VPN)

Name:

- Fill in the following fields:
 - Policy Container:** Leave the default value unchanged.
 - Policy:** Select *Create SSL VPN Default Policy* from the drop-down list. A policy pop-up appears. Click *Apply Changes* in the pop-up, then click *Close*.
 - Name:** Select *Create SSL VPN Default Protected Resource* from the drop-down list.
- Click *OK* to close the *Enable SSL VPN* pop-up.
- Select the *Web Servers* tab.
- Specify 8080 in the *Connect Port* field, then click *OK*

- 12 In the *Proxy Service List* section, click the name of the parent proxy service of the newly created SSL VPN proxy service. This host does not have a multi-homing value.
- 13 Select the *Protected Resources* tab.
- 14 Select *SSLVPN_Default* from *Protected Resources List*.
- 15 Select an authentication contract from the *Contract* drop-down list. Make sure you select *Name/Password - Form* as the authentication contract.
- 16 In the *URL Path List* section, ensure that the URL is */sslvpn/**.

Protected Resource: SSLVPN_Default

Description:

Contract:

URL Path List

New... | Delete 1 item(s)

<input type="checkbox"/> URL Path
<input type="checkbox"/> /sslvpn/*

IMPORTANT: Make sure that you configure the URL as given above. Any variation leads to the failure of SSL VPN service.

- 17 Click *Configuration Panel*, then click *OK*.
- 18 On the *Configuration* page, select *Apply Changes*, then click *OK*.
- 19 To update the Identity Server, click *Identity Servers > Setup > Update Servers*.
- 20 Click *Close*.

21.12 Moving SSL VPN Server to a Different Administration Console

If SSL VPN Gateway must be moved to a different Administration Console, use the following procedure:

- 1 In the Administration Console, delete the existing server
- 2 Enter the following command to stop the SSL VPN server:
`/etc/init.d/novell-sslvpn stop`
- 3 Enter the following command:
`sslvpnc --configure`
- 4 Specify the IP Address of the New Administration Console.
- 5 Specify the public IP Address of SSL VPN server.
- 6 Specify the Private IP Address of SSL VPN server.

- 7 At the console shell, access the `/opt/novell/devman/jcc` directory, then enter the following command:

```
conf/Configure.sh
```

- 8 Enter the following command to start the SSL VPN server:

```
/etc/init.d/novell-sslvpn start
```

- 9 Restart the Access Manager Server Communications using the following command:

```
/etc/init.d/novell-jcc restart
```

This imports the SSL VPN server to the new Administration Console. If you had configured multiple private IP addresses for the SSL VPN server, you can reconfigure it in the new Administration Console.

21.13 Accelerating SSL VPN with a New Access Gateway

If SSL VPN Gateway has to be accelerated by a different or new Access Gateway, auto-imported the Access Gateway to Administration Console and complete the basic configuration. For more information on accelerating the SSL VPN gateway with the new Access Gateway, see [Section 21.11, “Accelerating SSL VPN,”](#) on page 264.

Load Balancing and Fault Tolerance

22

You can import SSL VPN servers into the Administration Console. You can configure load balancing and fault tolerance on these servers by using the `config.txt` file.

- ♦ [Section 22.1, “Installing SSL VPN Servlet on a Separate Machine,” on page 269](#)
- ♦ [Section 22.2, “Load Balancing,” on page 270](#)
- ♦ [Section 22.3, “Fault Tolerance,” on page 270](#)

22.1 Installing SSL VPN Servlet on a Separate Machine

- 1 Download and install the SSL VPN servlet RPM.
- 2 To check if the SSL VPN Server is running, enter the following command:

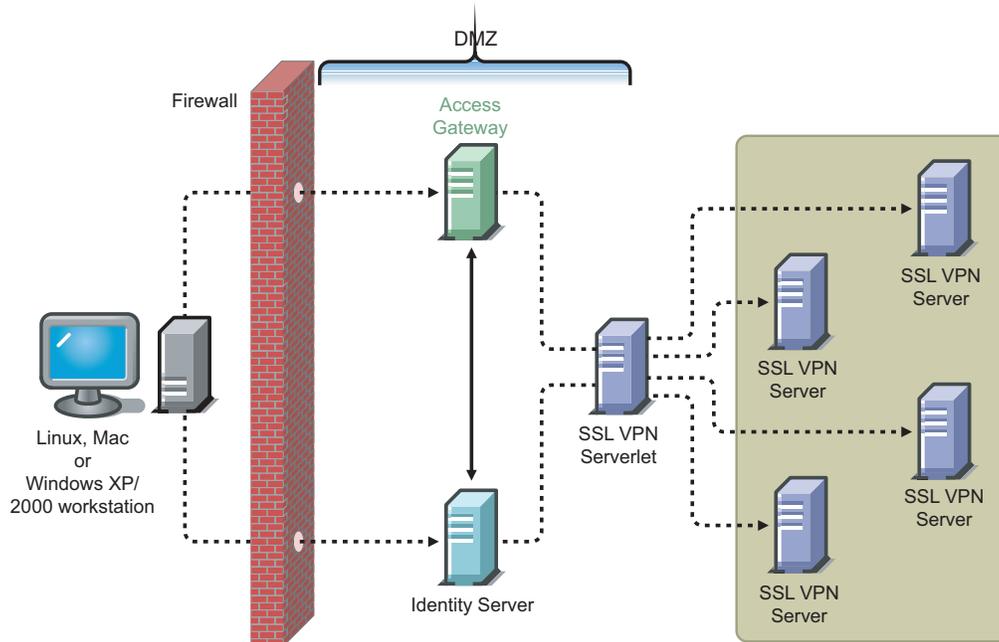
```
/etc/init.d/novell-sslvpn --status
```
- 3 To restart the SSL VPN Server, enter the following command:

```
/etc/init.d/novell-sslvpn --stop  
/etc/init.d/novell-sslvpn --start
```

22.2 Load Balancing

The SSL VPN server has load balancing capabilities so that more than one SSL VPN server can handle client connections.

Figure 22-1 Load Balancing SSL VPN servers



To configure load balancing:

- 1 Open `conf.txt` which is located in the following path:

```
/var/opt/novell/tomcat4/webapps/sslvpn/WEB-INF/
```

- 2 The first line of `conf.txt` contains the IP address and port number of the default server in the following format:

```
ServerIP=IPaddress:Port=Port number
```

NOTE: Add the IP address and port number of the servers in the same format in the next line. You can add a maximum of four servers to the failover group.

- 3 To enable load balancing among servers, set `RoundRobinCluster=true`
- 4 Save and close the file.
- 5 Restart the server by entering the following command:

```
/etc/init.d/novell-tomcat4 restart
```

22.3 Fault Tolerance

SSL VPN enables configuration of server failover groups, which enables fault tolerance. These groups ensure that when a server goes down, the other servers can service the clients. However, it is a passive fault tolerance because if a server goes down, all the client connections to that server are disconnected. When these clients try to reconnect, they are redirected to other servers in the failover group.

You can configure servers in two ways. In the first method, all the servers in the failover group receive connection. This way, client connections are distributed among the servers of the failover group, thereby balancing the load. To configure SSL VPN for both load balancing and fault tolerance, see [Section 22.2, “Load Balancing,” on page 270](#).

In the second method, all the client connections are received by one server and when that server goes down, all the connections are redirected to the next server. To configure failover groups:

- 1 Open `conf.txt`, which is located in the following path:

```
/var/opt/novell/tomcat4/webapps/sslvpn/WEB-INF/
```

- 2 The first line of `conf.txt` contains the IP address and port number of the default server in the following format:

```
ServerIP=IPaddress:Port=Port number
```

NOTE: Add the IP address and port number of the servers in the same format in the next line. You can add a maximum of four servers to the failover group.

- 3 Set `RoundRobinCluster=false`
- 4 Save and close the file.
- 5 Restart the server by entering the following command:

```
/etc/init.d/novell-tomcat4 restart
```


Accessing the SSL VPN User Portal

23

The Novell SSL VPN client can be launched through the web browser. The user can use a DNS name and the URL `http://<dns_name>/sslvpn/login` to log in to the SSL VPN server, which is behind a proxy server. This eliminates the need for VPN client to be installed and configured on remote users' machines, and gives the users a flexibility in accessing the protected network from their home computers or from Web browsing kiosks.

Java Applet or ActiveX applications are bundled with the product. These applications can encrypt traffic and send it to the SSL VPN server. When you access the SSL VPN client using the Firefox* or Safari* browsers, Java Applet is downloaded to your machine. If you use Internet Explorer to download the SSL VPN client, ActiveX is downloaded to your machine.

This section contains the following information accessing the SSL VPN client:

- ◆ [Section 23.1, “Accessing SSL VPN on Linux,” on page 273](#)
- ◆ [Section 23.2, “Accessing SSL VPN on Macintosh,” on page 275](#)
- ◆ [Section 23.3, “Accessing SSL VPN on Windows,” on page 276](#)

23.1 Accessing SSL VPN on Linux

Supported Operating System

- ◆ Novell Linux Desktop 9.0
- ◆ SUSE Linux Enterprise Desktop 10.0

Recommended Browser

- ◆ Mozilla Firefox 1.5 or higher

Mozilla Firefox Browser Requirements

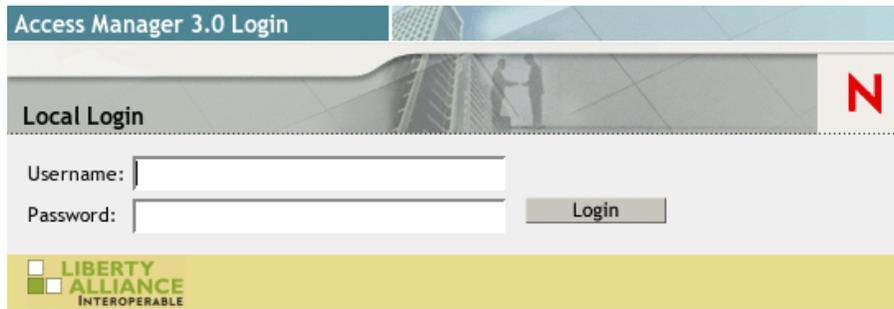
If you are using Mozilla Firefox browser, make sure the following:

- ◆ Sun JRE 1.4 or higher is installed in the workstation
- ◆ Browser is Java-enabled
- ◆ Browser is JavaScript-enabled

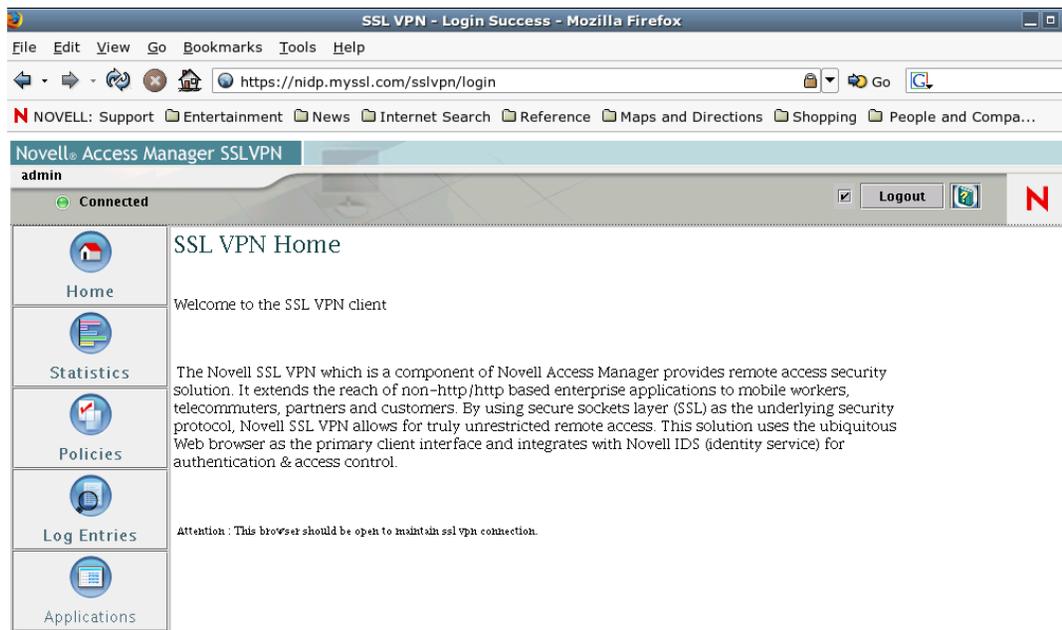
The following procedure describes a Linux user accessing SSL VPN, on a server behind a proxy server.

- 1 Log in to the SSL VPN server using the URL `http://<dns_name>/sslvpn/login`.

-
- 2 In the Access Manager page, specify the username and password. Click *OK*.



-
-
- 3 Click *Yes* in the warning message to accept and download the signed applet components required for SSL VPN client.
- 4 The Welcome Page of the SSL VPN service is displayed, allowing access to all the resources listed in the *Policy* tab.



-
-
-
-
- 5 Open a new terminal to launch applications that need to be SSLized.

NOTE: ♦To SSLize terminals that are already opened, run the `bash`, `tcsh` or `csch` depending on the shell being used.

♦If you have logged into the shell as a different user, use `source /tmp/sslize` to SSLize the shell.

-
-
-
-
-
- 6 Create desktop shortcuts for application that you want to SSLize.
- 7 Click *SSLize Application*.
- 8 Launch the application from the desktop shortcut.

23.2 Accessing SSL VPN on Macintosh

Supported Operating System

- ♦ Macintosh OS 10.4 Tiger

Recommended Browser

- ♦ Macintosh Safari

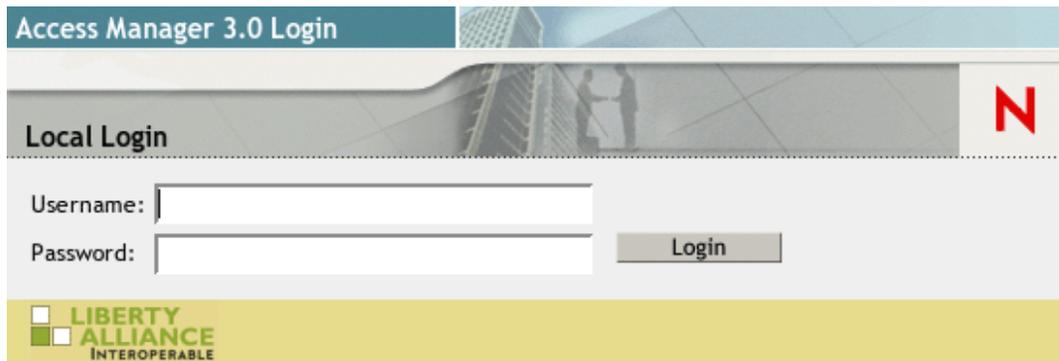
Macintosh Safari Browser Requirements

If you are using Macintosh Safari browser, ensure the following:

- ♦ Sun JRE 1.4 or higher is installed in the workstation
- ♦ Browser is Java-enabled
- ♦ Browser is JavaScript-enabled

The following procedure describes a Macintosh user accessing SSL VPN on a server behind a proxy server:

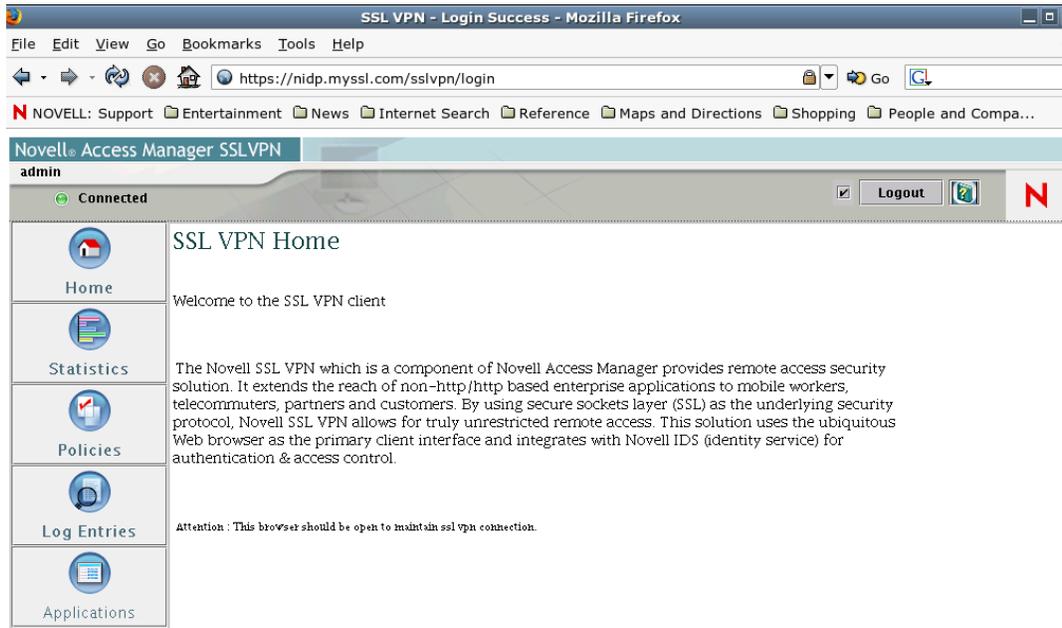
- 1 Log in to the SSL VPN server using the URL `https://<dns_name>/sslvpn/login`.
- 2 In the Access Manager page, specify the username and password. Click *OK*.



The screenshot shows the 'Access Manager 3.0 Login' interface. It features a teal header bar with the text 'Access Manager 3.0 Login'. Below this is a grey banner with the text 'Local Login' and a red 'N' icon. The main form area contains two input fields: 'Username:' and 'Password:', each followed by a text box. To the right of the password field is a 'Login' button. At the bottom of the page is a yellow banner with the 'LIBERTY ALLIANCE INTEROPERABLE' logo.

- 3 Click *Yes* in the warning message to accept and download the signed applet components required for SSL VPN client.

- 4 The Welcome Page of the SSL VPN service is displayed, allowing access to all the resources listed in the *Policy* tab.



- 5 Open a new terminal to launch applications that need to be SSLized.

NOTE: To SSLize terminals that are already opened, run the `bash`, `tcsh` or `csch` depending on the shell being used

- 6 Create alias for the application you want to SSLize by clicking the application and pressing Command+L.
- 7 Drag and drop the newly created alias into the SSL VPN folder on desktop.
- 8 Click *SSLize Application*.
- 9 Launch the application using the alias in the SSL VPN folder on desktop.

23.3 Accessing SSL VPN on Windows

Supported Operating Software

- ♦ Windows* 2000 SP 4
- ♦ Windows XP SP 2

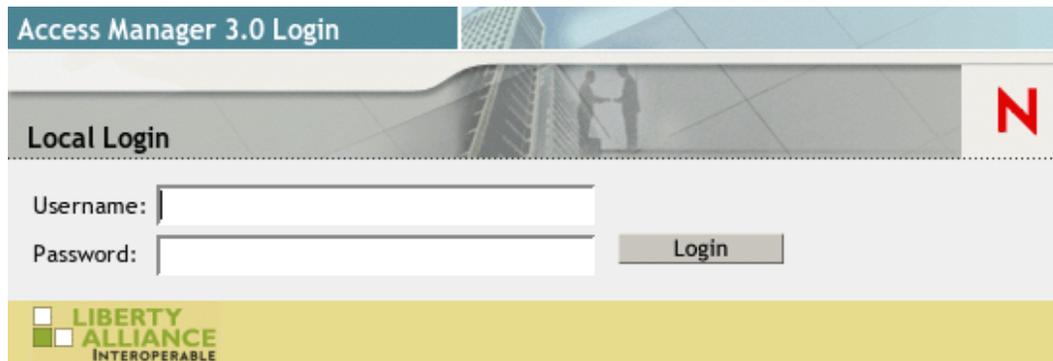
Recommended Browser

- ♦ Internet Explorer 6.0 SP 2
- ♦ Mozilla Firefox 1.5 or higher

Windows clients can use Mozilla Firefox or Internet Explorer to log in. The following procedure describes a Windows user accessing the SSL VPN on a server behind a proxy server.

- 1 Log in to the SSL VPN server using the URL `https://<dns_name>/sslvpn/login`.

- 2 In the Access Manager page, specify the username and password. Click *OK*.



A security alert message appears.

- 3 Click *Yes* to accept and download the signed ActiveX or Java Applet components required for SSL VPN client.
- 4 The Welcome Page of the SSL VPN service is displayed. This allows access to all the resources listed in the *Policy* table.

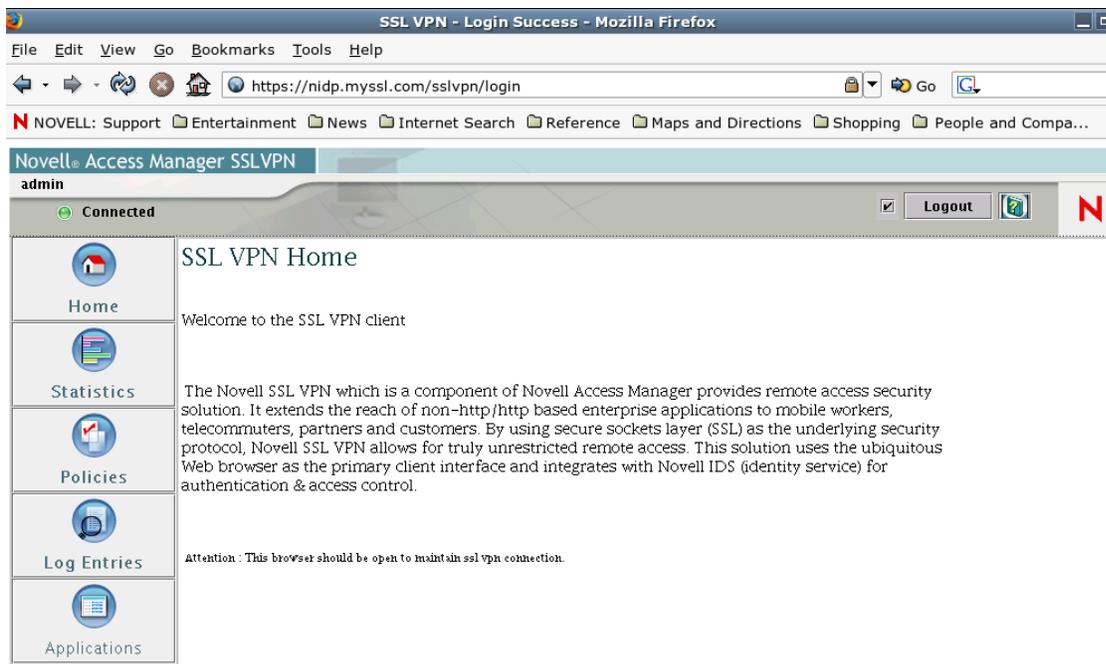


- 5 Launch applications to access your protected network.

Understanding the SSL VPN User Interface

24

When you access SSL VPN client through the browser, the following page is displayed after you authenticate to the server:



The SSL VPN client interface has the following information:

- ◆ **Username:** Specifies the name of the currently logged in user in the top left corner of the page.
- ◆ **Logout:** Click this to log out of the current session. Select the check box next to the *Logout* button to reduce the connection time when you log in again. If the check box is selected, some of the SSL VPN components are left behind in the client. This reduces the connection time, when the user logs in next time as these components need not be downloaded again.
- ◆ **Status of Connection:** Indicates the state of connection. For more information, see [Section 24.6, “Monitoring the Connection Status,” on page 284](#)

The following SSL VPN taps are displayed:

- ◆ **Home:** Displays the product information. For more information, see [Section 24.1, “Viewing the Home Page,” on page 280](#).
- ◆ **Statistics:** Displays the current statistics. [Section 24.2, “Viewing the Statistics Page,” on page 280](#).
- ◆ **Policies:** Displays the resources accessible by the user. [Section 24.3, “Viewing the Policies Page,” on page 281](#).
- ◆ **Log Entries:** Displays ActiveX or Applet logs. [Section 24.4, “Viewing the Log Entries Page,” on page 281](#)

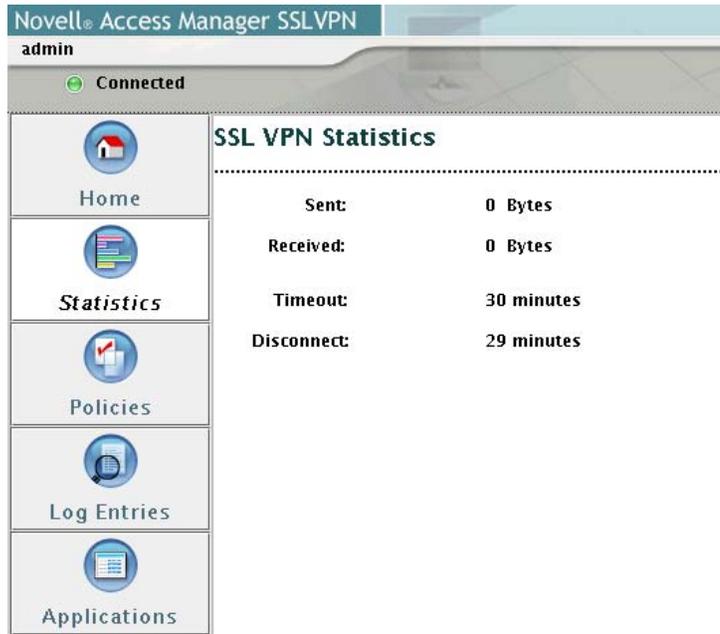
- ◆ **Applications:** Specifies steps to add applications to SSL VPN. [Section 24.5, “Viewing the Applications Page,” on page 283.](#)

24.1 Viewing the Home Page

This page displays the customer or the product information. This page can be customized for different organizations.

24.2 Viewing the Statistics Page

The Statistics page displays the client statistics.



The following table describes the information provided in the statistics page:

Column	Description
Sent	Displays bytes sent through the tunnel.
Received	Displays bytes received through the tunnel.
Timeout	Displays the time in minutes the client waits before disconnecting, if client is inactive. This time can be configured on the server.
Time to Disconnect	Displays the amount of time left before disconnecting. This counter is decremented only if the client is inactive or if there is no data transfer. The time is reset when there is a data transfer after the inactivity period.

24.3 Viewing the Policies Page

Displays the resources accessible by the user. Policies can be configured in the server.

The screenshot shows the Novell Access Manager SSLVPN interface. The user is logged in as 'admin' and is connected. The main content area is titled 'SSL VPN Policies' and contains a table with the following data:

Name	Destination	Port	Protocol	Action
Any_Role_TCP_...	0.0.0.0/255.0.0.0	Any	Tcp	Encrypt
Any_Role_UDP_...	0.0.0.0/255.0.0.0	Any	Udp	Encrypt

The following table describes the information provided in the Policies page:

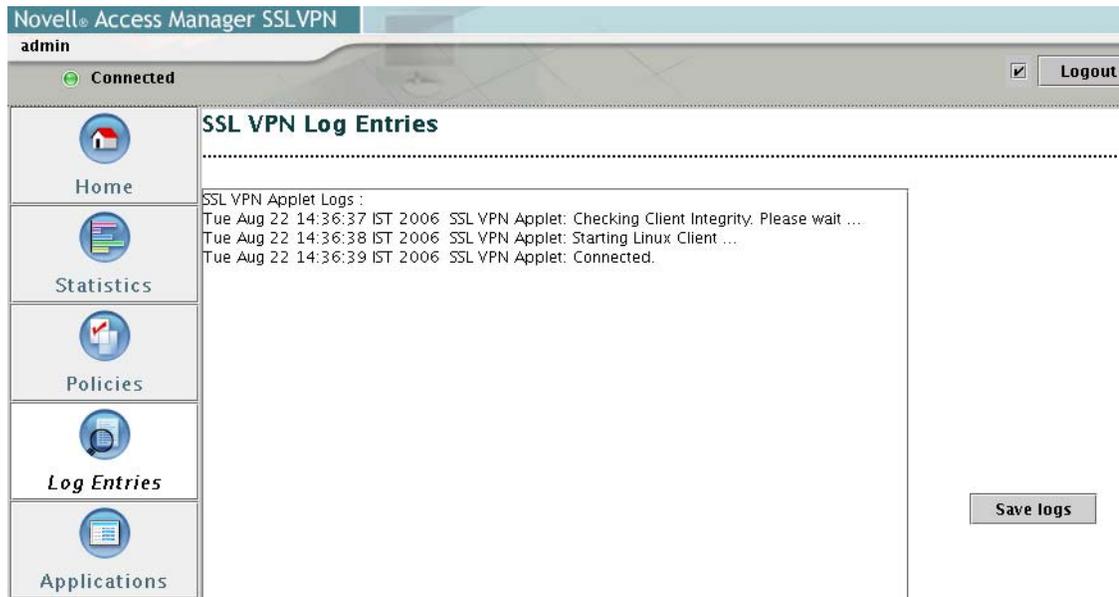
Column	Description
Destination	Displays the destination network.
Port	Displays the destination port.
Protocol	Displays the protocol.
Action	Action can be one of the following two options: <ul style="list-style-type: none">♦ Encrypt: Allows access to the protected resource.♦ Deny: Does not allow access to the protected resource.

24.4 Viewing the Log Entries Page

The Log Entries page specifies ActiveX or Java Applet log entries.

24.4.1 Saving Log Entries

You can save the log entries so that you can access them later if required.



- 1 In the SSL VPN user portal, access the *Log Entries* tab. The Log Entries page is displayed.
- 2 Click *Save Logs* to save the log entries.

If you are a Windows user, the log entries file is stored in the following location:

`[Your Home directory]\Novell\SSLVPN\log\nls\en`

If you are a Linux or Macintosh user, the log entries file is stored in the following location:

`~Userhome/[Your folder name]/sslvpnlogs`

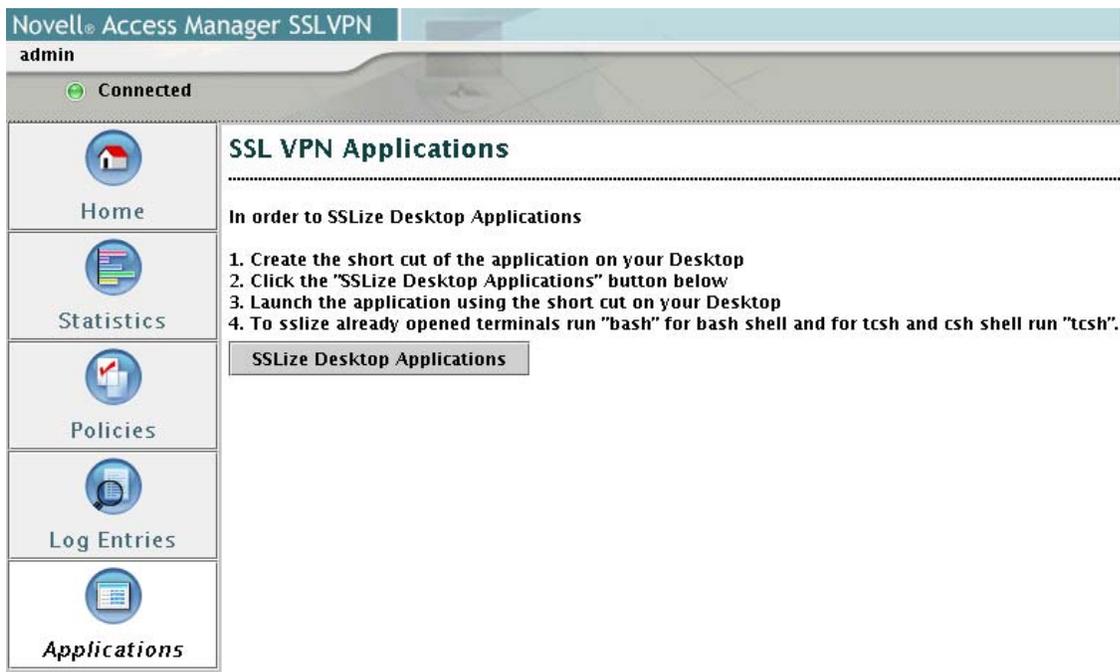
If you are a Linux or Macintosh user, the log entries file is stored temporarily in the following location:

`~Userhome/.sslvpn/log`

This directory is deleted when you log out of SSL VPN.

24.5 Viewing the Applications Page

The application page specifies steps to add applications to SSL VPN.



24.5.1 SSLizing Applications on Linux

To SSLize applications:

- 1 Start the SSL VPN services.
- 2 Create desktop shortcuts for application that you want to SSLize.
- 3 Click *SSLize Application*.
- 4 Launch the application from the desktop shortcut.

To SSLize terminals that were opened before the start of SSL VPN:

- 1 Run `bash` on bash shell.
- 2 Run `tcsh` on tcsh or csh shell.

Accessing Published Citrix Applications on Linux

To access published applications on Citrix using the SSL VPN client on Linux, you must change the browser settings as follows:

- 1 Start the SSL VPN services.
- 2 Manually connect to a Citrix server.
- 3 Click *SSLize Application*.
- 4 Click *OK* at the pop-up window, when prompted to change the browser setting.
- 5 Click *OK* to allow a Novell signed script to run.

You can now access the published applications by clicking the corresponding icons in the Citrix web page.

NOTE: The changes to browser settings are required only if you are using an SSL VPN Linux clients.

24.5.2 SSLizing Applications on Macintosh

- 1 Start the SSL VPN services.
- 2 Create alias of the application you want to SSLize by clicking on the application and pressing Command+L.
- 3 Drag and drop the newly created alias into the SSL VPN folder on desktop.
- 4 Click *SSLize Application*.
- 5 Launch the application using the alias in the SSL VPN folder on desktop.

To SSLize terminals that were opened either before or after the start of SSL VPN:

- 1 Run `bash` on bash shell.
- 2 Run `tcsh` on tcsh or csh shell.

24.6 Monitoring the Connection Status

The connection status is displayed at the top of the SSL VPN User Portal. The following table describes the different connection status:

Status	Description
Connected	Indicates that the Java applet or ActiveX has successfully established a connection to the SSL VPN server.
Disconnected	Indicates that the user has logged out of the SSL VPN server. This status is displayed when the user clicks the <i>Logout</i> button.
Connecting	Indicates that the connection is in progress. To avoid problems, user must wait until a successful connection status is displayed before pressing any other button.
Disconnecting	Indicates that the disconnection is in progress. To avoid problems, user must wait until a successful disconnection status is displayed before pressing any other button.
Error: Message	Indicates that the ActiveX or Java applet has an error. Check the ActiveX or Applet log for more information.

Security and Certificate Management



This section discusses the following topics:

- ♦ [Chapter 25, “Understanding How Access Manager Uses Certificates,” on page 287](#)
- ♦ [Chapter 26, “Managing Certificates,” on page 289](#)
- ♦ [Chapter 27, “Assigning Certificates to Access Manager Devices,” on page 301](#)

Understanding How Access Manager Uses Certificates

25

Access Manager allows you to manage centrally-stored certificates used for digital signatures and data encryption. eDirectory resides on the Administration Console is the main certificate store for all of the Access Manager components. If you use Novell® Certificate Server™, you can continue to create certificates there and import them into Access Manager.

By default, all Access Manager components (Identity Server, Access Gateway, SSL VPN, and J2EE agents) trust the local Access Manager CA. However, if the Identity Server is configured to use an SSL certificate signed externally, the trust store of the embedded service provider for each component must be configured to trust this new CA.

You can create and distribute certificates to the following components:

- ♦ **Identity Server:** Certificates allow you to provide secure authentication to the Identity Server and enable encrypted content from the Identity Server portal, via HTTPS. They also provide secure communications between trusted Identity Servers and user stores.

Liberty and SAML 2.0 protocol messages that are exchanged between identity and service providers often need to be digitally signed. The Identity Server uses the signing certificate included with the metadata of a trusted provider to validate signed messages from the trusted provider. For protocol messages to be exchanged between providers using SSL, each provider must trust the CA of the other provider. You must import the CA used by the other provider.

- ♦ **Access Gateway:** Access Gateway uses server certificates and trusted roots to protect Web servers, provide single sign-on, and enable the product's data confidentiality features, such as encryption.
- ♦ **SSL VPN:** SSL VPN uses server certificates and trusted roots to secure access to non-HTTP applications.
- ♦ **J2EE Agent:** The J2EE agent uses certificates to establish trust between the J2EE Agent and the Identity Server and for SSL between the J2EE server and the Identity Server.

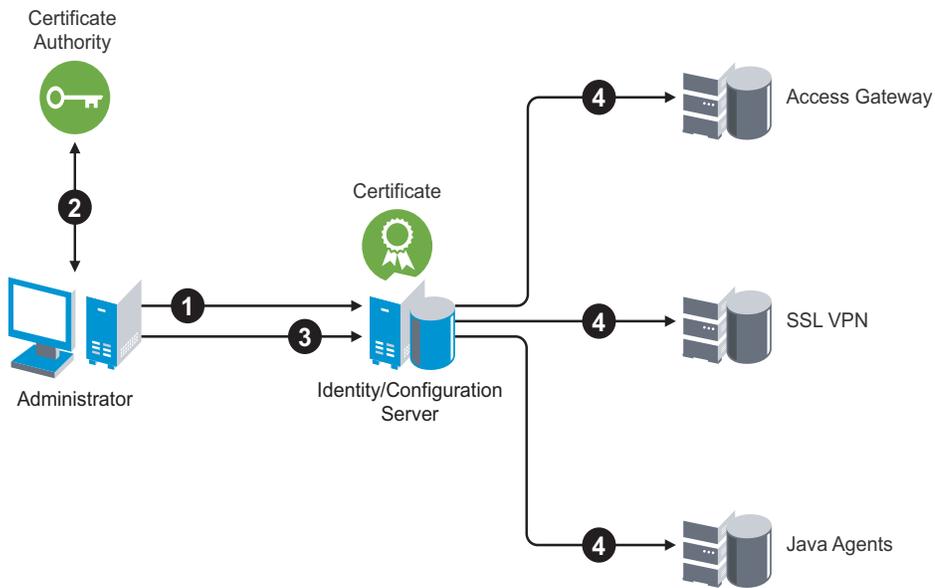
To ensure the validity of X.509 certificates, Access Manager supports both Certificate Revocation Lists (CRLs) and Online Certificate Status Protocol (OCSP) methods of verification.

Process Flow

You can install and distribute certificates to the Access Manager product components and configure how the components use certificates. This includes central storage, distribution, and expired

certificate renewal. **Figure 25-1** illustrates the primary administrative actions for certificate management in Access Manager:

Figure 25-1 Certificate Management



1. Create the certificate and generate a certificate signing request (CSR). See [Section 26.1, “Creating Certificates,”](#) on page 289.
2. Send the CSR to the external CA for signing.

A CA is a third party or network authority that issues and manages security credentials and public keys for message encryption. The CA’s certificate is held in the configuration store of the computers that trust the CA.

3. Import the signed certificate and CA chain into the configuration store. See [Section 26.5, “Importing Public Key Certificates \(Trusted Roots\),”](#) on page 297.
4. Assign certificates to devices. See [Chapter 27, “Assigning Certificates to Access Manager Devices,”](#) on page 301.

If you are unfamiliar with public key cryptography concepts, see [Public Key Cryptography Basics](#) (<http://www.novell.com/documentation/crt311/crtadmin/data/a2uqry.html#a2uqry>) in the *Novell Certificate Server 3.1.1 Guide* (<http://www.novell.com/documentation/crt311/treetitl.html>).

See [Appendix B, “Certificates Terminology,”](#) on page 521 for information about certificate terminology.

Access Manager comes with certificates for testing purposes. The test certificates are called test-signing, test-encryption, test-provider, test-consumer, and test-connector. At a minimum you must create two SSL certificates: one for Identity Server test-connector and one for the Access Gateway reverse proxy. Then you replace the pre-defined certificates with the new ones.

If you install a secondary Administration Console, the certificate authority (CA) is installed with the first instance of eDirectory, and the secondary consoles have eDirectory replicas, and therefore no CA software. All certificate management must be done from the primary Administration Console.

IMPORTANT: Before generating any certificates with the Administration Console CA, make sure time is synchronized within one minute among all of your Access Manager devices. If the time of the Administration Console has a time that is before the device for which you are creating the certificate, the device rejects the certificate.

The following sections contain detailed information about creating and managing certificates for Access Manager:

- ◆ [Section 26.1, “Creating Certificates,” on page 289](#)
- ◆ [Section 26.2, “Auto-Importing Certificates from Servers,” on page 295](#)
- ◆ [Section 26.3, “Importing a Private/Public Key Pair,” on page 296](#)
- ◆ [Section 26.4, “Exporting a Private/Public Key Pair,” on page 296](#)
- ◆ [Section 26.5, “Importing Public Key Certificates \(Trusted Roots\),” on page 297](#)
- ◆ [Section 26.6, “Renewing a Certificate,” on page 297](#)
- ◆ [Section 26.7, “Exporting a Public Certificate,” on page 298](#)
- ◆ [Section 26.8, “Enabling 4096k Keys,” on page 299](#)

26.1 Creating Certificates

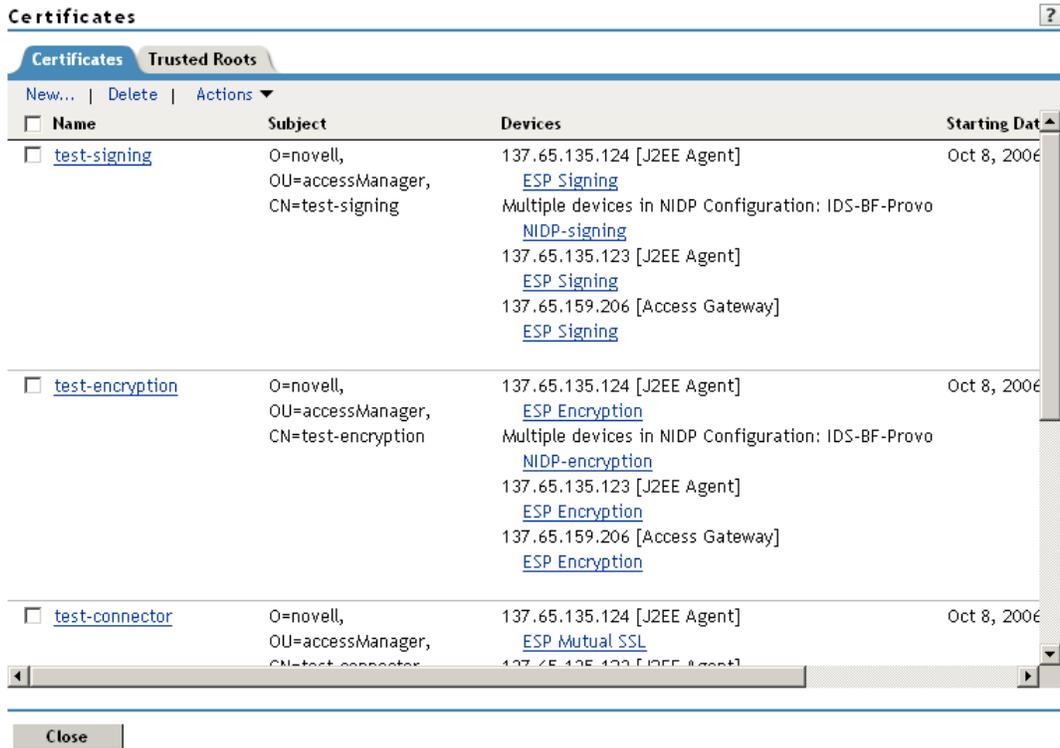
This task involves creating a certificate to be signed locally, or creating one that generates the CSR to be signed externally, which you later import after signing.

- ◆ [Section 26.1.1, “Creating a Locally Signed Certificate,” on page 290](#)
- ◆ [Section 26.1.2, “Generating a Certificate Signing Request,” on page 294](#)
- ◆ [Section 26.1.3, “Importing a Signed Certificate,” on page 295](#)

26.1.1 Creating a Locally Signed Certificate

By default, the Access Manager installation process creates the local CA for you. eDirectory contains a CA that can issue and sign certificates, and a certificate server that generates or imports certificates and keys, and generate CSRs

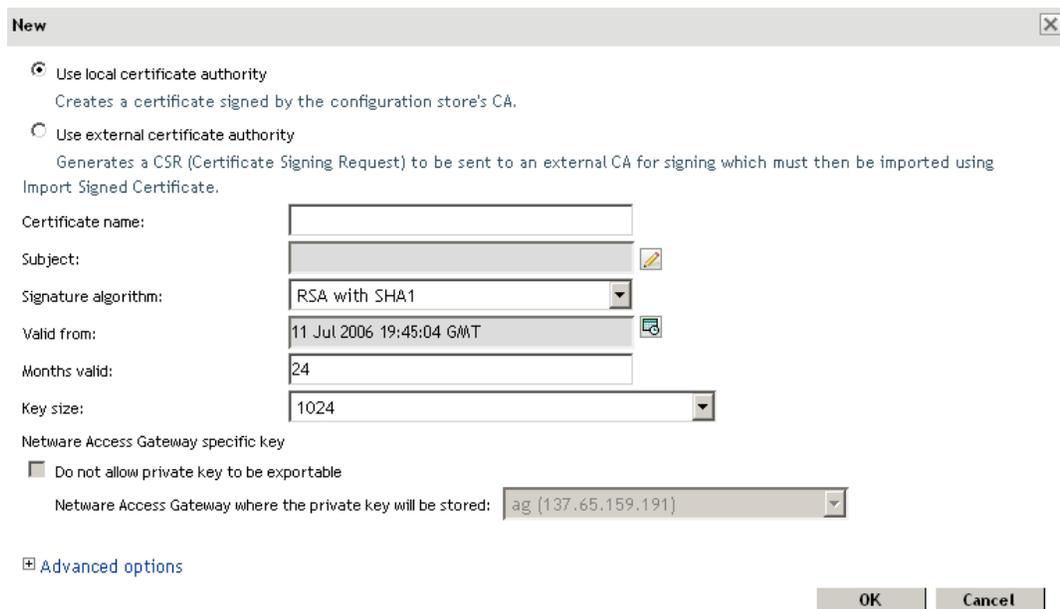
- 1 In the Administration Console, click *Certificates*.



The screenshot shows the 'Certificates' tab in the Administration Console. It displays a table with columns for Name, Subject, Devices, and Starting Date. Three certificates are listed: 'test-signing', 'test-encryption', and 'test-connector'. Each certificate entry includes a checkbox, a link to the certificate, its subject (O=novell, OU=accessManager, CN=test-...), and a list of devices with links to their respective configurations (e.g., ESP Signing, NIDP-signing).

Name	Subject	Devices	Starting Date
<input type="checkbox"/> test-signing	O=novell, OU=accessManager, CN=test-signing	137.65.135.124 [J2EE Agent] ESP Signing Multiple devices in NIDP Configuration: IDS-BF-Provo NIDP-signing 137.65.135.123 [J2EE Agent] ESP Signing 137.65.159.206 [Access Gateway] ESP Signing	Oct 8, 2006
<input type="checkbox"/> test-encryption	O=novell, OU=accessManager, CN=test-encryption	137.65.135.124 [J2EE Agent] ESP Encryption Multiple devices in NIDP Configuration: IDS-BF-Provo NIDP-encryption 137.65.135.123 [J2EE Agent] ESP Encryption 137.65.159.206 [Access Gateway] ESP Encryption	Oct 8, 2006
<input type="checkbox"/> test-connector	O=novell, OU=accessManager, CN=test-connector	137.65.135.124 [J2EE Agent] ESP Mutual SSL 137.65.135.123 [J2EE Agent]	Oct 8, 2006

- 2 Click *New*.



The 'New' dialog box is shown with the 'Use local certificate authority' option selected. It includes fields for Certificate name, Subject, Signature algorithm (RSA with SHA1), Valid from (11 Jul 2006 19:45:04 GMT), Months valid (24), and Key size (1024). There is also a checkbox for 'Do not allow private key to be exportable' and a dropdown for 'Network Access Gateway where the private key will be stored' (ag [137.65.159.191]).

Use local certificate authority
Creates a certificate signed by the configuration store's CA.

Use external certificate authority
Generates a CSR (Certificate Signing Request) to be sent to an external CA for signing which must then be imported using Import Signed Certificate.

Certificate name:

Subject:

Signature algorithm: RSA with SHA1

Valid from: 11 Jul 2006 19:45:04 GMT

Months valid: 24

Key size: 1024

Do not allow private key to be exportable

Network Access Gateway where the private key will be stored: ag [137.65.159.191]

[Advanced options](#)

OK Cancel

3 Select the following option:

Use local certificate authority: Creates a certificate signed by the local CA (or Organizational CA), and creates the private key. For information about creating a CSR, see [Section 26.1.2, “Generating a Certificate Signing Request,”](#) on page 294.

4 Fill in the following fields:

Certificate name: The name of the certificate. Pick a unique, system-wide name for the certificate that you can easily associate with the certificate’s purpose. The name must contain only alphanumeric characters and no spaces.

4a For *Subject*, click the *Edit* button to display a dialog box that lets you add the appropriate locality information types for the subject name.

Edit Subject ?

Commonly used attributes

Common name:

Organizational unit:

Organization:

City or town:

State or province:

Country:

Additional attributes

----- Select one ----- :

----- Select one ----- :

The subject is an X.500 formatted distinguished name that identifies the entity that is bound to the public key in an X.509 certificate. Choose the subject name that the browser expects to find in the certificate. The name you enter must be fully distinguished. Completing all the fields creates a fully distinguished name that will include the appropriate types (such as C for country, ST for state, L for location, O for organization, OU for organizational unit, and CN for common name). For example, `cn=AcmeWebServer.ou=Sales.o=Acme.c=US`.

The following attributes are the most common ones used in certificate subjects:

Common name: The name or IP of the Web server.

Enter just the value, for example `AcmeWebServer`. Do not include the type (`cn=`). The UI adds that for you.

For the Identity Server, this is the domain name of the base URL of the Identity Server configuration. This value cannot be an IP address or begin with a number, in order to ensure that trust does not fail between providers

Organizational unit: Describes departments or divisions.

Organization: Differentiates between organizational divisions.

City or town: Commonly referred to as the Locality.

State or province: Commonly referred to as the State.

Country: The country, such as US.

Use the *Additional Attributes* drop-down menus to add additional attributes. These values allow you to specify additional fields that are supported by eDirectory, and you can include them as part of the subject to further identify the entity represented by the certificate.

5 Click *OK*, then fill in the following fields:

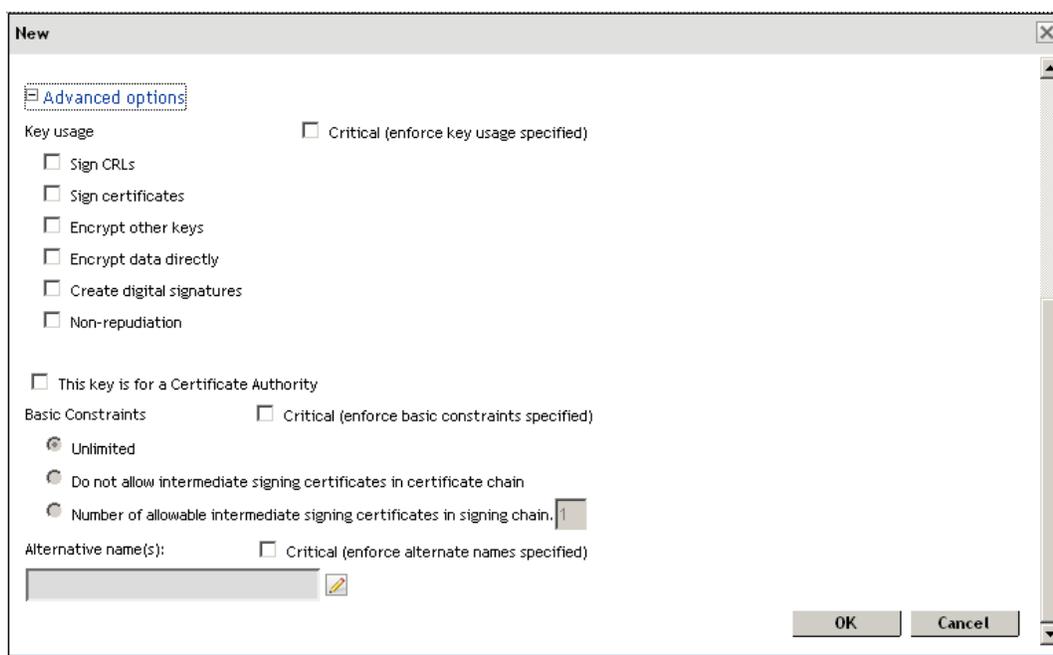
Signature algorithm: The algorithm you want to use (SHA-1, MD-2, or MD-5). SHA-1 is currently recommended.

Valid from: The date from which the certificate is valid. For externally signed certificates, the external certificate authority sets the validity period.

Months valid: The number of months that the certificate is valid.

Key size: The size of the key. Select 512, 1024, or 2048. 2048 bit is recommended. For 4096 key information, see [Section 26.8, “Enabling 4096k Keys,” on page 299](#).

6 (Optional) To configure advanced options, click *Advanced Options*.



7 Configure the following options as necessary for your organization:

Critical: Specifies that an application should reject the certificate if the application does not understand the key usage extensions.

Sign CRLs: Specifies whether the certificate is used to sign CRLs (Certificate Revocation Lists).

Sign certificates: Specifies that the certificate is used to sign other certificates.

Encrypt other keys: Specifies that the certificate is used to encrypt keys.

Encrypt data directly: Encrypts data for private transmission to the key pair owner. Only the intended receiver can read the data.

Create digital signatures: Specifies that the certificate is used to create digital signatures.

Non-repudiation: Links a digital signature to the signer and the data. This prevents others from duplicating the signature because no one else has the signer's private key. Additionally, the signer cannot deny having signed the data.

This key is for a Certificate Authority: Specifies that this certificate is for the local configuration (eDirectory) certificate authority.

If you create a new CA, all the keys signed by the CA being replaced no longer have a trusted CA. Thus, you might also need to reassign the new CA to all the trust stores that contained the old CA.

8 Under Basic Constraints, configure the following options as necessary:

Critical: Enforces the basic constraints you specify.

Unlimited: Specifies no restriction on the number of subordinate certificates that the CA can verify.

Do not allow intermediate signing certificates in certificate chain: Prevents the CA from creating other CAs, but it can create server or user certificates.

Number of allowable intermediate signing certificates in signing chain: Specifies how many subordinate certificates are allowed in the certificate chain. Values must be 1 or more. Entering 0 creates only entity objects.

Critical: Specifies that if an application does not understand the alternate name extensions, it should reject the certificate.

9 (Optional) To create subject alternative names used by the certificate, click the *Edit Subject Alternate Names* button.

Alternate names can represent the entity identified by the certificate. The certificate can identify the subject CN=www.OU=novell.O=com, but the subject can also be known by an IP address, such as 222.111.100.101, or a URI, such as www.novell.com, for example.

10 Click *New*.



Name Type: Names as specified by RFC 2459. Use the drop-down list to specify a name type, such as:

- ♦ **Directory name:** An X.500 directory name.
- ♦ **IP Address:** An IP address such as 222.123.123.123
- ♦ **URI:** A URI such as www.novell.com.
- ♦ **Registered ID:** An ASN.1 object identifier.
- ♦ **DNS Name:** A domain name such as novell.com.
- ♦ **RFC822 Name:** An e-mail address.

- ♦ **X400 Name:** The messaging and e-mail standard specified by the ITU-TS (International Telecommunications Union - Telecommunication Standard Sector). It is an alternative to the more prevalent Simple Mail Transfer Protocol (SMTP) e-mail protocol. X.400 is common in Europe and Canada.
- ♦ **EDI Party:** EDI (Electronic Data Interchange) is a standard format for exchanging business data.
- ♦ **Other:** A user-defined name.

Name: The display alternative name.

11 Click *OK*.

See Also

- ♦ [Section 26.3, “Importing a Private/Public Key Pair,” on page 296](#)
- ♦ [Section 26.4, “Exporting a Private/Public Key Pair,” on page 296](#)
- ♦ [Section 26.5, “Importing Public Key Certificates \(Trusted Roots\),” on page 297](#)

26.1.2 Generating a Certificate Signing Request

1 In the Administration Console, click *Certificates*, then click *New*.

2 Select the following option:

Use external certificate authority: Generates a Certificate Signing Request (CSR) for you to send to the CA for signing. A third-party CA is managed by a third party outside of the eDirectory tree. An example of a third party CA is VeriSign*. After the signed certificate is received, you need to import the certificate. See [Section 26.1.3, “Importing a Signed Certificate,” on page 295](#).

3 Fill in the following fields:

Certificate name: The name of the certificate. Pick a name unique, system-wide name for the certificate that you can easily associate with the certificate’s purpose. The name must contain only alphanumeric characters and no spaces.

Subject: An X.500 formatted distinguished name that identifies the entity that is bound to the public key in an X.509 certificate. Choose the subject name that the browser expects to find in the certificate. The name you enter must be fully distinguished. Completing all the fields creates a fully distinguished name that will include the appropriate types (such as C for country, ST for state, L for location, O for organization, OU for organizational unit, and CN for common name). For example, cn=AcmeWebServer.ou=Sales.o=Acme.c=US

4 Click the *Edit* button to display a dialog box that lets you add appropriate locality information types for the subject name.

The following attributes are the most common ones used in certificate subjects:

Common name: The name or IP of the Web server. Enter just the value. Do not enter the type (cn=). The UI adds it for you.

Organizational unit: Describes departments or divisions.

Organization: Differentiates between organizational divisions.

City or town: Commonly referred to as the Locality.

State or province: Commonly referred to as the State.

Country: The country, such as US.

Use the *Additional Attributes* drop-down lists to add additional attributes. These values allow you to specify additional fields that are supported by eDirectory, and you can include them as part of the subject to further identify the entity represented by the certificate.

5 Click *OK*, then fill in the following fields:

Signature algorithm: The algorithm you want to use (SHA-1, MD-2, or MD-5). SHA-1 is currently recommended.

Valid from: The date from which the certificate is valid. For externally signed certificates, the external certificate authority sets the validity period.

Months valid: The number of months that the certificate is valid.

Key size: The size of the key. Select 512, 1024, or 2048. 2048 bit is recommended. For 4096 key information, see [Section 26.8, “Enabling 4096k Keys,” on page 299](#).

6 If necessary, fill in the certificate fields, which are described in [Section 26.1.1, “Creating a Locally Signed Certificate,” on page 290](#).

7 Click *OK*.

8 On the Certificate Details page, copy the CSR data and send the information to the external CA.

The certificate status is CSR Pending until you import the signed certificate.

9 Click *Close*.

Continue with [Section 26.1.3, “Importing a Signed Certificate,” on page 295](#) after you receive the signed certificate and the trusted root (CA chain).

26.1.3 Importing a Signed Certificate

After you receive the signed certificate and the CA chain, you must import it. There are several ways in which the CA can return the certificate. Typically, the CA either returns one or more files each containing one certificate, or returns a file with multiple certificates in it.

1 In the Administration Console, click *Certificates*, then click the certificate name.

2 On the Certificate Details page, click *Import Signed Certificate*.

3 In the Import Signed Certificate dialog box, browse to locate the certificate data file, or paste the certificate data text into the *Certificate data text* field.

4 To import the CA chain, click *Add trusted root*, then locate the certificate data.

5 Click *Add intermediate certificate* if you need to continue adding certificates to the chain.

6 Click *OK*, then click *Close* on the Certificate Details page.

The certificate is now available for use by Access Manager devices.

26.2 Auto-Importing Certificates from Servers

You can import certificates from other servers, such as an LDAP server, and make them available for use in Access Manager. You must provide the IP address, port, and certificate name.

1 In the Administration Console, click *Access Manager > Certificates > Trusted Roots > Auto-Import from Server*.

2 Fill in the following fields:

Server IP Address: Specifies the server IP address. You can use a DNS name.

Server Port: Specifies the server port.

Certificate Name: Specifies a unique name of the certificate to be store in Access Manager.

3 Click *OK*.

26.3 Importing a Private/Public Key Pair

If you created a key pair that was exported from another certificate management system, you can import the key pair and then assign it to an Access Manager device. The file needs to be in PKCS12 (*.pfx) or (*.p12) format.

1 In the Administration Console, click *Certificates*.

2 Choose *Actions > Import Private/Public Keypair*.

3 Fill in the following fields:

Certificate name: The name of the certificate. This is a system-wide, unique name used by Access Manager.

Password: Type the encryption/decryption password established when exporting the certificate.

Certificate data file: The certificate file to import. You can browse to locate the .pfx or .p12 file.

Certificate data text: An editable field used to enter or paste certificate data text. This is valid if your PKCS12 file is in Base64-encoded format. The first line of the data is -----BEGIN PKCS12-----.

Overwrite an existing certificate with the same name: Specifies whether to replace any existing certificates with the same name as this one.

4 Click *OK*.

26.4 Exporting a Private/Public Key Pair

When you create a certificate, you can specify whether it is exportable. If a key is exportable, it can be extracted and put in a file along with the associated certificate. The file is written in an industry standard format, PKCS#12, which allows it to be transported to other platforms. It is encrypted with a user-specified password to protect the private key. You can export private certificates to obtain a backup copy of the key, to move the key to a different server, or to share the key between servers.

You cannot export a certificate if you created it for NetWare Access Gateway and enabled the *Do not allow private key to be exportable option* while creating the certificate.

1 In the Administration Console, click *Certificates*.

2 On the Certificates page, click the certificate.

- 3 On the Certificate Details page, click *Export Private/Public Keypair*.

Certificates ▸ Certificate Details

Certificate: test-signing



- 4 Specify the password in the *Encryption/decryption* password field, then click OK.

IMPORTANT: Remember this password because you need it to re-import the key.

26.5 Importing Public Key Certificates (Trusted Roots)

You import trusted roots so that the specific device can trust the certificate sent by other computers at runtime. After you import a trusted root, you can assign it to the proper trust store associated with a device, which allows the device to trust certificates signed by the trusted root.

- 1 In the Administration Console, click *Access Manager*, then click *Certificates*.
- 2 Click the *Trusted Roots* tab.
- 3 Click *Import*, then fill in the following fields:

Certificate name: The name of the certificate. This is a system-wide, unique name used by Access Manager.

Certificate data file: The certificate file to import. You can browse to locate the file or copy and paste text into the *Certificate data text* field.

Certificate data text: An editable field used to enter or paste Base64-encoded certificate data text.

- 4 Click *OK*.

26.6 Renewing a Certificate

The Certificate Details page lists the properties of a certificate, such as certificate type, name, subject, and assigned keystores. This page also includes the original CSR. If the certificate has expired, you can cut and paste its text to send it to the CA to get a renewed certificate, then import the newly signed certificate.

- 1 In the Administration Console, click *Certificates*.

2 Click the certificate name.

Certificates ▶ Certificate Details

Certificate: test-signing ?

Renew... | Export Private/Public Keypair... | Export Public Certificate | Add Certificate to Keystores...

Issuer: O=SPA_UNSTABLE_TREE, OU=Organizational CA
Serial number: 62535869181227109482020130595159804222133434715605398600291586264512628279869725
Subject: CN="test-signing,OU=accessManager,O=novell"
Valid from: Jun 2, 2006
Valid to: Jun 2, 2008
Devices: 151.155.167.53 [Access Gateway]
[ESP Signing](#)
Multiple devices in NIDP Configuration: IDP_A
[spa-unstable-signing](#)

Key size: 2048
Signature algorithm: RSA with SHA1

3 Click *Renew*.

Certificates ▶ Certificate Details

Certificate: test-encryption ?

Renew... | Export Private/Public Keypair... | Export Public Certificate | Add Certificate to Keystores...

264512628279869727

Renew X

Certificate data file
 Browse...

Certificate data text

Certificate Chain
 Add trusted root
 Add intermediate certificate
 Add intermediate certificate

Key Usage Critical

4 On the Renew page, browse to locate the certificate, then click *OK*.

26.7 Exporting a Public Certificate

You can export a trusted root or a public key certificate to a file so that a client can use it to verify the certificate chain sent by a cryptography-enabled application, or to have a backup copy of the file.

You can export the certificate in two file formats: DER-encoded (.der) and Base64-encoded (.b64). The .crt extension can also be used for DER-encoded certificates. You can also export to the system Clipboard in Base64 format so that the certificate can be pasted directly into a cryptography-enabled application.

- 1 In the Administration Console, click *Certificates*.
- 2 From either the *Certificates* tab or the *Trusted Roots* tab, click the certificate name.
- 3 On the Certificate Details page, click *Export Public Certificate*, then click the file type.
- 4 Save the output file to the location of your choosing.

26.8 Enabling 4096k Keys

The basic functionality for using cryptographic techniques in Java is provided by the Java Cryptography Architecture (JCA) and Java Cryptography Extension (JCE). This architecture is what is referred to as provider-based (pluggable) architecture. In this case, it means that the JCE and JCA provide a set of classes and interfaces that an application developer writes to, together with factories that enable the creation of the objects that conform to the interfaces and classes.

This key size is not available for the NetWare Access Gateway.

26.8.1 Jurisdiction Policy Files

Because of various export and import restrictions in various geographies, the Java Development Kit (JDK*) download ships with a set of policy files that place certain restrictions on the key sizes that can be used. Key sizes are limited in general to 128 bits (except for the symmetric cipher Triple-DES), and RSA key generation is limited to 2,048 bits. The easiest way to deal with this restriction if it need not apply to you is to download the unrestricted policy files.

Installing the Unrestricted Policy Files

You can find the unrestricted policy files on the same page as the JCE/JDK downloads are found. Normally it is a discrete link at the bottom of the download page, such as *Unlimited Strength Jurisdiction Policy Files*. If it is legal for you to do so, you should download the ZIP file and install the two JAR files it contains, according to the instructions in the Readme file contained in the ZIP file.

On Linux, make sure you install the policy files in the Java runtime that you are using; you need `root` access or the assistance of a `root` user to do so.

Typically these files are installed (copied) in the directory `JAVA_HOME/jre/lib/security`, where `JAVA_HOME` is the home directory path of JVM used by Access Manager components. For example, `/opt/novell/java`.

Installing and Enabling the Bouncy Castle Provider

The Bouncy Castle provider can be used to handle greater key sizes. You can configure the JRE* to statically preload the provider for you so that it is available in the environment.

- 1 Stop all Access Manager components.
- 2 Under your Java installation (`JAVA_HOME`), locate the directory `jre/lib/ext` (Linux), or if you are using Windows, `jre\lib\ext`.

The purpose of this directory is to provide a home for standard extensions at runtime. The extensions are not normally distributed with the default setup.

3 Place the JAR in this location.

For the Bouncy Castle provider, you can find the JAR file you need (bcprov-jdk14-128.jar) at [The Legion of the Bouncy Castle \(http://www.bouncycastle.org/\)](http://www.bouncycastle.org/), or from an installed Access Manager location (such as /var/opt/novell/tomcat4/webapps/nidp/WEB-INF/lib/bcprov-jdk14-128.jar). The naming convention used for provider JARs in Bouncy Castle is bcprov-JdkVersion-Version.jar.

4 Enable the provider by adding it to the java.security file in the directory *JAVA_HOME/jre/lib/security*.

This file contains a section which lists JCA/JCE providers with their precedence (among other things). Add the bouncy castle provider second in the list, as follows:

```
security.provider.1=sun.security.provider.Sun
security.provider.2=org.bouncycastle.jce.provider.BouncyCastleProvider
security.provider.3=com.sun.net.ssl.internal.ssl.Provider
security.provider.4=com.sun.rsa.jca.Provider
security.provider.5=com.sun.crypto.provider.SunJCE
```

The list might be different, depending on which version of Java you have installed. Ensure that the Bouncy Castle provider is in the second position as shown in the list above. This alters the cryptographic behavior of any process using this JVM. Do not put the Bouncy Castle provider at the top of the list, because some Java system software packages rely on the Sun providers being the first ones in the list, and processing can stop working if they aren't positioned correctly.

5 Restart Access Manager components, as well as the JVM.

Assigning Certificates to Access Manager Devices

27

This section discusses how you update, renew, and assign certificates to Access Manager devices.

- ♦ [Section 27.1, “Importing a Trusted Root to the LDAP User Store,” on page 301](#)
- ♦ [Section 27.2, “Replacing Identity Server SSL Certificates,” on page 303](#)
- ♦ [Section 27.3, “Changing a Trust Store Password,” on page 304](#)
- ♦ [Section 27.4, “Assigning Certificates to an Access Gateway,” on page 305](#)
- ♦ [Section 27.5, “Assigning Certificates to J2EE Agents,” on page 305](#)
- ♦ [Section 27.6, “Configuring SSL for Authentication between the Identity Server and Access Gateway,” on page 305](#)
- ♦ [Section 27.7, “Changing a Non-Secure \(HTTP\) Environment to a Secure \(HTTPS\) Environment,” on page 306](#)
- ♦ [Section 27.8, “Creating Keystores and Trust Stores,” on page 306](#)

27.1 Importing a Trusted Root to the LDAP User Store

When you specify the settings of a user store for an Identity Server configuration, or add a user store, you can import the trusted root certificate to the LDAP user store device.

- 1 In the Administration Console, click *Identity Servers > [Configuration] > Local > [User Store]*.

- 2 Under *Server Replicas*, click the name of the server replica.

Identity Servers ► IDP_A ►

Installed User Store

Name:

Admin name:
(Ex: cn=admin,o=novell)

Admin password:

Confirm password:

Directory type:

Server replicas

New | Delete

Name

[Installed User Store Replica](#)

Search Contexts

Specify server replica information

Name:

IP Address: :

Use secure LDAP connections

[Auto import trusted root](#)

Connection limit:

OK Cancel

- 3 Enable the *Use secure LDAP connections* option.

This option allows SSL communication to occur between the Identity Server and the user store.

- 4 Click *Auto import trusted root*.
- 5 Click *OK* to confirm the import.

Ensure that you have pop-ups enabled, or the browser cannot display the Confirm dialog box.

Select Certificate to Trust

Alias:

Server Certificate

Subject: O=.SPB_UNSTABLE_TREE., CN=spb-unstable.provo.novell.com
Issuer: O=SPB_UNSTABLE_TREE, OU=Organizational CA
Valid starting date: 30 May 2006 17:10:44 GMT
Valid ending date: 29 May 2008 17:10:44 GMT
Signature algorithm: SHA1withRSA
Finger print (MD5): 3C:7A:99:81:05:2F:40:23:0E:94:14:68:A5:D3:29:3D
Finger print (SHA1): 74:86:DD:23:F4:23:5B:95:8C:78:F7:86:6B:05:91:8C:8C:98:0D:99

Root CA Certificate

Subject: O=SPB_UNSTABLE_TREE, OU=Organizational CA
Issuer: O=SPB_UNSTABLE_TREE, OU=Organizational CA
Valid starting date: 28 May 2006 19:10:40 GMT
Valid ending date: 27 May 2016 19:10:40 GMT
Signature algorithm: SHA1withRSA
Finger print (MD5): F4:D9:FE:A5:F9:93:01:02:62:85:29:44:53:D4:5B:90
Finger print (SHA1): AF:EC:A7:1C:22:10:B7:35:91:FE:B9:6E:51:92:B8:9A:6C:0E:A1:5F

6 Select one of the certificates in the list.

You are prompted to choose either a server certificate or a root CA certificate. To trust one certificate, choose *Server Certificate*. Choose *Root CA Certificate* to trust any certificate signed by that certificate authority.

7 Specify an alias, then click *OK*.

You use the alias to identify the certificate in Access Manager.

8 On the User Store page, click *OK*.

9 Restart the Identity Server.

27.2 Replacing Identity Server SSL Certificates

This procedure allows you to replace a trusted root certificate that is stored in the trust store assigned to the Identity Server. You must create an SSL certificate for the Identity Server and then replace the predefined test-connector certificate that comes with Access Manager. You can also replace the test-provider and test-consumer certificates in the *NIDP-provider* and *NIDP-consumer* keystores. The steps for replacing the signing, encryption, provider, and consumer certificates are similar.

You can also add the trusted roots to the trust stores used by the Identity Server, or auto-import them from a server. The NIDP trust store is the certificate container for CA certificates associated with the Identity Server.

You can also access the OCSP trust store to add OCSP server certificates. Online Certificate Status Protocol is a method used for checking the revocation status of a certificate. For this feature, you must set up an OCSP server. The Identity Server sends an OCSP request to the OCSP server to determine if a certain certificate has been revoked. The OCSP server replies with the revocation

status. If this revocation checking protocol is used, the Identity Server does not cache or store the information in the reply, but sends a request every time it needs to check the revocation status of a certificate. The OCSP reply is signed by the OCSP server. To verify that it was signed by the correct OCSP server, the OCSP server certificate needs to be added to this trust store. The OCSP server certificate itself is added to the trust store, not the CA certificate

- 1 In the Administration Console, click *Identity Servers > [Configuration] > Security*.

- 2 Click the certificate link that you want to replace:

Encryption: Displays the encryption certificate keystore. The encryption certificate is used to encrypt specific fields or data in the assertions.

Signing: Displays the signing certificate keystore. Click this option to access the keystore and replace the signing certificate as necessary. The signing certificate is used to sign the assertion or specific parts of the assertion.

SSL: Displays the SSL connector keystore. Click this option to access the keystore and replace the SSL certificate as necessary. This certificate is used for SSL connections.

Provider: Displays the identity provider keystore. Click this option to access the keystore and replace the identity provider certificate.

Consumer: Displays the identity consumer keystore. Click this option to access the keystore and replace the identity consumer certificate as necessary.

- 3 Click *Replace*.

NOTE: A keystore stores only one certificate at a time. When you replace a certificate, you overwrite the existing one.

- 4 In the Replace dialog box, click the *Select Certificate* icon and browse to select the certificate you created in [Section 26.1, “Creating Certificates,” on page 289](#).

- 5 Click *OK*.

- 6 Click *OK* in the Replace dialog box.

- 7 Restart Tomcat, as prompted by the system.

The system restarts Tomcat for you if you click *Restart Now* at the prompt. If you want to restart at your convenience, select *Restart Later* and then manually restart Tomcat via ssh. Enter `/etc/init.d/novell tomcat4 restart`, then press Enter.

- 8 Update the Identity Server configuration on the Setup page, as prompted.

27.3 Changing a Trust Store Password

You can change the password for a trust store associated with the Identity Server configuration.

- 1 In the Administration Console, click *Identity Servers > [Configuration] > Security*.

- 2 Under *Trust Stores*, click a trust store name.

- 3 On the Trust Store page, select the trust store check box, then click *Change Password*.

- 4 Provide a password and retype the password to verify.

- 5 Click *OK*, then click *Close*.

27.4 Assigning Certificates to an Access Gateway

The Access Gateway can be configured to use certificates for SSL communication with three types of entities:

- ♦ **Identity Server:** The Access Gateway uses the embedded service provider to communicate with the Identity Server. The Access Manager CA automatically generates the required certificates for secure communication when you set up a trusted relationship with the Identity Server. To manage these certificates in the Administration Console, click *Access Gateways > [Configuration Link] > Service Provider Certificates*. For more information, see [Section 13.5, “Managing the Certificates of the Embedded Service Provider,” on page 178](#).
- ♦ **Client browsers:** You can enable SSL communication between the client browsers and the Access Gateway. When setting up this feature, you can either have the Access Manager CA automatically generate a certificate key or you can select a certificate key you have already imported (or created) for the reverse proxy. To manage this certificate in the administration console, click *Access Gateways > [Configuration Link] > [Name of Reverse Proxy]*. For more information, see [Section 12.1, “Creating a Reverse Proxy and Proxy Service,” on page 142](#).
- ♦ **Protected Web servers:** You can enable SSL communication between the Access Gateway and the Web servers it is protecting. This option is only available if you have enabled SSL communication between the browsers and the Access Gateway. You can enable SSL or mutual SSL. To manage these certificates in the Administration Console, click *Access Gateways > [Configuration Link] > [Name of Reverse Proxy] > [Name of Proxy Service] > Web Servers*. For more information, see [Section 12.3, “Configuring the Web Servers of a Proxy Service,” on page 148](#).

27.5 Assigning Certificates to J2EE Agents

To enable the J2EE agent for SSL, you must set up the following trust relationships:

- ♦ The J2EE server with the Identity Server
- ♦ The J2EE agent with the Identity Server

For instructions on setting up these certificates, see “[Configuring SSL Certificate Trust](#)” in the *Novell Access Manager 3.0 J2EE Agent Guide*.

27.6 Configuring SSL for Authentication between the Identity Server and Access Gateway

By default, all Access Manager components (Identity Server, Access Gateway, SSL VPN, and J2EE agents) trust the certificates signed by the local CA. However, if the Identity Server is configured to use an SSL certificate signed externally, the trusted store of the service provider for each component must be configured to trust this new CA. Import the public certificate of the CA into the following trust stores:

- ♦ For an Access Gateway, click *Access Gateways > [Configuration Link] > Service Provider Certificates > Trusted Roots*.
- ♦ For a J2EE agent, click *J2EE Agents > Edit > Trusted Roots*.
- ♦ For an SSL VPN server, click *SSL VPNs > Edit > SSL VPN Certificates > Trusted Root*.

If an Access Gateway, a J2EE agent, or an SSL VPN server is configured to use an SSL certificate signed externally, the trusted store of the Identity Server must be configured to trust this new CA. Import the public certificate of the CA into the Identity Server configuration that the component is using for authentication.

In the Administration Console, click *Identity Servers* > [*Configuration Assignment*] > *Security* > *NIDP Trust Store* and add the certificate to the Trusted Roots list.

NOTE: Whenever you replace certificates on a device, you must update the Identity Server configuration (by clicking *Update Servers* on the Setup page), or restart the Access Gateway ESP application.

27.7 Changing a Non-Secure (HTTP) Environment to a Secure (HTTPS) Environment

If you are running in a non-secure staging environment, and you're ready to move to production, you must perform the following steps to enable security.

- 1 Change the Identity Server configuration protocol to HTTPS.
- 2 Replace the test certificates with your own.
- 3 Reimport metadata for trusted service and identity providers.
- 4 Change the Access Gateway configuration to HTTPS.

27.8 Creating Keystores and Trust Stores

A keystore is storage file containing keys, certificates and trusted roots. Access Manager agents can access them to retrieve certificates, keys, and trusted roots as needed. A trust store is a keystore containing only trusted roots. Intermediate CAs and end entity public certificates can be part of a trust store.

Access Manager comes with predefined stores for certificate management. However, in certain situations you might need to create a keystore or trust store. For example, if you are using JBoss keystore certificates that you need to import into Access Manager, you must create a keystore and assign it to the JBoss agent. It is probable that the keystore already exists on the JBoss file system, as created and configured by JBoss. Creating it again through Access Manager does not delete the existing keystore. This does allow Access Manager to recognize the existing keystore and add or remove the certificates. Access Manager cannot manage certificates that were created before the keystore is created in Access Manager.

The easiest way to create a keystore is to do so when you are adding the certificate to the keystore. If you want to create a trust store, the steps are identical, except you select trusted roots from the Trusted Roots page, rather than the certificates from the Certificates page.

A keystore stores only one certificate at a time. When you replace a certificate, you overwrite the existing one.

- 1 In the Administration Console, click *Certificates*.
- 2 Import the certificate, if you have not done so already. See [Section 26.3, "Importing a Private/Public Key Pair," on page 296](#).
- 3 Click the certificate name.

- 4 In the Certificate Details page, click *Add Certificate to Keystores...*
- 5 On the Add Certificate to Keystores dialog box, click the *Select Keystore* button to browse for key stores.
- 6 On the Keystore page, click *New*.

Certificates

Keystores	
New, Delete, Actions	
	Device
	Multiple
	Multiple
	Multiple
	151.155
	151.155
	151.155
	store 151.155
	151.155

- 7 Fill in the following fields:

Keystore name: Specifies the name of the keystore. This maps to a name that the server communication recognizes to identify the keystore on the device.

Keystore type: Specifies whether to use Java, PEM, PKCS12.

Keystore password: Specifies the password to revise the keystore settings.

Device: Specifies the device (by IP) to which you assign the keystore. The device can be an Identity Server, Access Gateway, agent, or SSL VPN. You cannot assign one keystore to multiple devices.

Directory: Specifies the directory where PKCS12 or PEM files are stored.

For example, `/var/opt/novell/keystores/`.

File: Specifies the path and filename of the Java keystore (JKS).

For example, `/var/opt/novell/keystores/myKeystore.keystore`.

Description: Describes the keystore.

- 8 Click *OK*.

This creates the key store.

- 9 (Optional) On the Keystore page, you can assign a certificate to the new keystore by selecting the store's check box.

- 10 Click *OK* in the *Add Certificate to Keystores* dialog box.

Policy Management

VI

This section describes how Access Manager uses policies.

- ◆ [Chapter 28, “Managing Policies,” on page 311](#)
- ◆ [Chapter 29, “Creating Role Policies,” on page 317](#)
- ◆ [Chapter 30, “Creating Authorization Policies,” on page 339](#)
- ◆ [Chapter 31, “Creating Identity Injection Policies,” on page 373](#)
- ◆ [Chapter 32, “Creating Form Fill Policies,” on page 385](#)

Policies are logical and testable rules that you use to maintain order, security, and consistency within your Access Manager infrastructure. You can specify activation criteria, deactivation criteria, temporal constraints (such as time of day or subnet), identity constraints (such as user object attribute values), and additional separation-of-duty constraints. Identity information can come from any identity source (such as LDAP, an Identity Vault, or a directory) or from the Access Manager's Identity Server, which provides full Liberty Alliance specification support and SAML 2.0 support. Identity is available throughout the determination of rights and permissions.

- ◆ [Section 28.1, “Selecting a Policy Type,” on page 311](#)
- ◆ [Section 28.2, “Managing Policy Containers,” on page 312](#)
- ◆ [Section 28.3, “Managing Policies,” on page 312](#)
- ◆ [Section 28.4, “Managing a Rule List,” on page 313](#)
- ◆ [Section 28.5, “Enabling Policy Logging,” on page 315](#)

28.1 Selecting a Policy Type

Access Manager uses the policy type to define the context within which a policy is evaluated. Each type of policy differs in purpose, which in turn determines the conditions and actions that apply. For example, the conditions and actions of an authorization policy are going to differ from the conditions and actions of an Identity Injection policy.

When you click *New* on the Policies page, the system displays the predefined policy types in a drop-down list. Each policy type represents the set of conditions and actions that are available. You then configure rules to determine user roles, make decision requests, and enforce authorization decisions. You can also set up policies with no conditions, allowing actions to always take place. As policies and conditions become complex, it can be simpler and more manageable to design policies with conditions that deny or restrict access to large groups of users, rather than setting up policies that permit access to certain users.

The policy types available for Access Manager include:

- ◆ **Access Gateway: Authorization:** This policy type is used to permit or deny access to protected resources, such as Web servers. After you have set up the protected resource, you use the policy rules to define how you want to restrict access. For example, if a user is denied access to a resource, you can use the policy to redirect them to a URL where they can request access to the resource.
- ◆ **Access Gateway: Identity Injection:** This policy type evaluates the rules for Identity Injection, which retrieves identity data from a data source (user store) and forwards it to Web applications. Roles can also be used as data items for Identity Injection and Form Fill.
- ◆ **Access Gateway: Form Fill:** This policy type is used to create a policy that automatically fills in the information required in a form, after the user has accessed the form once.
- ◆ **Identity Server: Roles:** This policy type evaluates rules for establishing the roles of an authenticated user. Roles are generated based on policy statements each time a user authenticates. Roles are placed into an Authentication Profile, which can be used as input in policies for other Novell® Identity Server services.

- ◆ **J2EE Agent: EJB Authorization:** This policy type allows you to create policies that protect Enterprise JavaBeans. You can protect the entire bean or specific interfaces or methods.
- ◆ **J2EE Agent:Web Authorization:** This policy type allows you to create policies that protect the Web applications on a J2EE server.

28.2 Managing Policy Containers

You use policy containers to store and organize policies, similar to how you organize typical files in folders. The *Master_Container* is a permanent policy container, but you can use *Edit Policy Containers* to create new containers for purposes to suit your needs.

In the Administration Console, click *Access Manager > Policies*, then click the *Manage Policy Containers* icon by the *Policy Container* selection box.

On the Container List page, click *New*.

- 1 Name the policy container, then click *OK*.
- 2 Click *Close*.

After you add a policy container, the system displays it in the *Policy Container* drop-down list on the Policy List page.

If you have only one administrator configuring and managing policies, you can create additional policy containers to help you keep them organized. If you have multiple administrators managing and configuring policies, we suggest that you create a policy container for each administrator. This allows multiple administrators to modify policies at the same time. If two or more administrators are modifying policies in the same container at the same time, some modifications will be lost.

You must delete all the policies in a policy container before you can delete the policy container.

28.3 Managing Policies

- ◆ [Section 28.3.1, “Creating Policies,” on page 312](#)
- ◆ [Section 28.3.2, “Deleting Policies,” on page 313](#)
- ◆ [Section 28.3.3, “Importing and Exporting Policies,” on page 313](#)

28.3.1 Creating Policies

Before creating policies, you need to design your policy strategy. For example, if you are going to use role-based access, you need to decide which roles you need and which roles allow access to your protected resources. Roles, which are used by authorization policies that grant and deny access, need to be created first. If you have already created the roles and assigned them to users in your LDAP user store, you can use the values of your role attributes in the authorization policies rather than using Access Manager roles.

To create a policy, see the following sections:

- ◆ [Chapter 29, “Creating Role Policies,” on page 317](#)
- ◆ [Chapter 30, “Creating Authorization Policies,” on page 339](#)
- ◆ [Chapter 31, “Creating Identity Injection Policies,” on page 373](#)
- ◆ [Chapter 32, “Creating Form Fill Policies,” on page 385](#)

28.3.2 Deleting Policies

A policy cannot be deleted as long as a resource is configured to use the policy. For Access Gateway and J2EE Agent policies, this means that you must remove the policy you need to delete from all protected resources.

Roles can be used by Authorization, Form Fill, and Identity Injection policies. Before you can delete a role policy, you must remove any reference to the role from all other policies.

28.3.3 Importing and Exporting Policies

Policies which are created in the Administration Console can be exported and used in another Administration Console that is managing a different group of Access Gateways and other devices. Each policy type has slightly different import requirements. See the following:

- ◆ [Section 29.6, “Importing and Exporting Role Policies,” on page 337](#)
- ◆ [Section 30.7, “Importing and Exporting Authorization Policies,” on page 371](#)
- ◆ [Section 31.8, “Importing and Exporting Identity Injection Policies,” on page 382](#)
- ◆ [Section 32.5, “Importing and Exporting Form Fill Policies,” on page 398](#)

28.4 Managing a Rule List

You configure rules to create a policy. The rules collectively represent a desired course of action when the required conditions are met, such as denying entry-level employees access to a secure Web site, and permitting access for employees who have a role of Manager.

When the system evaluates the policy conditions, it begins with the rule with the highest priority and evaluates the conditions, starting with the first condition group in the rule. Each rule contains one or more conditions and one or more actions. If a rule’s conditions are met, the rule’s action is performed. For some policy types, the performance of any rule’s action terminates the policy evaluation. With authorization policies, for example, after the policy has determined that a user is either permitted or denied access to a resource, there is no reason to evaluate the policy further. However, a role policy may identify multiple roles to which a user belongs. In this case, each rule of the policy must be evaluated to determine all roles to which the user belongs.

IMPORTANT: The interface to the policy engine is designed for flexibility. It does not protect you from creating rules that do nothing because they are always true or always false. For example, you can set up a condition where Client IP is equal to Client IP, which is always true. You are responsible for defining the condition so that it does a meaningful comparison.

You use rules to coordinate how a policy operates, and the behavior varies according to the policy type:

- ◆ [Section 28.4.1, “Rule Evaluation for Role Policies,” on page 314](#)
- ◆ [Section 28.4.2, “Rule Evaluation for Authorization Policies,” on page 314](#)
- ◆ [Section 28.4.3, “Rule Evaluation for Identity Injection and Form Fill Policies,” on page 314](#)

28.4.1 Rule Evaluation for Role Policies

A role policy is used to determine which role or roles a user is assigned to. However, you can specify only one role per rule. Role policies are evaluated when a user authenticates. Role policies do not directly deny or allow access to any resource, nor do they determine if a user is authenticated. A user's role can be used in the evaluation of an authorization policy, but at that point the evaluation of the role policy has already occurred and is not directly part of the authorization process. The performance of an action (assigning a user to a role) does not terminate the evaluation of the policy; meaning, subsequent rules in the authorization policy continue to be evaluated.

28.4.2 Rule Evaluation for Authorization Policies

When the Access Gateway discovers a rule in an Authorization policy that either permits or denies a user access to a protected resource, it stops processing the rules in the policy. Use the following guidelines in determining whether your Authorization policy needs multiple rules:

- ◆ If the policy enforces multiple access requirements that can result in differing actions (either permit or deny), use separate rules to define the conditions and actions.
- ◆ If you want other conditions or actions processed when a rule fails, you must create a second rule for the failure actions or conditions.

If you create multiple rules, you can modify the order that the rules are processed. This allows you to create policies that contain a number of Permit rules that allow access if the user matches the rule. The lowest priority rule in such a policy is a Deny rule, which denies access to everyone who has not previously matched a Permit rule.

You can also create a number of policies and enable multiple policies for the same protected resource. Rule priority determines how the enabled policies interact with each other. The rules in the policies are gathered into one list, then sorted by priority. The same processing rules are applied as if the rules came from one policy. Thus, it is a personal design issue whether you create a policy with multiple rules or create multiple policies that you enable on a single protected resource. Either design produces a list of rules, sorted by priority, that is applied to the user requesting access to the protected resource.

28.4.3 Rule Evaluation for Identity Injection and Form Fill Policies

Rules in Identity Injection and Form Fill policies have actions, but no conditions. Because they have no conditions, all the rules are evaluated and the actions are performed. Identity Injection policies have two exceptions to this rule; they can insert only one authentication header and one cookie header. If you create multiple rules, each with an authentication header and a cookie header, the rule with the highest priority is processed and its actions performed. The actions in the second rule for injecting an authentication header and a cookie header are ignored.

You cannot create multiple rules for a Form Fill policy.

28.5 Enabling Policy Logging

Policy logging is expensive; it uses processing time and disk space. In a production environment, you should enable it only under the following types of conditions:

- ◆ You have created a new policy and need to verify its functionality
- ◆ You are troubleshooting a policy that is not behaving as expected.

To gather troubleshooting information, you should enable the *File Logging* and *Echo To Console* options in the Identity Server configuration and set the *Component File Logger Levels* for *Application* to at least *info*. Then you must update the Identity Server configuration and restart any Access Gateway ESPs, so that the ESPs read the logging options. See [Section 34.2, “Configuring Component Logging,” on page 410](#). When you have solved the problem, you should disable these options.

The log file on the component that executed the policy is where you should look for logging information. For example, if you have an Access Gateway: Authorization error, look at the log on the Access Gateway that executed the policy.

For additional policy troubleshooting procedures, see [Chapter 41, “Troubleshooting Access Manager Policies,” on page 473](#).

This section describes the following topics for Identity Server roles.

- ◆ [Section 29.1, “Understanding RBAC in Access Manager,” on page 317](#)
- ◆ [Section 29.2, “Creating Roles,” on page 320](#)
- ◆ [Section 29.3, “Creating Access Manager Roles from an Existing Role-Based Policy System,” on page 327](#)
- ◆ [Section 29.4, “Mapping Roles between Trusted Providers,” on page 335](#)
- ◆ [Section 29.5, “Enabling and Disabling Role Policies,” on page 336](#)
- ◆ [Section 29.6, “Importing and Exporting Role Policies,” on page 337](#)

29.1 Understanding RBAC in Access Manager

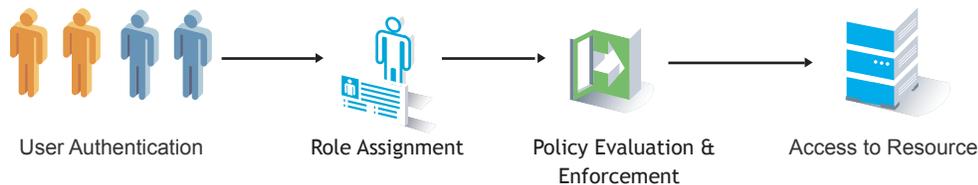
Role-based access control (RBAC) provides a convenient way to assign a user to a particular job function or set of permissions within an enterprise, in order to control access. As an administrator, you probably have defined a set of roles for your needs. Your roles might include Employee, Student, Administrator, Manager, and so on. You might have Web resources that you want available to all employees, or only to managers, as shown in [Figure 29-1](#).

Figure 29-1 Traditional RBAC



Access Manager supports core RBAC functionality by providing user role mapping and the mapping of roles to resource rights and permissions. User role mapping is a primary function of a role policy. Role mapping to resource rights is accomplished through [authorization policies](#) and role settings in J2EE and SSL VPN environments. When creating a role, you assign users to the role, based on attributes of their identity. You also specify the constraints to place on the role.

Figure 29-2 RBAC Using Policy



As shown in [Figure 29-2](#), during user authentication, the system checks the existing role policy to determine which roles that a user must be assigned to. After authentication, assigned roles can be used as evaluated conditions of an authorization policy.

Java applications and Web server applications can also be configured to use roles for access control. For these applications you can use Access Manager to assign the users to the required roles. You can then use the J2EE agent to forward the user’s assigned roles to the Java application, or use Access Gateway Identity Injection policies to inject the assigned roles into the HTTP header that is sent to the Web server.

The following examples describe ways to use roles in Access Manager.

- ◆ “Assigning all Authenticated Users to a Role” on page 318
- ◆ “Using a Role to Create an Authentication Policy” on page 318
- ◆ “Using Prioritized Rules in an Authorization Policy” on page 320

Assigning all Authenticated Users to a Role

The system assigns users to roles when they authenticate. The following example illustrates a role policy that creates an Employee role. All authenticated users are assigned to the role of Employee, because it does not include any conditions (see “Employee Role” on page 323).

Edit Policy: Employee - Rule 1 ?

Type: Identity Server: Roles

Description: Employee Activation Policy

Priority: 1

Conditions Condition structure: AND Conditions, OR group

Condition Group 1

New

No conditions in Rule 1. (Actions will always occur unconditionally.)

Actions

Activate Role

Do Activate Role

: Employee

Changes made on this panel must be applied from the [Policies](#) Panel.

OK Cancel

Role assignment audit events can be created during authentication to the Identity Server. You enabled this on the Logging page in the Identity Server configuration when you enable the *Login Provided* or *Login Consumed* options.

Using a Role to Create an Authentication Policy

The simplest implementation of RBAC policies is to include roles as evaluated conditions when creating authorization policies.

Suppose you belong to a company of 300 employees, and ten of them are managers. You can assign all employees to an Employee role, and make it a condition of an authorization policy with no

restrictions. Such a policy would permit access to Web resources intended for all employees, as shown in the following example:

Edit Policy: Authorize_All - Rule 1 ?

Type: Access Gateway: Authorization
Description: Allow All
Priority: 1

Conditions Condition structure: AND Conditions, OR group

If

Condition Group 1 ✕

New

If Roles for Current User ✕

Comparison: String : Equals
Mode: Case Sensitive
Value: Roles Employee
Result on Condition Error: False

Append New Group

Actions

Do Permit

Changes made on this panel must be applied from the [Policies](#) Panel.

OK **Cancel**

For more sensitive Web resources intended only for managers, you might create a role called Manager. (See “[Manager Role](#)” on page 325). The Manager role might be a condition of an authorization policy that denies access to any employee that has not been assigned to the Manager role when the user authenticated. The following example illustrates this. Notice that the operand for the governing condition logic is set to If Not.

Edit Policy: Deny_All_but_Manager - Rule 1 ?

Type: Access Gateway: Authorization
Description: Deny All but Manager to Web Resource
Priority: 1

Conditions Condition structure: AND Conditions, OR group

If

Condition Group 1 ✕

New

If Not Roles for Current User ✕

Comparison: String : Equals
Mode: Case Sensitive
Value: Roles Manager
Result on Condition Error: False

Append New Group

Actions

Do Deny Deny Message

You are not authorized to access this site.

Changes made on this panel must be applied from the [Policies](#) Panel.

OK **Cancel**

After you have created the authorization policies, you need to assign the policies to the resources they were designed to protect.

See [Section 12.4.2, “Assigning an Authorization Policy to a Protected Resource,” on page 152](#), [“Assigning a Web Authorization Policy to the Resource”](#), and [“Assigning an Enterprise JavaBean Authorization Policy to a Resource”](#).

Using Prioritized Rules in an Authorization Policy

In another policy example, you might create an authorization policy for the Sales Department and set up a list of rules that evaluate whether a user has been authenticated (assigned) to one of the roles associated with the department, and then deny access if the user has not been assigned to any of them, as shown in the Rule List page for the authorization policy below:

Edit Policy: Auth_Policy_for_Sales_Dept ?

Type: Access Gateway: Authorization
Description: Sales Department

Rule List				
New Enable Disable Delete 4 item(s)				
<input type="checkbox"/>	Rule	Priority	Enabled	Action Description
<input type="checkbox"/>	1	1	<input checked="" type="checkbox"/>	Permit Sales Representative
<input type="checkbox"/>	2	2	<input checked="" type="checkbox"/>	Permit Sales Manager
<input type="checkbox"/>	3	4	<input checked="" type="checkbox"/>	Permit Sales President
<input type="checkbox"/>	4	10	<input checked="" type="checkbox"/>	Deny Deny

Changes made on this panel must be applied from the [Policies](#) Panel.

In this example, you specify a first-priority rule with a condition that allows access if a user has been assigned to the role of Sales Representative. You add rules for users assigned to the a role of Sales Manager, Sales Vice President, and so on. You then create a lowest-priority rule that contains no conditions, and an action of Deny. This policy denies any user who has not authenticated to the Sales department roles. When users do not meet the conditions of the rules, the user is denied access by the lowest-priority rule.

For more information on using roles in authorization policies, see [Chapter 30, “Creating Authorization Policies,” on page 339](#).

29.2 Creating Roles

To implement RBAC, you must first define all of the roles within your organization and the permissions attached to each role. A collection of users requiring the same access can be assigned to a single role. Each user can also be assigned to one or more roles and receive the collective rights associated with the assigned roles. A role policy consists of one or more rules, and each rule consists of one or more conditions and an action.

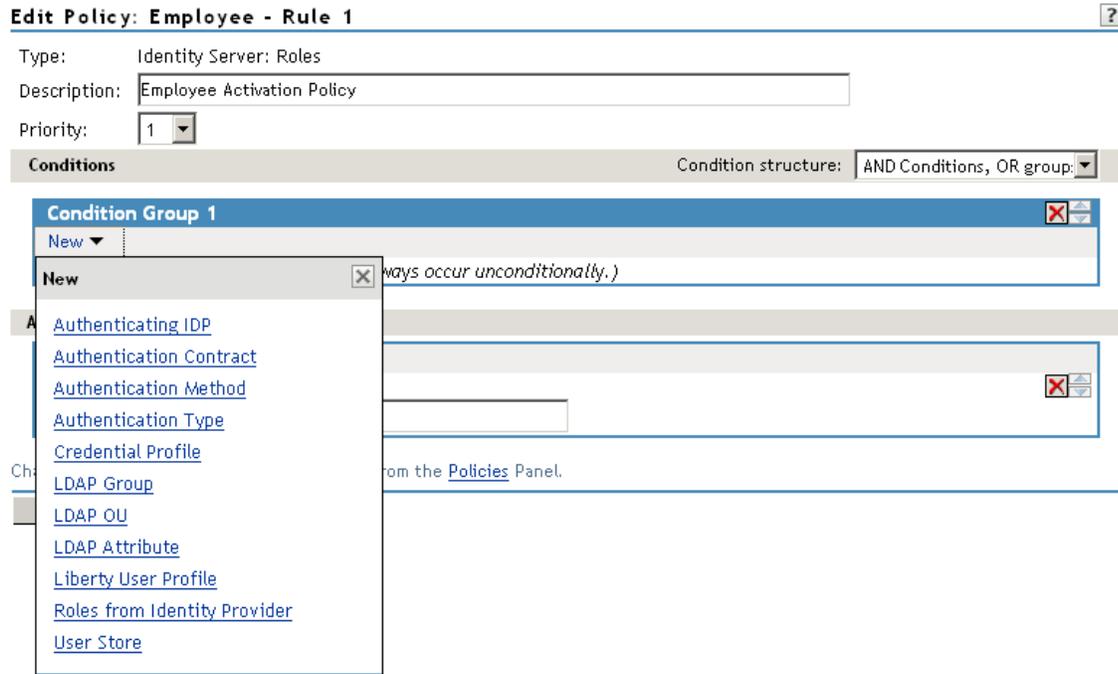
The following topics discuss how to create a role.

- ◆ [Section 29.2.1, “Selecting Conditions,” on page 321](#)
- ◆ [Section 29.2.2, “Selecting an Action,” on page 323](#)
- ◆ [Section 29.2.3, “Reviewing the Rules,” on page 323](#)

- ◆ [Section 29.2.4, “Example Role Policies,” on page 323](#)

29.2.1 Selecting Conditions

You create a role by selecting the appropriate conditions that qualify a user to be assigned to a role, as shown in the following screen.



The following table describes each condition:

Role Condition	Description
Authenticating IDP	Specifies the identity provider that authenticated the current user. This condition is significant when you have implemented identity federation.
Authentication Contract	Specifies the contract used to authenticate the current user. The selections in this list are defined in the Identity Server configuration (<i>Local > Contracts</i>). The <i>Comparison</i> value may be an exact string, the start, the end, or a substring.
Authentication Method	Specifies the method used to authenticate the current user.
Authentication Type	Compares a selected authentication type to the authentication types used to authenticate the current user. <i>[Current]</i> represents the current set of authentication types used to authenticate the user. The other selections represent specific authentication types that can be used to compare with <i>[Current]</i> . The <i>Authentication Type</i> condition returns True if the selected authentication type is contained in the set of authentication types for <i>[Current]</i> . For example, if the current user had been required to satisfy the authentication types of Basic and SmartCard, then a selected authentication type of either Basic or SmartCard would match.

Role Condition	Description
Credential Profile	<p>Specifies the credentials used by the user during authentication. Only values used at authentication time are available for this comparison. The <i>Comparison</i> value can be an exact string, the start, the end, or a substring.</p> <p>If your user store is an Active Directory server, you need to specify the cn attribute for matching the Credential profile in a policy. Even though the user login is chosen from the SAMAccountName, the credential profile is chosen from the cn attribute.</p> <p>The default contracts assign the cn attribute to the Credential profile. If you create your own authentication contract, you can assign a different attribute to the Credential profile.</p>
LDAP Group	<p>Specifies a group in which the authenticating user is evaluated for membership. The value, an LDAP DN, must be a fully distinguished name of a group.</p>
LDAP OU	<p>Specifies an OU for which the authenticating user's DN (distinguished name) is evaluated for containment. The value, an LDAP DN, must be a fully distinguished name of an organizational unit.</p>
LDAP Attribute	<p>Specifies an attribute from the user object of an authenticated user. By default the selection values include those defined by InetOrgPerson.</p>
Liberty User Profile	<p>Specifies any one of a number of data values that have been mapped to a Liberty Profile attribute. To check the mapping of attribute values, click <i>Identity Servers > Liberty > LDAP Attribute Mapping</i>.</p>
Roles from Identity Provider	<p>Roles that are passed from an identity provider to another trusted provider. For example, a service provider requests authentication from an identity provider, which returns a role. This role might be passed to the protected resource as a condition for authorization.</p> <p>This condition uses the mapped attribute All Roles. All roles that are assigned to the user can be mapped to Liberty and SAML 2 attributes and assigned to a trusted identity provider. (See Section 8.7, "Selecting Attributes for a Trusted Provider," on page 97 for information about enabling All Roles.)</p> <p>For an example of how to use <i>Roles from Identity Provider</i>, see Section 29.4, "Mapping Roles between Trusted Providers," on page 335.</p>
User Store	<p>Compares a selected user store with the user store from which the current user is authenticated. The <i>[Current]</i> selection represents the user store from which the user was authenticated. The other selections represent all of the configured user stores that can be used to compare with <i>[Current]</i>.</p> <p>For example, if the configured user stores are eDir1 and AD1 and the current user is authenticated from eDir1, then a selected user store of eDir1 would match and a selected user store of AD1 would not match.</p>

29.2.2 Selecting an Action

The policy action specifies the role to which the user belongs. Roles are activated at the time the role policy is evaluated. In the following screen, the role Employee is specified to be assigned to the user.

Actions

Activate Role

Do Activate Role

: Employee

Changes made on this panel must be applied from the [Policies](#) Panel.

OK Cancel

29.2.3 Reviewing the Rules

After you create roles, they are displayed as rules on the Edit Policy page, where you can review the priority, action, and a description of the role, as shown in the following screen.

Edit Policy: Employee

Type: Identity Server: Roles

Description: Employee Activation

Rule List

New | Enable | Disable | Delete 1 item(s)

<input type="checkbox"/>	Rule	Priority	Enabled	Action	Description
<input type="checkbox"/>	1	1	<input checked="" type="checkbox"/>	Activate Role	Employee Activation Policy

Changes made on this panel must be applied from the [Policies](#) Panel.

OK Cancel

29.2.4 Example Role Policies

The following instructions describe how to create two types of roles: a general Employee role and a restrictive Manager role. These roles can be used in Access Gateway Identity Injection policies and in Access Gateway, J2EE Web applications, and J2EE Enterprise JavaBeans.

- ◆ “Employee Role” on page 323
- ◆ “Manager Role” on page 325

Employee Role

This role policy creates an Employee role. All authenticated users are assigned to this role when they log in (because it does not include conditions) and activates it during authentication. This role can then be used to grant access to resources to all users in your user stores.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Roles > Manage Policies*.

- 2 On the Policies page, click *New*.

Policies ?

View/Edit the policies in the specified container.

Policy Container:

Policy List 3 item(s)

New... | Delete | Rename... | Import... | Export...

New X

Name:

Type:

OK Cancel

- 3 Select a policy type of *Identity Server: Roles* and specify a display name, such as *Employee*.
If this role needs to match the name of a role required by a Java or Web application, ensure that the case of the name matches the application's name.
- 4 Click *OK*.
- 5 On the Edit Policy page, enter a description in the *Description* field.
It is important to use this field to keep track of your roles and policies. The policy feature is powerful, and setup can be as large and complex as you want it to be, with a potentially unlimited number of conditions and choices. This description is useful to help keep track of various role and policy configurations.
- 6 Make sure the *Condition Group 1* section has no conditions, so that all users who authenticate match the condition.

Edit Policy: employee - Rule 1 ?

Type:

Description:

Priority:

Conditions Condition structure:

Condition Group 1 X

New

Actions

Activate Role

Do

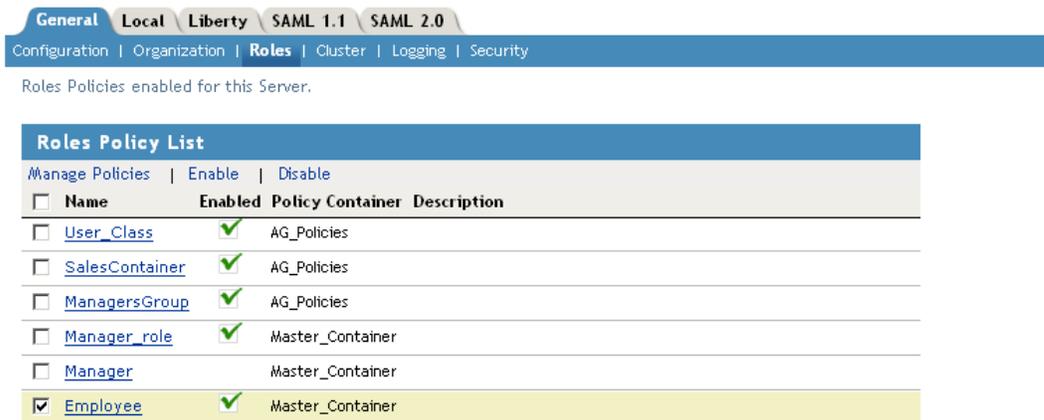
:

Changes made on this panel must be applied from the [Policies](#) Panel.

OK Cancel

- 7 In the *Actions* section, click *Activate Role*.
- 8 In the *Activate Role* box, type *Employee*, then click *OK*.
This entry is case sensitive and must match the name used in other applications that require the role.
- 9 On the Edit Policy (Rule List) page, click *OK*.

- 10 On the Policies page, click *Apply Changes*, then click *Close*.

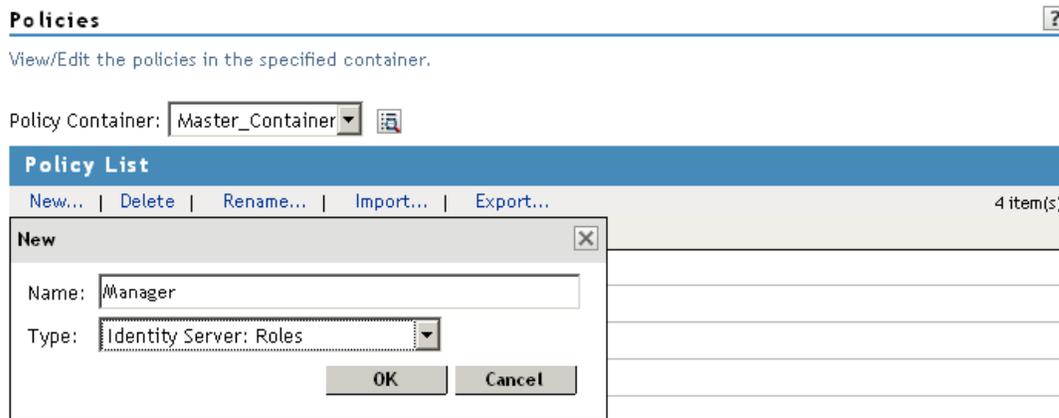


- 11 On the Role Policy page, select the Employee role, then click *Enable*
- 12 On the *Setup* tab, click *Update Servers*.
This step updates the Identity Server configuration, which is required after you create a role.
- 13 To create a Manager role, continue with [Section , “Manager Role,” on page 325](#).

Manager Role

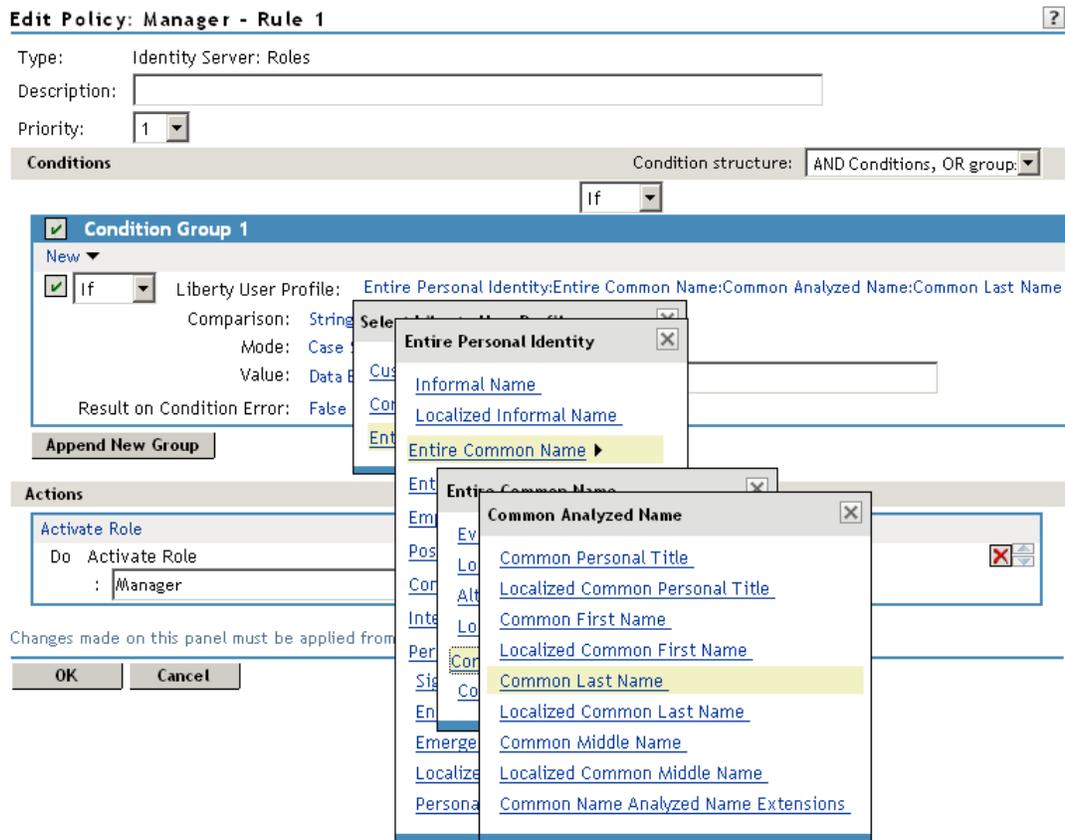
Because the Manager role is restrictive, role policy conditions must be specified. The Manager role is assigned only to the users who meet the conditions.

- 1 Click *Identity Servers > Setup > [Configuration] > Roles > Manage Policies*.
- 2 On the Policies page, click *New*.



- 3 Select a policy type of *Identity Server: Roles* and specify a display name (for this example, Manager.)
- 4 Click *OK*.

5 In the *Conditions* section, click *New > Liberty User Profile*.



6 In the *Condition Group 1*, select the conditions the user must meet:

Liberty User Profile: Select *Entire Personal Identity > Entire Common Name > Common Analyzed Name > Common Last Name*.

Comparison: Select how you want the attribute values to be compared. For the *employeeType* attribute, select *String > Equals*.

Mode: Select *Case Insensitive*.

Value: Select *Data Entry Field* and type the person's name in the box (Smith, in this example). This sets up the condition that if the user has the name Smith, his or her role as Manager is activated at authentication.

Result on Condition Error: This sets up the results that are returned if an error occurs while evaluating the condition (for example, the LDAP server goes down). This rule is set up to grant the user the role of Manager if the condition evaluates to *True*. If an error occurs, you do not want random users assigned the role of Manager. Therefore, for this rule, you need to select *False*.

7 In the Actions section, click *Activate Role*.

Edit Policy: Manager - Rule 1 ?

Type: Identity Server: Roles

Description:

Priority:

Conditions Condition structure: AND Conditions, OR group:

Condition Group 1

New

Liberty User Profile: Entire Personal Identity:Entire Common Name:Common Analyzed Name:Common Last Name

Comparison:

Mode:

Value: :

Result on Condition Error:

Append New Group

Actions

Activate Role

Do Activate Role ✖

:

Changes made on this panel must be applied from the [Policies](#) Panel.

8 In the *Activate Role* box, type *Manager*, then click *OK* twice.

9 On the *Policies* page, click *Apply Changes*.

General Local Liberty SAML 1.1 SAML 2.0

Configuration | Organization | **Roles** | Cluster | Logging | Security

Roles Policies enabled for this Server.

Roles Policy List			
Manage Policies Enable Disable			
<input type="checkbox"/> Name	Enabled	Policy Container	Description
<input type="checkbox"/> User_Class	<input checked="" type="checkbox"/>	AG_Policies	
<input type="checkbox"/> SalesContainer	<input checked="" type="checkbox"/>	AG_Policies	
<input type="checkbox"/> ManagersGroup	<input checked="" type="checkbox"/>	AG_Policies	
<input type="checkbox"/> Manager_role	<input checked="" type="checkbox"/>	Master_Container	
<input checked="" type="checkbox"/> Manager	<input checked="" type="checkbox"/>	Master_Container	
<input type="checkbox"/> Employee	<input checked="" type="checkbox"/>	Master_Container	

10 Select the *Manager* role, then click *Enable*

11 On the *Setup* tab, click *Update Servers*.

29.3 Creating Access Manager Roles from an Existing Role-Based Policy System

If you have already implement a role-based administration policy for granting access to print, file, and LDAP resources, you can leverage your role definitions and use Access Manager policies to

control access to Web resources. If your role definitions use the following types of LDAP features, you can create Access Manager Role policies that use them:

- ◆ The values found in LDAP attributes
- ◆ The location of the user objects in the directory tree
- ◆ Membership in groups or roles

The Access Manager Role policies that you create using these features can then be used to control access to protected Web resources.

- ◆ [Section 29.3.1, “Creating a Role Using an LDAP Attribute,” on page 328](#)
- ◆ [Section 29.3.2, “Creating a Role Using the Location of the User Objects,” on page 330](#)
- ◆ [Section 29.3.3, “Creating a Role Using Role or Group Membership,” on page 332](#)

29.3.1 Creating a Role Using an LDAP Attribute

You can assign a user to a role by using a value found in any LDAP attribute in your directory. The following example uses the `objectClass` attribute because every object in an LDAP directory has an `objectClass` attribute that contains the object classes to which the object belongs. This attribute contains the name of the object class that was used to create the object as well as the names of the superior object classes of this class. All you need to know is the name of the object class you used to create your users in the LDAP directory. For example, the following instructions create a Role policy for users who were created with the User object class.

- 1 In the Administration Console, click *Access Manager > Policies*.
- 2 Click *New*, specify a name for the Role policy, select *Identity Server: Roles* for the type, then click *OK*.
- 3 In Condition Group 1, click *New*, then select *LDAP Attribute*.
- 4 In the *Condition Group 1*, select the conditions the user must meet:

LDAP Attribute: Select the `objectClass` attribute. If you have not added this attribute, it won't appear in the list. Scroll to the bottom of the list, click *New*, specify `objectClass` for the name, then click *OK*.

If you are using eDirectory for your LDAP directory, you need to specify standard LDAP names for the attributes. Access Manager does not support spaces or colons in attribute names.

Comparison: Select how you want the attribute values to be compared. For the `objectClass` attribute, select *String > Contains Substring*.

The `objectClass` attribute is a multi-valued attribute and, for most objects, contains multiple values. For example in eDirectory, users created with the User object class have `User`, `organizationalPerson`, `person`, `ndsLoginProperties`, and `top` as values in the `objectClass` attribute.

Mode: Select *Case Insensitive*.

Value: Select *Data Entry Field* and specify `User` as the value.

Result on Condition Error: This sets up the results that are returned if an error occurs while evaluating the condition (for example, the LDAP server goes down). This rule is set up to grant the user the role of `UserClass` if the condition evaluates to *True*. If an error occurs, you do not want random users assigned the role of `UserClass`. Therefore, for this rule, you need to select *False*.

- 5 In the *Actions* section, click *Activate Role*.
- 6 In the *Activate Role* box, enter `User`, then click *OK*

The name you enter in the box is the role you want assigned to the users who match the condition.

Your rule should look similar to the following:

Type: Identity Server: Roles

Description: Object class rule for UserClass role

Priority: 1

Conditions Condition structure: AND Conditions, OR group:

If

Condition Group 1

New

If LDAP Attribute: objectClass Comparison: String : Equals Mode: Case Insensitive Value: Data Entry Field : user

Result on Condition Error: False

Append New Group

Actions

Activate Role

Do Activate Role : UserClass

Changes made on this panel must be applied from the [Policies](#) Panel.

OK Cancel

- 7 Click *OK* twice, then click *Apply Changes*.
- 8 To enable the role so that it can be used in Authorization and Identity Injection policies, click *Identity Servers > [Name of Configuration] > Roles*.
- 9 Select the checkbox by the name of the role, then click *Enable*.
- 10 Click *OK*.
- 11 To update the Identity Server, click *Setup > Update Servers*.

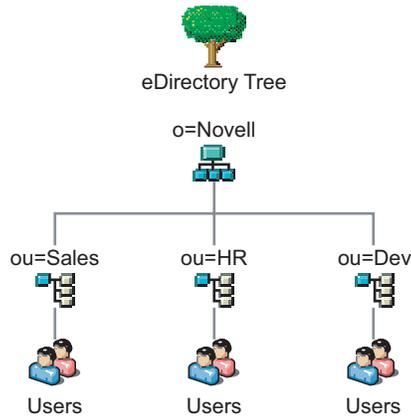
You can now use this role when creating Authorization and Identity Injection policies, which control access to protected Web resources. For more information, see

- ♦ [Chapter 30, “Creating Authorization Policies,” on page 339](#)
- ♦ [Chapter 31, “Creating Identity Injection Policies,” on page 373](#)

29.3.2 Creating a Role Using the Location of the User Objects

If you have created your users in specific containers in your LDAP tree, you can use these container objects to assign users to roles. For example, suppose your LDAP tree looks similar to the following tree.

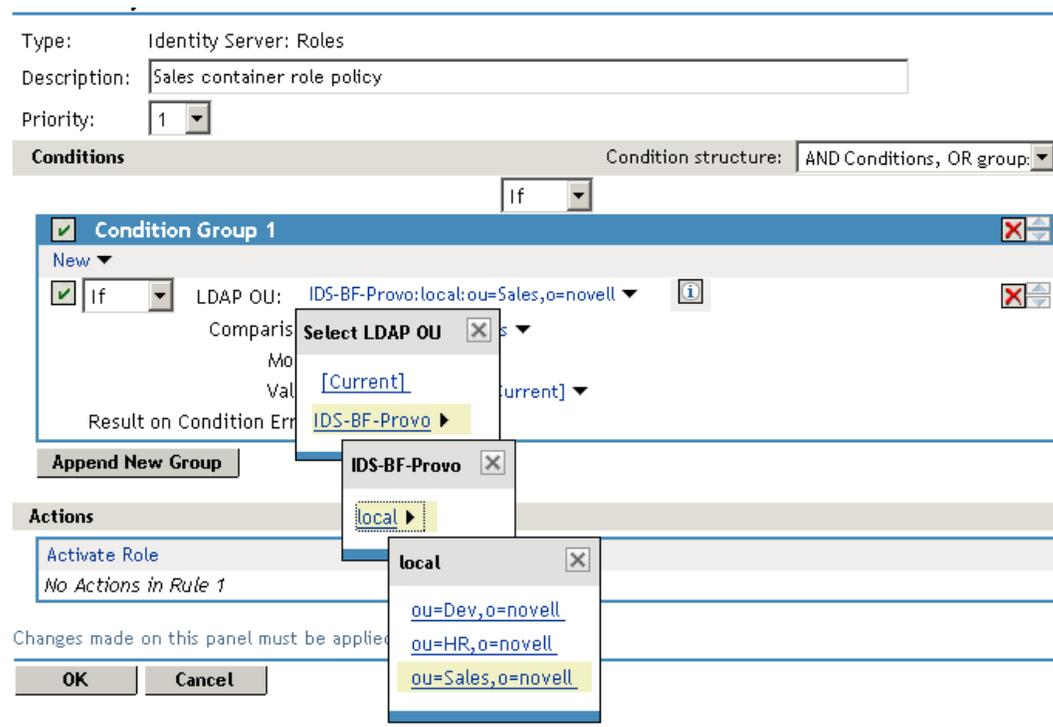
Figure 29-3 Using an eDirectory Tree for access control



Such a tree organization can be used to control access to resources. The following instructions explain how to create a Role policy for the users created under the Sales container.

- 1 In the Administration Console, click *Access Manager > Policies*.
- 2 Click *New*, specify a name for the Role policy, select *Identity Server: Roles* for the type, then click *OK*.
- 3 In *Condition Group 1*, click *New*, and select *LDAP OU > [Identity Server Configuration] > [Replica] > [DN of OU]*.

The following example illustrates how to make these selections:



Comparison: Select how you want the attribute values to be compared. For LDAP OU, select *Contains*.

Mode: Select *One Level* if all your users are created in ou=Sales. Select *Subtree* if your users are created in various containers under the ou=Sales container.

Value: Select *LDAP OU*, then select *[Current]*.

The DN of the authenticated user is compared with the value specified in LDAP OU. If the DN of the user contains the LDAP OU value, the user matches the condition. For example, if the DN of the user is cn=bsmith,ou=sales,o=novell and the LDAP OU value is ou=sales,o=novell, the user matches the condition. If you selected *Subtree* for the Mode, a user with the following DN also matches the condition: cn=djones,ou=provo,ou=sales,o=novell.

Result on Condition Error: This sets up the results that are returned if an error occurs while evaluating the condition (for example, the LDAP server goes down). This rule is set up to grant the user the role of Sales if the condition evaluates to *True*. If an error occurs, you do not want random users assigned the role of Sales. Therefore, for this rule, you need to select *False*.

- 4 In *Condition Group 1*, select the conditions the user must meet:
- 5 In the *Actions* section, click *Activate Role*.
- 6 In the *Activate Role* box, enter `Sales`, then click *OK*

The name you enter in the box is the role you want assigned to the users who match the condition.

Your rule should look similar to the following:

Type: Identity Server: Roles
Description: Sales container role policy
Priority: 1

Conditions Condition structure: AND Conditions, OR group: If

Condition Group 1

New

If LDAP OU: IDS-BF-Provo:local:ou=Sales,o=novell
Comparison: LDAP OU : Contains
Mode: One Level
Value: LDAP OU [Current]
Result on Condition Error: False

Append New Group

Actions

Activate Role
Do: Activate Role
: Sales

Changes made on this panel must be applied from the Policies Panel.

OK Cancel

- 7 Click *OK* twice, then click *Apply Changes*.
- 8 To enable the role so that it can be used in Authorization and Identity Injection policies, click *Identity Servers > [Configuration] > Roles*.
- 9 Select the checkbox by the name of the role, then click *Enable*.
- 10 Click *OK*.
- 11 To update the Identity Server, click *Setup > Update Servers*.

You can now use this role when creating Authorization and Identity Injection policies, which control access to protected Web resources. For more information, see

- ♦ [Chapter 30, “Creating Authorization Policies,” on page 339](#)
- ♦ [Chapter 31, “Creating Identity Injection Policies,” on page 373](#)

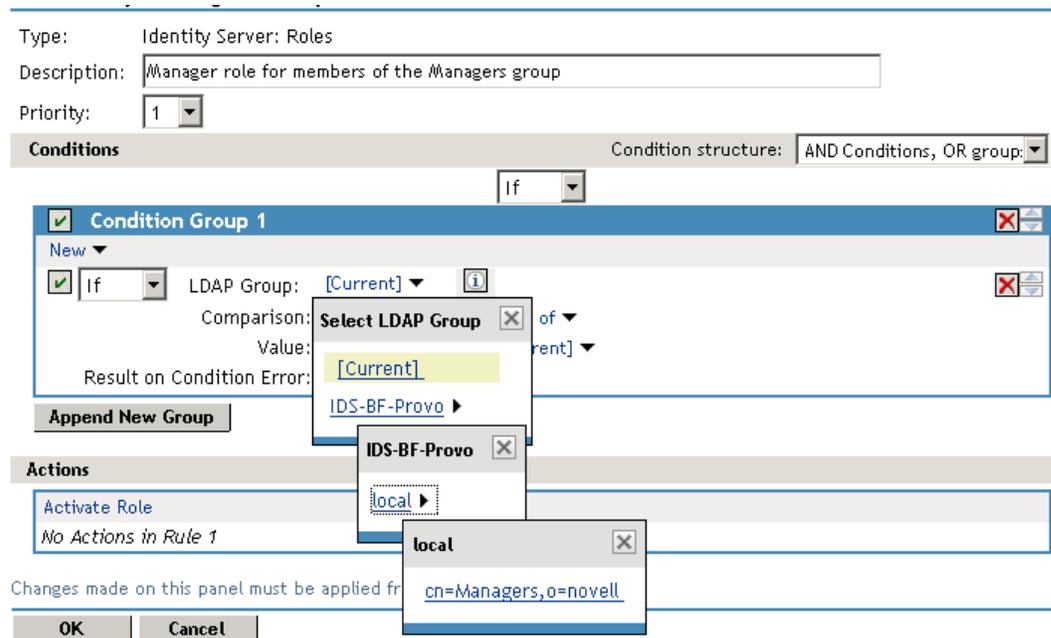
29.3.3 Creating a Role Using Role or Group Membership

If you have created an LDAP group and assigned users to the group, you can use whether a user is a member of a group to assign a role to the user. For example, you might have created a first level managers group and made all your first level managers a member of this group. You would have other groups for your upper level managers. You can create a Role policy that assigns the user a role if the user is a member of a specific group. The Role policy can then be used in an Authorization or Identity Injection policy to protect a Web resource.

- 1 In the Administration Console, click *Access Manager > Policies*.
- 2 Click *New*, specify a name for the Role policy, select *Identity Server: Roles* for the type, then click *OK*.

- 3 In *Condition Group 1*, click *New*, then select *LDAP Group*.
- 4 In *Condition Group 1*, select the conditions the user must meet:

LDAP Group: Select the Identity Server Configuration, the replica, then the Group. The following figure illustrates this selection process.



Comparison: Select how you want the attribute values to be compared. For LDAP Group, select *Is Member of*.

Value: Select *LDAP Group*, then select *[Current]*.

The DN of the authenticated user is compared with the members of the LDAP Group. If the DN of the user matches one of the members, the user matches the condition.

Result on Condition Error: This sets up the results that are returned if an error occurs while evaluating the condition (for example, the LDAP server goes down). This rule is set up to grant the user the role of ManagersGroup if the condition evaluates to *True*. If an error occurs, you do not want random users assigned the role of ManagersGroup. Therefore, for this rule, you need to select *False*.

- 5 In the *Actions* section, click *Activate Role*.
- 6 In the *Activate Role* box, enter `Managers Group`, then click *OK*.

The name you enter in the box is the role you want assigned to the users who match the condition.

Your rule should look similar to the following:

Type: Identity Server: Roles

Description: Manager role for members of the Managers group

Priority: 1

Conditions Condition structure: AND Conditions, OR group

If

Condition Group 1

New

If LDAP Group: IDS-BF-Provo:local:cn=Managers,o=novell

Comparison: LDAP Group : Is Member of

Value: LDAP Group [Current]

Result on Condition Error: False

Append New Group

Actions

Activate Role

Do Activate Role

: ManagersGroup

Changes made on this panel must be applied from the Policies Panel.

OK Cancel

- 7 Click *OK* twice, then click *Apply Changes*.
- 8 To enable the role so that it can be used in Authorization and Identity Injection policies, click *Identity Servers > [Name of Configuration] > Roles*.
- 9 Select the checkbox by the name of the role, then click *Enable*.
- 10 Click *OK*.
- 11 To update the Identity Server, click *Setup > Update Servers*.

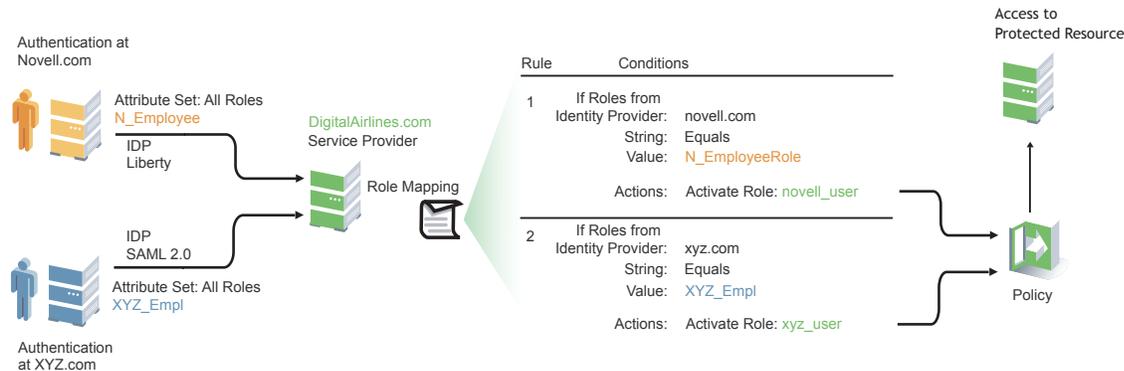
You can now use this role when creating Authorization and Identity Injection policies, which control access to protected Web resources. For more information, see:

- ♦ Chapter 30, “Creating Authorization Policies,” on page 339
- ♦ Chapter 31, “Creating Identity Injection Policies,” on page 373

29.4 Mapping Roles between Trusted Providers

The Identity Server can send roles in an authentication assertion. You can map these roles that are received from trusted providers to your own roles. [Figure 29-4](#) illustrates this process.

Figure 29-4 Role Mapping



In this example, employees authenticate to identity providers novell.com (Liberty) or xyz.com (SAML 2.0). Each user is assigned to a role (such as N_EmployeeRole or XYZ_Empl, respectively). Attribute sets at each of the identity providers are configured to exchange the *All Roles* attribute with the trusted service provider, DigitalAirlines.com. DigitalAirlines.com consumes the authentication assertions, then maps the incoming roles to local roles. The mapped roles at DigitalAirlines.com can be used as evaluated conditions in authorization or J2EE policies, which can provide access to resources intended for the authenticated employees.

Prerequisites

- ◆ Configure trust between trusted providers, using the Liberty or SAML 2.0 protocol.
You should be familiar with [Chapter 8, “Configuring Trusted Providers,”](#) on page 87.
- ◆ Configure local authentication.
You must create an external contract at the service provider that matches the contract of the identity provider. See [Chapter 7, “Configuring Local Authentication,”](#) on page 71.
- ◆ Create an attribute set and select the local attribute *All Roles* in the set. This must be done at the identity provider and service provider.
This attribute set is used to pass roles from an identity provider to an external service provider in authentication assertions. See [Section 6.1, “Configuring Attribute Sets,”](#) on page 65.

The following procedure describes how the service provider configures this type of role policy for novell.com, mapping N_EmployeeRole to:

- 1 In the Administration Console, click *Access Manager > Policies*.
- 2 Click *New*, then specify a name for the Role policy.
- 3 Select *Identity Server: Roles* for the type, then click *OK*.

4 Configure the role policy as shown on the following page.

Edit Policy: Novell_Employees - Rule 1 ?

Type: Identity Server: Roles
Description: Novell
Priority: 1

Conditions Condition structure: AND Conditions, OR group: ▾

If ▾

Condition Group 1 ✕

New ▾

If ▾ Roles from Identity Provider: Novell IDP Liberty i ✕

Comparison: String : Equals ▾

Mode: Case Sensitive ▾

Value: Data Entry Field ▾ : N_EmployeeRole

Result on Condition Error: False ▾

Append New Group

Actions

Activate Role

Do Activate Role

: novell_user ✕

Changes made on this panel must be applied from the [Policies](#) Panel.

OK **Cancel**

- 5 In the *Conditions* section, click *New > Roles from Identity Provider*.
- 6 Select the trusted identity provider in the drop-down menu.
- 7 For *Comparison*, choose *String > Equals*.
- 8 Choose *Value > Data Entry Field*.
- 9 Type the name of the role used by the trusted identity provider.
- 10 Under the *Actions* section, click *Activate Role*.
- 11 Type the name of the role you want to activate at the trusted service provider.
- 12 Click *OK*.
- 13 On the *Policies* page, click *Apply Changes*.
- 14 Click *Identity Servers > Setup > Update Servers*.

29.5 Enabling and Disabling Role Policies

In order for a role policy to function, you must enable it for the Identity Server configuration.

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Roles*.
- 2 Click the role policy's check box, then click *Enable*.
- 3 To disable the role policy, click the role policy's check box, then click *Disable*.
- 4 After enabling or disabling role policies, update the Identity Server configuration on the Setup tab.

29.6 Importing and Exporting Role Policies

You can import and export role policies in order to run them in other Identity Server configurations. When you import a role, ensure that you have enabled any Liberty profile that is referenced in the role policy, in order to correctly display the policy in the interface. However, the policy still evaluates if you have not enabled the profile.

You must also enable roles after importing them to an Identity Server configuration. See [Section 29.5, “Enabling and Disabling Role Policies,” on page 336](#).

When you export a role policy, the system saves it as a `.txt` file at the location of your choosing. After you import a role policy, you must update the Identity Server configuration.

To export a role policy:

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Roles > Manage Policies*.
- 2 Select a policy, then click *Export*.
- 3 (Optional) Modify the name suggested for the file.
- 4 Click *OK*.
- 5 Using the features of your browser, specify where the file is copied.

To import a role policy:

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > [Configuration] > Roles > Manage Policies*.
- 2 Click *Import*, then browse to the location of the file.
- 3 Click *OK*.
- 4 When the policy appears in the list, click *Apply Changes*.

Authorization policies are used when you want to protect a resource based on criteria other than authentication, and you want Access Manager to enforce the access restrictions. Authorization policies are enforced when a user requests data from a resource.

The Access Manager supports three types of Authorization policies:

- ♦ **Access Gateway Authorization policies** for protecting resources of the Access Gateway
- ♦ **Web Authorization policies** for protecting Java applications on a J2EE server
- ♦ **Enterprise JavaBean Authorization policies** for protecting the Enterprise JavaBeans of a J2EE application

The first step in creating an Authorization policy is determining the criteria for restricting access. The second step is translating those criteria into rules and conditions for a policy. This section describes the policy elements, but your resource and your security requirements determine which elements to use when creating the policy.

- ♦ **Section 30.1, “Designing an Authorization Policy,” on page 339**
- ♦ **Section 30.2, “Creating Access Gateway Authorization Policies,” on page 344**
- ♦ **Section 30.3, “Creating Web Authorization Policies for J2EE Agents,” on page 350**
- ♦ **Section 30.4, “Creating Enterprise JavaBean Authorization Policies for J2EE Agents,” on page 351**
- ♦ **Section 30.5, “Conditions,” on page 352**
- ♦ **Section 30.6, “Using Multiple Conditions,” on page 370**
- ♦ **Section 30.7, “Importing and Exporting Authorization Policies,” on page 371**

30.1 Designing an Authorization Policy

When creating an Authorization policy, you need to configure one or more rules. Each rule consists of two parts: (1) one or more conditions the user must meet and (2) the action to perform when the user meets the conditions or doesn't meet the conditions. The action can be to either allow or deny access to the resource. This section describes how to use the following elements when creating a policy:

- ♦ **Section 30.1.1, “Many Rules or Many Conditions,” on page 340**
- ♦ **Section 30.1.2, “Controlling Access with a Deny Rule,” on page 340**
- ♦ **Section 30.1.3, “Controlling Access with Multiple Conditions,” on page 341**
- ♦ **Section 30.1.4, “Using Permit Rules with a Deny Rule,” on page 342**
- ♦ **Section 30.1.5, “Public Policies,” on page 344**
- ♦ **Section 30.1.6, “General Design Principles,” on page 344**
- ♦ **Section 30.1.7, “Assigning Policies to Resources,” on page 344**

30.1.1 Many Rules or Many Conditions

You can design your policy to have many rules with a single condition and action, or you can design your policy to have fewer rules, with each rule containing many conditions.

For example, suppose you have a resource that you don't want users accessing on Monday, Wednesday, and Friday between 1 am and 2 am. You could set up three rules, one for each day, or you could set up one rule with three conditions. If all the conditions have the same action (for example, deny access with the same reason), it is simpler to put them in the same rule. However, if you have a customized message to return for each day, you need to put them in separate rules.

Each rule contains the following:

- ♦ Zero or more conditions. A condition specifies how the request data is evaluated for a true or false match. Conditions are evaluated in the order in which they are listed.
- ♦ One or more condition groups. Conditions are placed in condition groups, which gives you the flexibility of creating a policy that allows the user to match the conditions in one group but not the conditions in the other condition groups. Or you can set up the condition groups to require that the user matches at least one condition in each condition group.
- ♦ An action, which either denies or grants access to the users who match the conditions.

Conditions, conditions groups, and the interaction among them allow you to create very simple rules (if A, then grant access) to very complex rules (if A, B, and C, but not D and E, then grant access). For more information on how multiple conditions and condition groups can interact with each other, see [Section 30.6, "Using Multiple Conditions," on page 370](#).

30.1.2 Controlling Access with a Deny Rule

As you design your policies, remember that authenticated users are allowed access to protected resources unless the policy denies access. To deny access to the correct set of users, you need to know the characteristics of the users you don't want accessing the resource, as well as the characteristics of the users you want accessing the resource.

Some very simple policies can be created using a deny access action. For example, suppose you have an application that you only want managers to access. If you have set up a role that assigns all

managers to the Manager role, you can use this characteristic for an Authorization policy. The policy requires only one rule, which would be similar to the following:

Figure 30-1 Simple Rule

The screenshot displays a configuration window for a policy rule. At the top, the 'Conditions' section is active, showing a 'Condition structure' dropdown set to 'AND Conditions, OR group'. Below this, a 'Condition Group 1' is defined with a structure of 'If'. The condition itself is 'If Not Roles for Current User', with a comparison of 'String : Equals', a mode of 'Case Sensitive', and a value of 'Roles' and 'Manager'. The 'Result on Condition Error' is set to 'False'. An 'Append New Group' button is visible below the condition group. The 'Actions' section below shows a 'Do Deny' action with a message box containing the text 'You must be a manager to access this resource.' At the bottom of the window, there are 'OK' and 'Cancel' buttons, and a note stating 'Changes made on this panel must be applied from the Policies Panel.'

This rule evaluates the user, and if the user does not belong to the Manager role, the user matches the condition. The action for matching the condition is to deny access. The managers, who belong to the Manager role, do not match the condition. Because the rule does not apply to them, they are allowed access to the resource.

30.1.3 Controlling Access with Multiple Conditions

Your policy becomes more complicated if you want only some of the managers to have access to the application. With the *Condition structure* set to *OR Conditions, AND groups*, you'll need to add another condition group. This second condition should match the group of managers you don't want accessing the application. If this application is a sales management application, you could set up this second condition to match if the user has the Manager role but does not belong to the Sales

department. This rule would allow managers who belong to the Sales department to have access to the application. Such a rule would look similar to the following:

Figure 30-2 A Rule with Two Condition Groups

The screenshot displays a rule configuration window with the following elements:

- Conditions Section:**
 - Condition structure:** AND Conditions, OR group
 - Condition Group 1:**
 - Operator: If Not
 - Field: Roles for Current User
 - Comparison: String : Equals
 - Mode: Case Sensitive
 - Value: Roles / Manager
 - Result on Condition Error: True
 - Condition Group 2:**
 - Operator: Or
 - Condition 1:
 - Operator: If
 - Field: Roles for Current User
 - Comparison: String : Equals
 - Mode: Case Sensitive
 - Value: Roles / Manager
 - Result on Condition Error: True
 - Condition 2:
 - Operator: And If Not
 - Field: Liberty User Profile: Corporate Employment Identity:Department
 - Comparison: String : Equals
 - Mode: Case Insensitive
 - Value: Data Entry Field : sales
 - Result on Condition Error: True
 - Append New Group** button
- Actions Section:**
 - Operator: Do
 - Action: Deny
 - Display Default Deny Page

The second condition group matches all of the users who are managers but don't belong to the sales department, and the match causes all of these users to be denied access to the resource.

This second condition group could be implemented as the second rule of the policy. If so, it should be set as a lower priority than the first rule. Because most systems would have more users than managers, the user rule would be used more frequently, so it should come first.

30.1.4 Using Permit Rules with a Deny Rule

You can also create policies that contain one or more Permit rules and then as the lowest priority rule in the policy, a Deny rule with no conditions. In such a policy, as soon as an allow match is processed, the rest of the rules are not processed and the user is granted access to the resource. The Deny rule is only processed if the user does not match one of the allow rules, and because all users

match a rule with no conditions, the user is denied access to the resource. The first rule in such a policy for the sales application would look similar to the following.

Figure 30-3 Rule 1 Granting Access

The screenshot displays the configuration for Rule 1. At the top, the 'Type' is 'Access Gateway: Authorization', the 'Description' is 'Sales Department Permit Rule', and the 'Priority' is set to 1. The 'Conditions' section is set to 'AND Conditions, OR group' and contains a single 'If' condition group. This group has two conditions: 'Roles for Current User' with comparison 'String : Equals', mode 'Case Sensitive', and value 'Roles' (not equal to) 'Manager'; and 'Liberty User Profile: Corporate Employment Identity:Department' with comparison 'String : Equals', mode 'Case Insensitive', and value 'Data Entry Field' (equals) 'Sales'. The 'Actions' section is set to 'Do Permit'.

The conditions in Rule 1 are ANDed, which requires the user to match both conditions before they are granted access to the resource. The priority is set to 1, so this rule is the first rule that the Access Gateway processes. The J2EE authorization policies use the same logic.

The second rule would look similar to the following.

Figure 30-4 Rule 2 Denying Access

The screenshot displays the configuration for Rule 2. The 'Type' is 'Access Gateway: Authorization', the 'Description' field is empty, and the 'Priority' is set to 4. The 'Conditions' section is set to 'AND Conditions, OR group' and contains no conditions, with a message stating 'No conditions in Rule 2. (Actions will always occur unconditionally.)'. The 'Actions' section is set to 'Do Deny' with the 'Deny Message' 'Access is restricted to Sales Managers.'.

Changes made on this panel must be applied from the [Policies](#) Panel.

OK Cancel

Because this rule has no conditions, any user who did not match the first rule, matches this rule and is denied access. The priority of this rule is set lower than the Permit rule so that the Permit rule is processed first.

30.1.5 Public Policies

You can create public authorization policies, which are policies that apply to everyone, by leaving the *Condition* section empty. In the *Action* section, you specify either to deny or to permit access to the resource. Then you assign the policy to the protected resource.

30.1.6 General Design Principles

When designing a policy, remember the following principles:

1. Logged-in users are allowed access to a protected resource unless the policy denies access.
2. Priority determines the order in which rules are applied.
3. Rules are only processed until user matches the conditions in a rule and its action is applied. If a user matches the first rule in a policy, that action is applied, and the rest of the rules in the policy are ignored.
4. If two rules have the same priority, Deny rules are applied before Permit rules.

After you have designed your policy, created it, and assigned it to a resource, you need to test the policy. You need to log in as the type of user who should be granted access and as the type of user who should not be granted access.

30.1.7 Assigning Policies to Resources

For information on how to assign the policy to a resource, see the following:

- ♦ For an Access Gateway policy, see [Section 12.4.2, “Assigning an Authorization Policy to a Protected Resource,” on page 152.](#)
- ♦ For a Web Authorization policy, see [“Assigning a Web Authorization Policy to the Resource” in the *Novell Access Manager 3.0 J2EE Agent Guide*.](#)
- ♦ For a Enterprise JavaBean Authorization policy, see [“Assigning an Enterprise JavaBean Authorization Policy to a Resource” in the *Novell Access Manager 3.0 J2EE Agent Guide*.](#)

30.2 Creating Access Gateway Authorization Policies

An Authorization policy specifies conditions that a user must meet in order to access a resource. The Access Gateway enforces these conditions. The policy specifies the criteria a user must meet to either allow access or deny access. This section describes the following:

- ♦ [Section 30.2.1, “The Process,” on page 345](#)
- ♦ [Section 30.2.2, “Sample Policies,” on page 346](#)

30.2.1 The Process

To create an Authorization policy:

- 1 In the Administration Console, click *Access Manager > Policies > New*.
- 2 Specify a name for the policy, then select *Access Gateway: Authorization* for the type of policy.
- 3 Configure the following fields:
 - Description:** (Optional) Describe the purpose of this policy.
 - Priority:** Specify the order in which a rule is applied in the policy, when the policy has multiple rules. The highest priority is 1 and 10 is the lowest. If two rules have the same priority, a Deny rule is applied before a Permit rule.
- 4 In the *Condition Group 1* section, click *New*, then select one of the following:
 - ♦ **Authentication Contract:** Allows you to control access based on the contract the user used for login. For configuration information, see [Section 30.5.1, “Authentication Contract Condition,”](#) on page 353.
 - ♦ **Client IP:** Allows you to control access based on the IP address of the client making the request. For configuration information, see [Section 30.5.2, “Client IP Condition,”](#) on page 355.
 - ♦ **Credential Profile:** Allows you to control access based on the credentials the user specified during authentication. For configuration information, see [Section 30.5.3, “Credential Profile Condition,”](#) on page 356.
 - ♦ **Current Date:** Allows you to control access based on the date of the request. For more information, see [Section 30.5.4, “Current Date Condition,”](#) on page 357.
 - ♦ **Current Day of Month:** Allows you to control access based on the month the request is made. For configuration information, see [Section 30.5.6, “Current Day of Month Condition,”](#) on page 359.
 - ♦ **Current Day of Week:** Allows you to control access based on the day the request is made. For configuration information, see [Current Day of Week Condition.](#)
 - ♦ **Current Time of Day:** Allows you to control access based on the time the request was made. For configuration information, see [Section 30.5.7, “Current Time of Day Condition,”](#) on page 359.
 - ♦ **HTTP Request Method:** Allows you to control access based on the request method. For configuration information, see [Section 30.5.8, “HTTP Request Method Condition,”](#) on page 360.
 - ♦ **LDAP Attribute:** Allows you to control access based on the value of an LDAP attribute. For configuration information, see [Section 30.5.9, “LDAP Attribute Condition,”](#) on page 361.
 - ♦ **Liberty User Profile:** Allows you to control access based on the value of a profile attribute. For configuration information, see [Section 30.5.10, “Liberty User Profile Condition,”](#) on page 362.
 - ♦ **Roles for Current User:** Allows you to control access based on the roles a user has been assigned. For configuration information, see [Section 30.5.11, “Roles for Current User Condition,”](#) on page 363.
 - ♦ **URL:** Allows you to control access based on the URL in the request. For configuration information, see [Section 30.5.12, “URL Condition,”](#) on page 363.

- ♦ **URL Scheme:** Allows you to control access based on the scheme in the URL of the request (for example, http or https). For configuration information, see [Section 30.5.13, “URL Scheme Condition,” on page 364.](#)
 - ♦ **URL Host:** Allows you to control access based on the hostname in the URL of the request. For configuration information, see [Section 30.5.14, “URL Host Condition,” on page 365.](#)
 - ♦ **URL Path:** Allows you to control access based on the path in the URL of the request. For configuration information, see [Section 30.5.15, “URL Path Condition,” on page 366.](#)
 - ♦ **URL File Name:** Allows you to control access based on the filename in the URL of the request. For configuration information, see [Section 30.5.16, “URL File Name Condition,” on page 367.](#)
 - ♦ **URL File Extension:** Allows you to control access based on the file extension in the URL of the request. For configuration information, see [Section 30.5.17, “URL File Extension Condition,” on page 368.](#)
 - ♦ **X-Forwarded-For IP:** Allows you to control access based on the value in the X-Forwarded-For IP header of the HTTP request. For configuration information, see [Section 30.5.18, “X-Forward-For IP Condition,” on page 369.](#)
- 5** To add multiple conditions to the same rule, either add a condition to the same condition group or create a new condition group. For information on how conditions and condition groups interact with each other, see [Section 30.6, “Using Multiple Conditions,” on page 370.](#)
- 6** In the *Actions* section, select either *Permit* or *Deny*. If you select *Deny*, select one of the following:
- ♦ **Display Default Deny Page:** Displays a generic message, indicating that users have insufficient rights to access the resource.
 - ♦ **Deny Message:** Allows you to enter a customized message that is displayed to users who are denied access.
 - ♦ **Redirect to URL:** Allows you to specify a URL that users are redirected to when they are denied access. For example:
http://www.novell.com
- 7** To save the rule, click *OK* twice, then click *Apply Changes*.
- 8** For information on how to assign the policy to a protected resource, see [Section 12.4.2, “Assigning an Authorization Policy to a Protected Resource,” on page 152.](#)

30.2.2 Sample Policies

The following sections describe a scenario and then describe two types of policies that enforce the requirements of the scenario:

- ♦ [“Company Scenario” on page 346](#)
- ♦ [“LDAP Context Policies” on page 347](#)
- ♦ [“Role Policies with Authorization Policies” on page 348](#)

Company Scenario

Suppose that the company LDAP directory has the following organization.

ou=sales,o=acme

ou=dev,o=acme

ou=hr,o=acme

Suppose that this company has the following configuration and requirements:

- ◆ Under each branch of the tree, the system administrator has created the users who work in these departments.
- ◆ Each department has its own Web resources, and other departments must be denied access to these resources.

With this type of configuration, you can use the LDAP context condition to create authorization policies or you can create role policies that are used in conjunction with authorization policies.

LDAP Context Policies

With such an organization, you can create a policy that either allows or denies access based on the LDAP context of the user's DN. You can use the LDAP context of the user DN to separate the users into their departments and then grant access based on the context match. You need to create protected resources for the Web resources of the department, create a policy for each protected resource, and assign a policy to the protected resources.

The following procedure explains how to configure such a policy for the sales department.

- 1 Click *Policies > New*, specify a name for the policy, select *Access Gateway: Authorization* as the type, then click *OK*.
- 2 For *Condition Group 1*, click *New*, then select *Credential Profile*.
- 3 Fill in the following fields:
 - LDAP Credentials:** Select *LDAP User DN*.
 - If/If Not:** Select *If Not*.
 - Comparison:** Select *Contains Substring*.
 - Mode:** Select *Case Insensitive*.
 - Value:** Select *Data Entry Field*. In the text box, type the following value:
ou=sales,o=acme
 - Result on Condition Error:** Select *True*.
- 4 In the *Actions* section, select *Deny*.

Your policy should look similar to the following:

Type: Access Gateway: Authorization
Description: LDAP context policy
Priority: 1

Conditions Condition structure: AND Conditions, OR group:
If

Condition Group 1
New
 If Not Credential Profile: LDAP Credentials:LDAP User DN
Comparison: String : Contains Substring
Mode: Case Sensitive
Value: Data Entry Field : ou=sales,o=acme
Result on Condition Error: True

Append New Group

Actions

Do Deny Deny Message
You do not belong to the sales department.

This sets up the condition so that the following occurs:

- ◆ When the user does not belong to the sales department, the user is denied access.
 - ◆ When the user belongs to the sales department, the user is granted access.
 - ◆ When an error occurs evaluating the conditions in the rule, the user is denied access.
- 5 Assign the policy to the protected Web resources of the sales department (see [Section 12.4.2, “Assigning an Authorization Policy to a Protected Resource,”](#) on page 152).
 - 6 Repeat these steps for the other two departments, changing the *Value* field to match the appropriate department.

Role Policies with Authorization Policies

With the company’s organization, you need to create three role policies, one for the sales users, one for the development users, and one for the human resource users. You can then use these roles as conditions in authorization policies to allow and deny access. The first time you use roles in an authorization policy, there is extra set up because you have to create the role policies. But after the role policies are created, you can use them in multiple authorization policies.

The following instructions explain how to use the Sales role to create a policy that controls access to a protected resource. For instructions on how to create the Sales role, see [Section 29.3.2, “Creating a Role Using the Location of the User Objects,”](#) on page 330.

You need to decide on the type of Authorization policy you want to create. For example, you can create a Deny policy, which denies access to everyone who does not match the condition (in this case, the Sales role). Or you can create a two-rule policy which allows access to everyone that matches the condition. The first rule grants access to everyone who has the Sales role, and the second rule denies access to everyone who did not match the conditions of the first rule. (Other

methods are also possible.) Because the proposed Deny policy is very similar to the LDAP context policy example, the following procedures explain how to create the two-rule policy.

- 1 In the Administration Console, click *Access Manager > Policies > New*.
- 2 Specify a name for the policy, select *Access Gateway: Authorization* as the type, then click *OK*.
- 3 (Optional) Enter a description for the rule.
- 4 In *Condition Group 1*, click *New*, and select *Roles for Current User*.

5 Fill in the following fields:

If/If Not: Select *If*.

Comparison: Select *String: Equals*.

Mode: Select *Case Insensitive*

Value: Select *Roles*, then select *Sales*.

Result on Condition Error: Select *False*.

- 6 Under *Actions*, select *Permit*, then click *OK*.

These steps create the Permit rule and sets up the condition so that the following occurs:

- ♦ When the user does not match the condition because the user does not belong to the Sales role, the policy engine moves to the next rule in the policy.
- ♦ When the user does match the condition because the user belongs to the Sales role, the user is granted access.
- ♦ If an error occurs evaluating the condition of the policy, the user does not match the condition and the policy engine moves to the next rule in the policy.

- 7 In the *Rule List*, click *New*.

This second rule is for denying access to everyone who does not match the condition in Rule 1. Processing of the policy stops when a user matches a rule; therefore all users who match Rule 1 are granted access and the policy engine does not evaluate the second rule.

- 8 Set the *Priority* to be 2 or greater.

You want the Permit rule to be processed first, so it should have a priority of 1. The Deny rule needs to be processed last, so it needs a lower priority than the Permit rule.

- 9 Leave the *Condition Group 1* empty.

The *Conditions* section is left empty so that everyone who does not match the conditions of the Permit rule are denied access to the resource

- 10 In the *Actions* section, select *Deny* and either accept the default action or select one of the other actions.
- 11 Click *OK* twice.
- 12 Click *Apply Changes* on the Policies page.
- 13 Assign the policy to the protected Web resources of the sales department (see [Section 12.4.2, "Assigning an Authorization Policy to a Protected Resource,"](#) on page 152).

30.3 Creating Web Authorization Policies for J2EE Agents

A Web Authorization policy specifies conditions that a user must meet in order to access a resource on a J2EE server. The Web Authorization policy specifies the criteria a user must meet to either allow access or deny access. For example, if you create a sales role and assign it to the users, the role can be used to allow access to the sales applications and to deny access to resource management applications. For information about designing a policy, see [Section 30.1, “Designing an Authorization Policy,” on page 339](#).

To create a Web Authorization policy:

- 1 In the Administration Console, click *Access Manager > Policies > New*.
- 2 Specify a name for the policy, select *J2EE Agent: Web Authorization* as the type, then click *OK*.
- 3 Fill in the following fields:
 - Description:** (Optional) Specify a description for the rule.
 - Priority:** Specify the order in which a rule is applied in the policy, when the policy has multiple rules. The highest priority is 1 and 10 is the lowest. If two rules have the same priority, a Deny rule is applied before a Permit rule.
- 4 In the *Condition Group 1* section, click *New*, then select one of the following:
 - ♦ **Client IP Address:** Allows you to control access based on the IP address of the client making the request. For configuration information, see [Section 30.5.2, “Client IP Condition,” on page 355](#).
 - ♦ **Credential Profile:** Allows you to control access based on the credentials the user specified during authentication. For configuration information, see [Section 30.5.3, “Credential Profile Condition,” on page 356](#).
 - ♦ **Current Date:** Allows you to control access based on the date of the request. For more information, see [Section 30.5.4, “Current Date Condition,” on page 357](#).
 - ♦ **Current Day of Month:** Allows you to control access based on the month the request is made. For configuration information, see [Section 30.5.6, “Current Day of Month Condition,” on page 359](#).
 - ♦ **Current Day of Week:** Allows you to control access based on the day the request is made. For configuration information, see [Section 30.5.5, “Current Day of Week Condition,” on page 358](#).
 - ♦ **Current Time of Day:** Allows you to control access based on the time the request was made. For configuration information, see [Section 30.5.7, “Current Time of Day Condition,” on page 359](#).
 - ♦ **HTTP Request Method:** Allows you to control access based on the request method. For configuration information, see [Section 30.5.8, “HTTP Request Method Condition,” on page 360](#).
 - ♦ **LDAP Attribute:** Allows you to control access based on the value of an LDAP attribute. For configuration information, see [Section 30.5.9, “LDAP Attribute Condition,” on page 361](#).
 - ♦ **Liberty User Profile:** Allows you to control access based on the value of a profile attribute. For configuration information, see [Section 30.5.10, “Liberty User Profile Condition,” on page 362](#).

- ♦ **Roles for Current User:** Allows you to control access based on the roles a user has been assigned. For configuration information, see [Section 30.5.11, “Roles for Current User Condition,” on page 363.](#)
 - ♦ **URL:** Allows you to control access based on the URL in the request. For configuration information, see [Section 30.5.12, “URL Condition,” on page 363.](#)
 - ♦ **URL Scheme:** Allows you to control access based on the scheme in the URL of the request (for example, http or https). For configuration information, see [Section 30.5.13, “URL Scheme Condition,” on page 364.](#)
 - ♦ **URL Host:** Allows you to control access based on the hostname in the URL of the request. For configuration information, see [Section 30.5.14, “URL Host Condition,” on page 365.](#)
 - ♦ **URL Path:** Allows you to control access based on the path in the URL of the request. For configuration information, see [Section 30.5.15, “URL Path Condition,” on page 366.](#)
 - ♦ **URL File Name:** Allows you to control access based on the filename in the URL of the request. For configuration information, see [Section 30.5.16, “URL File Name Condition,” on page 367.](#)
 - ♦ **URL File Extension:** Allows you to control access based on the file extension in the URL of the request. For configuration information, see [Section 30.5.17, “URL File Extension Condition,” on page 368.](#)
 - ♦ **X-Forwarded-For IP:** Allows you to control access based on the value in the X-Forwarded-For IP header of the HTTP request. For configuration information, see [Section 30.5.18, “X-Forward-For IP Condition,” on page 369.](#)
- 5 To add multiple conditions to the same rule, either add a condition to the same condition group or create a new condition group. For information on how conditions and condition groups interact with each other, see [Section 30.6, “Using Multiple Conditions,” on page 370.](#)
 - 6 In the *Actions* section, select either *Permit* or *Deny*.
 - 7 To save the rule, click *OK* twice, then *Apply Changes*.
 - 8 Assign the policy to a Web resource. See “[Assigning a Web Authorization Policy to the Resource](#)” in the *Novell Access Manager 3.0 J2EE Agent Guide*

30.4 Creating Enterprise JavaBean Authorization Policies for J2EE Agents

An Enterprise JavaBean (EJB) Authorization policy allows you to protect the entire bean or specific interfaces or methods. For information about designing a policy, see [Section 30.1, “Designing an Authorization Policy,” on page 339.](#)

To create an EJB Authorization policy:

- 1 In the Administration Console, click *Access Manager > Policies > New*.
- 2 Specify a name for the policy, select *J2EE Agent: EJB Authorization* as the type, then click *OK*.
- 3 Fill in the following fields:

Description: (Optional) Specify a description for the rule.

Priority: Specify the order in which a rule is applied in the policy, when the policy has multiple rules. The highest priority is 1 and 10 is the lowest. If two rules have the same priority, a Deny rule is applied before a Permit rule.

- 4 In the *Condition Group 1* section, click *New*, then select one of the following:
 - ♦ **Credential Profile:** Allows you to control access based on the credentials the user specified during authentication. For configuration information, see [Section 30.5.3, “Credential Profile Condition,”](#) on page 356.
 - ♦ **Current Date:** Allows you to control access based on the date of the request. For more information, see [Section 30.5.4, “Current Date Condition,”](#) on page 357.
 - ♦ **Current Day of Month:** Allows you to control access based on the month the request is made. For configuration information, see [Section 30.5.6, “Current Day of Month Condition,”](#) on page 359.
 - ♦ **Current Day of Week:** Allows you to control access based on the day the request is made. For configuration information, see [Section 30.5.5, “Current Day of Week Condition,”](#) on page 358.
 - ♦ **Current Time of Day:** Allows you to control access based on the time the request was made. For configuration information, see [Section 30.5.7, “Current Time of Day Condition,”](#) on page 359.
 - ♦ **LDAP Attribute:** Allows you to control access based on the value of an LDAP attribute. For configuration information, see [Section 30.5.9, “LDAP Attribute Condition,”](#) on page 361.
 - ♦ **Liberty User Profile:** Allows you to control access based on the value of a profile attribute. For configuration information, see [Section 30.5.10, “Liberty User Profile Condition,”](#) on page 362.
 - ♦ **Roles for Current User:** Allows you to control access based on the roles a user has been assigned. For configuration information, see [Section 30.5.11, “Roles for Current User Condition,”](#) on page 363.
- 5 To add multiple conditions to the same rule, either add a condition to the same condition group or create a new condition group. For information on how conditions and condition groups interact with each other, see [Section 30.6, “Using Multiple Conditions,”](#) on page 370.
- 6 In the *Actions* section, select either *Permit* or *Deny*.
- 7 To save the rule, click *OK*, then *Apply Changes*.
- 8 Assign the policy to a EJB resource. See “[Assigning an Enterprise JavaBean Authorization Policy to a Resource](#)” in the *Novell Access Manager 3.0 J2EE Agent Guide*

30.5 Conditions

This section describes the possible conditions for an Authorization policy. For the specific policies they can be used in, see the following:

- ♦ [Section 30.2, “Creating Access Gateway Authorization Policies,”](#) on page 344
- ♦ [Section 30.3, “Creating Web Authorization Policies for J2EE Agents,”](#) on page 350
- ♦ [Section 30.4, “Creating Enterprise JavaBean Authorization Policies for J2EE Agents,”](#) on page 351

Some conditions can be set up so that the current values in the request are compared against static values (A to B), or you can compare static values to current values in the request (B to A). Within one policy, you should probably decide which direction to set up the comparisons and remain consistent unless there is a compelling reason to switch the direction for a particular condition.

For example, suppose you set up a rule to allow access to a resource only during the weekdays (Monday through Friday). You set up four of these conditions to compare if the current date of when the request is made matches with Monday, Tuesday, Wednesday, or Thursday. You set up the fifth condition to compare whether Friday matches the current date of when the request is made. This works, but the maintenance of the this policy is more difficult because each new policy manager will ponder about the Friday condition and wonder why it is configured differently.

Many conditions, when used as the sole condition of a rule, do not make very useful rules. For example, you can create a rule that grants access if the user specifies a specific URL in the request. Such a rule has limited application. But a rule that requires that the request contain a specific URL and that the user have a specific role has greater application because it can be used to limit access to the URL based on the user's role. For information about how conditions can be ANDed or ORed together or placed in different condition groups, see [Section 30.6, "Using Multiple Conditions," on page 370](#).

Authorization policies use the following conditions:

- ◆ [Section 30.5.1, "Authentication Contract Condition," on page 353](#)
- ◆ [Section 30.5.2, "Client IP Condition," on page 355](#)
- ◆ [Section 30.5.3, "Credential Profile Condition," on page 356](#)
- ◆ [Section 30.5.4, "Current Date Condition," on page 357](#)
- ◆ [Section 30.5.5, "Current Day of Week Condition," on page 358](#)
- ◆ [Section 30.5.6, "Current Day of Month Condition," on page 359](#)
- ◆ [Section 30.5.7, "Current Time of Day Condition," on page 359](#)
- ◆ [Section 30.5.8, "HTTP Request Method Condition," on page 360](#)
- ◆ [Section 30.5.9, "LDAP Attribute Condition," on page 361](#)
- ◆ [Section 30.5.10, "Liberty User Profile Condition," on page 362](#)
- ◆ [Section 30.5.11, "Roles for Current User Condition," on page 363](#)
- ◆ [Section 30.5.12, "URL Condition," on page 363](#)
- ◆ [Section 30.5.13, "URL Scheme Condition," on page 364](#)
- ◆ [Section 30.5.14, "URL Host Condition," on page 365](#)
- ◆ [Section 30.5.15, "URL Path Condition," on page 366](#)
- ◆ [Section 30.5.16, "URL File Name Condition," on page 367](#)
- ◆ [Section 30.5.17, "URL File Extension Condition," on page 368](#)
- ◆ [Section 30.5.18, "X-Forward-For IP Condition," on page 369](#)

30.5.1 Authentication Contract Condition

The *Authentication Contract* condition matches the URI of the contract the user logged in with to the URI of the contract specified in this condition. The Identity Server has the following default contracts:

URI	Contract
basic/name/password/uri	Name/Password - Basic

URI	Contract
name/password/uri	Name/Password - Form
secure/basic/name/password/uri	Secure Name/Password - Basic
secure/name/password/uri	Secure Name/Password - Form

To configure other contracts for your system, click *Identity Servers > Setup > [Name of Configuration] > Local > Contracts*.

To specify an *Authentication Contract* condition, fill in the following fields:

Authentication Contract: To compare the contract that the user used with a static value, select *Current*. To compare a static value with what the user used, select a contract from the list.

Comparison: Select one of the following types:

- ◆ **Comparison: String:** Specifies that you want the values compared as strings and how you want the string values compared. Select one of the following:
 - ◆ **Equals:** Indicates that the values must match, letter for letter.
 - ◆ **Starts with:** Indicates that the *Authentication Contract* value must begin with the letters specified in the *Value* field.
 - ◆ **Ends with:** Indicates that the *Authentication Contract* value must end with the letters specified in the *Value* field.
 - ◆ **Contains Substring:** Indicates that the *Authentication Contract* value must contain the letters, in the same sequence, as specified in the *Value* field.
- ◆ **Comparison: Regular Expression: Matches:** Specifies that you want the values compared as regular expressions.

Mode: Select the mode appropriate for the comparison type:

- ◆ **Comparison: String:** Specify whether case is important by selecting *Case Sensitive* or *Case Insensitive*.
- ◆ **Comparison: Regular Expression: Matches:** Select one or more of the following:
 - Canonical Equivalence
 - Case Insensitive
 - Comments
 - Dot All
 - Multi-Line
 - Unicode
 - Unix Lines

Value: Specify the value you want to compare with the *Authentication Contract* value. If you selected a static value for the *Authentication Contract* value, select *Authentication Contract* and *Current*. If you selected *Current* for the *Authentication Contract* value, select *Authentication Contract*, then the name of a contract.

Other value types are possible if you selected *Current* for the *Authentication Contract* value. For example:

- ◆ You can select *Data Entry Field* and specify the name of the contract in the text box.
- ◆ If you have defined a Liberty User Profile attribute for authentication contracts, you can select *Liberty User Profile* and your defined attribute.
- ◆ If you have defined an LDAP attribute for authentication contracts, you can select *LDAP Attribute* and your defined attribute.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.2 Client IP Condition

The *Client IP* condition allows you to use the IP address of the user making the request to determine whether the user is allowed access to a resource. Fill in the following fields:

Comparison: Specify how you want the client's IP address compared by selecting one of the following:

- ◆ **Equals:** Allows you to specify an IP address that the client must match. You can specify more than one.
- ◆ **In Range:** Allows you to specify a range of IP addresses that the client's address must fall within. You can specify more than one range.
- ◆ **In Subnet:** Allows you to specify the subnet that the client's address must belong to. You can specify more than one subnet

Value: Select *Data Entry Field* and specify a value appropriate for your comparison type. Use the Edit button to access a text box where you can enter multiple values, each on a separate line. Use the Add button to add values one at a time. For example:

Comparison Type	Value
Equals	10.10.10.10 10.10.10.11
In Range	10.10.10.10 - 10.10.10.100 10.10.20.10 - 10.10.20.100
In Subnet	10.10.10.12 / 22 10.10.20.30 / 22

Other values types are possible. For example, if your user store contains an LDAP attribute with the IP address of your users, you could select to compare the client's current IP address with the stored value by using an LDAP attribute or a Liberty User Profile value.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.3 Credential Profile Condition

The Credential Profile condition allows you to control access based on the credentials the user entered when authenticating to the system.

Fill in the following fields:

Credential Profile: Specify the type of credential your users are using for authentication. Select one of the following:

- ◆ **LDAP Credentials:** If you prompt the user for a user name, select this option, then *LDAP User Name* (the cn attribute of the user), *LDAP User DN* (the fully distinguished name of the user), or *LDAP Password*.

If your user store is an Active Directory server, you cannot use an LDAP credential for the user name. Active Directory does not use the cn or the dn attributes for login, but instead uses the SAMAccountName attribute. You need to use the LDAP Attribute condition and specify the SAMAccountName attribute. The SAMAccountName is not a member of the default set of attribute names derived from the inetOrgPerson class; therefore it must be added before it can be used. See [Section 30.5.9, “LDAP Attribute Condition,” on page 361](#).

- ◆ **X509 Credentials:** If you prompt the user for a certificate, select this option, then select one of the following:
 - ◆ **X509 Public Certificate Subject:** Retrieves the subject field from the certificate, which can match the DN of the user, depending upon who issued the certificate.
 - ◆ **X509 Public Certificate Issuer:** Retrieves the issuer field from the certificate, which is the name of the CA that issued the certificate.

Comparison: Select one of the following types:

- ◆ **Comparison: String:** Specifies that you want the values compared as strings and how you want the string values compared. Select one of the following:
 - ◆ **Equals:** Indicates that the values must match, letter for letter.
 - ◆ **Starts with:** Indicates that the *Credential Profile* value must begin with the letters specified in the *Value* field.
 - ◆ **Ends with:** Indicates that the *Credential Profile* value must end with the letters specified in the *Value* field.
 - ◆ **Contains Substring:** Indicates that the *Credential Profile* value must contain the letters, in the same sequence, as specified in the *Value* field.
- ◆ **Comparison: Regular Expression: Matches:** Specifies that you want the values compared as regular expressions.

Mode: Select the mode appropriate for the comparison type:

- ◆ **Comparison: String:** Specify whether case is important by selecting *Case Sensitive* or *Case Insensitive*.
- ◆ **Comparison: Regular Expression: Matches:** Select one or more of the following:
 - Canonical Equivalence
 - Case Insensitive
 - Comments
 - Dot All

Multi-Line
Unicode
Unix Lines

Value: Specify the second value for the comparison. Select one of the following data types:

- ♦ **LDAP Attribute:** If you have an LDAP attribute that corresponds to the Credential Profile you have specified, select this option and the attribute.
- ♦ **Liberty User Profile:** If you have a Liberty User Profile attribute that corresponds to the Credential Profile you have specified, select this option and the attribute.
- ♦ **Data Entry Field:** Specify the string you want matched. If you have selected LDAP User DN or X509 Public Certificate Subject, you can specify the name of an object that you expect to be in the DN of a group of users, and use this match to control access for a group of users. For example you could specify

`ou=sales`

This allows all users with `ou=sales` as part of their DN to match the condition.

Other values are possible. Your policy requirements determine whether they are useful.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.4 Current Date Condition

The *Current Date* condition allows you to use the date to determine whether the user is allowed access to a resource. Fill in the following fields:

Comparison: Specify how the current date is compared to the data in the *Value* field. Select one of the following:

- ♦ **Equals:** Requires that the current date must equal the specified value.
- ♦ **Greater Than:** Requires that the current date after than the specified value.
- ♦ **Greater Than or Equal to:** Requires that the current date be after or equal to the specified value.
- ♦ **Less Than:** Requires that the current date be before the specified value.
- ♦ **Less Than or Equal to:** Requires that the current date be before or equal to the specified value.

Date Format: Specify the format of the *Value* field for the *Data Entry Field* value type. D specifies a number from 1 to 31. M specifies a number from 1 to 12 or the name of the month in three letters (Sep) or complete (September). Y specifies the year in a four digit format. Select one of the following formats:

- ♦ **D/M/Y** = 1/Jul/2006 or 1/7/2006
- ♦ **D-M-Y** = 1-Jul-2006 or 1-7-2006
- ♦ **D.M.Y** = 1.Jul.2006 or 1.7.2006
- ♦ **M/D/Y** = Jul/1/2006 or 7/1/2006

- ♦ **M-D-Y** = Jul-1-2006 or 7-1-2006
- ♦ **M.D.Y** = Jul.1.2006 or 7.1.2006
- ♦ **YYYY-MM-DD** = 2006-07-01
- ♦ **YYYY.MM.DD** = 2006.07.01

Value: Specify the second value for the comparison. If you select *Data Entry Field* as the value type, specify the date in the format you select in the *Date Format* field.

Other value types are possible. If you have set up a Liberty User Profile attribute that corresponds to the day of the month, you can use this option and select your attribute. The *Date Format* field does not apply to these value types.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.5 Current Day of Week Condition

The *Current Day of Week* condition allows you to restrict access based on which day of the week the request is made. Fill in the following fields:

Current Day of Week: Select the name of the day from the list. To compare the day specified in the current request with a static value, select *Current*. To compare a static value with the day specified in the current request, select the name of a day from the list

Comparison: Specify how the current day of the week is compared to the data in the *Value* field. Select one of the following.

- ♦ **Equals:** Allows you to specify a day that the client must match.
- ♦ **In Range:** Allows you to specify a range of days that the client's request must fall within, for example, Monday to Friday.

Value: Specify the second value for the comparison. If you selected *Current* for the *Current Day of Week* field, you need to specify a static value. If you selected a static value for the *Current Day of the Week* field, you need to selected *Current* for the *Value* field. If you select *Data Entry Field* as the value type, days of the week are specified in the following format:

Sun or Sunday
 Mon or Monday
 Tue or Tuesday
 Wed or Wednesday
 Thu or Thursday
 Fri or Friday
 Sat or Saturday

If you selected *In Range* as the *Comparison* type, specify the first day of the range in the left text box and the end day of the range in the right text box.

Other value types are possible. If you have set up a Liberty User Profile or an LDAP attribute that corresponds to a day or a day of the week, you can use this option and select your attribute.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you

want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.6 Current Day of Month Condition

The *Current Day of Month* condition allows you to restrict access based on the day of the month the request is made. Fill in the following fields:

Comparison: Specify how the current day of the month is compared to the data in the *Value* field. Select one of the following:

- ◆ **Equals:** Allows you to specify a day that the client must match.
- ◆ **In Range:** Allows you to specify a range of days that the client's request must fall within.

Value: Specify the second value for the comparison:

- ◆ If you selected *Equals* for the *Comparison* type, normally you would select *Data Entry Field* for the *Value* field and specify a number from 1 to 31 in the text box.
- ◆ If you selected *In Range* for the *Comparison* type, normally you would select *Data Entry Field* for the *Value* field and specify the first value of the range in the first text box and the second value of the range in the second text box. If you specified 1 in the first box and 15 in the second box, you can use this condition to restrict access between the first day of the month and the 15th day.

Other value types are possible. If you have set up a Liberty User Profile or and LDAP attribute that corresponds to a day or a day of the month, you can use this option and select your attribute.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.7 Current Time of Day Condition

The *Current Time of Day* condition allows you to restrict access based on the time the request is made. Fill in the following fields:

Comparison: Specify how you want the current time compared to the data in the *Value* field. Select one of the following:

- ◆ **Greater Than:** Requires that the current time is greater than the specified value.
- ◆ **Greater Than or Equal to:** Requires that the current time is greater than or equal to the specified value.
- ◆ **Less Than:** Requires that the current time is less than the specified value.
- ◆ **Less Than or Equal to:** Requires that the current time is less than or equal to the specified value.
- ◆ **In Range:** Requires that the current time must fall within the specified range, such as 08:00 and 17:00.

Time Zone: Select the time zone, either *Local* or *GMT*.

Value: Specify the second value for the comparison. If you select *Data Entry Field* as the value type, hours and minutes are specified in the following format:

hour:minute

Hour is a number from 00 to 23, minute is a number from 00 to 59.

Time can only be specified in a 24-hour clock format. For example, 8 am is 08:00 and 5:30 pm is 17:30.

Other value types are possible. If you have set up a Liberty User Profile or an LDAP attribute that corresponds to a date or time of day, you can use this option and select your attribute.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.8 HTTP Request Method Condition

The HTTP Request Method condition allows you to restrict accessed based on the request method in the current request.

HTTP Request Method: Select the request method from the list or select *Current* to specify the method in the current request.

Comparison: Select one of the following types:

- ◆ **Comparison: String:** Specifies that you want the values compared as strings and how you want the string values compared. Select one of the following:
 - ◆ **Equals:** Indicates that the values must match, letter for letter.
 - ◆ **Starts with:** Indicates that the *HTTP Request Method* value must begin with the letters specified in the *Value* field.
 - ◆ **Ends with:** Indicates that the *HTTP Request Method* value must end with the letters specified in the *Value* field.
 - ◆ **Contains Substring:** Indicates that the *HTTP Request Method* value must contain the letters, in the same sequence, as specified in the *Value* field.
- ◆ **Comparison: Regular Expression: Matches:** Specifies that you want the values compared as regular expressions.

Mode: Select the mode appropriate for the comparison type:

- ◆ **Comparison: String:** Specify whether case is important by selecting *Case Sensitive* or *Case Insensitive*.
- ◆ **Comparison: Regular Expression: Matches:** Select one or more of the following:
 - Canonical Equivalence
 - Case Insensitive
 - Comments
 - Dot All
 - Multi-Line
 - Unicode

Unix Lines

Value: Specify the value you want compared to the *HTTP Request Method* value. If you selected a method from the list for the *HTTP Request Method* value, select *HTTP Request Method > Current*. If you selected *Current* for the *HTTP Request Method* value, select a request method from the list.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.9 LDAP Attribute Condition

The *LDAP Attribute* condition allows you to restrict access based on a value in an LDAP attribute defined for the *inetOrgPerson* class or any other LDAP attribute that you have added. You can have the user's attribute value retrieved from your LDAP directory and compared to a value of the following type:

- ◆ Roles
- ◆ Date and time and its various elements
- ◆ URL and its various elements
- ◆ IP address
- ◆ Authentication contract
- ◆ Credential profile
- ◆ HTTP request method
- ◆ Liberty User Profile attribute
- ◆ Static value

To set up the matching for this condition, fill in the following fields:

LDAP Attribute: Specify the LDAP attribute you want to use in the comparison. Select from the listed LDAP attributes. To add an attribute that isn't in the list, click *New*, then specify the name of the attribute or add the desired attribute as a Custom Attribute. See [Section 6.4.2, "Creating LDAP Attribute Names,"](#) on page 69.

Comparison: Specify how you want the values compared. All data types are available. Select one that matches the value type of your attribute.

Mode: Select the mode, if available, that matches the data type. For example, if you select to compare the values as strings, you can select between a *Case Sensitive* mode and a *Case Insensitive* mode.

Value: Specify the second value for the comparison. All data types are available. For example, you can select to compare the value of one LDAP attribute with the value of another LDAP attribute. Only you can determine if such a comparison is meaningful.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.10 Liberty User Profile Condition

The *Liberty User Profile* condition allows you to restrict access based on a value in a Liberty User Profile attribute. These attributes can be mapped to LDAP attributes (see [Section 11.9, “Mapping LDAP and Liberty Attributes,” on page 127](#)). When mapped, the actual value comes from your user store. If you are using multiple user stores with different LDAP schemas, mapping similar attributes to the same Liberty User Profile attribute allows you to create one policy with the Liberty User Profile attribute rather than multiple policies for each LDAP attribute.

If no profile attributes are available, you have not enabled their use in the Identity Server configuration. See [Section 11.2, “Enabling Web Services and Profiles,” on page 116](#).

The selected attribute is compared to a value of the following type:

- ◆ Roles
- ◆ Date and time and its various elements
- ◆ URL and its various elements
- ◆ IP address
- ◆ Authentication contract
- ◆ Credential profile
- ◆ HTTP request method
- ◆ Liberty User Profile attribute
- ◆ LDAP attribute
- ◆ Static value

To set up the matching for this condition, fill in the following fields:

Liberty User Profile: Select the Liberty User Profile attribute. These attributes are organized into four main groups: Custom Profile: Customizable Strings, Custom Profile Extensions, Corporate Employee Identity, and Entire Personal Identity. By default, the Common Last Name attribute for Liberty User Profile is mapped to the sn attribute for LDAP. To select this attribute for comparison, click *Entire Personal Identity > Entire Common Name > Common Analyzed Name > Common Last Name*.

Comparison: Select the comparison type that matches the data type of the selected attribute and the value.

Mode: Select the mode, if available, that matches the data type. For example, if you select to compare the values as strings, you can select between a *Case Sensitive* mode and a *Case Insensitive* mode.

Value: Select one of the values that is available from the current request or select *Data Entry Field* to enter a static value. The static value that you can enter is dependent upon the comparison type you selected.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.11 Roles for Current User Condition

If you have configured some Access Manager role policies (see [Chapter 29, “Creating Role Policies,” on page 317](#)), you can use these roles as conditions to control access. Fill in the following fields:

Comparison: Select one of the following types:

- ♦ **Comparison: String:** Specifies that you want the values compared as strings and how you want the string values compared. Select one of the following:
 - ♦ **Equals:** Indicates that the values must match, letter for letter.
 - ♦ **Starts with:** Indicates that the *Roles for Current User* value must begin with the letters specified in the *Value* field.
 - ♦ **Ends with:** Indicates that the *Roles for Current User* value must end with the letters specified in the *Value* field.
 - ♦ **Contains Substring:** Indicates that the *Roles for Current User* value must contain the letters, in the same sequence, as specified in the *Value* field.
- ♦ **Comparison: Regular Expression: Matches:** Specifies that you want the values compared as regular expressions.

Mode: Select the mode appropriate for the comparison type:

- ♦ **Comparison: String:** Specify whether case is important by selecting *Case Sensitive* or *Case Insensitive*.
- ♦ **Comparison: Regular Expression: Matches:** Select one or more of the following:

Canonical Equivalence
Case Insensitive
Comments
Dot All
Multi-Line
Unicode
Unix Lines

Value: If you have created Identity Server roles policies, select *Roles*, then select the role. If you have defined a Liberty User Profile or an LDAP attribute for roles, select this option, then your attribute.

You can use the *Data Entry Field* option to enter the name of the role you know you want to test for.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.12 URL Condition

The *URL* condition allows you to restrict access based on the URL specified in the request. In an Access Gateway Authorization policy if you have users requesting a resource with a URL you don't want them to use, you can use this condition to deny them access to this URL, and use the Actions

section to redirect the request to the URL you want them to use. In a J2EE agent policy, you can only deny or allow; you cannot redirect.

To set up matching for this condition, fill in the following fields:

Value: If you have defined a Liberty User Profile or an LDAP attribute for a URL, you can select this option, then your attribute. To enter a static value to compare to the URL in the current request, select *Data Entry Field* and specify the URL. This needs to be the complete URL, starting with the URL scheme (`http://` or `https://`) and including the domain name. If the URL contains a path, you need to include it.

Use the Edit button to access a text box where you can enter multiple values, each on a separate line. Use the Add button to add values one at a time. All entered URLs are compared to the request URL until a match is found or the list is exhausted.

The wildcard characters, (?) or (*) can be specified as the last element of the URL path to aid in matching basic URL patterns. These wildcard characters are interpreted as follows:

- ◆ ? matches all files at the specified directory level.
- ◆ * matches all files and directories at and beyond the specified directory level.

For example, if the request URL is `http://www.resourcehost.com/path/resource.gif`, then the following entered URLs would match the request URL:

```
http://www.resourcehost.com/path/resource.gif
http://www.resourcehost.com/path/?
http://www.resourcehost.com/path/*
http://www.resourcehost.com/*
```

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.13 URL Scheme Condition

The *URL Scheme* condition allows you to restrict access based on the scheme specified in the URL of the request. For example in an Access Gateway Authorization policy, if the request contains HTTP as the scheme in the URL and you require them to use HTTPS, you can use this condition to deny access and redirect them to another URL. In a J2EE Agent policy, you can only deny or allow; you cannot redirect.

This condition allows you to compare A to B or B to A. You need to decide whether you want to compare a static value to the current value in the HTTP request, or whether you want to compare the current value in the HTTP request to a specified value. The comparison type you use depends upon the value you want to specify. If you want more flexibility in specifying the value, you should select to compare the current value in the HTTP request with a specified value.

To set up matching for this condition, fill in the following fields:

URL Scheme: Specify the scheme that you want compared. You can select *Current* for the current value in the HTTP request or specify a static value of *http* or *https*.

Comparison: URL Scheme: Specify how you want the values compared. Select one of the following:

- ♦ **Equals:** Indicates that the URL scheme must contain the same letters, in the same order, as specified in the value.
- ♦ **Starts with:** Indicates that the URL scheme must begin with the letters specified in the value.
- ♦ **Ends with:** Indicates that the URL scheme must end with the letters specified in the value.
- ♦ **Contains Substring:** Indicates that the URL scheme must contain the letters specified in the value.

Mode: Specify whether case is important by selecting *Case Sensitive* or *Case Insensitive*.

Value: Specify the value you want to compare with the *URL Scheme* value. If you selected a static value for the *URL Scheme* value, select *URL Scheme* and *Current*. If you selected *Current* for the *URL Scheme* value, select one of the following value types:

- ♦ **Data Entry Field:** Allows you to specify the scheme value you want to use in the comparison. The scheme can be specified with or without a trailing colon (:) character. Use the Edit button to access a text box where you can enter multiple values, each on a separate line. Use the Add button to add values one at a time.

All specified URL schemes are compared to the request URL scheme until a match is found or the list is exhausted.

- ♦ **LDAP Attribute:** If you have defined an LDAP attribute containing a URL or URL scheme, you can select this option, then select your attribute.
- ♦ **Liberty User Profile:** If you have defined a Liberty User Profile attribute containing a URL or URL scheme, you can select this option, then select your attribute.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.14 URL Host Condition

The *URL Host* condition allows you to restrict access based on the hostname specified in the URL of the request. For example, you can use this condition to create rules that allow access if the URL contains one hostname, but denies access if the URL contains another hostname. The URL Host condition compares the hostname in the URL of the current request to the URL hostname specified in the *Value* field.

To set up matching for this condition, fill in the following fields:

Value: Specify the value type and value for the comparison. Select one of the following:

- ♦ **Data Entry Field:** To specify a static value to compare to the URL host in the current request, select this value type and specify the DNS name or the IP address of the host.

Use the Edit button to access a text box where you can enter multiple values, each on a separate line. Use the Add button to add values one at a time. All listed hostnames and addresses are compared to the requested URL until a match is found or the list is exhausted.

For example, if the request URL is `http://www.resourcehost.com/path/resource.gif` and the IP address of the resource host is `10.10.10.10`, then the following entered hostname or IP address would match the resource URL:

```
www.resourcehost.com
10.10.10.10
```

- ◆ **LDAP Attribute:** If you have defined an LDAP attribute containing a URL or URL host, you can select this option, then select your attribute.
- ◆ **Liberty User Profile:** If you have defined a Liberty User Profile attribute containing a URL or URL host, you can select this option, then select your attribute.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.15 URL Path Condition

The *URL Path* condition allows you to restrict access based on the path specified in the URL of the request. This condition compares the path of the URL in the current request to the path specified in the *Value* field.

To set up matching for this condition, fill in the following fields:

Comparison: URL Path: Specify how you want the values compared. Select one of the following:

- ◆ **Equals:** Indicates that the URL path must contain the same letters, in the same order as specified in the value.
- ◆ **Starts with:** Indicates that the URL path must begin with the letters specified in the value.
- ◆ **Ends with:** Indicates that the URL path must end with the letters specified in the value.
- ◆ **Contains Substring:** Indicates that the URL path must contain the letters specified in the value.

Mode: Indicate whether case is important by selecting *Case Sensitive* or *Case Insensitive*.

Value: Specify the value type and value for the comparison. Select one of the following:

- ◆ **Data Entry Field:** To specify a static value to compare to the URL path in the current request, select this value type and specify the path. Start the path with a forward slash. The path can end with a filename or a wildcard. An asterisk (*) matches all files and directories at and beyond the specified directory level. A question mark (?) matches all files at the specified directory level. For example:

Path	Match Description
<code>/path1/path2/</code>	Requires an exact match of the URL path. It matches if the URL does not contain anything beyond <i>path2</i> .
<code>/path1/file.ext</code>	Requires an exact match of the URL path, including the extension on the filename.
<code>/path1/path2/?</code>	Matches everything that immediately follows <i>path2</i> . It does not match anything if the path contains another directory such as <code>/path1/path2/path3/file3.ext</code> .

Path	Match Description
/path1/path2/*	Matches everything that follows path2, including a filename or another directory such as /path1/path2/path3/file3.ext.

Use the Edit button to access a text box where you can enter multiple values, one on each line. Use the Add button to add values one at a time. All specified URL paths are compared to the request URL path until a match is found or the list is exhausted.

- ◆ **LDAP Attribute:** If you have defined an LDAP attribute containing a URL or URL path, you can select this option, then select your attribute.
- ◆ **Liberty User Profile:** If you have defined a Liberty User Profile attribute containing a URL or URL path, you can select this option, then select your attribute.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.16 URL File Name Condition

The *URL File Name* condition allows you to restrict access based on the filename specified in the URL. It compares the filename in the URL of the current request to the filename specified in the *Value* field.

To set up matching for this condition, fill in the following fields:

Comparison: URL File: Specify how you want the values compared. Select one of the following:

- ◆ **Equals:** Indicates that the filenames must contain the same letters, in the same order, as specified in the value.
- ◆ **Starts with:** Indicates that the filenames must begin with the letters specified in the value.
- ◆ **Ends with:** Indicates that the filenames must end with the letters specified in the value.
- ◆ **Contains Substring:** Indicates that the filenames must contain the letters specified in the value.

Mode: Indicate whether case is important by selecting *Case Sensitive* or *Case Insensitive*.

Value: Specify the value type and value for the comparison. Select one of the following:

- ◆ **Data Entry Field:** To specify a static value to compare to the filename in the current request, select this value type and specify the filename. You can specify the filename with or without an extension. For example:

```
filename
filename.ext
```

This condition does not support wildcards. Use the Edit button to access a text box where you can enter multiple values, each on a separate line. Use the Add button to add values one at a time. All listed filenames are compared to the requested URL filename until a match is found or the list is exhausted.

- ◆ **LDAP Attribute:** If you have defined an LDAP attribute containing a URL or filename, you can select this option, then select your attribute.

- ♦ **Liberty User Profile:** If you have defined a Liberty User Profile attribute containing a URL or filename, you can select this option, then select your attribute.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.17 URL File Extension Condition

The *URL File Extension* condition allows you to restrict access based on the file extension specified in the URL of the request. It compares the file extension in the URL of the current request to the extension specified in the *Value* field

To set up matching for this condition, fill in the following fields:

Comparison: URL File: Specify how you want the values compared. Select one of the following:

- ♦ **Equals:** Indicates that the file extensions must contain the same letters, in the same order, as specified in the value.
- ♦ **Starts with:** Indicates that the file extensions must begin with the letters specified in the value.
- ♦ **Ends with:** Indicates that the file extensions must end with the letters specified in the value.
- ♦ **Contains Substring:** Indicates that the file extensions must contain the letters specified in the value.

Mode: Indicate whether case is important by selecting *Case Sensitive* or *Case Insensitive*.

Value: Specify the value type and value for the comparison. Select one of the following:

- ♦ **Data Entry Field:** To enter a static value to compare to the file extension in the current request, select this value type and specify the file extension. You can enter the extension or the period and the extension. For example:

```
.ext  
ext
```

This condition does not support wildcards. Use the Edit button to access a text box where you can enter a list of URL file extensions, each on a separate line. Use the Add button to add values one at a time.

All listed URL file extensions are compared to the request URL file extension until a match is found or the list is exhausted.

- ♦ **LDAP Attribute:** If you have defined an LDAP attribute containing a URL or file extension, you can select this option, then select your attribute.
- ♦ **Liberty User Profile:** If you have defined a Liberty User Profile attribute containing a URL or file extension, you can select this option, then select your attribute.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.5.18 X-Forward-For IP Condition

For added security, you can add the IP address of the reverse proxy as a condition to check before granting access. One way to implement this is to create a rule that requires the X-Forwarded-For IP address in the HTTP header to match the configured IP address of the reverse proxy that is using the policy. The X-Forwarded-For IP condition matches the first IP address in the X-Forwarded-For header with the IP address specified in the *Value* field.

To set up matching for this condition, fill in the following fields:

Comparison: Specify how you want the X-Forwarded-For IP address compared by selecting one of the following:

- ♦ **Equals:** Allows you to specify an IP address that the X-Forwarded-For IP address must match. You can specify more than one.
- ♦ **In Range:** Allows you to specify a range of IP addresses that the X-Forwarded-For IP address must fall within. You can specify more than one range.
- ♦ **In Subnet:** Allows you to specify the subnet that the X-Forwarded-For IP address must belong to. You can specify more than one subnet

Value: Specify the value type and value for the comparison. Select one of the following:

- ♦ **Data Entry Field:** To specify a static value, select *Data Entry Field* and specify a value appropriate for your comparison type. For example:

Comparison Type	Value
Equals	10.10.10.10
	10.10.10.11
In Range	10.10.10.10 - 10.10.10.100
	10.10.20.10 - 10.10.20.100
In Subnet	10.10.10.12 / 22
	10.10.20.30 / 22

Use the Edit button to access a text box where you can enter multiple values, each on a separate line. Use the Add button to add values one at a time. All specified values are compared to the IP address in the header until a match is found or the list is exhausted.

- ♦ **Client IP:** If you want the first IP address in the X-Forwarded-For header compared to the IP address of the client making the request, select this option.
- ♦ **LDAP Attribute:** If you have defined an LDAP attribute for an IP address, you can select this option, then select your attribute.
- ♦ **Liberty User Profile:** If you have defined a Liberty User Profile attribute for an IP address, you can select this option, then select your attribute.

Result on Condition Error: Specify what the condition returns when the comparison of the two values returns an error rather than the results of the comparison. Select either *False* or *True*. If you want the rule to fail if the condition does not match, select *False*. If you want the rule to fail if the condition matches, select *True*.

30.6 Using Multiple Conditions

In the bar of the *Conditions* section, the *Condition structure* option controls how conditions within a condition group interact with each other and how condition groups interact with each other. Select one of the following:

- ♦ **AND Conditions, OR groups:** If the conditions are ANDed, the user must meet all the conditions in a condition group to match the profile. If the condition groups are ORed, the user must meet all of the conditions of one group to match the profile. This option allows you to set up two or more profiles into which a user could fit and be considered a match. For example, you could create the following Permit rule:

The first condition group could contain the following conditions:

- a. The user's department must be Engineering.
- b. The request must come on a weekday.

The second condition group could contain the following conditions:

- a. The user's department must be Information Services and Technology (IS&T).
- b. The request must come on a weekend.

With this rule, the engineers who match the first condition group have access to the resource during the week, and the IS&T users who match the second condition group have access to the resource on the weekend.

- ♦ **OR Conditions, AND groups:** If the conditions are ORed, the user must meet at least one condition in the condition group to match the profile. If the conditions groups are ANDed, the user must meet at least one condition in each condition group to match the profile. For example, suppose you created the following allow rule:

The first condition group could contain the following conditions:

- a. The user's department is Engineering.
- b. The user's department is Sales.

The second condition group could contain the following conditions:

- a. The user has been assigned the Party Planning role.
- b. The user has been assigned the Vice President role.

With this rule, the Vice Presidents of both the Engineering and Sales departments can access the resource, and the users from the Engineering and Sales department who have been assigned to the Party Planning role can access the resource.

At the top of each condition group, there is an option that allows you to control whether the user must match the conditions to match the profile or whether the user matches the profile if the user doesn't match any of the conditions. Depending upon your selection for the *Condition structure*, you can select from the following:

- ♦ If/If Not
- ♦ Or/Or Not
- ♦ And/And Not

Conditions also have similar Not options, so that a user can match a condition by not matching the specified value.

The check box by each condition allows you to enable the condition or disable it. Usually you use the disable option when testing a new rule, and if you decide the condition is not needed, you then use the Delete button to delete the condition from the rule. Use the Move buttons by the Delete button to move a condition up or down within its group.

30.7 Importing and Exporting Authorization Policies

You can import and export authorization policies in order to run them in other Access Manager configurations and to analyze the authorization logic. The policy is exported as a text file with XML tags. We do not recommend editing the exported file with a text editor. Any changes you want to make to a policy ought to be done through the Administration Console.

To export an Authorization policy:

- 1 In the Administration Console, click *Access Manager > Policies*.
- 2 Select an Authorization policy, then click *Export*.
- 3 (Optional) Modify the name suggested for the file.
- 4 Click *OK*
- 5 Using the features of your browser, specify where the file is copied.

To import a policy:

- 1 Make sure any referenced role policies have been imported.
See [Section 29.6, “Importing and Exporting Role Policies,”](#) on page 337.
- 2 If the policy uses LDAP or Liberty Profile attributes, make sure the Identity Server has been configured for these same attributes.
- 3 In the Administration Console, click *Policies*.
- 4 Click *Import*, then browse to the location of the file.
- 5 Click *OK*.
- 6 When the policy appears in the list, click *Apply Changes*.

Identity injection allows you to add information to the URL or to the HTML page before it is posted to the Web server. The Web server uses this information to determine whether the user should have access to the resource, so it is the Web server that determines the information that you need to inject to allow access to the resource.

When the policy is configured correctly, the user is unaware that additional information is required to access the Web server.

IMPORTANT: Identity Injection policies allow you to inject the user's password into the HTTP header. If you set up such a policy, you should also configure the Access Gateway to use SSL between itself and the back-end Web server. This is the only way to ensure that the password is encrypted on the wire.

This section describes the elements available for an Identity Injection policy, but your Web servers determine which elements you use.

- ◆ [Section 31.1, “Designing an Identity Injection Policy,” on page 373](#)
- ◆ [Section 31.2, “Configuring an Identity Injection Policy,” on page 374](#)
- ◆ [Section 31.3, “Configuring an Authentication Header Policy,” on page 375](#)
- ◆ [Section 31.4, “Configuring a Custom Header Policy,” on page 378](#)
- ◆ [Section 31.5, “Configuring a Custom Header with Tags,” on page 379](#)
- ◆ [Section 31.6, “Specifying a Query String for Injection,” on page 381](#)
- ◆ [Section 31.7, “Injecting into the Cookie Header,” on page 382](#)
- ◆ [Section 31.8, “Importing and Exporting Identity Injection Policies,” on page 382](#)

31.1 Designing an Identity Injection Policy

Before setting up an Identity Injection policy, you need to know the following about your Web application:

- ◆ Does it require an authentication header? Does this header need just the user name or does it also need the password?
- ◆ Does it use a custom header with custom names (x-names)? If so, you need to know their names and their expected values.
- ◆ Does the custom header require any custom names (x-names) with tags? If so, gather this information.
- ◆ Does the application expect specific values in the query string of the URL? If so, gather this information.

After gathering the information, you need to determine whether you need to create one policy with one rule, one policy with multiple rules, or multiple policies. If you have multiple applications that require the same type of authentication header, you might want to create an authentication header policy and separate policies for the application-specific information. You can then enable both the authentication header policy and the application-specific policy for the resource that is protecting

the application. Everything defined in a policy is injected into the header, even if the values are empty because the Access Manager could not obtain the value for the item. For some applications, this is still useful information and the application uses it to make access decisions.

You should design your policies so that the application receives just what it needs. It should not inject custom names and values it does not use.

Whether you create a policy with one rule or multiple rules is a personal design decision. If you put all the actions in one rule, you have only one description field to describe the function of the policy. If you put each action type in a separate rule, you have multiple description fields to describe the function of the policy. Select the method that is easiest for you. Your policy can inject only one authentication header and one cookie header.

31.2 Configuring an Identity Injection Policy

- 1 In the Administration Console, click *Access Manager > Policies > New*.
- 2 Specify a name for the policy, then select *Access Gateway: Identity Injection* for the type of policy.

The screenshot shows the configuration interface for a policy. The 'Type' field is set to 'Access Gateway: Identity Injection'. The 'Description' field is empty. The 'Priority' is set to '1'. The 'Actions' section is expanded, and a 'New' dropdown menu is open, showing the following options: 'Inject into Authentication Header', 'Inject into Custom Header', 'Inject into Custom Header with Tags', 'Inject into Cookie Header', and 'Inject into Query String'. The background shows the 'Policies Panel'.

- 3 Fill in the following fields:
 - Description:** (Optional) Describe the purpose of this policy. Because Identity Injection policies are customized to match the content of a specific Web server, you might want to include the name of the Web server as part of the description.
 - Priority:** Specify the order in which a rule is applied in the policy, when the policy has multiple rules. The highest priority is 1 and 10 is the lowest.
- 4 In the *Actions* section, click *New* and select one of the following.
 - Repeat this process to add multiple actions to the same rule. If a particular action is allowed only once per rule, then the action does not appear in the *New* menu if that action has already been defined in the rule.
 - ♦ **Inject into Authentication Header:** Inserts the user name and password into the header. For information about how to configure this type of policy, see [Section 31.3, “Configuring an Authentication Header Policy,”](#) on page 375.
 - ♦ **Inject into Custom Header:** Inserts custom names with values into the custom header. For information about how to configure this type of policy, see [Section 31.4, “Configuring a Custom Header Policy,”](#) on page 378.

- ♦ **Inject into Custom Header with Tags:** Inserts custom tags with name/value content into the custom header. For information about how to configure this type of policy, see [Section 31.5, “Configuring a Custom Header with Tags,” on page 379.](#)
 - ♦ **Inject into Query String:** Inserts a query string into the URL for the page. For information about how to configure this type of policy, see [Section 31.6, “Specifying a Query String for Injection,” on page 381.](#)
 - ♦ **Inject into Cookie Header:** Inserts the session cookie into the cookie header. For information about how to configure this type of policy, see [Section 31.7, “Injecting into the Cookie Header,” on page 382.](#)
- 5 To save the policy, click *OK* twice, then click *Apply Changes*.
 - 6 For information on how to assign the policy to a protected resource, see [Section 12.4.2, “Assigning an Authorization Policy to a Protected Resource,” on page 152.](#)

31.3 Configuring an Authentication Header Policy

To inject values into the authentication header, you need to know what the Web server requires. For basic authentication, you need to inject the user name and password. For a sample policy for a Web server that requires the LDAP username and password to be injected into the header, see [“Setting Up an Identity Injection Policy”](#) in the *Novell Access Manager 3.0 Setup Guide*.

To create and configure an authentication header policy:

- 1 In the Administration Console, click *Access Manager > Policies > New*.
- 2 Specify a name for the policy, select *Access Gateway: Identity Injection* for the type, then click *OK*.
- 3 (Optional) Specify a description for the injection policy. This is useful if you plan to create multiple custom header policies to be used for multiple resources.
- 4 In the *Actions* section, click *New*, then select *Inject into Authentication Header*.
- 5 Fill in the *User Name* field.

Select *Credential Profile* to insert the name the user entered when the user authenticated. This is the most common value type to use for user name. However, if your user store is an Active Directory server, you cannot use an LDAP credential for the user name. Active Directory does not use the *cn* or the *dn* attributes for login, but instead uses the *SAMAccountName* attribute. You need to use LDAP attribute as the value type and specify the *SAMAccountName* attribute.

Depending upon what the user must supply for authentication, select one of the following:

- ♦ **LDAP Credentials:** If you prompt the user for a user name, select this option, then select either *LDAP User Name* (the *cn* attribute of the user) or *LDAP User DN* (the fully distinguished name of the user). Your Web server requirements determine which one you use.
- ♦ **X509 Credentials:** If you prompt the user for a certificate, select this option, then select one of the following. Your Web server requirements determine which one you use.
 - ♦ **X509 Public Certificate Subject:** Injects just the subject field from the certificate, which can match the DN of the user, depending upon who issued the certificate.
 - ♦ **X509 Public Certificate Issuer:** Injects just the issuer field from the certificate, which is the name of the CA that issued the certificate.

You can also select one of the following values to insert into the header as the user name:

- ♦ **Authentication Contract:** Injects the URI of the local contract for the protected resource:

URI	Contract
basic/name/password/uri	Name/Password - Basic
name/password/uri	Name/Password - Form
secure/basic/name/password/uri	Secure Name/Password - Basic
secure/name/password/uri	Secure Name/Password - Form

To view other possible values configured for your system, click *Identity Servers > Setup > [Name of Configuration] > Local > Contracts*.

- ♦ **Client IP:** Injects the IP address associated with the user.
- ♦ **LDAP Attribute:** Injects the values of the selected attribute. For Active Directory servers, specify the SAMAccountName attribute for the user name.
- ♦ **Liberty User Profile:** Injects the value of the selected attribute. If no profile attributes are available, you have not enabled their use in the Identity Server configuration. See [Section 11.2, “Enabling Web Services and Profiles,” on page 116](#).
- ♦ **Proxy Session Cookie:** Injects the session cookie associated with the user.
- ♦ **Roles for Current User:** Injects the roles that have been assigned to the user.
- ♦ **Shared Secret:** Injects the username that has been stored in the selected shared secret store. For more information, see [Section 32.4, “Creating and Managing Shared Secrets,” on page 396](#).
- ♦ **String Constant:** Injects a static value that you specify in the text box. This name is used by all users who access the resources assigned to this policy.
- ♦ **Java Data Injection Module:** Specifies the name of a custom Java plug-in, which injects custom values into the header. Usually, you can use either the *LDAP Attribute* or *Liberty User Profile* option to supply custom values, because both are extensible. For more information about creating a custom plug-in, see [Novell Access Manager Developer Tools and Examples \(http://developer.novell.com/wiki/index.php/Nacm\)](http://developer.novell.com/wiki/index.php/Nacm).

The value type you use depends upon how you have set up the application.

6 Fill in the *Password* field.

Select *Credential Profile* to insert the password the user entered when the user authenticated. This is the most common value type to use for the password. Depending upon what the user must supply for authentication, select one of the following:

- ♦ **LDAP Credentials:** If you prompt the user for a password, select this option, then *LDAP Password*. If the user’s password is the same as the name of the user, you can select either *LDAP User Name* (the cn attribute of the user) or *LDAP User DN* (the fully distinguished name of the user).
- ♦ **X509 Credentials:** If you use a certificate for the password, select this option, then select one of the following:
 - ♦ **X509 Public Certificate Subject:** Injects just the subject from the certificate, which can match the DN of the user, depending upon who issued the certificate.

- ◆ **X509 Public Certificate Issuer:** Injects just the issuer from the certificate, which is the name of the CA that issued the certificate.

Your Web server requirements determine which one you use. If your application requires it, you can select value types other than *LDAP Credentials* or *X509 Credentials*.

You can also select one of the following values to insert into the header as the password:

- ◆ **Authentication Contract:** Injects the URI of the local contract for the protected resource:

URI	Contract
basic/name/password/uri	Name/Password - Basic
name/password/uri	Name/Password - Form
secure/basic/name/password/uri	Secure Name/Password - Basic
secure/name/password/uri	Secure Name/Password - Form

To view other possible values configured for your system, click *Identity Servers > Setup > [Name of Configuration] > Local > Contracts*.

- ◆ **Client IP:** Injects the IP address associated with the user.
- ◆ **LDAP Attribute:** Injects the values of the selected attribute.
- ◆ **Liberty User Profile:** Injects the value of the selected attribute.
- ◆ **Proxy Session Cookie:** Injects the session cookie associated with the user.
- ◆ **Roles for Current User:** Injects the roles that have been assigned to the user.
- ◆ **Shared Secret:** Injects the password that has been stored in the selected shared secret store. For more information, see [Section 32.4, “Creating and Managing Shared Secrets,” on page 396](#).
- ◆ **String Constant:** Injects a static value that you specify in the text box. This name is used by all users who access the resources assigned to this policy.
- ◆ **Java Data Injection Module:** Specifies the name of a custom Java plug-in, which injects custom values into the header. Usually, you can use either the *LDAP Attribute* or *Liberty User Profile* option to supply custom values, because both are extensible. For more information about creating a custom plug-in, see [Novell Access Manager Developer Tools and Examples \(http://developer.novell.com/wiki/index.php/Nacm\)](http://developer.novell.com/wiki/index.php/Nacm).

The value type you use depends upon how you have set up the application.

- 7 Select a value for the *Multi-Value Separator*. If the value type you have selected for User Name or Password can have multiple values, select the character to use when separating the values. For example, the *cn* attribute, which you can select for User Name, is multi-valued.

Multi-Value Separator: Select a value separator, if the value type you have select is multi-valued. For example, *Roles for Current User* can contain multiple values.

- 8 Click *OK*.
- 9 (Optional) To add a second rule, click *New* in the Rule List.

You can inject only one authentication header into an Identity Injection rule. However, your policy can have multiple rules. If you inject two authentication headers, each in a separate rule, the authentication header in the rule with the highest priority is applied, and the authentication header action in the second rule is ignored.

10 To save the policy, click *OK*, then click *Apply Changes*.

31.4 Configuring a Custom Header Policy

To inject values into a custom header, you need to know the name of the tag and its expected value type. The names are specific to the application. The names might be case sensitive. They might require an X- prefix. Because the requirements vary, you need to enter them in the same format as specified by the application. For example, an application might require the following to be in the custom header:

Name/Value Pair	Description
X-First_Name=givenName	A first name tag with an LDAP attribute value
X-Last_Name=sn	A last name tag with an LDAP attribute value
Role=sales_role	A role tag with the role name as the value.

If you create a custom header policy with these name/value pairs, the policy injects these names with their values into a custom header, before sending the request to the Web server.

To create such a policy:

- 1 In the Administration Console, click *Access Manager > Policies > New*.
- 2 Specify a name for the policy, select *Access Gateway: Identity Injection* for the type, then click *OK*.
- 3 (Optional) Specify a description for the injection policy. This is useful if you plan to create multiple custom header policies to be used for multiple resources.
- 4 In the *Actions* section, click *New*, then select *Inject into Custom Header*.
- 5 Fill in the following fields:

Custom Header Name: Specify the name to be inserted into the custom header. These are the names required by your application. If your application requires the X- prefix, make sure you include the prefix in this field.

Value: Select the value required by the name. Select one of the following:

- ♦ **Authentication Contract:** Injects the URI of the local contract for the protected resource:

URI	Contract
basic/name/password/uri	Name/Password - Basic
name/password/uri	Name/Password - Form
secure/basic/name/password/uri	Secure Name/Password - Basic
secure/name/password/uri	Secure Name/Password - Form

To view other possible values configured for your system, click *Identity Servers > Setup > [Name of Configuration] > Local > Contracts*.

- ♦ **Client IP:** Injects the IP address associated with the user.

- ♦ **Credential Profile:** Injects the credentials that the user specified at login. You can select between *LDAP Credentials* and *X509 Credentials*. For more information, see [Section 31.3, “Configuring an Authentication Header Policy,” on page 375](#).
- ♦ **LDAP Attribute:** Injects the value of the selected attribute.
- ♦ **Liberty User Profile:** Injects the value of the selected attribute. If no profile attributes are available, you have not enabled their use in the Identity Server configuration. See [Section 11.2, “Enabling Web Services and Profiles,” on page 116](#).
- ♦ **Proxy Session Cookie:** Injects the session cookie associated with the user.
- ♦ **Roles for Current User:** Injects the roles that have been assigned to the user.
- ♦ **Shared Secret:** Injects a value that has been stored in the selected shared secret store. Select the shared secret store and the name of the value you want injected. For more information, see [Section 32.4, “Creating and Managing Shared Secrets,” on page 396](#).
- ♦ **String Constant:** Injects a static value that you specify in the text box. This value is used by all users who access the resources assigned to this policy.
- ♦ **Java Data Injection Module:** Specifies the name of a custom Java plug-in, which injects custom values into the header. Usually, you can use either the *LDAP Attribute* or *Liberty User Profile* option to supply custom values, because both are extensible. For more information, see [Novell Access Manager Developer Tools and Examples \(http://developer.novell.com/wiki/index.php/Nacm\)](http://developer.novell.com/wiki/index.php/Nacm).

Multi-Value Separator: Select a value separator, if the value type you have select is multi-valued. For example, *Roles for Current User* can contain multiple values.

- 6 In the *Actions* section, continue selecting *New > Inject into Custom Header* and configuring the fields until you have added all the name/value pairs that the application requires.
- 7 To save the policy, click *OK* twice, then click *Apply Changes*.

31.5 Configuring a Custom Header with Tags

Some Web applications require more than a name and a value to be injected into the custom header. Sometimes they require a custom name, a tag, and a value. Sometimes the application requires a custom name with multiple tags and values. The *Inject into Custom Header with Tags* option provides you with the flexibility to add such values to the custom header. For example, your application could be expecting the following custom header with tag:

```
X-Custom_Role Role=Manager
```

You can inject this information by setting the *Custom Header Name* to X-Custom, the *Tag Name* to Role, and the *Tag Value* to Manager. The value can be set as a static variable or you can retrieve it from various sources such as a Liberty User Profile attribute or the roles assigned to the current user.

- 1 In the Administration Console, click *Access Manager > Policies > New*.
- 2 Specify a name for the policy, select *Access Gateway: Identity Injection* for the type, then click *OK*.
- 3 (Optional) Specify a description for the injection policy. This is useful if you plan to create multiple custom header policies to be used for multiple resources.
- 4 In the *Actions* section, click *New*, then select *Inject into Custom Header with Tags*.
- 5 Fill in the following fields:

Custom Header Name: Specify the name that the application expects. If your application requires the X- prefix, make sure you include the prefix in this field.

Tag Name: Specify the tag name that the application expects.

Tag Value: Specify the value. Select from the following data types:

- ♦ **Authentication Contract:** Injects the URI of the local contract for the protected resource:

URI	Contract
basic/name/password/uri	Name/Password - Basic
name/password/uri	Name/Password - Form
secure/basic/name/password/uri	Secure Name/Password - Basic
secure/name/password/uri	Secure Name/Password - Form

To view other possible values configured for your system, click *Identity Servers > Setup > [Name of Configuration] > Local > Contracts*.

- ♦ **Client IP:** Injects the IP address associated with the user.
- ♦ **Credential Profile:** Injects the credentials that the user specified at login. You can select between *LDAP Credentials* and *X509 Credentials*. For more information, see [Section 31.3, “Configuring an Authentication Header Policy,” on page 375](#).
- ♦ **LDAP Attribute:** Injects the value of the selected attribute.
- ♦ **Liberty User Profile:** Injects the value of the selected attribute. If no profile attributes are available, you have not enabled their use in the Identity Server configuration. See [Section 11.2, “Enabling Web Services and Profiles,” on page 116](#).
- ♦ **Proxy Session Cookie:** Injects the session cookie associated with the user.
- ♦ **Roles for Current User:** Injects the roles that have been assigned to the user.
- ♦ **Shared Secret:** Injects a value that has been stored in the selected shared secret store. The name specified as the Tag Name must match the name of a name/value pair stored in the shared secret. For more information, see [Section 32.4, “Creating and Managing Shared Secrets,” on page 396](#).
- ♦ **String Constant:** Injects a static value that you specify in the text box. This value is used by all users who access the resources assigned to this policy.
- ♦ **Java Data Injection Module:** Specifies the name of a custom Java plug-in, which injects custom values into the header. Usually, you can use either the *LDAP Attribute* or *Liberty User Profile* option to supply custom values, because both are extensible. For more information about creating a custom plug-in, see [Novell Access Manager Developer Tools and Examples \(http://developer.novell.com/wiki/index.php/Nacm\)](http://developer.novell.com/wiki/index.php/Nacm).

Multi-Value Separator: Select a value separator, if the value type you have select is multi-valued. For example, *Roles for Current User* can contain multiple values.

- 6 To add multiple tag and value pairs to the custom name, click *New* in the *Tags* section.
- 7 In the *Actions* section, continue selecting *New > Inject into Custom Header with Tags* and configuring the fields until you have added all the custom name with tag and value pairs that the application requires.
- 8 To save the policy, click *OK* twice, then click *Apply Changes*.

31.6 Specifying a Query String for Injection

Some applications require custom information in a query string of the URL. The *Inject into Query String* option allows you to inject this information without prompting the user for it. To inject the information, you must specify a tag name and a tag value. The tag name is what your application requires. For example, suppose your application expects the following query string for user jsmith:

```
?name=jsmith
```

You can inject this information into the URL by specifying name for the *Tag Name* and *Credential Profile* for the *Tag Value*. The *Credential Profile* value type inserts the name the current user specified when authenticating to the Access Gateway.

- 1 In the Administration Console, click *Access Manager > Policies > New*.
- 2 Specify a name for the policy, select *Access Gateway: Identity Injection* for the type, then click *OK*.
- 3 (Optional) Specify a description for the injection policy. This is useful if you plan to create multiple custom header policies to be used for multiple resources.
- 4 In the *Actions* section, click *New*, then select *Inject into Query String*.
- 5 Fill in the following fields:

Tag Name: Specify the tag name that the application expects.

Tag Value: Specify the value. Select from the following data types:

- ♦ **Authentication Contract:** Injects the URI of the local contract for the protected resource:

URI	Contract
basic/name/password/uri	Name/Password - Basic
name/password/uri	Name/Password - Form
secure/basic/name/password/uri	Secure Name/Password - Basic
secure/name/password/uri	Secure Name/Password - Form

To view other possible values configured for your system, click *Identity Servers > Setup > [Name of Configuration] > Local > Contracts*.

- ♦ **Client IP:** Injects the IP address associated with the user.
- ♦ **Credential Profile:** Injects the credentials that the user specified at login. You can select between *LDAP Credentials* and *X509 Credentials*. For more information, see [Section 31.3, “Configuring an Authentication Header Policy,” on page 375](#).
- ♦ **LDAP Attribute:** Injects the value of the selected attribute.
- ♦ **Liberty User Profile:** Injects the value of the selected attribute. If no profile attributes are available, you have not enabled their use in the Identity Server configuration. See [Section 11.2, “Enabling Web Services and Profiles,” on page 116](#).
- ♦ **Proxy Session Cookie:** Injects the session cookie associated with the user.
- ♦ **Roles for Current User:** Injects the roles that have been assigned to the user.
- ♦ **Shared Secret:** Injects a value that has been stored in the selected shared secret store. The name specified as the Tag Name must match the name of a name/value pair stored in the shared secret.

- ♦ **String Constant:** Injects a static value that you specify in the text box. This value is used by all users who access the resources assigned to this policy.
- ♦ **Java Data Injection Module:** Specifies the name of a custom Java plug-in, which injects custom values into the header. Usually, you can use either the *LDAP Attribute* or *Liberty User Profile* option to supply custom values, because both are extensible. For more information about creating a custom plug-in, see [Novell Access Manager Developer Tools and Examples \(http://developer.novell.com/wiki/index.php/Nacm\)](http://developer.novell.com/wiki/index.php/Nacm).

Multi-Value Separator: Select a value separator, if the value type you have select is multi-valued. For example, *Roles for Current User* can contain multiple values.

6 To add multiple tag and value pairs, click *New* in the *Tags* section.

You can inject only one query string into a rule, but you can inject multiple tag name and tag value pairs in the single query string.

7 To save the policy, click *OK* twice, then click *Apply Changes*.

31.7 Injecting into the Cookie Header

Some applications require access to the Access Gateway session cookie and expect to find it in the cookie header. You can create an Identity Injection policy that adds this cookie to the cookie header.

1 In the Administration Console, click *Access Manager > Policies > New*.

2 Specify a name for the policy, select *Access Gateway: Identity Injection* for the type, then click *OK*.

3 (Optional) Specify a description for the injection policy. This is useful if you plan to create multiple custom header policies to be used for multiple resources.

4 In the *Actions* section, click *New*, then select *Inject into Cookie Header*.

This action allows only one value, so the value is configured automatically.

5 To save the policy, click *OK* twice, then click *Apply Changes*.

31.8 Importing and Exporting Identity Injection Policies

You can import and export Identity Injection policies in order to run them in other Access Manager configurations. The policy is exported as a text file with XML tags. We do not recommend editing the exported file with a text editor. Any changes you want to make to a policy should be done through the Administration Console.

To export an Identity Injection policy:

1 In the Administration Console, click *Access Manager > Policies*.

2 Select an Identity Injection policy, and click *Export*.

3 (Optional) Modify the name suggested for the file (optional).

4 Click *OK*.

5 Using the features of your browser, specify where the file is to be copied.

To import a policy:

- 1** Make sure any referenced shared secret stores have been created. See [Section 32.4, “Creating and Managing Shared Secrets,”](#) on page 396.
- 2** If the policy uses LDAP or Liberty Profile attributes, make sure the Identity Server has been configured for these same attributes.
- 3** Make sure any referenced role policies have been imported.
See [Section 29.6, “Importing and Exporting Role Policies,”](#) on page 337.
- 4** In the Administration Console, click *Access Manager > Policies*.
- 5** Click *Import*, then browse to the location of the file.
- 6** Click *OK*.
- 7** When the policy appears in the list, click *Apply Changes*.

A Form Fill policy allows you to prepopulate fields in a form on first login and then save the information in the completed form to a secret store for subsequent logins. The user is prompted to reenter the information only when something changes such as an expired password. The HTML page determines the requirements for the Form Fill policy. This section describes the following:

- ♦ Section 32.1, “Understanding an HTML Form,” on page 385
- ♦ Section 32.2, “Creating a Form Fill Policy for the Sample Form,” on page 388
- ♦ Section 32.3, “Implementing Form Fill Policies,” on page 391
- ♦ Section 32.4, “Creating and Managing Shared Secrets,” on page 396
- ♦ Section 32.5, “Importing and Exporting Form Fill Policies,” on page 398

32.1 Understanding an HTML Form

The following figure is an example of a Web page containing an HTML form.

Figure 32-1 Sample HTML Form

Novell Services Login

Username:

Password:

City of Employment:

Web 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.

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 \(http://www.w3.org/TR/html401/interact/forms.html\)](http://www.w3.org/TR/html401/interact/forms.html).

The following HTML data corresponds to the sample form (see [Figure 32-1](#)). 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>Form Fill 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><b><input type="text" name="username" size="30"></b></td>
            </tr>
            <tr align="left">
              <td>Password:</td>
              <td><b><input type="password" name="password" size="30"></b>
            </td>
            </tr>
            <tr align="left">
              <td>City of<br>Employment:</td>
              <td><b><input type="text" name="city" size="30"></b>
            </td>
            </tr>
          </table>
        </td>
      </tr>
    </table>
  </form>
</body>
</html>
```

```

<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.

32.2 Creating a Form Fill Policy for the Sample Form

The sample form has ten input fields and five select options that need to be configured in the Form Fill policy. The following steps explain how to create a shared secret to store the values and use that shared secret to create a Form Fill policy for this sample form.

- 1 To create the policy, click *Policies > New*.
- 2 Specify a display name for the policy and select *Access Gateway: Form Fill* for its type.

The screenshot shows a configuration panel for a policy. It includes the following elements:

- Type:** A dropdown menu with the selected option "Access Gateway: Form Fill".
- Description:** An empty text input field.
- Priority:** A dropdown menu with the selected option "1".
- Actions:** A section with a "New" dropdown menu and the text "No Actions in Policy".

Below the configuration fields, there is a note: "Changes made on this panel must be applied from the [Policies](#) Panel." At the bottom of the panel are two buttons: "OK" and "Cancel".

- 3 (Optional) Specify a description for the Form Fill policy. This is useful if you plan to create multiple Form Fill policies.

You might want to specify the name of the HTML page that contains the form this policy is designed to fill.

- In the *Actions* section, click *New*, then select *Form Fill*.

Changes made on this panel must be applied from the [Policies](#) Panel.

OK Cancel

- In the *Form Selection* section, select *Form Name* and specify *mylogin* in the text box. The form name comes from the HTML page. See the following line in the source for the page:


```
<form name="mylogin" action="validatepassword.php" method="post" id="mylogin">
```
- In the *Fill Options* section, specify all the input fields and select options. For each new field, click *New*. Specify the fields in the order in which they appear on the form. The following table displays the Fill Options selected for each input field.

Form Name	Fill Options
username	<p>Input Field Name: username</p> <p>Input Field Type: Text</p> <p>Input Field Value: Credential Profile: LDAP Credentials: LDAP User Name</p> <p>If your user store is an Active Directory server, you cannot use an LDAP credential for the user name. Active Directory does not use the cn or the dn attributes for login, but instead uses the SAMAccountName attribute. You need to use the LDAP Attribute value type and specify the SAMAccountName attribute.</p>

Form Name	Fill Options
password	Input Field Name: password Input Field Type: Password Input Field Value: Credential Profile: LDAP Credentials: LDAP Password
webserv	Input Field Name: webserv Input Field Type: Select Input Field Value: Shared Secret: sampleLogin: webServer
role	Input Field Name: role Input Field Type: Radio Button Input Field Value: Shared Secret: sampleLogin: role
mail	Input Field Name: mail Input Field Type: Checkbox Input Field Value: Shared Secret: sampleLogin: mail
payroll	Input Field Name: payroll Input Field Type: Checkbox Input Field Value: Shared Secret: sampleLogin: payroll
selfservice	Input Field Name: selfservice Input Field Type: Checkbox Input Field Value: Shared Secret: sampleLogin: selfService

7 In the *Submit Options* section, fill in the following fields:

Auto Submit: Select this option to submit the form as soon as all the values are filled in. If this option is not selected, even though all the values are filled in for the user, the user must click the Submit button.

Debug Mode: Select the *Debug Mode* option, which allows you to verify that the information is correct before submitting the form. If values must be filled in, you first see the form to add the values. When the form is submitted, you are presented with a JavaScript that contains all of the name/value pairs. To submit the form, you need to click the Submit button.

Insert Text in Post Header: Select this option so you can add a static value. In the *Post Header Text* box, specify the city value. Enter:

```
city = Provo
```

- 8 To create a login failure policy, click *New* in the *Actions* section, then select *Form Login Failure*.

The screenshot shows a configuration window with two main sections. The top section is titled 'Form Selection' and contains a dropdown menu for 'Form Name' with a colon and an empty text box to its right. Below this are two links: 'CGI Matching Criteria...' and 'Form Matching Criteria...'. The bottom section is titled 'Login Failure Processing' and contains a 'Redirect to URL:' label with an empty text box. Below that is a checked checkbox labeled 'Clear Shared Secret Data Values from Policy:' followed by a dropdown menu showing 'simple_ff'.

- 9 In the *Form Selection* section, select *Form Name* and specify *mylogin* in the text box. The form name comes from the HTML page.
- 10 In the *Login Failure Processing* section, fill in the following field:
Clear Shared Secret Data Values from Policy: Select this option to clear the data stored in the Shared Secret object when log in fails. Select the name you have given to this policy.
- 11 Click *OK*.
- 12 On the Policies page, click *Apply Changes*.

32.3 Implementing Form Fill Policies

Section 32.2, “Creating a Form Fill Policy for the Sample Form,” on page 388 section describes how to create a simple Form Fill policy for a few input fields. This section describes all available options and explains how to use them to create a Form Fill policy and a login failure policy.

- ♦ Section 32.3.1, “Creating a Form Fill Policy,” on page 391
- ♦ Section 32.3.2, “Creating a Login Failure Policy,” on page 395
- ♦ Section 32.3.3, “Troubleshooting a Form Fill Policy,” on page 396

32.3.1 Creating a Form Fill Policy

- 1 Examine the source code for the HTML form and determine what data the form requires and where that data is stored (LDAP attributes, Liberty User Profile attributes, shared secrets, credential profiles, etc.)

Ideally, the form should be its own HTML page, and page should be as small as possible. Form Fill has to parse the entire file and assemble the body in contiguous memory before the first byte of the form is displayed to the user. On a large file, this can take enough time that your users will think the system has a problem.

If it isn’t possible to have the form on its own HTML page, make sure the form is easily identifiable on the page. For example, give the form a name or use CGI data (the text that the follows the question mark in the URL) to identity the page and form.

- 2 In the Administration Console, click *Access Manager > Policies > New*.
- 3 Specify a name for the policy, select *Access Gateway: Form Fill* as its *Type*, then click *OK*.

- In the *Actions* section, click *New* and select *Form Fill*.

Changes made on this panel must be applied from the [Policies](#) Panel.

If you are converting an iChain[®] Form Fill policy written in XML to an Access Gateway policy, see “[URLs Requiring Form Fill](#)” in the *Novell Access Manager 3.0 Installation Guide*.

- In the *Form Selection* section, specify how the Access Gateway can identify the form on the page. Select one or more of the following methods. The goal should be to be as specific as possible.

Form Name/Form Number: Allows the Access Gateway to identify the form by name or number.

If your form contains a line similar to the following, select to use *Form Name* and specify the name in the text box. In this example, the form name is mylogin.

```
<form name="mylogin" action="validatepassword.php" method="post"
      id="mylogin">
```

The Access Gateway numbers forms sequentially from the top of the HTML page. If your page has multiple forms, you can use Form Number option and specify the form’s sequential location in the text box.

CGI Matching Criteria: Allows the Access Gateway to evaluate the query string in the URL (the portion after the question mark) to differentiate pages that have the same URL. Consider the following URL:

```
http://webaccess.novell.com/servlet/webacc?Action=User.login
```

For this URL, enter the following string in the text box for *CGI Matching Criteria*:

```
Action=User.login
```

If possible, copy the text from the form and paste it into the *CGI Matching Criteria* text box.

Form Matching Criteria: Causes the Access Gateway to search the HTML page for the specified text. If the specified text is found on the page, the page is a match for the policy. If it isn’t found, the page is not a match for the policy and the policy is not applied. For example, suppose your HTML page has the following string:

<title>Form Fill Test Page</title>

If you enter this string in the *Form Matching Criteria* box, the Access Gateway searches the form for this string. If it finds the string, it knows it has a match.

White space is significant. If the text in the text box is left-justified, the text can be found any where on the HTML page. If the text contains leading white space, such as ten spaces, the text must be found with ten leading spaces. If possible, copy the text as it appears on the form and paste it into *Form Matching Criteria* text box.

The more specific your information is, the faster Access Gateway can match the form. Parsing form matching criteria is a very intensive process. If possible, use *Form Name*, *Form Number*, or *CGI Matching Criteria* to identify the form.

- 6 In the *Fill Options* section, create an entry for all the input fields and select options in the form. For each input field or select option, you need to specify the following information:

Input Field Name: Specifies the name of the field or option. This is the name attribute of the element on the form.

Input Field Type: Specifies the type attribute for the input field or select option in the form. Select one of the following data types for the field:

- ♦ **Text:** Indicates that the field is a text field on the form.
- ♦ **Password:** Indicates that the field is a password field on the form.
- ♦ **Checkbox:** Indicates that the field is a check box on the form.
- ♦ **Radio Button:** Indicates that the field is a radio button on the form.
- ♦ **Select:** Indicates that the field is a select option on the form.
- ♦ **Hidden:** Indicates that the field is an input field, but that this field is hidden from the user.
- ♦ **Not Specified:** Indicates that the field is an input field, but the data type is not specified in the Form Fill policy.

Input Field Value: Specify the value for the field. You must specify the data type, then enter the value. Select one of the following data types:

- ♦ **Credential Profile:** Specifies that the value should be retrieved from the credentials the user specified during authentication.
 - ♦ **LDAP Credentials:** If you prompt the user for a username and password, select this option, then either *LDAP User Name* (the cn of the user) or *LDAP User DN* (the fully distinguished name of the user). Your Web server requirements determine which one you use.

If your user store is an Active Directory server, you cannot use an LDAP credential for the user name. Active Directory does not use the cn or dn attributes for login, but instead uses the SAMAccountName attribute. You need to use the LDAP Attribute condition and specify the SAMAccountName attribute.

- ♦ **X509 Credentials:** If you prompt the user for a certificate, select this option, then one of the following. Your Web server requirements determine which one you use.

X509 Public Certificate Subject: Specifies that the subject field from the certificate should be the value, which can match the DN of the user, depending upon who issued the certificate.

X509 Public Certificate Issuer: Specifies that the issuer field from the certificate should be the value, which is the name of the CA that issued the certificate.

- ♦ **LDAP Attribute:** Indicates that the value should be retrieved from the specified LDAP attribute. If the attribute you require does not appear in the list, click *New* to add the attribute.
- ♦ **Liberty User Profile:** Indicates that the input field contains a Liberty User Profile attribute. In the value field, select the attribute. The attribute you select must be mapped to an LDAP attribute, and the Access Gateway retrieves its value from the LDAP directory.
- ♦ **Shared Secret:** Indicates that the input field contains a user-entered value that is to be stored in the specified shared secret store.

You can create your own. Click *New*, specify a display name for the store, and the Access Manager creates the store. Select the store, click *New*, specify a name for the attribute, then click *OK*. The store can contain one name/value pair or a collection of name/value pairs. For more information, see [Section 32.4, “Creating and Managing Shared Secrets,” on page 396](#).

Data Conversion: Specify whether the case of the value entered by the user should be converted. Select one of the following options:

- ♦ **None:** Indicates that no conversion should be performed on the value.
- ♦ **To Upper Case:** Indicates that the value should be converted to uppercase.
- ♦ **To Lower Case:** Indicates that the value should be converted to lowercase.

- 7 In the *Submit Options* section, specify how you want the information in the form submitted to the Web server. (The HTML form page determines whether the post or the get method is used for the submission) Select one or more of the following options:

Auto Submit: Indicates that you want the form submitted to the Web server without having the user confirm the submission by clicking a Submit button. If this option is not selected, Form Fill can fill in the data, but the user must click the Submit button before the data is sent to the Web server. If you select *Auto Submit*, you can select one or more of the following options:

- ♦ **Debug Mode:** Allows you to verify that the information in the filled in form is valid before it is posted to the Web server. You can right-click and view the source that is being submitted to the Web server. If it is correct, click Submit to send it to the Web server.

This is a troubleshooting option. We recommend that you use it when creating a new Form Fill policy, and that you remove it when you have determined that the policy is behaving as expected.

- ♦ **Mask Data:** Replaces text input field values (username, password, etc.) with nov-ss-ff-masked instead of the value specified by the value parameter when the form is sent to the browser. The Access Gateway replaces these masked values with the real values when the Access Gateway submits the form to the Web server. The user’s browser never sees the actual values for these fields.

Insert Text in Post Header: If this option is selected, you can use the *Post Header Text* option to specify text to add to the post header. Use this option to insert static values into the form.

Enable JavaScript Handling: Retains JavaScript from the original page. Use the following fields to specify how you want the JavaScript handled.

- ♦ **Functions to Keep:** Specifies the functions you want executed from the JavaScript on the original page. In the text box, use the following format:

```
function setCookie
```

where *function* is a key word, followed by a space, and then the name of the function. Each function should be entered on a separate line, but you need only one function per

script block. Everything must match exactly (name, capitalization, white space.) If possible, copy the function name from the HTML page.

- ♦ **Statements to Execute on Post:** Specifies the functions you want executed just before the form is posted. Copy the JavaScript from the HTML page into this text box or add a Java function that you want called that is not one the HTML page. This allows you to modify the behavior of the form when you can't modify the form.

If the text box is empty, the JavaScript function specified in the submit field of the HTML page executes before the form is posted.

- 8 In the *Error Handling* section, specify how you want errors handled.

Redirect to URL: When an LDAP or NSS error occurs, the user is redirected to the URL you specify in the text box. This is optional and allows you to customize the error handling process. If you do not customize it, a standard error page is displayed.

- 9 Click *OK*, then click *Apply Changes*.

- 10 Continue with [Section 12.4.4, “Assigning a Form Fill Policy to a Protected Resource,” on page 154](#) or [Section 32.3.2, “Creating a Login Failure Policy,” on page 395](#).

32.3.2 Creating a Login Failure Policy

The login failure policy can be part of the same policy as the Form Fill policy, if both share the same URL. If the user is redirected to a different page when login fails, it is best to create a separate policy for that page, create a protected resource that includes just that page, and assign your login failure policy to that resource.

To create a login failure policy:

- 1 In the Administration Console, click *Access Manager > Policies > New*.
- 2 Specify a name for the policy, select *Access Gateway: Form Fill* as its *Type*, then click *OK*.
- 3 In the *Actions* section, click *New*, then select *Form Login Failure*.
- 4 In the *Form Selection* section, identify the form. This section uses the same criteria for identifying a form as the Form Fill policy. See [Step 5 on page 392](#).
- 5 In the *Login Failure Processing* section, define the actions you want executed when a user fails to log in. Fill in the following fields:

Redirect to URL: When a user's login attempt fails, use this option with its text box to specify the URL you want the user redirected to. This is optional and allows you to customize what happens on login failures.

Clear Data Value: Select this field to delete the user's stored data for a Form Fill policy. If the user has the ability (and perhaps the requirement) to periodically change his or her password or any other information on the form, you need to select this field. Otherwise, the wrong data can be stored for the user, and the Access Gateway has no way of updating the information.

- 6 Click *OK*, then click *Apply Changes*.
- 7 Continue with [Section 12.4.4, “Assigning a Form Fill Policy to a Protected Resource,” on page 154](#).

32.3.3 Troubleshooting a Form Fill Policy

When creating a new Form Fill policy, you should always select the *Debug Mode* option. This option prepares the form for submission, but doesn't submit the form until you click the *Submit* button. This allows you to view the source, and determine if the policy is generating the required data.

Check to ensure that all input fields have valid names, that the fields are being filled in the correct order, and that any JavaScript commands have been entered correctly.

On the NetWare[®] Access Gateway, you can add a command to the startup NCF file that increases the detail generated from a Form Fill event.

Edit the `sys:system\ap_start.ncf` file, and add the following to the `load sso` line.
`load sso /D<number> L<number>`

Replace `<number>` with a value of 1 to 5, with 5 specifying the most detail.

The following table describes these options:

Switch	Purpose
<code>/D<number></code>	Enables Form Fill debugging at the specified level: <ul style="list-style-type: none">◆ 0 enables standard output◆ 1-5 enables debugging output, with each higher level providing more information. Default: 0
<code>/L<number></code>	Enables or disables the logging of Form Fill debugging information. <ul style="list-style-type: none">◆ 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

For this setting to take effect, you need to restart the Access Gateway.

32.4 Creating and Managing Shared Secrets

A shared secret is an object that holds name and value pairs for Form Fill and Identity Injection policies.

- ◆ If your HTML form prompts the user for more than credential information, you need to create a shared secret to store the values.
- ◆ If your Web servers requires some name/value pairs to be injected and these are not available from the HTTP request, you need to create a shared secret to store these name/value pairs so that they can be injected into the header before it is sent to the Web server.

Access Manager supports the creation and use of local secret stores and the use of remote secret stores if you are using eDirectory™ as your User Store and have set up a secret store there.

- ♦ [Section 32.4.1, “Local Shared Secret,” on page 397](#)
- ♦ [Section 32.4.2, “Remote Shared Secret,” on page 397](#)

32.4.1 Local Shared Secret

You can create a shared secret store as part of the process of creating a Form Fill or Identity Injection policy. You can also create and manage a local shared secret independent of a policy:

- 1 In the Administration Console, click *Access Manager > Identity Servers > Setup > Custom Attributes*.
- 2 To create a new shared secret, click *New* in the *Shared Secret Names* section and specify a display name.
The Identity Server creates and encrypts the object, and provides a way to store any name and value pair specified in the Form Fill policy.
- 3 To delete a shared secret, select the display for the shared secret and click *Delete*.
- 4 Click *Apply*.

A local shared secret does not contain any name/value pairs until you configure a Form Fill policy to add name/value pairs or enable the *Allow End Users to See Credential Profile* option. This option allows the username and password to be stored in the local secret store. For more information, see [Section 11.4, “Configuring Credential Profile Security and Display Settings,” on page 118](#).

Your applications, how you use them, and your personal preferences determine whether you create one shared secret store and use it for all your applications or whether you create a shared secret store for each application. If the applications use some of the same secrets, then it makes sense to use the same secret store for these applications. If the application does not use the same secrets as another application and you want the freedom to remove the application and its secrets without effecting other applications, then you should create a separate secret store for this application.

32.4.2 Remote Shared Secret

If you have installed a shared secret store on an eDirectory server, you can configure Access Manager to use that secret store. The following conditions must be met before the Access Manager can use the secret store:

- ♦ The eDirectory server must have Novell® SecretStore® installed and must be configured for one or more shared secret stores.
- ♦ The Identity Server must be configured to use the eDirectory server as a user store
- ♦ Secure communications must be set up between the Identity Server and the user store.

To use a remote shared secret:

- 1 In the Administration Console, click *Access Manager > Identity Servers > [Name of Configuration] > Liberty > Web Service Provider > Credential Profile*.
- 2 In the *Novell Secret Store User Store References* section, click *New*.
- 3 Select the user store that has Novell SecretStore installed, then click *OK*.
- 4 Apply the changes.

- 5 Click *Identity Servers > Setup > Custom Attributes*.
- 6 In the *Shared Secret Names* section, click *New*, specify the name of a shared secret store on the eDirectory server, then click OK.
This name must match the eDirectory name
- 7 Click *Setup*, then click *Update Servers*.

32.5 Importing and Exporting Form Fill Policies

You can import and export Form Fill policies in order to run them in other Access Manager configurations and to analyze the policy. The policy is exported as a text file with XML tags. We do not recommend editing the exported file with a text editor. Any changes you want to make to a policy ought to be done through the Administration Console.

To export a Form Fill policy:

- 1 In the Administration Console, click *Access Manager > Policies*.
- 2 Select a Form Fill policy, then click *Export*.
- 3 (Optional) Modify the name suggested for the file.
- 4 Click *OK*.
- 5 Using the features of your browser, specify where the file is to be copied.

To import a policy:

- 1 Make sure any referenced shared secret stores have been created. See [Section 32.4, “Creating and Managing Shared Secrets,” on page 396](#).
- 2 If the policy uses LDAP or Liberty Profile attributes, make sure the Identity Server has been configured for these same attributes.
- 3 In the Administration Console, click *Access Manager > Policies*.
- 4 If the policy uses LDAP or Liberty Profile attributes, make sure the Identity Server has been configured for these same attributes.
- 5 Click *Import*, then browse to the location of the file.
- 6 Click *OK*.
- 7 When the policy appears in the list, click *Apply Changes*.

Monitoring Access Manager Components

VII

This section describes the various ways you can determine whether the Access Manager is functioning normally and whether an Internet attack is in progress. This section discusses the following topics:

- ◆ Chapter 33, “Enabling Auditing,” on page 401
- ◆ Chapter 34, “Configuring Logging,” on page 409
- ◆ Chapter 35, “Viewing Statistics,” on page 427
- ◆ Chapter 36, “Managing Server Health,” on page 443
- ◆ Chapter 37, “Reviewing Command Status,” on page 449
- ◆ Chapter 38, “Reviewing Alerts,” on page 455

Access Manager includes a licensed version of Novell Audit to provide compliance assurance logging and to maintain audit log entries that can be subsequently included in reports. In addition to selectable events, device generated alerts are automatically sent to the audit server.

Audit logs record events that have occurred in the identity and access management system and are primarily intended for auditing and compliance purposes. The types of events that are logged include the following:

- ◆ Starting, stopping, and configuring a component
- ◆ Success or failure of user authentication
- ◆ Role assignment
- ◆ Allowed or denied access to a protected resource
- ◆ Error events
- ◆ Denial of service attacks
- ◆ Security violations and other events necessary for verifying the correct and expected operation of the identity and access management system.

Audit logging does not track the operational processing of the Access Manager components, that is, the processing and interactions between the Access Manager components required to fulfill a user request (for this type of logging, see [Section 34.2, “Configuring Component Logging,” on page 410](#)). Audit logs record the results of user and administrator requests and other system events. While the primary purpose for audit logging is for auditing and compliance, the types of events logged can also be useful for detecting abnormal and error conditions and can be used as a first alert mechanism for system support. You can configure the audit log entries to generate alerts by leveraging the Novell Audit Notification feature. You can select to generate e-mail, syslog, and SNMP notifications.

Access Manager has been assigned the Novell Audit server-alert event code 0x002E0605. The Novell Audit Platform Agent is responsible for packaging and forwarding the audit log entries to the configured Novell Audit server. If the Novell Audit server is not available, the platform agent caches log entries until the server is operational and can accept audit log data.

For additional information about Novell Audit, see [Novell Audit 2.0.2 \(http://www.novell.com/documentation/novellaudit20/index.html\)](http://www.novell.com/documentation/novellaudit20/index.html) at the Novell Documentation Web site.

This section describes the following Access Manager features of auditing:

- ◆ [Section 33.1, “Configuring Access Manager for Novell Auditing,” on page 402](#)
- ◆ [Section 33.2, “Enabling Identity Server Audit Events,” on page 403](#)
- ◆ [Section 33.3, “Enabling Access Gateway Audit Events,” on page 405](#)
- ◆ [Section 33.4, “Enabling SSL VPN Audit Events,” on page 406](#)

33.1 Configuring Access Manager for Novell Auditing

By default, Access Manager is preconfigured to use the Novell Audit server it installs on the first instance of the Administration Console. If you install more than one instance of the Administration Console for failover, Novell Audit is installed with each instance. However, if you already use Novell Audit, you can continue using your existing installation with Access Manager. You'll need to configure Access Manager to use your audit servers. You'll also need to register the Access Manager with your Audit servers by importing the `nids_en.lsc` and `sslvpn_en.lsc` files.

This section includes the following topics:

- ♦ [Section 33.1.1, “Specifying the Logging Server and Events,” on page 402](#)
- ♦ [Section 33.1.2, “Generating Queries,” on page 403](#)

33.1.1 Specifying the Logging Server and Events

The Secure Logging Server manages the flow of information to and from the Novell auditing system. It receives incoming events and requests from the Platform Agents, logs information to the data store, monitors designated events, and provides filtering and notification services. It can also be configured to automatically reset critical system attributes according to a specified policy.

- 1 To specify the logging server, click *Access Manager > Auditing*.
- 2 Enter the IP address or DNS name of your auditing server (your Administration Console machine).

Server: The audit logging server you want to use. For failover protection, you can configure up to three servers. By default, the system uses the primary Administration Console IP address. If you want to use a different server, specify that server's IP address here.

Port: The port where the Platform Agents connect to the Secure Logging Server.
- 3 Under *Management Console Audit Events*, specify the system-wide events you want to audit:

Select All: Selects all of the audit events.

Health Changes: Generated whenever the health of a server changes.

Server Imports: Generated whenever a server is imported into the Administration Console.

Server Deletes: Generated whenever a server is deleted from the Administration Console.

Configuration Changes: Generated whenever you change a server configuration.
- 4 Click *OK*.

If you did not change the address or port of the Secure Logging Server, this completes the process. It may take up to fifteen minutes for the events you selected to start appearing in the audit files.

If you changed the address or the port of the Secure Logging Server, complete the following the steps.
- 5 If the Administration Console is the only Access Manager component installed on the machine, edit the Novell Audit Configuration file.

For security reasons, this file cannot be edited from the Administration Console when it is the only Access Manager component on the machine.

Edit the `/etc/logevent.conf` file and specify the new address and port of the Secure Logging Server.

6 Restart the Administration Console. From a terminal window, enter the following command:
`/etc/init.d/novell-tomcat4 restart`

7 Restart every device imported into the Administration Console.

The devices (Identity Server, Access Gateway, SSL VPN, J2EE Agents) will not start reporting events until they have been restarted.

33.1.2 Generating Queries

Queries let you create, run, edit and delete queries and event verifications. You can create two kinds of queries in Access Manager: manual queries and saved queries. Manual queries are simply queries that are not saved; they only run one time. All verification queries are saved. Saved queries and verifications are listed in the Queries list and can be run again and again against different databases.

Access Manager uses queries to request information from MySQL* and Oracle* databases. All queries are defined in SQL*. Although you must be familiar with the SQL language to create SQL query statements, this is the most powerful and flexible query method.

For information about queries, see [Novell Audit 2.0.2 \(http://www.novell.com/documentation/novellaudit20/index.html\)](http://www.novell.com/documentation/novellaudit20/index.html).

33.2 Enabling Identity Server Audit Events

All user and administrator actions can be logged to Novell Audit. You can generate a Novell Audit logging event to indicate whether authentications are successful or unsuccessful. The following steps assume that you have already set up Novell Audit on your network. For more information, see [Section 33.1, “Configuring Access Manager for Novell Auditing,” on page 402](#)

- 1** In the Administration Console, click *Access Manager > Identity Server > Setup > [Configuration] > Logging*.
- 2** In the *Novell Audit Logging* section, select *Enabled*.
- 3** Select the events for notification.

Select All. Select this option for all events. Otherwise, select one or more of the following:

Event	Description
Login Provided	Generated when an identity provider sends authentication to a service provider. Role assignment audit events are included in authentication audit events for the identity server.
Login Provided Failure	Generated when an identity provider attempts to send authentication to a service provider but fails.

Event	Description
Login Consumed	Generated when the Identity Server is authenticated either locally or by an external identity provider. Role assignment audit events are included in authentication audit events for the identity server.
Login Consumed Failure	Generated when the Identity Server initiates authentication, but the process fails.
Logout Provided	Generated when an identity provider sends a logout request to a service provider that it has authenticated.
Logout Local	Generated when the Identity Server receives a command to logout from the user.
Federation Request Sent	Generated when a service provider attempts to federate with an identity provider.
Federation Request Handled	Generated by the Identity Server when processing a request for federation.
Defederation Request Sent	Generated by the identity provider when a request for defederation is sent to another provider.
Defederation Request Handled	Generated when the Identity Server processes a request for defederation.
Register Name Request Handled	Generated when the Identity Server processes a request for changing a name identifier.
Attribute Query Request Handled	Generated by the Identity Server when processing an attribute request from a service provider.
Web Service Query Handled	Causes a Web service query request to be sent to an identity provider.
Web Service Modify Handled	Causes a web service modify request to be sent to an identity provider.
User Account Provisioned	Generated by the Identity Server when functioning as an identity consumer and when an account has been provisioned.
User Account Provisioned Failure	Generated by the Identity Server when functioning as an identity consumer and when account provisioning has failed.
Ldap Connection Lost	Generated when the LDAP connection is lost.
Ldap Connection Reestablished	Generated when the LDAP connection is reestablished.
Server Started	Generated when the server gets a start command from the server communications module.

Event	Description
Server Stopped	Generated when the server gets a stop command from the server communications module.
Server Refreshed	Generated when the server gets a refresh command from the server communications module.
Intruder Lockout Detected	Generated when an attempt to login as a particular user with an invalid password has occurred more times than is allowed by the directory.
Component Log Severe Messages	Logged for all component messages with level of Severe.
Component Log Warning Messages	Logged for all component messages with level of Warning.

4 Click *Apply*, then *OK*.

5 Click *Setup > Update Servers*.

After changing settings for Novell Audit logging for an Identity Server configuration, you must click Update Servers on the Setup page. You must also restart the Novell Audit server.

33.3 Enabling Access Gateway Audit Events

The *Novell Audit* option in the Access Gateway allows you to configure the events you want audited. The following steps assume that you have already set up Novell Audit on your network. For more information, see [Section 33.1, “Configuring Access Manager for Novell Auditing,” on page 402](#).

1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Novell Audit*.

The screenshot shows a configuration panel titled "Events". At the top, there is a "Select All" checkbox. Below it, there is a grid of 12 checkboxes for different event types:

<input type="checkbox"/> Access Denied	<input type="checkbox"/> Access Allowed	<input type="checkbox"/> Identity Injection Failed	<input type="checkbox"/> Identity Injection Parameters
<input type="checkbox"/> System Started	<input type="checkbox"/> System Shutdown	<input type="checkbox"/> Form Fill Success	<input type="checkbox"/> Form Fill Failed
<input type="checkbox"/> URL Accessed	<input type="checkbox"/> URL Not Found	<input type="checkbox"/> IP Access Attempted	

Below the grid, a note states: "Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel." At the bottom, there are "OK" and "Cancel" buttons.

2 Select the events for notification.

Select All. Select this option for all events. Otherwise, select one or more of the following:

Event	Description
Access Denied	Generated when a requested action is denied because the requester has insufficient access rights to a URL.

Event	Description
System Started	Generated when the Access Gateway is started.
URL Accessed	Generated when a user accesses a URL.
Access Allowed	Generated when a requested action is allowed because the requester has the correct access rights to a URL.
System Shutdown	Generated when the Access Gateway is stopped.
URL Not Found	Generated when a requested URL cannot be found.
Identity Injection Failed	Generated when an Identity Injection policy fails to obtain a requested value to inject into the HTTP header.
Form Fill Success	Generated when a Form Fill policy successfully fills in a form.
IP Access Attempted	Generated when a user attempts to access a URL with an IP address instead of the published DNS name configured in the Access Gateway.
Identity Injection Parameters	Generated when the Identity Injection policy successfully injects data into the HTTP header. Some of the data might be injected with the value field empty. When this happens, this event should also produce an <i>Identity Injection Failed</i> event.
Form Fill Failed	Generated when a Form Fill policy fails to successfully fill in a form.

3 To save your modifications, click *OK*, then on the Configuration page, click *Apply Changes*.

33.4 Enabling SSL VPN Audit Events

The *Novell Audit Settings* option allows you to configure the events you want audited. The following steps assume that you have already set up Novell Audit on your network. For more information, see [Section 33.1, “Configuring Access Manager for Novell Auditing,” on page 402](#).

- 1 In the Administration Console, click *Access Manager > SSL VPNs > Edit*.
- 2 Select *Novell Audit Settings* from the *Novell Audit and Alerts* section. The Novell Audit Settings for SSL VPN page is displayed.

Servers > Configuration > Novell Audit Settings

Novell Audit Settings for SSL VPN: 12.12.12.124 ?

Events	
<input type="checkbox"/> Select All	
<input checked="" type="checkbox"/> Authentication Logs	<input type="checkbox"/> Command Line Interface Logs
<input checked="" type="checkbox"/> Command Line Interface Debug Logs	<input checked="" type="checkbox"/> Servlet Communications Logs
<input checked="" type="checkbox"/> Connection Manager Logs	<input checked="" type="checkbox"/> Certificate Management Logs
<input checked="" type="checkbox"/> Certificate Management Debug Logs	<input checked="" type="checkbox"/> SSL VPN Incoming Connections Logs
<input checked="" type="checkbox"/> SSL VPN Incoming Connections Debug Logs	<input checked="" type="checkbox"/> Other SSL VPN Gateway Logs

Note: This configuration is currently being edited by cn=admin,o=novell (164.99.170.122) since 6/21/06 6:36:11 PM. Click unlock on the main page to override.

- 3 Select the *Select All* option to receive logs for all the events. Otherwise, select one or more of the following:

Event	Description
Authentication Logs	Generates a log file with the authentication details.
Command Line Interface Logs	Generates a log file with command line actions.
Command Line Interface Debug Logs	Generates a log file with command line actions.
Servlet Communications Logs	Generates a log file with information on servlet communication.
Connection Manager Logs	Generates a log file containing information on the connection activity.
Certificate Management Logs	Generates a log file with certificate management information.
Certificate Management Debug Logs	Generates a log file with certificate management information.
SSL VPN Incoming Connections Logs	Generates a log file containing information on the incoming connection.
SSL VPN Incoming Connections Debug Logs	Generates a log file containing debug information on the incoming connection.
Other SSL VPN Gateway Logs	Generates a log file containing miscellaneous information.

4 To save your modifications, click *OK*, then click *Apply Changes* on the Configuration page.

- ♦ [Section 34.1, “Understanding the Types of Logging,” on page 409](#)
- ♦ [Section 34.2, “Configuring Component Logging,” on page 410](#)
- ♦ [Section 34.3, “Configuring Debug Trace Logging,” on page 415](#)
- ♦ [Section 34.4, “Configuring Access Gateway Logging,” on page 416](#)

34.1 Understanding the Types of Logging

Access Manager supports three types of logging:

- ♦ [Section 34.1.1, “Component Logging for Troubleshooting Configuration or Network Problems,” on page 409](#)
- ♦ [Section 34.1.2, “Debug Trace Logging to Discover Software Problems,” on page 410](#)
- ♦ [Section 34.1.3, “HTTP Transaction Logging for Proxy Services,” on page 410](#)

34.1.1 Component Logging for Troubleshooting Configuration or Network Problems

Each Access Manager component maintains log files that contain entries documenting the operation of the component. Component file logging records the processing and interactions between the Access Manager components that occur while satisfying user and administrative requests and during general system processing. By enabling the correct levels of logging for the various Access Manager components, an administrator can monitor how the Access Manager processes user and administrative requests. Transaction flows have been defined to help the administrator identify the processing steps that occur during the execution of specific types of user or administrative requests. All component file logs include tags and values that allow the administrator to identify and correlate which component file log entries pertain to a given transaction and user.

Component file logs are not primarily intended for debugging the software itself, although they can be used to detect software that is not behaving properly. Rather the intent of component file logging is to document the operational processing of the Access Manager components so that system administrators and support personnel can identify and isolate problems caused by configuration errors, invalid user data, or network problems such as broken connection. However, component file logging will typically be the first step in identifying software bugs.

Component file logging is more verbose than audit logging. It increases processing load, and on a day-to-day basis, it should be enabled only to log error conditions and system warnings. If a specific problem occurs, component file logging can be set to *info* or *config* to gather the information needed to isolate and repair the detected problem. Once the problem is resolved, component file logging should be reconfigured to log only error conditions and system warnings.

Log files can be configured to include entries for the following events:

- ♦ Initialization and shutdown
- ♦ Configuration

- ◆ Events processed by the component, such as authentication, role assignment, resource access, and policy evaluation
- ◆ Error conditions

See [Section 34.2, “Configuring Component Logging,” on page 410](#).

34.1.2 Debug Trace Logging to Discover Software Problems

Debug trace logging is used to debug the software execution flow of an Access Manager component. Debug trace logging is the most verbose of the Access Manager logging categories and by its nature includes data that generally encompasses the information provided by both audit logging and component file logging. The information contained in debug trace logs can generally only be interpreted by those with access to the source code such as Novell support personnel or software engineers. System administrators might be required to enable debug trace logging in order to provide support personnel with the information necessary to resolve a software bug. Debug trace logging should not be enabled during normal operation of Access Manager.

See [Section 34.3, “Configuring Debug Trace Logging,” on page 415](#).

34.1.3 HTTP Transaction Logging for Proxy Services

The Access Gateway allows you to log HTTP transactions. You can log what happens with an HTTP request and response

- ◆ Between the browser and the Access Gateway
- ◆ Between the Access Gateway and the back-end Web server

You select fields from the HTTP header of a request and these fields are logged. You can then use these logged transactions to bill customers for Web services or to troubleshoot whether a request is refused because the browser didn't send the required information or because the Access Gateway didn't send the Web server the required information. This type of logging conforms to the W3C specification for proxy server logging in the common and extended log formats. This type of logging provides no information about the exchanges between the Access Gateway and the Identity Server. If you need to discover whether the Access Gateway is obtaining the correct information from the Identity Server for an Identity Injection or Form Fill policy, you need to turn on Component logging. See [Section 34.2, “Configuring Component Logging,” on page 410](#).

For HTTP transaction logging, see [Section 34.4, “Configuring Access Gateway Logging,” on page 416](#).

34.2 Configuring Component Logging

You can enable and configure how the system performs logging. Logging is the main tool you use for debugging the Identity Server configuration. All administrative and end-user actions and events are logged to a central event log. This allows easy access to this information for security and operational purposes. Additionally, the log system provides the ability to monitor ongoing activities (such as identity provider authentication activity, up-time of the system, and so on) using this page. File logging is not enabled by default.

Identity Servers, Access Gateways (Linux and NetWare[®]), and embedded service providers use these logging features. If you change or enable logging, you must update the Identity Server configuration (using Update Servers on the Setup page) and restart the service providers on the

Access Gateways, in order to apply the changes. When you disable logging, you must also restart the Access Gateway embedded service provider. See [Section 2.1.6, “Restarting the Access Gateway,” on page 29](#).

This section describes the following about component logging:

- ♦ [Section 34.2.1, “Enabling Component Logging,” on page 411](#)
- ♦ [Section 34.2.2, “Downloading the Log Files,” on page 412](#)
- ♦ [Section 34.2.3, “Understanding the Log File Format,” on page 414](#)

34.2.1 Enabling Component Logging

File logging records the actions that have occurred. For example, Web servers maintain log files listing every request made to the server. With log file analysis tools, it’s possible to get a good idea of where visitors are coming from, how often they return, and how they navigate through a site. The content logged to file logging can be controlled by specifying logger levels and by enabling Statistics Logging.

- 1 In the Administration Console, click *Access Manager > Identity Server > Setup > [Configuration] > Logging*.
- 2 The following options are available for component logging in the *File Logging* section:
 - ♦ **Enabled:** Enables file logging for this server and its associated Embedded Service Providers.
 - ♦ **Echo To Console:** Copies the Identity Server log file to `/var/opt/novell/tomcat4/logs/catalina.out`. You can download the file from *Access Manager > Auditing > General Logging*. If you want to view Identity Server logs mixed with logs from other application devices, you use `catalina.out`.

For the Embedded Service Providers, it depends upon the platform:

- ♦ For a Linux Access Gateway, this sends the messages to the `catalina.out` file of the Access Gateway.
- ♦ For a NetWare Access Gateway, this sends the messages to the NetWare console.
- ♦ For a SSL VPN, this sends the messages to the `catalina.out` file of the SSL VPN.
- ♦ **Log File Path:** Specifies the path that the system uses to save the Identity Server XML log file. The default path is `tomcat application directory/web-inf/logs`. If you change this path, you must ensure that the user associated with configuring the identity or service provider has administrative rights to the Tomcat application directory in the new path.

If you have a mixed platform environment (for example, the Identity Server is installed on Linux and the Access Gateway is on NetWare), do not specify a path. In a mixed platform environment, you must use the default path.
- ♦ **File Wrap:** Specifies the frequency (hour, day week, month) for the system to use when closing a log file and creating a new one. The system saves each file based on the time you specify and attaches the date and/or time to the filename.

- 3 In the Component File Logger Levels, you can specify the logging sensitivity for the following:
 - Application:** Logs system-wide events, except events that belong to a specific subsystem.

Liberty: Logs events specific to the Liberty IDFF Protocol and Profiles.

SAML 1: Logs events specific to the SAML1 Protocol and Profiles.

SAML 2: Logs events specific to the SAML2 Protocol and Profiles.

Web Service Provider: (Liberty) Logs events specific to fulfilling Web service requests from other Web Service Consumers.

Web Service Consumer: (Liberty) Logs all events specific to requesting Web services from a Web Service Provider.

Use the drop-down menu to categorize logging sensitivity. Higher logging levels include the lower levels in the log.

- ♦ **Severe:** Logs serious failures that can cause system processing to not proceed.
- ♦ **Warning:** Logs potential failures, but the impact on execution is minimal. Warnings indicate that you should be aware that this event is happening and might want to make a configuration change to avoid it.
- ♦ **Info:** Logs informational events. No execution or data impact occurred.
- ♦ **Config:** Logs static configuration information. The system logs any configuration errors under one of the primary three levels: Severe, Warning, and Info.
- ♦ **Off:** Turns off logging.

4 (Optional) Enable statistics logging.

When statistics logging is enabled, the system periodically sends the system statistics, in string format, to the current file logger. Statistical data (such as counts, levels, and so on) are included in the file log.

4a In the *Statistics Logging* section, select *Enabled*.

4b In the *Log Interval* field, specify the time interval in seconds that statistics are logged.

5 Click *OK*.

6 Update the Identity Server configuration (using *Update Servers* on the *Setup* page).

7 Restart the embedded service providers on the Access Gateways, in order to apply the changes.

When you disable component logging, you need to update the Identity Server configuration and restart the embedded service providers.

34.2.2 Downloading the Log Files

The *General Logging* page displays the location of the files that the Access Manager components use for logging system messages. The J2EE agent is the one exception. The J2EE agent uses the J2EE global logger, and the location of this file is customizable. For information about J2EE agent log files, see “[Viewing Log Files](#)” in the *Novell Access Manager 3.0 J2EE Agent Guide*.

To view or download the log file, click the filename, then use your browser prompts to either view or save the file. You can use a text editor to view the logs.

1 In the Administration Console, click *Auditing > General Logging*.

2 Click the link for the log file name, then save it to disk.

Table 34-1 Access Manager Log Files

Component	Filename	Description
Administration Console		
	<code>/var/opt/novell/tomcat4/logs/catalina.out</code>	Contains Tomcat related errors.
	<code>/opt/novell/devman/share/logs/app_sc.0.log</code>	Contains events related to importing devices, device configuration changes, health status changes, statistics reporting, and communication problems.
	<code>/opt/novell/devman/share/logs/app_cc.0.log</code>	Contains events related to policy configuration.
	<code>/opt/novell/devman/share/logs/platform.0.log</code>	Contains XML related events for configuration changes. This log file contains very little useful information for system administrators.
Identity Server		
	<code>/var/opt/novell/tomcat4/logs/catalina.out</code>	Logging to this file only occurs if you have selected the <i>Echo to Console</i> option from the <i>Identity Servers > Setup > [Configuration] > Logging</i> page. When component logging has been set to info for Applications, it contains entries tracing user authentication and role assignment.
	<code>/opt/novell/devman/jcc/logs/jcc-0.log.0</code>	Contains the log entries for the server communications module related to interaction of the Identity Server with the Administration Console such as imports, certificates, and configuration.
Linux Access Gateway		
	<code>/var/log/novell/reverse/common</code>	If logging is enabled on one or more reverse proxies (see Section 34.4, "Configuring Access Gateway Logging," on page 416), this directory contains the log files. A directory is listed for each reverse proxy on which you have enabled logging.
	<code>/var/log/ics_dyn.log</code>	Contains all the log entries generated by the Linux Access Gateway. Use syslog to control file rolling and log file distribution.

Component	Filename	Description
	/opt/novell/devman/jcc/logs/jcc-0.log.0	Contains the log entries for the server communications module related to interaction of the Access Gateway with the Administration Console such as imports, certificates, and configuration.
	/var/opt/novell/tomcat4/logs/catalina.out	Logging to this file only occurs if you have selected the <i>Echo to Console</i> option from the <i>Identity Servers > Setup > [Configuration] > Logging</i> page. Check this file for entries tracing the evaluation of authorization, identity injection, and form fill policies.
NetWare Access Gateway		
	log:\etc\proxy\data\logs\reverse\common\	If logging is enabled on one or more reverse proxies (see Section 34.4, "Configuring Access Gateway Logging," on page 416), this directory contains the log files. A directory is listed for each reverse proxy on which you have enabled logging.
	SYS:\etc\proxy\data\debug.log	Contains the abend messages.
	SYS:\jcc\logs\jcc-0.log.0	Contains the log entries for the server communications module.
SSL VPN		
	/var/opt/novell/tomcat4/logs/catalina.out	Logging to this file only occurs if you have selected the <i>Echo to Console</i> option from the <i>Identity Servers > Setup > [Configuration] > Logging</i> page.
	/opt/novell/devman/jcc/logs/jcc-0.log.0	Contains the log entries for the server communications module related to interaction of the SSL VPN with the Administration Console such as imports, certificates, and configuration.

34.2.3 Understanding the Log File Format

There is not a fixed field format for log file entries. However, because most requests handled by Access Manager are processed by multiple Access Manager components, there is a mechanism defined that facilitates the correlation of log entries for a single Access Manager request in the various component log files. This mechanism defines common tag values for the following information:

- ◆ error code, tag: AM#error-code:

error-code is an error number defined in [Error Codes \(http://www.novell.com/documentation/beta/novellaccessmanager/pdfdoc/qsc/configissues.pdf\)](http://www.novell.com/documentation/beta/novellaccessmanager/pdfdoc/qsc/configissues.pdf).

- ◆ authentication identifier, tag: AMAUTHID#auth-id:
auth-id is a unique identifier assigned by the IDP to each user authentication
- ◆ event identifier, tag: AMEVENTID#event-id:
event-id is a unique identifier assigned to an Access Manager operation or request
- ◆ device identifier, tag: AMDEVICE#device-id:
device-id is the JCC device identifier of the device generating the log entry

For policy logging information, see [Section 41.3, “Understanding Policy Evaluation Traces,” on page 475](#).

34.3 Configuring Debug Trace Logging

Novell® recommends that you use the tracing feature only for software debugging. Sensitivity levels do not apply to trace logging. Therefore, you would not activate this feature during production, because it impacts processing speed. This feature is filterable by Java* class or package.

To enable debug trace logging:

1 In the Administration Console, click *Access Manager > Identity Server > Setup > [Configuration] > Logging*.

2 In the *File Logging* section, select *Enabled*.

It is assumed that you have set up the Echo To Console, Log File Path, and File Wrap options when you set up Component File Logging. If you need help with these options, see [Step 2 in Section 34.2.1, “Enabling Component Logging,” on page 411](#).

3 In the *Trace Logging* section, select *Enabled*.

This option enables trace logging and the *Custom Content Filter* link.

4 (Optional) Click *Custom Content Filter* to display the *Edit custom trace logging content filter* text box.

The *Custom Content Filter* allows you to focus trace content on a specific section of the system where you suspect a problem exists. The filter is an XML document that specifies which trace logging content to send to the trace logger. You can limit the trace logging to one or more Java class files, or to one or more Java packages, or to one or more Novell-defined thread identifiers.

4a Click *Default* to insert the default XML text.

4b To validate this XML, the Java class or package must be completed.

Knowledge of the Java class structure of the Access Manager product is required to create a Custom Content Filter. Therefore, it is recommended that this feature be used only with help from Novell Customer Support.

For information about using the filter, see [Appendix D, “Logging: Using the Custom Content Filter,” on page 529](#).

5 To quickly trace content for specific parts of the system, select one of the following filters. The results are written to the file logger.

Application: Logs system-wide trace content, except content that belongs to a specific protocol subsystem.

Liberty: Logs trace content specific to the Liberty IDFF protocol and profiles.

SAML 1: Logs trace content specific to the SAML 1.1 protocol and profiles.

SAML 2: Logs trace content specific to the SAML 2 protocol and profiles.

Web Service Provider: Logs trace content specific to fulfilling Web service requests from other Web service consumers.

Web Service Consumer: Logs trace content specific to requesting Web services from a Web service provider.

Request/Response: Logs trace content specific to sending and receiving requests on all protocols, such as Liberty, SAML 1.1, and SAML 2.

User Stores: Logs trace content specific to accessing user stores. During a health check, the system includes all user stores in the configuration store.

Configuration: Logs trace content specific to configuring the system.

- 6 Click OK.
- 7 Update the Identity Server configuration (using Update Servers on the Setup page).
- 8 Restart the embedded service providers on the Access Gateways, in order to apply the changes.
When you disable trace logging, you need to update the Identity Server configuration and restart the embedded service provides.

34.4 Configuring Access Gateway Logging

Logging HTTP transactions has associated costs. The Access Gateway is capable of handling thousands of transactions per second. If transaction volume is high and each log entry consumes a few hundred bytes, the Access Gateway can fill up the available disk space in a matter of minutes. HTTP logging also increases system overhead, which causes some degradation in performance. By default, the logging of HTTP transactions is turned off. Before enabling logging, you need to determine what needs to be logged and then plan a logging strategy.

- ◆ [Section 34.4.1, “Determining Logging Requirements,” on page 416](#)
- ◆ [Section 34.4.2, “Calculating Roll Over Requirements,” on page 417](#)
- ◆ [Section 34.4.3, “Enabling Logging,” on page 420](#)
- ◆ [Section 34.4.4, “Configuring Common Log Options,” on page 420](#)
- ◆ [Section 34.4.5, “Configuring Extended Log Options,” on page 422](#)
- ◆ [Section 34.4.6, “Configuring Log Pushing,” on page 424](#)
- ◆ [Section 34.4.7, “Configuring the Size of the Log Partition,” on page 426](#)

34.4.1 Determining Logging Requirements

Because logging requirements and transaction volume vary widely, Novell cannot make recommendations regarding a specific logging strategy. The following tasks guide you through the process of creating a strategy that fits your business needs.

- 1 Identify the reasons for tracking transactions such as customer billing, statistical analysis, or growth planning.
- 2 Determine which resources need logging.

You enable logging at the proxy service level. If you have a proxy service protecting resources whose transaction do not need to be logged, reconfigure your proxy services so that the proxy

service you configure for logging contains only the resources for which you want to log transactions.

3 Determine what information you need in each log entry.

The common configuration for a log entry contains minimal information: the date, time, and client IP address for each entry. If you need more information, you can select the extended log configuration. Do not select all available fields, but carefully select what you really need. For example, you can include cookie information, but cookie information can consume a large amount of space and might not include any critical information you need.

You should log only the essential data because a few bytes can add up quickly when the Access Gateway is tracking thousands of hits every second. For information about what is available in an extended log profile, see [Section 34.4.5, “Configuring Extended Log Options,” on page 422.](#)

4 Design a roll over strategy.

A log must be closed before it can be downloaded to another server for analysis or deleted. You specify either by time or size when the Access Gateway closes a log file and creates a new one. For each proxy service that you enable for logging, you need to reserve enough space for at least two files: one for logging and one for roll over. To calculate the best procedure, see [Section 34.4.2, “Calculating Roll Over Requirements,” on page 417.](#)

5 Design a log deletion strategy

The Access Gateway has a limited amount of disk space allocated for logging, and you need to decide how you are going to manage this space. You can limit the number of roll over files by number or age. You can also select to copy the files to another server and then delete them. To calculate the best procedure, see [Section 34.4.2, “Calculating Roll Over Requirements,” on page 417.](#)

34.4.2 Calculating Roll Over Requirements

You can have the Access Gateway roll over log files based on time or on size, but not both. If you already know which option you want to use, scan this section and then complete only the calculations pertinent to your choice. If you don't know which option best matches your situation, completing the calculations in this section should help you decide.

The following variables are used in the formulas:

- ◆ **logpartition_size:** The total disk capacity reserved for log files on the Access Gateway.

The Access Gateway reserves 4 GB to share between logging and system files. The system files do not grow significantly, so you can assume that you have about 2 GB for logging. To increase this size, see [Section 34.4.7, “Configuring the Size of the Log Partition,” on page 426.](#)

- ◆ **logentry_size:** The average log entry size.

You can determine this by configuring a proxy service to track the required information, generating traffic to the proxy service, 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 Access Gateway in service and get more accurate data by accessing generated statistics. See [Section 35.2, “Monitoring Access Gateway Statistics,” on page 429.](#)

- ◆ **num_services:** The number of proxy services for which you plan to enable logging.

- ◆ **logs_per_service:** The number of log files, both active and closed, that you want the Access Gateway to generate for each proxy service before the disk fills.

You must plan to have at least two logs per proxy service, but you can have three or more.

The following formulas can help you estimate when the system would run out of resources:

- ◆ [“Calculating diskfull_time” on page 418](#)
- ◆ [“Calculating max_roll_time” on page 418](#)
- ◆ [“Calculating max_log_roll_size” on page 419](#)

Calculating diskfull_time

Using the following formula, you can calculate how long it will take the Access Gateway to fill your logging disk space:

$$\text{diskfull_time in seconds} = \frac{\text{logpartition_size}}{(\text{request_rate} * \text{logentry_size} * \text{num_services})}$$

For example, assume the following:

```
logpartition_size = 1 GB (1,073,741,824 bytes)
request_rate = 1000 requests per second
logentry_size = 1 KB (1,024 bytes)
num_services = 1
diskfull_time = (1 GB) / (1000 * 1 KB * 1) = 1048 seconds (17.47
minutes)
```

The logging disk space will fill up every 17.47 minutes.

To calculate the diskfull_time for your Access Gateway:

- 1 Determine the values of the four variables listed above.
- 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](#).

If your diskfull_time interval is too short to be practical for your roll over schedule, the easiest option is to reduce the log entry size by configuring the proxy services to log less information per transaction.

Calculating max_roll_time

Using the following formula, you can calculate the maximum roll-over time value you should specify in the *Roll over every* field

$$\text{max_roll_time} = \text{diskfull_time} / \text{logs_per_service}$$

For example, assume the following:

```
diskfull_time = 12 hours
logs_per_service = 2
max_roll_time = 12 / 2 = 6 hours
```

If you roll your logs over by time intervals, the maximum time should be less than six hours. Otherwise, scheduling the download and deletion of log files is much more complicated and the window in which this can be done is narrower.

To calculate the `max_roll_time` for your Access Gateway:

- 1 Determine how many log files you want the Access Gateway 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 418, calculate how often you should have the cache device 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 *Maximum File Size* field:

```
max_log_roll_size = logpartition_size / (num_services *  
    logs_per_service)
```

For example, assume the following:

```
logpartition_size = 600 MB
```

```
num_services = 2
```

```
logs_per_service = 3
```

```
max_log_roll_size = 600 MB / (2 * 3) = 100 MB
```

If you roll your logs over when they reach a specific size, the file size must be no more than 100 MB. Otherwise, the system 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.

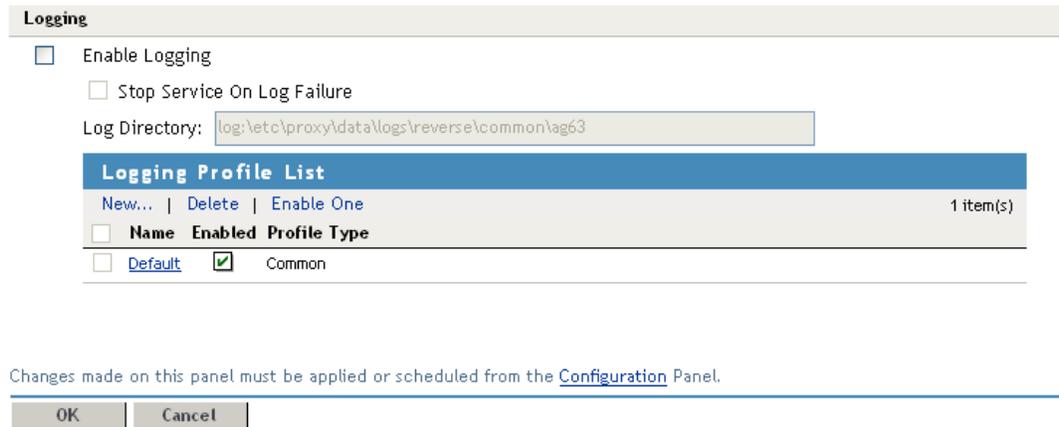
To calculate the `max_log_roll_size` for your Access Gateway, complete the following steps:

- 1 Determine the values of the three variables listed above.
- 2 Using the `max_log_roll_size` formula, calculate the maximum size a log file should reach before the cache device rolls it over.

34.4.3 Enabling Logging

Do not enable logging until you have designed a logging strategy. See [Section 34.4.1, “Determining Logging Requirements,”](#) on page 416.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > Logging*.



Logging

Enable Logging

Stop Service On Log Failure

Log Directory:

Logging Profile List		
New... Delete Enable One		
<input type="checkbox"/>	Name	Enabled Profile Type
<input type="checkbox"/>	Default	<input checked="" type="checkbox"/> Common

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

OK Cancel

- 2 Fill in the following fields:

Enable Logging: Select this field to enable logging.

Stop Service On Log Failure: Select this field if you want the Access Gateway to deny requests to this proxy service because the Access Gateway cannot log entries for it.

Log Directory: Displays the default location for the log files for this proxy service.

- 3 In the *Logging Profile List*, click one of the following options:

- ♦ **New:** Click this option to create a new logging profile. Then specify a name and select either *Common* or *Extended*.
- ♦ **Default:** Click *Default* to modify or view the settings for the Default profile. The Default profile uses the common log options.

- 4 Continue with one of the following:

- ♦ [Section 34.4.4, “Configuring Common Log Options,”](#) on page 420
- ♦ [Section 34.4.5, “Configuring Extended Log Options,”](#) on page 422

34.4.4 Configuring Common Log Options

Use the common log options page to control log roll over and old file options. The data included in a log entry is controlled by a default configuration that includes the following:

- ♦ Date and time of the request
- ♦ Username of the client
- ♦ Remote host name
- ♦ The request line as it came from the client
- ♦ The HTTP status code returned to the client

- ◆ The number of bytes in the document transferred to the client

The Access Gateway does not allow active log files to be deleted. Only log files that have been closed can be deleted. The roll over options allow you to control when a file is rolled over and closed, and a new file is created. The old file options allow you to control when the rolled-over log files are deleted.

To configure a default log file for a selected proxy service:

- 1 Click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > Logging > [Name of Common Log Profile]*.

Rollover Options

Maximum File Size: MB

Roll over every Hour(s) beginning at

Old File Options

Maximum Number of Files:

Delete Files Older Than:

Do Not Delete

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

OK
Cancel

- 2 Select one of the following roll over options:

Maximum File Size: Rolls the file when it reaches the specified number of megabytes.

Rollover every: Rolls the file at the specified interval. You can specify the interval in minutes or hours.

- ◆ **beginning:** Specifies the day that the interval should begin. You can select a day of the week or the first of the month.
- ◆ **at:** Select the hour of the day that the interval should begin and the time zone (either the local time zone or GMT).

- 3 Select one of the following old file options.

Maximum Number of Files: Allows you to limit the number of old log files on the system to the number specified in this option. The oldest file is automatically deleted when this number is reached. All logging data in deleted files is lost. If you configure the *Log Push* option, you can set the system up so that the files are copied to another server before they are deleted.

Delete Files Older Than: Allows you to configure the Access Gateway to delete files when they are older than the time you specify. All logging data in deleted files is lost. If you configure the *Log Push* option, you can set the system up so that the files are copied to another server before they are deleted.

Do Not Delete: Prevents the system from automatically deleting the log files. You can use the *Log Push* option to copy the files to another server and then either manually delete them or have the *Log Push* option delete them after they are copied to another server.

For information about the *Log Push* option, see [Section 34.4.6, “Configuring Log Pushing,” on page 424](#).

4 Click *Configuration Panel*, then click *Apply Changes*.

34.4.5 Configuring Extended Log Options

Use the extended log options page to control log entry content, log roll over, and old file options. A log entry always includes the date, time, and client IP address for each entry, but with the log data options, you can add other fields such as the IP address of the server and the username of the client.

The Access Gateway does not allow active log files to be deleted. Only log files that have been closed can be deleted. The roll over options allow you to control when a file is rolled over and closed, and a new file is created. The old file options allow you to control when the rolled-over log files are deleted.

To configure an extended log file for a selected proxy service:

- 1 Click *Access Manager > Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > Logging > [Name of Extended Log Profile]*.

Log Data

Date, Time and Client IP are always provided.

Select All

<input type="checkbox"/> User Name	<input type="checkbox"/> Server IP	<input type="checkbox"/> Site Name	<input type="checkbox"/> Method	<input type="checkbox"/> URI
<input type="checkbox"/> URI Stem	<input type="checkbox"/> URI Query	<input type="checkbox"/> Version	<input type="checkbox"/> Status	<input type="checkbox"/> Bytes Sent
<input type="checkbox"/> Bytes Recieved	<input type="checkbox"/> Time Taken	<input type="checkbox"/> User Agent	<input type="checkbox"/> Cookie	<input type="checkbox"/> Referrer
<input type="checkbox"/> Cached Status	<input type="checkbox"/> Fill Proxy	<input type="checkbox"/> Origin Server	<input checked="" type="checkbox"/> X-Forward-For	

Rollover Options

Maximum File Size: MB

Roll over every Hour(s) beginning at Local

Old File Options

Maximum Number of Files:

Delete Files Older Than: Week(s)

Do Not Delete

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

OK
Cancel

- 2 Select one or more of the log data options:

Name	Description
User Name	The name of the user sending the request.
URI Stem	The stem portion of the HTTP URL the browser sent to the Access Gateway. The stem is everything in the URL up to the first question mark. If the URL has no question mark, the <i>URI Stem</i> field is the same as the <i>URI</i> field. It is redundant if <i>URI</i> is selected.
Bytes Received	The number of bytes of HTTP request data the proxy service received from the browser.

Name	Description
Cached Status	The value indicates whether the request was filled from cache. 1 = filled from cache 0 = not filled from cache
Server IP	The IP address of the Access Gateway.
URI Query	The query portion of the HTTP URL the browser sent to the Access Gateway. The query is everything from the first question mark through the end of the URL. If the URL has no question mark, this field has no value. It is redundant if URI is selected.
Time Taken	The time in seconds it took the Access Gateway resources to deal with the request.
Fill Proxy	The IP address of the upstream proxy.
Site Name	The name of the reverse proxy.
Version	The HTTP version specified in the URL the browser sent to the Access Gateway.
User Agent	The User-Agent HTTP request header value the browser sent to the Access Gateway.
Origin Server	The IP address of the Web server. This assumes the Access Gateway retrieved the requested information directly from the Web server.
Method	The HTTP method the browser sent to the Access Gateway.
Status	The HTTP status code the Access Gateway sent to the browser.
Cookie	The Cookie HTTP request header value the browser sent to the Access Gateway. The Access Gateway doesn't cache cookie information. Cookies can consume a lot of space. If you select this option, make sure it contains the critical information that you need.
X-Forward-For	The X-Forwarded-For HTTP request header value the browser sent to the Access Gateway. Do not confuse this with the X-Forwarded-For option that causes the Access Gateway to generate or forward headers to upstream proxies or Web servers.
URI	The HTTP URL the browser sent to the Access Gateway.
Bytes Sent	The number of bytes of HTTP response data the Access Gateway sent to the browser.
Referer	The Referer HTTP request header value the browser sent to the Access Gateway.

3 Select one of the following roll over options:

Maximum File Size: Rolls the file when it reaches the specified number of megabytes.

Roll over every: Rolls the file at the specified interval. You can specify the interval in minutes or hours.

- ♦ **beginning:** Specifies the day that the interval should be begin. You can select a day of the week or the first of the month.

- ◆ **at:** Select the hour of the day that the interval should begin and the time zone (either the local time zone or GMT).
- 4 Select one of the following old file options.
- Maximum Number of Files:** Allows you to limit the number of old log files on the system to the number specified in this option. The oldest file is automatically deleted when this number is reached. All logging data in deleted files is lost. If you configure the Log Push option, you can set the system up so that the files are copied to another server before they are deleted.
- Delete Files Older Than:** Allows you to configure the Access Gateway to delete files when they are older than the time you specify. All logging data in deleted files is lost. If you configure the *Log Push* option, you can set the system up so that the files are copied to another server before they are deleted.
- Do Not Delete:** Prevents the system from automatically deleting the log files. You can use the *Log Push* option to copy the files to another server and then either delete them manually or have the *Log Push* option delete them when they have been copied to another server.
- For information about the *Log Push* option, see [Section 34.4.6, “Configuring Log Pushing,” on page 424](#).
- 5 Click *Configuration Panel*, then click *Apply Changes*.

34.4.6 Configuring Log Pushing

(NetWare only) The *Log Push* option allows you to configure the NetWare Access Gateway to copy log files to an FTP server at specified intervals. The *Log Push* option is configured for all log files on the Access Gateway. If you have enabled logging on multiple proxy services, the Access Gateway uses the same configuration to push the log files of each proxy service.

This feature works within the following parameters:

- ◆ The Access Gateway tries as many times as necessary to establish a connection with the FTP server during the hour of the scheduled push. When the hour changes, the Access Gateway stops trying until the next interval you have specified.
- ◆ When the connection is established, the Access Gateway assumes that pushing the log files was successful. The Access Gateway does not detect any errors that prevent the successful pushing of the files.

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, the Access Gateway begins trying to establish a connection with the FTP server.

- ◆ If a connection cannot be established before the hour changes to 1 a.m., the Access Gateway 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 the Access Gateway doesn't try to connect again until 12 midnight the next day.

To configure log pushing:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Log Push*.

Enable Log Push

FTP Settings

DNS or IP Address: *

Default Directory: *

Login Name: *

Password: *

Rollover Options

Push Logs when the Logs Roll Over

Group Member:

Push Using Address:

Days to Push the Logs: 1st of Month Sunday Monday Tuesday
 Wednesday Thursday Friday Saturday

Time to Push the Logs:

Delete Log Files from Server after Push

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

- 2 To enable log pushing, select *Enable Log Push*.
- 3 Configure the following FTP settings. All of them are required settings.
 - DNS or IP Address:** Specify the DNS name or the IP address of your FTP server.
 - Default Directory:** Specify the directory on the FTP server to which the Access Gateway should copy the log files.
 - Login Name:** Specify the name that the Access Gateway should use to log in to the FTP server.
 - Password:** Specify the password that the Access Gateway should use for logging in.
- 4 To schedule when the log files are copied to the FTP server, fill in the following fields:
 - Push Logs when the Logs Roll Over:** Select this option to have the Access Gateway try to push a log as soon as it rolls over. This is the recommended option, because it ensures that log files are copied as soon as possible.
 - Group Member:** (Available only if the Access Gateway is a member of a group.) Select the server you want to configure from the list of servers. The modifications made to the *Push Using Address* option apply only to the selected group member. Modifications made to any other options on the page apply to all members of the group.
 - Push Using Address:** Select the IP address you want to use for sending the log files to the FTP server.

Days to Push the Logs: Specify the days when the log push should occur. You can select multiple days for pushing.

Time to Push the Logs: Specify the time of day when the log files are pushed.

- 5 Specify what you want done with the log files after they have been copied to the FTP server.

Select the *Delete Log Files from Server after Push* option to have the Access Gateway delete the log files after they have been copied to the FTP server. This is the recommended method. If you do not select this option, you must manually delete them or use the old file options on the Logging page (see [Section 34.4.4, “Configuring Common Log Options,”](#) on page 420).

- 6 Click *Configuration Panel*, then click *Apply Changes*.

34.4.7 Configuring the Size of the Log Partition

The size of the log partition should be configured as part of the installation process. See one of the following in the *Novell Access Manager 3.0 Installation Guide*:

- ♦ The NetWare Access Gateway creates a 2 GB log: volume. To increase its size, see “[Configuring the Log Partition on the NetWare Access Gateway](#)”
- ♦ Linux Access Gateway logs are stored in `/root` partition by default. You can create a `/var` partition to store the logs. The size of this partition depends on your requirements. For more information on creating the `/var` partition, see “[Customizing the Partitions](#)”.

Statistics can indicate that the system is functioning optimally or that it has some bottlenecks.

- ◆ [Section 35.1, “Monitoring Identity Server Statistics,” on page 427](#)
- ◆ [Section 35.2, “Monitoring Access Gateway Statistics,” on page 429](#)
- ◆ [Section 35.3, “Viewing SSL VPN Statistics,” on page 438](#)

35.1 Monitoring Identity Server Statistics

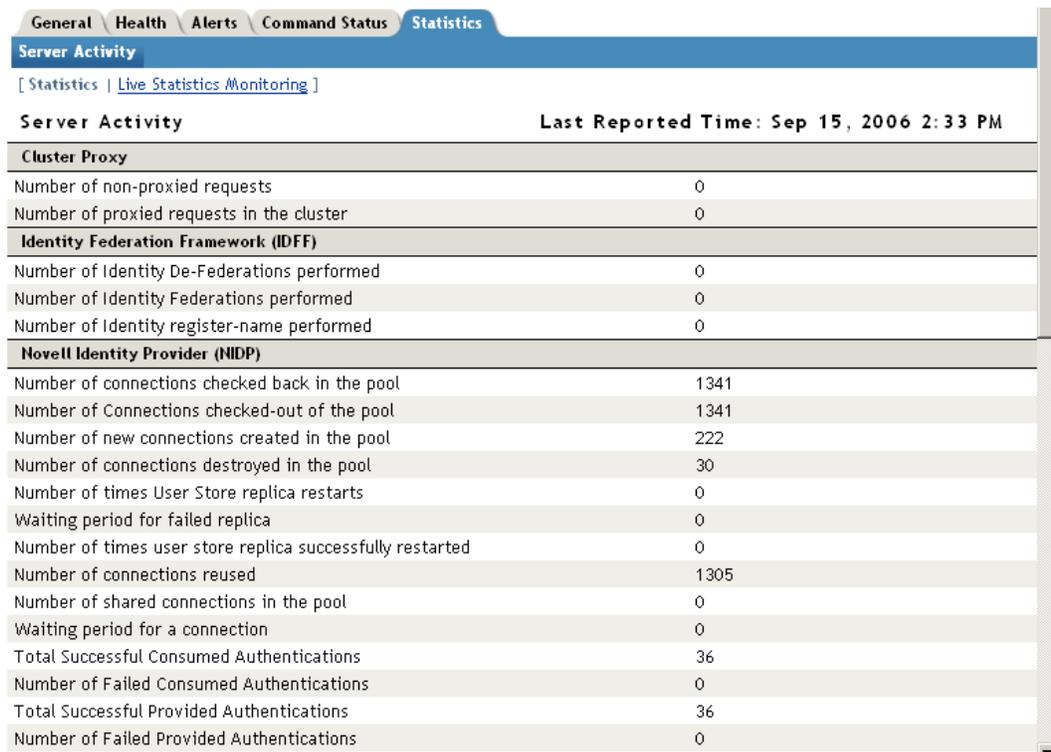
Activity for the following Identity Server components is provided:

- ◆ Cluster proxy
- ◆ IDFF (Identity federation)
- ◆ NIDP (Identity provider)
- ◆ SAML
- ◆ SAML 2.0
- ◆ WSF (Web Services Framework)

You can specify the intervals for the refresh rate and, where allowed, view graphic representations of the activity.

- 1 In the Administration Console, choose *Access Manager > Identity Servers*.

2 In the Statistics column, click *View*.



Server Activity		Last Reported Time: Sep 15, 2006 2:33 PM
Cluster Proxy		
Number of non-proxied requests		0
Number of proxied requests in the cluster		0
Identity Federation Framework (IDFF)		
Number of Identity De-Federations performed		0
Number of Identity Federations performed		0
Number of Identity register-name performed		0
Novell Identity Provider (NIDP)		
Number of connections checked back in the pool		1341
Number of Connections checked-out of the pool		1341
Number of new connections created in the pool		222
Number of connections destroyed in the pool		30
Number of times User Store replica restarts		0
Waiting period for failed replica		0
Number of times user store replica successfully restarted		0
Number of connections reused		1305
Number of shared connections in the pool		0
Waiting period for a connection		0
Total Successful Consumed Authentications		36
Number of Failed Consumed Authentications		0
Total Successful Provided Authentications		36
Number of Failed Provided Authentications		0

3 Click *Live Statistics Monitoring* if you want the display refreshed at certain time intervals.

4 Click *Close* to return to the Servers page.

35.2 Monitoring Access Gateway Statistics

The Statistics page allows you to monitor how much data and the type of data the Access Gateway is processing.

- 1 In the Administration Console, click *Access Manager > Access Gateways > [Name of Server] > Statistics*.

Server Activity	Last Reported Time: Jul 28, 2006 4:11 PM
CPU Utilization	9.0 % Graphs
Cache Hit	10.0 % Graphs
Mounted Partitions Disk Space	6.005 GB
Mounted Partitions Disk Space Used	870.2 MB
Mounted Partitions Disk Space Free	5.155 GB
Boot Partition Disk Space	0.0 Bytes
Boot Partition Disk Space Used	0.0 Bytes
Boot Partition Disk Space Free	0.0 Bytes
Swap Partition Disk Space	1.003 GB
Swap Partition Disk Space Used	0.0 Bytes
Swap Partition Disk Space Free	1.003 GB
Cache Disk Space	148817920 KB
Cache Disk Space Utilization	0.0 %
Total Installed Memory	1 MB
Start Up Time	Fri, 21 Jul 2006 16:20:57 GMT
Up Time	7 Days, 5 Hours, 49 Minutes, 40 Seconds
Number of Objects Cached	34.0

- 2 Select from the following types:
 - ♦ [Section 35.2.1, “Server Activity,” on page 429](#)
 - ♦ [Section 35.2.2, “Server Benefits,” on page 433](#)
 - ♦ [Section 35.2.3, “Server SSL Activity,” on page 434](#) (NetWare only)
 - ♦ [Section 35.2.4, “Top 20 Sites,” on page 434](#)
 - ♦ [Section 35.2.5, “Service Provider Activity,” on page 435](#)
 - ♦ [Section 35.2.6, “Configured Services,” on page 438](#) (NetWare only)
- 3 Click *Close*.

35.2.1 Server Activity

Select whether to view live or static statistics:

- ♦ **Statistics:** Select this option to view the statistics as currently gathered. The data is not updated.
- ♦ **Live Statistics Monitoring:** Select this option to view the statistics as currently gathered and to have them refreshed at the rate specified in the *Refresh Rate* field.

These general statistics are grouped into the following categories:

- ◆ “Server Activity” on page 430
- ◆ “Connections” on page 431
- ◆ “Bytes” on page 431
- ◆ “Requests” on page 432
- ◆ “Cache Freshness” on page 433

Server Activity

The Server Activity section displays general server utilization statistics.

Column	Description
CPU Utilization	Displays the current CPU utilization rate. Use the available graph for capacity planning.
Cache Hit	Displays the current cache hit rate. A high cache hit rate indicates that the caching system is off-loading significant request processing from the Web servers whose objects have been cached. Use the available graph for capacity planning.
Mounted Partitions Disk Space	(Linux only) Displays the total disk space configured for mounted partitions.
Mounted Partitions Disk Space Used	(Linux only) Displays the disk space in use on mounted partitions.
Mounted Partitions Disk Space Free	(Linux only) Displays the disk space available on mounted partitions.
Boot Partition Disk Space	(Linux only) Displays the total disk space configured for the boot partition.
Boot Partition Disk Space Used	(Linux only) Displays the disk space in use on the boot partition.
Boot Partition Disk Space Free	(Linux only) Displays the disk space available on the boot partition.
Swap Partition Disk Space	(Linux only) Displays the total disk space configured for the swap partition.
Swap Partition Disk Space Used	(Linux only) Displays the disk space in use on the swap partition.
Swap Partition Disk Space Free	(Linux only) Displays the disk space available on the swap partition.
Cache Disk Space	Displays the total disk space available for caching. The amount shown is smaller than the total disk space available on the Access Gateway because it doesn't include the disk space reserved for the operating system and for log files.
Cache Disk Space Utilization	Displays the percentage of caching disk space currently in use.
Total Installed Memory	Displays the amount of memory that is installed on the Access Gateway.

Column	Description
Start Up Time	Displays the last time the Access Gateway was started.
Up Time	Displays the total time the Access Gateway has been running since it was last started.
Number of Objects Cached	Displays the total number of Web objects that have been cached.

Connections

The connection statistics show the current and peak levels of usage in terms of TCP connections.

Column	Description
Current Connections to Origin Server	Displays the current number of connections that the Access Gateway has established with Web servers.
Current Connections to Browsers	Displays the current number of connections that the Access Gateway has established with browsers.
Current Total Connections	Displays the current total of all connections that the Access Gateway has established.
Connections to Origin Server	Displays the total number of connections that the Access Gateway has established with Web servers since it was last started.
Peak Connections from Origin server	Displays the peak number of connections that the Access Gateway has established with Web servers.
Connections to Browsers	Displays the total number of connections that the Access Gateway has established with browsers since it was last started.
Peak Connections to Browsers	Displays the peak number of connections that the Access Gateway has established with browsers.
Total Connections through SOCKS	Displays the total number of connections the Access Gateway has established through a firewall.
Failed Connection Attempts	Displays the total number of failed connection attempts the Access Gateway has made while attempting to fill its object cache.

Bytes

The bytes statistics show how fast information is being sent in response to the following types of requests:

- ◆ Browser requests to the Access Gateway
- ◆ Access Gateway requests to the Web servers

Column	Description
Bytes per second from Origin Server	Displays the number of bytes of data being sent each second from the Web servers to the Access Gateway.
Bytes per second to Browsers	Displays the number of bytes of data being sent each second from the Access Gateway to the browsers.

Column	Description
Total bytes per second	Displays the total number of bytes being sent each second, from the Access Gateway and from the Web servers.
Bytes Received from Origin Server	Displays the total number of bytes of data sent to the Access Gateway from the Web servers since the Access Gateway last started.
Bytes Sent to Browser	Displays the total number of bytes of data sent to the browsers from the Access Gateway since the Access Gateway last started.
Total Bytes	Displays the total number of bytes sent from the Access Gateway and from the Web servers since the Access Gateway last started.

Requests

The request statistics show the number of requests that are being sent from the browsers to the Access Gateway and from the Access Gateway to the Web servers.

Column	Description
Current Requests to Origin Server	Displays the current number of requests that the Access Gateway has made to the Web servers.
Current Requests from Browsers	Displays the current number of requests that the browsers have made to the Access Gateway
Total Current Requests	Displays the total number of current requests that the Access Gateway has received from the browsers and that the Access Gateway has sent to the Web servers.
Requests to Origin Server from Download	If read ahead is enabled, displays the number of requests that have generated read ahead downloads.
Successful Requests to Origin Server	Displays the total number of successful requests that the Access Gateway has sent to the Web servers since the Access Gateway last started.
Failed Requests to Origin Server	Displays the total number of failed requests that the Access Gateway has sent to the Web servers since the Access Gateway last started.
Cumulative Requests to Origin Server	Displays the total number of requests that the Access Gateway has sent to the Web servers since the Access Gateway last started.
Cumulative Requests to Browsers	Displays the total number of requests that the browsers have sent to the Access Gateway since the Access Gateway last started
Total Cumulative Requests	Displays the total number of cumulative requests that the Access Gateway has processed since the Access Gateway last started.
Requests per second from Origin Server	Displays the number of requests that are being sent each second from the Access Gateway to the Web servers.
Requests per second from Browsers	Displays the number of requests that are being sent each second from the browsers to the Access Gateway
Total Requests per sec.	Displays the total number of requests that are being sent each second from the Access Gateway and from the browsers
Peak request per second from Origin Server	Displays the peak number of requests that have been sent in one second from the Access Gateway to the Web servers.

Column	Description
Peak requests per second to Browsers	Displays the peak number of requests that have been sent in one second from the browsers to the Access Gateway.

Cache Freshness

The cache freshness statistics display information about the cache refresh process.

Column	Description
Total "Get If Modified Since" Request	Displays the total number of Get If Modified Since Requests that the Access Gateway has received from browsers.
Total Not Modified Replies	Displays the total number of 304 Not Modified replies that the Access Gateway has received from the Web servers for updated content.
Cache Freshness	Displays the percentage of objects in cache that the Access Gateway considers fresh.
Oldest Object in Memory	Displays how long the oldest cache object has been cached.

35.2.2 Server Benefits

Select whether to monitor live or static statistics:

- ◆ **Statistics:** Select this option to view the statistics as currently gathered. The data is not updated.
- ◆ **Live Statistics Monitoring:** Select this option to view the statistics as currently gathered and to have them refreshed at the rate specified in the *Refresh Rate* field.

The Server Benefits page displays information about bandwidth and DNS lookups.

Column	Description
Total Bandwidth Saved	(Linux only) Displays the amount of bandwidth saved by using data cached by the Access Gateway rather than requesting the data from the Web servers.
Bytes Saved per Second	Displays how many bytes of data the Access Gateway has able to send from cache rather than requesting it from the Web servers.
Bandwidth Saved	Displays the amount of bandwidth saved by using data cached by the Access Gateway rather than requesting the data from the Web servers.
Total DNS Lookups Saved	(Linux only) Displays the number of DNS requests that the Access Gateway could solve locally without performing a DNS lookup.
DNS "Modified Since" Queries Returning False	(Linux only) Displays the number of DNS "Modified Since" queries that the Access Gateway was able to service with a false.
Total Number of Connections Saved	Displays the number of connections that the Access Gateway has with clients minus the number of connections that the Access Gateway has with Web servers. This statistic indicates the number of connections that the Access Gateway is off loading from the Web servers.

35.2.3 Server SSL Activity

(NetWare only) Select whether to monitor live or static statistics:

- ♦ **Statistics:** Select this option to view the statistics as currently gathered. The data is not updated.
- ♦ **Live Statistics Monitoring:** Select this option to view the statistics as currently gathered and to have them refreshed at the rate specified in the *Refresh Rate* field.

The Server SSL Activity page displays information about the SSL communication process between the browsers and the Access Gateway and between the Access Gateway and the Web servers.

Column	Description
Current SSL Connection	Displays the current number of connections that are SSL connections.
Peak SSL Connections	Displays the peak number of SSL connections that the Access Gateway has established since it was last started.
Browser Full SSL Handshakes	Displays the number of browser connection requests that required a full handshake because the handshake information was not available from cache.
Browser Abbreviated Handshakes	Displays the number of browser connections requests that could perform an abbreviated handshake because the handshake information was still available from cache.
Browser SSL Alerts	Displays the number of SSL alerts that the browsers sent to the Access Gateway.
Server SSL Full Handshakes	Displays the number of Web server connection requests that required a full handshake because the handshake information was not available from cache.
Server SSL Abbreviated Handshakes	Displays the number of Web server connection requests that could perform an abbreviated handshake because the handshake information was still available from cache.
Server SSL Alerts	Displays the number of SSL alerts that the Access Gateway sent to the Web servers.

35.2.4 Top 20 Sites

Select whether to monitor live or static statistics:

- ♦ **Statistics:** Select this option to view the statistics as currently gathered. The data is not updated.
- ♦ **Live Statistics Monitoring:** Select this option to view the statistics as currently gathered and to have them refreshed at the rate specified in the *Refresh Rate* field.

The Top 20 Sites page displays the DNS name or IP address of the sites that are receiving the most traffic. These sites are sorted by the number of hits they received and by the number of bytes of data they sent.

35.2.5 Service Provider Activity

Select whether to monitor live or static statistics:

- ◆ **Statistics:** Select this option to view the statistics as currently gathered. The data is not updated.
- ◆ **Live Statistics Monitoring:** Select this option to view the statistics as currently gathered and to have them refreshed at the rate specified in the *Refresh Rate* field.

The Service Provider Activity page displays information about the communication process between the Access Gateway module (the embedded service provider) and the Identity Server. These statistics are grouped into the following categories:

- ◆ [“Cluster Proxy Statistics” on page 435](#)
- ◆ [“Identity Federation Framework \(IDFF\) Statistics” on page 435](#)
- ◆ [“Novell Identity Provider \(NIDP\) Statistics” on page 436](#)
- ◆ [“SAML Statistics” on page 437](#)
- ◆ [“SAML 2 Statistics” on page 437](#)
- ◆ [“Web Service Framework \(WSF\) Statistics” on page 437](#)

Cluster Proxy Statistics

Column	Description
Number of non-proxied requests	The total number of times the L4 switch sent the request to the server that established the session for user making the request. When this happens, the server does not need to proxy the request to a peer cluster member.
Number of proxied request in the cluster	The total number of times a cluster member has determined that it did not establish the session for the user making the request and then proxied the request to the peer cluster member that did establish the session.

Identity Federation Framework (IDFF) Statistics

Column	Description
Number of Identity De-Federations performed	The number of requests to defederated user accounts that the Identity Server has processed.
Number of Identity Federations performed	The number of requests to federated user accounts that the Identity Server has processed.
Number of Identity register-name performed	The total number of register name requests that the Identity Server has processed.

Novell Identity Provider (NIDP) Statistics

Column	Description
Number of connections checked back in the pool	The total number of times a user store connection has been checked into a connection pool after being checked out and used.
Number of Connections checked-out of the pool	The total number of times a user store connection has been checked out of a connection pool and used.
Number of new connections created in the pool	The total number of times the Identity Server has created a new connection to a user store.
Number of connections destroyed in the pool	The total number of times the Identity Server has destroyed a connection to a user store.
Number of times User Store replica restarts	When the Identity Server loses a connection to an LDAP user store, that user store is placed on a restart thread. After a period of time, the restart thread attempts to reconnect to the user store. This count is the total number of times that user stores have been placed on the restart thread.
Waiting period for failed replica	The total number of times the restart thread has failed to regain a connection with a user store and has had to wait the given time period before trying again.
Number of times user store replica successfully restarted	The total number of times the restart thread has successfully regained a connection with a user store.
Number of connections reused	The total number of times a user store connection has been reused. This means that the Identity Server was able to check out a connection from the pool and use an existing connection.
Number of shared connections in the pool	Each user store has two connection pools: a user pool and an admin pool. As connections are checked out of each of these pools, it might become apparent to the Identity Server that one pool is overworked and the other pool has connections doing nothing. When this situation is detected, a connection is shared from one pool to the other. So, the admin pool might gain a connection and the user pool might lose one. This is the total number of times that connections have been shared (over all user stores).
Waiting period for a connection	The total number of times that all connections have been checked out, and the requesting thread has waited for a connection to become available.
Total Successful Authentications	The number of successful logins that the Identity Server has processed.
Number of Failed Authentications	The total number of failed logins that the Identity Server has processed (for any reason).
Number of Logouts	The total number of logout requests that the Identity Server has processed.
% of free memory	The current percentage of system memory that Java* considers free.
Number of users currently logged-in	The number of sessions that are currently active, which equates with the number of currently logged-in users.
Total Requests	The total number of requests that have passed through the Identity Server.

SAML Statistics

Column	Description
Number of Saml requests	The total number of SAML1.1 query attribute requests that the Identity Server has processed.

SAML 2 Statistics

Column	Description
Number of Saml-2 De-Federations	The total number of SAML2 DeFederate requests that the Identity Server has processed.
Number of Saml-2 Federations	The total number of SAML2 Federate requests that the Identity Server has processed.
Number of Saml-2 requests	The total number of SAML2 Query Attribute requests that the Identity Server has processed.
Number of Saml-2 register-name	The total number of SAML2 Register Name requests that the Identity Server has processed.

Web Service Framework (WSF) Statistics

Column	Description
Number of credential-profile service 'modify'	The total number of modify requests made to the Novell® Credential Profile Web Service.
Number of credential-profile service 'query'	The total number of query requests made to the Novell Credential Profile Web Service.
Number of discovery service 'modify'	The total number of modify requests made to the Discovery Web Service.
Number of discovery service 'query'	The total number of query requests made to the Discovery Web Service.
Number of employee-profile service 'modify'	The total number of modify requests made to the Employee Profile Web Service.
Number of employee-profile service 'query'	The total number of query requests made to the Employee Profile Web Service.
Number of custom-profile service 'modify'	The total number of modify requests made to the Novell Custom Profile Web Service.
Number of custom-profile service 'query'	The total number of query requests made to the Novell Custom Profile Web Service.
Number of personal-profile service 'modify'	The total number of modify requests made to the Personal Profile Web Service.
Number of personal-profile service 'query'	The total number of query requests made to the Personal Profile Web Service.

Column	Description
Number of role-profile service 'modify'	The total number of modify requests made to the Novell Role Profile Web Service.
Number of role-profile service 'query'	The total number of query requests made to the Novell Role Profile Web Service.
Number of interaction service redirects by web services consumer (client)	The total number of times the Identity Server has been redirected to perform user interaction using the User Interaction Redirection profile.
Number of interaction service redirects to server	The total number of times the Identity Server has handled a user interaction request that it received through the User Interaction Redirection profile.
Number of interaction service redirects initiated by web services consumer (client)	The total number of times the Identity Server has called a Trusted User Interaction Service using Trusted User Interaction Service profile.
Number of interaction service redirects handled by trusted server	The total number of times the Identity Server has handled a user interaction request that it received through the Trusted User Interaction Service profile.

35.2.6 Configured Services

(NetWare only) The Configured Services page displays information about all the services that have been configured on the selected Access Gateway. It includes information about the proxy services that you have configured as well as the system proxy services (the nesp and soapbc path-based proxy services).

35.3 Viewing SSL VPN Statistics

The Statistics page allows you to view such information as the number of active client connections and the time when the SSL VPN server was started.

- 1 In Administration Console, click *Access Manager > SSL VPNs > [Server Name] > Statistics*.

The Server Statistics page is displayed.

General Health Alerts Command Status **Statistics**

Server Activity

[Statistics | [Live Statistics Monitoring](#)]

Server Activity

Server Status	
Up Time:	1 days 20 hours 37 minutes 01 seconds
Sockd status:	Sockd is running
Stunnel status:	Stunnel is running
Connections	
Active SSL VPN Connections	2
User/Role/Uptime:	US / / 0 days 00 hours 20 m
User/Role/Uptime:	AK / / 1 days 07 hours 48 m
Bytes	
Bytes Received:	5.82909e+09 Graphs
Bytes Sent:	6.36377e+09 Graphs
Received Byte Rate:	0.71
Sent Byte Rate:	2.19
Total Byte Rate:	269337.47

Close

Server Status information is gathered in the following sections:

Column	Description
Up Time	Specifies the duration for which the server has been up and running.
Sockd Status	Specifies if the sockd is running or not.
Stunnel Status	Specifies if the Stunnel is running or not.

Connection information is gathered in the following sections:

Column	Description
Active SSL VPN Connections	Specifies the number of active SSL VPN connections. Username, role of the user and uptime of each user is specified for each active connection.

Bytes information is gathered in the following sections:

Column	Description
Bytes Received	Specifies the number of bytes received. You can also view a graph, which lists the number of bytes sent for fixed intervals. For more information, see Viewing the Bytes Graphs .
Bytes Sent	Specifies the number of bytes sent. You can also view a graph, which lists the number of bytes sent for fixed intervals. For more information see Viewing the Bytes Graphs .
Received Byte Rate	Specifies the percentage of bytes received.
Sent Byte Rate	Specifies the percentage of bytes sent.
Total Byte Rate	Specifies the total percentage of bytes transferred.

2 Select one of the following options:

- ♦ **Statistics:** To display the number of active client connections and the time when the server was started, click *Statistics*.
- ♦ **Live Statistics Monitoring:** To refresh the above information for a specified interval, click *Live Statistics Monitoring*. You can select the refresh interval from the *Refresh Rate* drop-down list.

3 Click *Close* to close the Statistics tab.

35.3.1 Viewing the Bytes Graphs

The number of bytes sent and bytes received can be viewed in the form of graphs. You can view graphs for the following time frames:

- ♦ **1 Hour:** The number of bytes sent or received every ten minutes.
- ♦ **1 Day:** The number of bytes sent or received every four hours.
- ♦ **1 Week:** The number of bytes sent or received every day.
- ♦ **1 Month:** The number of bytes sent or received every week.
- ♦ **6 Months:** The number of bytes sent or received every month for six weeks.
- ♦ **12 Months:** The number of bytes sent or received every month for one year.

To view graphs:

- 1 In Access Manager select, click *Access Manager > SSL VPNs > [Server Name] > Statistics*.

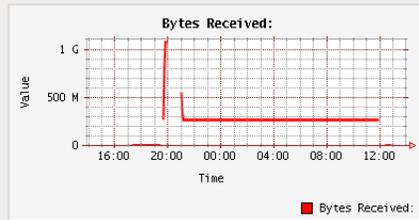
2 Select *Graphs* from either the *Bytes Received* or *Bytes Sent* section, depending on your needs.

Servers ▸ Statistics ▸ Statistics

Server Statistics: 12.12.12.124



Bytes Received:



[[1 hour](#) | [1 day](#) | [1 week](#) | [1 month](#) | [6 months](#) | [12 months](#)]

Close

3 Click *Close* to close the Graphs page.

You can monitor all of the components hosted by a server and quickly isolate and correct server issues. The system displays statuses (green, yellow, white, or red) for the Access Manager components. Health information can be accessed at the following places:

- ◆ *Access Manager > Overview*

The Overview page shows the health status at the component-level.

- ◆ *Access Manager > Auditing > Device Health*

The Device Health page shows the health status for all devices in one list.

- ◆ *Access Manager > [Component] > Servers*

The Servers page for each component provides a health status for each device.

This section discusses the following topics:

- ◆ [Section 36.1, “Health States,” on page 443](#)
- ◆ [Section 36.2, “Monitoring the Health of an Identity Server,” on page 444](#)
- ◆ [Section 36.3, “Monitoring the Health of an Access Gateway,” on page 446](#)
- ◆ [Section 36.4, “Monitoring the Health of an SSL VPN Server,” on page 448](#)

36.1 Health States

The Health page displays the current status of the server. The following states are possible:

Icon	Description
	A green status indicates that the server has not detected any problems
	A red status with a bar indicates that the server has been stopped.
	A white status with disconnected bars indicates that the server is not communicating with the Administration Console.
	A yellow status indicates that the server might be functioning sub-optimally because of configuration discrepancies.
	A yellow status with a question mark indicates that the server has not been configured.
	A red status with an x indicates that the server configuration might be incomplete or wrong, that a dependent service is not running or functional, or that the server is having a runtime problem.

36.2 Monitoring the Health of an Identity Server

To view detailed health status information for an Identity Server:

- 1 In the Administration Console, click *Access Manager > Identity Servers > [Name of Server] > Health*.

General	Health	Alerts	Command Status	Statistics
Refresh Update from Server				
Status	Description			
	Server is operational (Passed)			
Services Detail				
Type	Status	Message		
Services		Identity Server Configuration Configuration Datastore User Datastores		
Identity Server Configuration		Fully configured		
Configuration Datastore		Operating properly		
User Datastores		Operating properly		
Close				

The status icon is followed by a description that explains the significance of the current state.

- 2 To ensure that the information is current, select one of the following:
 - ♦ Click *Refresh* to refresh the page with the latest health available from the Administration Console.
 - ♦ Click *Update from Server* to send a request to the Identity Server to update its status information. This can take a few minutes.
- 3 Examine the *Services Detail* section which displays the status of each service. For an Identity Server, this includes information such as the following:

Status Category	If not healthy
Status: Indicates whether the Identity Server is online and operational.	Verify whether the Identity Server has been stopped or is not configured. Also verify that network problems are not interfering with communications between the Identity Server and the Administration Console.
Services: Indicates the general health of all configured services.	If one service is unhealthy, this category reflects that status. See the particular service that also displays an unhealthy status.

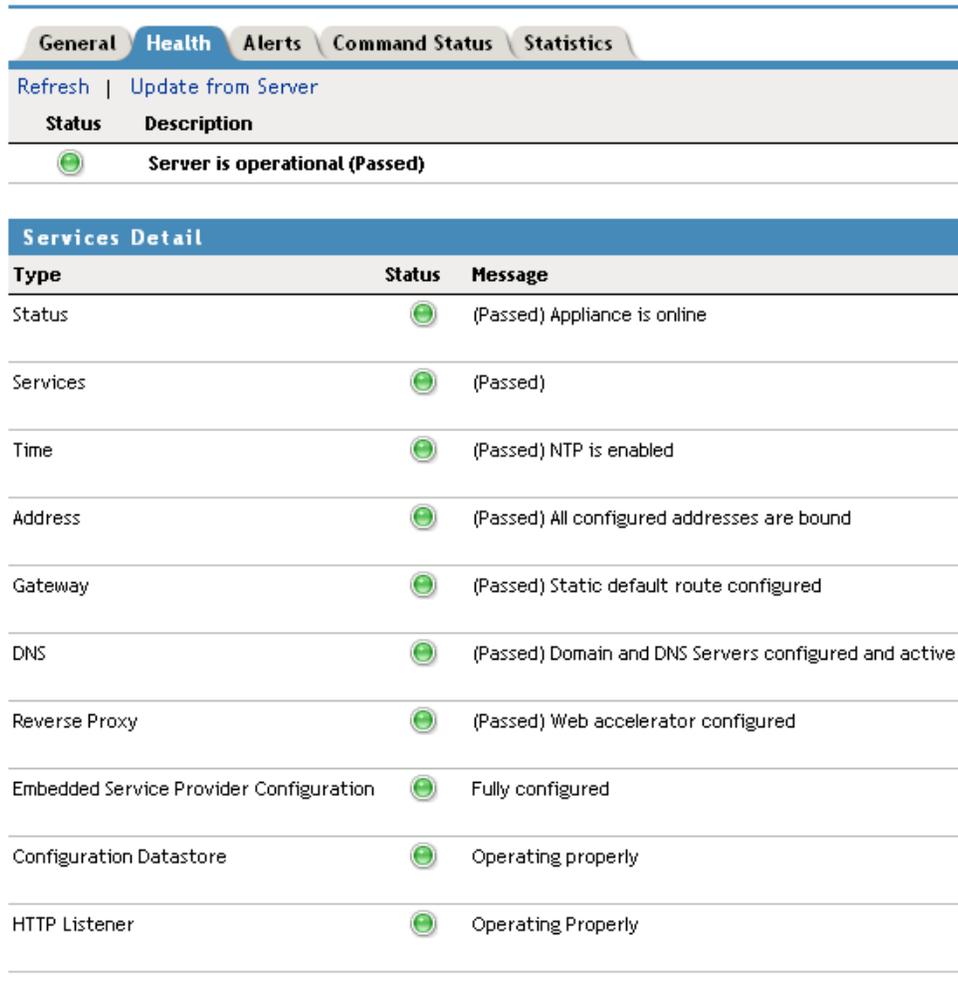
Status Category	If not healthy
Identity Server Configuration: Indicates the status of the configuration.	Configure the Identity Server or assign the server to a configuration. See Chapter 5, "Configuring an Identity Server," on page 53.
Configuration Datastore: Indicates the status of the installed configuration datastore.	<p>You might need to restart Tomcat or reinstall the Administration Console. If you have a backup Administration Console, you can restore it.</p> <p>See Section 2.2, "Backing Up and Restoring the Configuration Store," on page 31.</p> <p>If you want to convert a secondary console to your primary console, see Section 2.4, "Converting your Secondary Console into your Primary Console," on page 35.</p>
User Datastores: Specifies the user store that is configured for the server.	Ensure that the user store is operating correctly. You might need to import the SSL certificate for communication with the Identity Server. See Section 7.1, "Configuring Identity User Stores," on page 71.

4 Click *Close*.

36.3 Monitoring the Health of an Access Gateway

To view detailed health status information of an Access Gateway:

- 1 In the Administration Console, click *Access Manager > Access Gateways > [Name of Server] > Health*.



The screenshot displays the Health page for an Access Gateway. At the top, there are tabs for General, Health (selected), Alerts, Command Status, and Statistics. Below the tabs, there are links for Refresh and Update from Server. The main status is 'Server is operational (Passed)'. Below this is a 'Services Detail' table with the following data:

Type	Status	Message
Status		(Passed) Appliance is online
Services		(Passed)
Time		(Passed) NTP is enabled
Address		(Passed) All configured addresses are bound
Gateway		(Passed) Static default route configured
DNS		(Passed) Domain and DNS Servers configured and active
Reverse Proxy		(Passed) Web accelerator configured
Embedded Service Provider Configuration		Fully configured
Configuration Datastore		Operating properly
HTTP Listener		Operating Properly

The status icon is followed by a description that explains the significance of the current state.

- 2 To ensure that the information is current, select one of the following:
 - ♦ Click *Refresh* to refresh the page with the latest health available from the Administration Console.
 - ♦ Click *Update from Server* to send a request to the Access Gateway to update its status information. This can take a few minutes.
- 3 Examine the *Services Detail* section which displays the status of each service. For an Access Gateway, this includes information such as the following:

Status Category	If not healthy
Status: Indicates whether the Access Gateway is online.	<p>Check the status of the Enterprise Service Provider Configuration. If its status does not appear in the list of services, you need to start the service provider. In the Administration Console, click <i>Access Manager > Access Gateways > [Name of Server] > Actions > Start Service Provider</i>.</p> <p>Also verify that network problems are not interfering with communications between the Access Gateway and the Administration Console.</p>
Services: Indicates the general health of all configured services.	If one service is unhealthy, this category reflects that status. See the particular service that also displays an unhealthy status.
Time: Indicates the type of time configuration. Time must be configured so that it remains synchronized with the other servers in the configuration (the Identity Server, SSL VPN server, J2EE agents, Web servers, etc.).	See Section 14.3, "Setting Date and Time," on page 182
Address: Indicates whether an IP address has been configured for the reverse proxy to listen on. This is required for the Access Gateway to function.	See Section 12.1, "Creating a Reverse Proxy and Proxy Service," on page 142 .
Gateway: Specifies the type of routing that is configured for the gateway.	See Section 14.7.2, "Viewing and Modifying Gateway Settings," on page 191 .
DNS: Specifies whether a domain name server has been configured	.See Section 14.7.3, "Viewing and Modifying DNS Settings," on page 194 .
Reverse Proxy: Specifies whether a reverse proxy has been configured. An Access Gateway must have at least one reverse proxy configured.	See Section 12.1, "Creating a Reverse Proxy and Proxy Service," on page 142 .
Embedded Service Provider Configuration: Specifies whether the Access Gateway has a trusted relationship with an Identity Server. At least one Identity Server must be configured and set up as a trusted authentication source for the Access Gateway.	See Chapter 5, "Configuring an Identity Server," on page 53 for information on configuring an Identity Server. See Section 12.1, "Creating a Reverse Proxy and Proxy Service," on page 142 for information on assigning an Identity Server configuration to the Access Gateway.
Configuration Data store: Indicates whether the configuration data store is functioning correctly.	See Section 2.2, "Backing Up and Restoring the Configuration Store," on page 31 .

4 Click *Close*.

36.4 Monitoring the Health of an SSL VPN Server

You can monitor the health of an SSL VPN Server through the Health page, which displays the current status of the server.

- 1 In Administration Console, click *Access Manager > SSL VPNs > [Server Name] > Alerts*.

Servers ▸ **Health**

Server Health: 12.12.12.123

General **Health** Alerts Command Status Statistics

Refresh | Update from Server

Status	Description
	Server is operational (Passed)

Services Detail

Type	Status	Message
Socks		(Passed) Socks Server is up and running.
Stunnel		(Passed) Stunnel Server is running properly
Servlet		(Passed) Servlet is running and registered with Connection Manager.

Close

The *Status* column displays the current state, and the *Description* column explains the significance of the current state.

The *Services Details* section provides the following information:

Type: Specifies the type of service.

Status: Specifies the status of the service.

Message: Specifies a description of the status of the service.

- 2 To reload the current page with the latest status, click *Refresh*.
- 3 To send a request to the agent to update its status information, click *Update from Server*. Click *OK* in the confirmation dialog box. This can take a few minutes.
- 4 To close the Health tab, click *Close*.

Commands are issued to a device when you make configuration changes and when you select an action such as stopping or starting a device.

- ◆ [Section 37.1, “Viewing the Command Status of the Identity Server,” on page 449](#)
- ◆ [Section 37.2, “Viewing the Command Status of the Access Gateway,” on page 449](#)
- ◆ [Section 37.3, “Viewing Command Status of the SSL VPN Server,” on page 453](#)

37.1 Viewing the Command Status of the Identity Server

The Command Status page lists scheduled events and the current status of each event. A new command appears in the list each time you change a configuration. The commands remain listed until you delete them.

- 1 In the Administration Console, click *Access Manager > Identity Servers*.
- 2 Click the *Command Status* link for the server.
- 3 To delete an event, select it and click *Delete*.
- 4 Click *Refresh* to refresh the display.

The following table describes the columns on the Command Status page:

Column Name	Description
Name	Lists the Identity Server name.
Status	Lists the status of each server.
Type	Displays type of command issued to the server.
Admin	Displays the credentials of the administrator who performed the command.
Date & Time	The date and time that the command was issued. Date and time entries are specified in the local time.

37.2 Viewing the Command Status of the Access Gateway

You can view the status of the commands that have been sent to the Access Gateway for execution. The *Apply Changes* button on the configuration page issue a command, and the results appear on this page. The Actions options, such as restarting the embedded service provider or purging the cache, also appear on this page.

This section describes the following tasks related to commands:

- ◆ [Section 37.2.1, “Viewing the Status of Current Commands,” on page 450](#)
- ◆ [Section 37.2.2, “Viewing Detailed Command Information,” on page 450](#)

- ◆ [Section 37.2.3, “Issuing Commands,” on page 451](#)
- ◆ [Section 37.2.4, “Scheduling a Command,” on page 452](#)

37.2.1 Viewing the Status of Current Commands

- 1 In the Administration Console, click *Access Manager > Access Gateways > [Name of Server] > Command Status*.

General Health Alerts Command Status Statistics					
Delete Refresh					
<input type="checkbox"/>	Name	Status	Type	Admin	Date & Time (Note)
<input type="checkbox"/>	10.10.15.206 Start	EXECUTING	Service Provider Start	cn=admin,o=novell	Feb 27, 2007 3:12 PM
<input type="checkbox"/>	10.10.15.206 Stop	SUCCEEDED	Service Provider Stop	cn=admin,o=novell	Feb 27, 2007 3:12 PM

This page lists the current commands and the following information about the commands:

Column Name	Description
Name	Contains the display name of the command. Click the link to view additional details about the command. For more information, see Section 37.2.2, “Viewing Detailed Command Information,” on page 450 .
Status	Specifies the status of the command. Some of the possible states of the command include Pending, Incomplete, Executing, and Succeeded.
Type	Specifies the type of command.
Admin	Specifies if the system or a user issued the command. If a user issued the command, the DN of the user is displayed.
Date & Time	Specifies the local date and time the command was issued.

- 2 Select one of the following actions:
 - ◆ To view information about a particular command, click the name of a command.
 - ◆ To delete a command from the list, select the command, then click *Delete*.
 - ◆ To refresh the status of the listed commands, click *Refresh*.
- 3 Click *Close*.

37.2.2 Viewing Detailed Command Information

To view information about an individual command:

- 1 In Administration Console, click *Access Manager > Access Gateways > [Name of Server] > Command Status*.

- 2 Click the name of a command to get detailed information.

Note: Date and time entries are specified in local time.

Command Information

Refresh | Delete

Name: 10.10.15.206 Start
Type: Service Provider Start
Admin: cn=admin,o=novell
Status: SUCCEEDED
Last Executed On: Feb 27, 2007 3:12 PM

Command Execution Details

Command	Command Result
start	start successful

Close

To determine if any problems occurred, view the Command Execution Details section.

- 3 Select one of the following actions:
 - ♦ **Delete:** To delete a command, click *Delete*. Click *OK* in the confirmation dialog box.
 - ♦ **Refresh:** To update the current cache of recently executed commands, click *Refresh*.
- 4 Click *Close* to return to the command status page.

37.2.3 Issuing Commands

You can issue commands to a single Access Gateway, or if the Access Gateway is a member of a group, to a group of Access Gateways. For the group commands, see [Section 16.4.2, “Issuing Commands to a Group of Access Gateways,” on page 228](#).

- 1 In the Administration Console, click *Access Manager > Access Gateways > [Name of Server] > General > Actions*.
- 2 Select one of the following actions, then click either *OK* or *Cancel*.
 - ♦ **Schedule Shutdown**
 - ♦ **Schedule Restart**
 - ♦ **New NIC** (Linux only)
 - ♦ **Upgrade** (NetWare only)
 - ♦ **Purge List Now**
 - ♦ **Purge All Cache**
 - ♦ **Start Service Provider**
 - ♦ **Stop Service Provider**
 - ♦ **Restart Service Provider**

New NIC

Click this option to perform a server scan for new network interfaces. This might take some time because the server reboots.

37.2.4 Scheduling a Command

The same form is used for scheduling a shutdown, restart, or upgrade command.

Note: Date and time entries are specified in local time.

Name Scheduled Command: *

Type: Access Gateway Restart

Description:

Date & Time: 20 October 2006 at 2 pm 10 mins

To schedule a command, fill in the following fields:

Name Scheduled Command: (Required) Specifies a name for this scheduled command. This name is used in log and trace files.

Type: Displays the type of command that is being scheduled, such as *Access Gateway Shutdown*, *Access Gateway Restart*, or *Access Gateway Upgrade*.

Description: (Optional) Provides a field to describe the reason for the command.

Date & Time: The drop-down menus allow you to select the day, month, year, hour, and minute when the command should execute.

37.3 Viewing Command Status of the SSL VPN Server

Use the Command Status page to view the command status of the selected SSL VPN server.

- 1 In Administration Console, click *Access Manager > SSL VPNs > [Server Name] > Command Status*.

Servers ▸ Command Status

SSL VPNs: 12.12.12.124

General Health Alerts Command Status Statistics

Delete | Refresh

<input type="checkbox"/>	Name	Status	Type	Admin	Date & Time (Note)
<input type="checkbox"/>	12.12.12.124 Configuration	SUCCEEDED	Device Configuration	cn=admin,o=novell	Jun 19, 2006 5:34 PM
<input type="checkbox"/>	12.12.12.124 Configuration	SUCCEEDED	Device Configuration	cn=admin,o=novell	Jun 19, 2006 5:19 PM
<input type="checkbox"/>	12.12.12.124 Configuration	SUCCEEDED	Device Configuration	cn=admin,o=novell	Jun 19, 2006 4:26 PM
<input type="checkbox"/>	12.12.12.124 Configuration	SUCCEEDED	Device Configuration	cn=admin,o=novell	Jun 19, 2006 3:43 PM
<input type="checkbox"/>	12.12.12.124 Configuration	SUCCEEDED	Device Configuration	cn=admin,o=novell	Jun 19, 2006 3:42 PM
<input type="checkbox"/>	12.12.12.124 Configuration	SUCCEEDED	Device Configuration	cn=admin,o=novell	Jun 19, 2006 3:41 PM
<input type="checkbox"/>	12.12.12.124 Start	SUCCEEDED	SSL VPN Start	cn=admin,o=novell	Jun 19, 2006 3:40 PM
<input type="checkbox"/>	12.12.12.124 Configuration	SUCCEEDED	Device Configuration	cn=admin,o=novell	Jun 19, 2006 3:40 PM
<input type="checkbox"/>	12.12.12.124 Start	SUCCEEDED	SSL VPN Start	cn=admin,o=novell	Jun 19, 2006 3:38 PM
<input type="checkbox"/>	12.12.12.124 Configuration	EXECUTING	Device Configuration	cn=admin,o=novell	Jun 19, 2006 3:28 PM

This page lists the command and the following information about the command:

- ♦ **Name:** Contains the display name of the command. Click the link to view additional details about the command. For more information, see [Section 37.3.1, “Viewing Command Information,” on page 453](#).
 - ♦ **Status:** Specifies the status of the command. Some of the possible states of the command include Pending, Incomplete, Executing, and Succeeded.
 - ♦ **Type:** Specifies the type of command.
 - ♦ **Admin:** Specifies if the system or a user issued the command. If a user issued the command, the DN of the user is displayed.
 - ♦ **Date & Time:** Specifies the local date and time the command was issued.
- 2 To delete a command, select the check box for the command, then click *Delete*. The selected command is cleared.
 - 3 To update the current cache of recently executed commands, click *Refresh*.
 - 4 Click *Close* to close the Command Status tab.

37.3.1 Viewing Command Information

To view configuration of individual commands:

- 1 In Administration Console, click *Access Manager > SSL VPNs > [Server Name] > Command Status > [Individual Command]*. The command status page is displayed.

- 2 Click the command to get a detailed information on the command. The Server Configuration scheduled command page is displayed.

[Servers](#) ▶ **Server Scheduled Command**

Server Details Edit: Server Configuration Scheduled Command

Note: Date and time entries are specified in local time.

Command Information	
Delete Refresh	
Name:	12.12.12.124 Configuration
Type:	Device Configuration
Admin:	cn=admin,o=novell
Description:	12.12.12.124 Configuration
Status:	SUCCEEDED
Last Executed On:	Jun 19, 2006 5:34 PM
Aggregate Command Result:	Success
Command Execution Details	
Command	Command Result
<input type="button" value="Cancel"/>	

You can perform the following actions:

- ♦ **Delete:** To delete a command, click *Delete*. Click *OK* in the confirmation dialog box.
 - ♦ **Refresh:** To update the current cache of recently executed commands, click *Refresh*.
- 3 Click *Close* to return to the command status page.

- ♦ [Section 38.1, “Monitoring Identity Server Alerts,” on page 455](#)
- ♦ [Section 38.2, “Monitoring Access Gateway Alerts,” on page 455](#)

38.1 Monitoring Identity Server Alerts

The Alerts page allows you to view information about current Java alerts and to clear them. An alert is generated whenever the Identity Server detects a condition that prevents it from performing normal system services.

- 1 In the Administration Console, click *Access Manager > Identity Servers > [Name of Server] > Alerts* tab.
- 2 To acknowledge an alert, select the check box for the alert, then click *Acknowledge Alert(s)*. When you acknowledge an alert, you clear the alert from the list.
- 3 Click *Close*.
- 4 (Optional) To verify that the problem has been solved, *Identity Servers > [Name of Server] > Health > Update from Server*.

38.2 Monitoring Access Gateway Alerts

The Access Gateway has been programmed to issue events to various types of systems (such as a Novell® Audit server or a Syslog server) so that the administrator can be informed when significant changes occur that modify how the Access Gateway is performing. For information about auditing and audit events, see [Chapter 33, “Enabling Auditing,” on page 401](#). This section describes how to use the following types of alerts:

- ♦ [Section 38.2.1, “Reviewing Java Alerts,” on page 455](#)
- ♦ [Section 38.2.2, “Configuring Access Gateway Alerts,” on page 456](#)
- ♦ [Section 38.2.3, “Enabling SNMP,” on page 461](#)

38.2.1 Reviewing Java Alerts

The Alerts page allows you to view information about current Java alerts and to clear them. An alert is generated whenever the Access Gateway detects a condition that prevents it from performing normal system services.

- 1 In the Administration Console, click *Access Manager > Access Gateways > [Name of Server] > Alerts* tab.

Severity	Date & Time	Message
<input type="checkbox"/> Severe	Oct 16, 2006 4:21 PM	Access Gateway Embedded Service Provider failed to initialize after 300 seconds.

- 2 To acknowledge an alert, select the check box for the alert, then click *Acknowledge Alert(s)*. When you acknowledge an alert, you clear the alert from the list.
- 3 Click *Close*.
- 4 (Optional) To verify that the problem has been solved, *Access Gateways > [Server Name] > Health > Update from Server*.

The NetWare Access Gateway currently sends the following severe alerts when it is not functioning correctly:

Alert Message	Solution
Access Gateway Embedded Service Provider failed to initialize	Click <i>Access Gateways > [Name of Access Gateway] > General > Actions > Start Service Provider</i> .
Access Gateway Server communication channel failed to start	

38.2.2 Configuring Access Gateway Alerts

The configuration steps for Access Gateway Alerts are platform-specific, although both platforms support similar options. To set up notification for these types of alerts, see the following sections:

- ♦ [“NetWare Access Gateway Alerts” on page 457](#)
- ♦ [“Linux Access Gateway Alerts” on page 459](#)

NetWare Access Gateway Alerts

For a NetWare[®] Access Gateway, the *Legacy Alerts* option allows you to send notification of generated system alerts to a Syslog server, to a list of e-mail recipients, or to both.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Legacy Alerts*.

Syslog

Enable Syslog

Port: * 514

Identifier: *

Syslog Server List

New... | Delete

Syslog Server

No items

Email

Enable Email

Email Server List

New... | Delete

Email Server

No items

Email Address List

New... | Delete

Email Address

No items

- 2 Enable the Syslog services by configuring the following fields:

Enable Syslog: Selecting this option enables syslog alerts. You must also configure a Syslog server and select some alerts. (See [Step 3](#) and [Step 5](#).)

Port: Specifies the port where the Syslog server listens for Syslog messages. The default value for the UDP port is 514. Make sure this port value matches the port configuration of your Syslog server.

Identifier: Specifies a string that identifies the Access Gateway as the generator of the alert.

- 3 If you enabled Syslog services, configure a Syslog server.

3a Click *New* under the *Syslog Server List*.

3b Specify the DNS name or IP address of the Syslog server and click *OK*.

- 4 To enable e-mail notification, select *Enable Email*.

You must also set up an e-mail server and a list of recipients and select some alerts before any alert notifications are sent

4a Click *New* under the *Email Server List* section and specify the DNS name or IP address of your e-mail server.

Repeat this step if you have more than one e-mail server.

- 4b** Click *Email Address* to activate all servers in the list, or click the box by individual servers to select only some servers in the list.
 - 4c** Click *New* under the *Email Address List* section and specify the e-mail address of the user you want to receive alert notifications.
Repeat this step to add others to the list.
 - 4d** Click *Email Address* to activate notification for all users in the list, or click the box by individual users to select only some users in the list.
- 5** Select the alerts for notification.

Alerts

Select All

<input type="checkbox"/> Disk Space Shortage	<input type="checkbox"/> ECB Shortage	<input type="checkbox"/> Ping Flooding
<input type="checkbox"/> TCP Synchronization Flooding	<input type="checkbox"/> UDP Flooding	<input type="checkbox"/> Login Failure

Note: The settings below are applied globally.

<input type="checkbox"/> System Up	<input type="checkbox"/> System Down	<input type="checkbox"/> Configuration Change
------------------------------------	--------------------------------------	---

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

OK
Cancel

Select All: Select this option for all alerts. Otherwise, select one or more of the following:

Alert	Description
Disk Space Shortage	Generated when disk space is low on the OS (sys:) or Log (log:) volumes.
TCP Synchronization Flooding	Generated when TCP/IP detects a flooding of synchronization packets. This often happens during a denial-of-service attack.
ECB Shortage	Generated when network receive buffers are low.
UDP Flooding	Generated when TCP/IP detects a flooding of UDP packets. This often happens during a denial-of-service attack.
Ping Flooding	Generated when TCP/IP detects a flooding of ping packets. This often happens during a denial-of-service attack.
Login Failure	Generated each time a login failure occurs from the management tool or from FTP. The alert contains the IP address of the client making the unsuccessful attempt.
System Up	Generated each time the Access Gateway is started.
System Down	Generated each time the Access Gateway is stopped.
Configuration Change	Generated each time the configuration of the Access Gateway is modified.

- 6** To save your modifications, click *OK*, then on the Configuration page, click *Apply Changes*.

Linux Access Gateway Alerts

For a Linux Access Gateway, this option allows you to send notification of generated system alerts to a Syslog server, to SNMP, to a system controller, to a log file, or to a list of e-mail recipients.

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > Alerts*.

Servers ▶ Configuration ▶ Alerts

Alert Profiles: 10.10.171.140

New... | Enable Disable Delete

Profile Enabled

default Yes

Changes made on this panel must be applied or scheduled from the [Configuration Panel](#).

OK Cancel

- 2 To add a new profile, click *New*.
- 3 Specify a name for the profile, then click *OK*.
- 4 Click the new profile to configure alert events.

Alert Events

Enable All Disable All 12 item(s)

Enabled	Event	Enabled	Event
<input type="checkbox"/>	Connection Refused	<input type="checkbox"/>	Login Failed
<input type="checkbox"/>	Proxy Initialization Failure	<input type="checkbox"/>	DNS Server Not Responding
<input checked="" type="checkbox"/>	System Up	<input type="checkbox"/>	DNS Server is Now Responding
<input checked="" type="checkbox"/>	System Down	<input type="checkbox"/>	DNS Parent Address Invalid
<input checked="" type="checkbox"/>	Configuration Changed	<input type="checkbox"/>	DNS Resolver Initialization Failure (10 Seconds)
<input type="checkbox"/>	Login Succeeded	<input type="checkbox"/>	DNS Resolver Initialization Failure (2 minutes)

Alert Actions

Enable All Disable All

Enabled	Action
<input checked="" type="checkbox"/>	Send to System Controller
<input type="checkbox"/>	Send to SNMP

Send to Log File

New Enable All Disable All

Enabled	Action
No items	

Send Email Notifications

New Enable All Disable All

Enabled	Action
No items	

Send to Syslog

New Enable All Disable All

Enabled	Action
No items	

- 5 To select the alerts for notification, select one or more of the following:

Alert	Description
Connection Refused	Generated when the connection is refused.
Proxy Initialization Failure	Generated when the embedded service provider fails to initialize.
System Up	Generated each time the Access Gateway is started.
System Down	Generated each time the Access Gateway is stopped.
Configuration Changed	Generated each time the configuration of the Access Gateway is modified.
Login Succeeded	Generated each time login succeeds.
Login Failed	Generated each time login fails.
DNS Server Not Responding	Generated each time the DNS server fails to respond.
DNS Server Is Now Responding	Generated each time the DNS server comes up.
DNS Parent Address Invalid	Generated when the IP address of DNS parent is invalid.
DNS Resolver Initialization Failure (10 seconds)	Generated when the DNS resolver initialization fails.
DNS Resolver Initialization Failure (2 minutes)	Generated when the DNS resolver initialization fails.

- 6** To send alerts to all destinations, click *Enable All*. Otherwise, select the action for each destination.
- 7** To send alerts to the Administration Console select the *Send to System Controller* check box.
- 8** To send alerts to an SNMP server, click the *Send to SNMP* link.
 - 8a** Specify the IP address of the SNMP server in the *IP address* field, then click *Insert*. You can add multiple IP addresses similarly.
 - 8b** To delete an IP address, click the *Delete* button next to the IP address that you want to delete. Click *OK* in the confirmation dialog box.
- 9** To send alerts to a log file, click *New*, then specify a name for the log profile.
 - 9a** Click the newly added profile to configure the following Log File details:
 - ♦ **Send to Log File:** This field displays the log profile name entered in the previous dialog box.
 - ♦ **Log File Name:** Specify a name for the log file and path where the file would be stored.
 - ♦ **Max File Size:** Specify a maximum size in KB for the log file. The size can be in the range of 50 to 100000 KB. If this field is left blank, it indicates that the maximum file size is 100000 KB. Specify 0 to indicate that there is no maximum file size.
 - 9b** Click *OK*.
- 10** To enable e-mail notification click *New*, then specify a name for the e-mail profile.

10a Click the newly added profile to configure the following e-mail details:

- ♦ **Send E-Mail Notifications:** This field displays the e-mail profile name entered in the previous dialog box.
- ♦ **E-mail Recipients:** Specify the e-mail address of the recipient, then click *Insert*. You can add multiple e-mail addresses. Click *Delete* to delete any of the e-mail addresses, then click *OK* at the confirmation dialog box.
- ♦ **Mail Exchange Servers:** Specify the IP address or the DNS name of the mail exchange server. Click *Delete* to delete any of the mail exchange servers addresses, then click *OK* at the confirmation dialog box.

10b Click *OK*.

11 To enable Syslog alerts click *New*, then specify a name for the Syslog profile.

11a Click the newly added profile to configure the following Syslog details:

- ♦ **Send to Syslog:** This field displays the profile name entered in the previous dialog box.
- ♦ **Facility Name:** Specify a facility name for the Syslog server. It can be any name between local0 to local7. If you specify local0 as your facility name, the alerts are stored at `\var\logs\ics_dyn.log`. The Linux Access Gateway uses local0 for normal logging information. Therefore, it is not recommended to specify local0 as your facility name.

11b Click *OK*.

11c To delete a syslog profile, click *Delete*. Click *OK* in the confirmation dialog box.

12 To delete an Alert Profile, select the profile, then click *Delete*. Click *OK* in the confirmation dialog box.

13 To save your modifications, click *OK*, then on the Configuration page, click *Apply Changes*.

38.2.3 Enabling SNMP

The SNMP page allows you configure the Access Gateway with basic SNMP information so the Access Gateway can communicate with your SNMP management workstations.

This 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) (<http://www.faqs.org/rfcs/rfc1067.html>).

When SNMP-enabled components of Access Gateway 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.

The Access Gateway has 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 on the SNMP page, all alerts you select in the Legacy Alerts page (see “[NetWare Access Gateway Alerts](#)” on [page 457](#)) are automatically sent as SNMP traps even if you have not configured syslog or e-mail alert notification on the Legacy Alerts page.

To set up SNMP:

- 1 In the Administration Console, click *Access Manager > Access Gateways > Edit > SNMP*.

Monitor State

No Community May Read

Specified Community May Read:

Control State

No Community May Write

Specified Community May Write:

Trap State

Do Not Send Traps

Trap Community Name:

Node Name for SNMP:

SNMP Management Server IP Addresses

[New...](#) | [Delete](#)

IP Address

No items

Appliance Information

Hardware	Location	Contact
<input type="text" value="null"/>	<input type="text" value="null"/>	<input type="text" value="iChain"/>

Changes made on this panel must be applied or scheduled from the [Configuration](#) Panel.

- 2 Configure the following:

Monitor State: Specifies whether the community has Read access to monitor the Access Gateway. If it does, you need to specify the community name. Community names must contain only ASCII characters and must not have spaces.

Control State: Specifies whether the community has Write access to the control states of the Access Gateway. If it does, you need to specify the community name. Community names must contain only ASCII characters and must not have spaces.

Trap State: Specifies whether traps are sent. If they are sent, you can 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. You can also specify a *Node Name for SNMP* for management of the Access Gateway through SNMP.

- 3 Add an SNMP server.

- 3a In the *SNMP Management Server IP Addresses* section, click *New*.

- 3b Specify the IP address of the SNMP server, then click *OK*.

3c Repeat to add additional servers.

4 (Optional) Configure appliance information.

The *Appliance Information* fields allow you to enter additional information about the Access Gateway. You can describe the Access Gateway hardware and its location, and provide the name of the person responsible for the Access Gateway.

5 To save your modifications, click *OK*, then on the Configuration page, click *Apply Changes*.

Troubleshooting

VIII

The following section contains information about troubleshooting the components of Access Manager:

- ◆ [Chapter 39, “Troubleshooting the Administration Console,” on page 467](#)
- ◆ [Chapter 40, “Troubleshooting for the Identity Server and Authentication,” on page 469](#)
- ◆ [Chapter 41, “Troubleshooting Access Manager Policies,” on page 473](#)
- ◆ [Chapter 42, “Troubleshooting the Access Gateway,” on page 491](#)
- ◆ [Chapter 43, “Troubleshooting SSL VPN,” on page 511](#)

Troubleshooting the Administration Console

39

This section discusses general troubleshooting issues found in the Administration Console:

- ♦ [Section 39.1, “Session Conflicts,” on page 467](#)
- ♦ [Section 39.2, “Unable to Login to the Administration Console,” on page 467](#)
- ♦ [Section 39.3, “Exception Processing IdentityService_ServerPage.JSP,” on page 467](#)
- ♦ [Section 39.4, “Logging,” on page 467](#)

39.1 Session Conflicts

Do not use two instances of the same browser to simultaneously access the same Administration Console. Browser sessions share settings which can result in problems when you apply changes to configuration settings. However, if you are using two different brands of browsers simultaneously, such as Internet Explorer and Firefox, it is possible to avoid the session conflicts.

39.2 Unable to Login to the Administration Console

If you experience problems logging in to the Administration Console, you might need to restart Tomcat. In a terminal window on the console machine, enter the following command:

```
/etc/init.d/novell-tomcat restart
```

39.3 Exception Processing IdentityService_ServerPage.JSP

If you see the message `Exception processing IdentityService_ServerPage.jsp`, it is an indication that the system has run out of available file handles. You need to use the command line to increase the `ulimit` value (`ulimit -n [new limit]`), which sets the number of open file descriptors allowed.

To set this value permanently, you can create the file `/etc/profile.local` with the `ulimit` value, such as:

```
ulimit -n 4096
```

You can make changes to `/etc/security/limits.conf` with a line just to change the limit for a specific user, in this case the `novlwww` user. You would do this by adding the following line:

```
novlwww soft nofile [new limit]
```

39.4 Logging

You can troubleshoot by accessing the server logging page.

In the Administration Console, click *Identity Server > Setup > [Configuration] > Logging*.

See [Section 34.2, “Configuring Component Logging,”](#) on page 410.

Troubleshooting for the Identity Server and Authentication

40

This section discusses the following topics:

- ◆ [Section 40.1, “Authentication,” on page 469](#)
- ◆ [Section 40.2, “Translating the Identity Server Configuration Port,” on page 470](#)
- ◆ [Section 40.3, “Problems Reading Keystores after Identity Server Re-installation,” on page 471](#)

Identity Server logging information can be found in [Appendix D, “Logging: Using the Custom Content Filter,” on page 529](#).

40.1 Authentication

This section discusses the following issues that occur during authentication:

- ◆ [Section 40.1.1, “Browser Hangs in an Authentication Redirect,” on page 469](#)
- ◆ [Section 40.1.2, “Prompting Users to Refresh a Session before Expiration,” on page 469](#)

40.1.1 Browser Hangs in an Authentication Redirect

If the browser hangs when the user attempts to authenticate at an identity provider, determine whether a new authentication contract was created and set as the default contract on the Identity Server. If true, and you have an Access Gateway resource set to accept any contract from the identity provider, you should navigate to the *Overview* tab for the protected resource and specify *Any* again in the *Contract* drop-down menu. Then click *Apply Changes* on the Configuration page.

40.1.2 Prompting Users to Refresh a Session before Expiration

This section describes how you prompt users to refresh a session before it expires by enabling a pop-up box that warns users about session timeouts.

Use the following JavaScript code in a custom login page. The script checks whether the session is approaching the idle session timeout parameter and, if so, prompts the user to extend the session.

The following script should be inserted in the script portion of your *head* tag. Also, call the `timeoutClock` function in your body onload (`<body onload=timeoutClock();>`).

```
var x = 600; //20 minutes
var timerID = null;

function timeoutClock()
{
timerID = setTimeout(timeoutClock(), 1000); //run every second
if(x==30)
{
newwindow = window.open(path/to/
timeout_page.html,toWindow,toolbar=no,
```

```

menubar=no, resizable=no, scrollbars=no, status=no, location=no, width=300,
height=200);
    }
    if (x==0)
    {
    window.location.href = path/to/time_expired.html;
    }
    x=x-1;
    return;}

// Function to reset the timeout
function resetClock()
{
    clearTimeout(timerID);
    x = 600;
    timeoutClock();
    return;
}

```

Then you create the expired page (such as `time_expired.html`) and the timeout page (such as `timeout_page.html`) with the messaging you choose. Be sure to have a link on the timeout page similar to:

```

Click <a href=javascript:window.opener.location.href
=window.opener.location.href;window.close();>[here]</a> to refresh
now.

```

Also, you might want to have the timeout page close itself after 29 seconds by calling this `closeMe` function in the body onload:

```

var howLong = 29000;
t = null;
function closeMe()
{
t = setTimeout(self.close(), howLong);
}

```

See [Section 7.6, “Creating Custom Login Pages,” on page 85](#) for information about custom login pages.

40.2 Translating the Identity Server Configuration Port

If your Identity Server must communicate with an external Identity Server through a firewall, you must either set up a hole in your firewall for TCP ports 8080 or 8443 (default ports used respectively for non secure and secure communication with Identity Server), or configure the Identity Server service to use TCP port 80 or 443.

The Identity Server service (hosted on Tomcat) runs as a non-privileged user and cannot therefore bind to ports below 1024. In order to allow requests to port 80/443 while Tomcat is listening on 8080/8443, the preferred approach is to use iptables to perform a port translation. Assuming HTTPS

on port 443 is used, perform the following procedure. Similar steps apply to using HTTP on port 80 if a non secure channel is required.

1 In the Administration Console, click *Identity Server > Setup > [Configuration]*, and configure the base URL with HTTPS as protocol, and the TCP Port as 443.

2 Log in to SLES 9 as the root user.

3 Create a file to hold the iptables rule and place it in /etc/init.d.

For example, /etc/init.d/Redirect. Ensure it has execute rights. You can use CHMOD as appropriate.

4 Edit the file by adding the iptables rule to perform the translation, as follows:

```
iptables -t nat -A PREROUTING -i eth0 -p tcp --dport 443 -j DNAT
--to 10.0.0.0:8443
```

0.0.0.0 indicates the Identity Server IP. The iptables RPM is installed with SLES 9.

5 Ensure that the iptables rule is active after rebooting.

In Yast2 > System > Runlevel Editor, select the Redirect service and enable it.

40.3 Problems Reading Keystores after Identity Server Re-installation

This can occur if you replace a hard drive and incorrectly reinstall the Identity Server. See “[Reinstalling an Identity Sever onto a New Hard Drive](#)” in the *Novell Access Manager 3.0 Installation Guide* for the correct procedure.

Troubleshooting Access Manager Policies

41

This section discusses troubleshooting topics for the following:

- ◆ [Section 41.1, “Roles and Authorization Policies,” on page 473](#)
- ◆ [Section 41.2, “Turning on Logging for Policy Evaluation,” on page 474](#)
- ◆ [Section 41.3, “Understanding Policy Evaluation Traces,” on page 475](#)
- ◆ [Section 41.4, “Policy Page Timeout,” on page 483](#)
- ◆ [Section 41.5, “Policy Changes Are Not Saved,” on page 483](#)
- ◆ [Section 41.6, “Policy Creation and Storage,” on page 483](#)
- ◆ [Section 41.7, “Policy Distribution,” on page 484](#)
- ◆ [Section 41.8, “Policy Evaluation: Access Gateway Devices,” on page 485](#)

41.1 Roles and Authorization Policies

If you are using roles in your authorization policies, you need to make sure that the role is enabled for the Identity Server configuration. You can create roles and authorization policies independently of assigning them to protect a resource or to an Identity Server configuration. So when you assign an authorization policy to protect a resource, make sure you have enabled the role for the Identity Server configuration. (Click *Access Manager > Identity Servers > Setup > [Configuration] > Roles.*)

If you haven't enabled the role, users are not assigned the role when they log in, even when they meet all the criteria for the role.

- ◆ If the Authorization Policy is an Allow policy, the users might be denied access because they haven't been assigned the role.
- ◆ If the Authorization Policy is a Deny policy, the users might be allowed access because they haven't been assigned the role.

Whenever an Authorization Policy is not producing the expected results and the policy contains a role, the first troubleshooting step should always be to check whether the role has been enabled for the Identity Server configuration.

If the first step does not solve the problem, the second step should be to turn on logging of policy events and then trace the policy evaluation. See the following:

- ◆ [Section 41.2, “Turning on Logging for Policy Evaluation,” on page 474](#)
- ◆ [Section 41.3, “Understanding Policy Evaluation Traces,” on page 475](#)

If you have enabled logging, the trace for determining whether a user has been assigned a role is in the Identity Server log file.

41.2 Turning on Logging for Policy Evaluation

Policy evaluation for Roles occurs at the Identity Server. For Authorization, Form Fill, and Identity Injection policies, policy evaluation occurs on the Embedded Service Provider where the policy is enabled. This can be an Access Gateway, or for Authorization policies, this can be a J2EE Agent. Logging for policy evaluation is controlled by the log settings of the Identity Server configuration.

To enable logging for policy evaluation:

- 1 Click *Access Manager > Identity Servers > Setup > [Configuration] > Logging*.

If you have set up more than one Identity Server configuration, make sure you select the configuration to which the other Access Manager components have been assigned.

- 2 Select *Enabled* for *File Logging*.

- 3 Select to echo the trace messages to the console.

- ♦ For a Linux Access Gateway, this sends the messages to the `catalina.out` file.
- ♦ For a NetWare Access Gateway, this sends the messages to the NetWare console.
- ♦ For the Linux Identity Server, this sends the messages to the `catalina.out` file.

- 4 (Optional) Specify a path for the Identity Server log files.

If you have a mixed platform environment (for example, the Identity Server is installed on Linux and the Access Gateway is on NetWare), do not specify a path.

- 5 For policy evaluation tracing, set the *Application* level to *info* in the *Component File Logger Levels* section.

If you are only troubleshooting policies at this time, do not select any other options. This reduces the amount of information recorded in the log files.

To see the policy SOAP messages, you need to set the *Application* level to *config*.

- 6 Update the Identity Server.

- 7 Click *Auditing > General Logging*.

- ♦ For role evaluation traces, view the Identity Server `catalina.out` file.

If your Identity Servers are clustered, you need to look at the file from each Identity Server.

- ♦ For authorization, form fill, and identity injection evaluation traces, view the log file of the Embedded Service Provider of the device that is protecting the resource.
 - ♦ For a Linux Access Gateway, this is the `catalina.out` file of the Access Gateway where the protected resource is defined. If the Linux Access Gateway is part of a group, you need to look at this file from each Access Gateway in the group. The actual ESP log file is not displayed in the list. To view this file which contains only ESP log messages, see the `nidp.*.xml` files in the `/var/ops/novell/tomcat4/logs` directory (or the directory you specified in [Step 4](#)). Depending upon how you have configured File Wrap, the `*` portion of the filename contains the month, the week, the day, and the hour.
 - ♦ For a NetWare Access Gateway, the file is not displayed in the list. To view the trace messages, you need to go to the system console or view the `nipd.*.xml` file in the `sys:\tomcat\4\webapps\nesp\WEB-INF\logs` directory. Depending upon how you have configured File Wrap, the `*` portion of the filename contains the month, the week, the day, and the hour.

To view the nipd.*.xml file, you need to enable FTP or SSH and copy the file.

- ♦ For a J2EE Agent, see “[Viewing Log Files](#)” in the *Novell Access Manager 3.0 J2EE Agent Guide*.

8 To understand what you are looking for in the log file, continue with one of the following:

- ♦ [Section 41.3, “Understanding Policy Evaluation Traces,” on page 475](#) if you set *Application* level to *info*.
- ♦ [Section 41.8, “Policy Evaluation: Access Gateway Devices,” on page 485](#) if you set *Application* level to *config*.

41.3 Understanding Policy Evaluation Traces

- ♦ [Section 41.3.1, “Format,” on page 475](#)
- ♦ [Section 41.3.2, “Policy Result Values,” on page 480](#)
- ♦ [Section 41.3.3, “Sample Policies,” on page 481](#)

41.3.1 Format

All policy trace messages start with the following elements:

```
<Address of Information Context>~<TimeStamp>~<Type>~<InterfaceId>~
```

Element	Description
<code><Address of Information Context></code>	A hexadecimal numeric string. In the sample policy, this number is 1af33d6.
<code><TimeStamp></code>	The number of seconds since 1970, expressed as day of week, month, day, hour, minutes, seconds, time zone, year.
<code><Type></code>	A two letter string indicating the type of trace string. The following types are supported: <ul style="list-style-type: none">♦ CO - Condition Evaluation Result♦ CS - Condition Set Evaluation Result♦ PA - Policy Action Initiation♦ PC - Policy Action Completion♦ RU - Rule Evaluation Result♦ RL - Rule List Evaluation Result
<code><InterfaceID></code>	The unique ID of the policy, rule, condition, or action. This identifier is generated by the Administration Console when the policy is created.

Elements are separated from each other with the tilde (~) character. If an element does not have a value, no value is inserted, which results in two or more tildes between values. Two tildes means one element didn't have a value, three tildes means that two elements didn't have values, and so forth.

Condition Evaluation Result

After the initial four fields, a CO trace has the following fields:

~<LHSOperand>~<Operator>~<RHS>~<Values>~<NOT>~<Result>[~<ResultOnError>]

A CO trace looks similar to the following:

```
com.novell.nxpe.NxpePolicyEvaluation$InformationContext@1af33d6~Thu
Jan 25 09:20:30 MST 2007~CO~0~1004:no-param~date-equal~1004:no-
param~~~69
```

Element	Description
<LHSOperand>	<p>The enumerative value and parameter list of the left-hand operand. This data can be displayed in one of two formats:</p> <p>Embedded: <LHS parameter> or <User Data ID>: <LHS parameter></p> <p>Embedded indicates that the information was entered by the administrator when the condition was configured.</p> <p><User Data ID> is a numerical value assigned to user data that might be sensitive.</p> <p><LHS parameter> is the string <code>no-param</code> if no parameters are specified for this operand or the value of the parameter if the value was specified by the administrator when the condition was configured.</p> <p>In the sample CO trace, this is <code>1004:no-param</code>. The 1004 is a <User Data ID>.</p>
<Operator>	<p>The display name of the condition operator.</p> <p>In the sample CO trace, this is <code>date-equal</code>.</p>
<RHSOperand>	<p>The enumerative value and parameter list of the right-hand operand. This data can be displayed in one of two formats:</p> <p>Embedded: <RHS parameter> or <User Data ID>: <RHS parameter></p> <p>Embedded indicates that the information was entered by the administrator when the condition was configured.</p> <p><User Data ID> is a numerical value assigned to user data that might be sensitive.</p> <p><RHS parameter> is the string <code>no-param</code> if no parameters are specified for this operand or the value of the parameter if the value was specified by the administrator when the condition was configured.</p> <p>In the sample CO trace, this is <code>1004:no-param</code>. The 1004 is a <User Data ID>.</p>

Element	Description
<Values>	<p>The values of the left-hand and right-hand operands, separated by the colon (:) character. User data values are never displayed to ensure privacy. If both the left-hand and the right-hand operands reference user data, this element is represented by a tilde.</p> <p>In the sample CO trace, this element is represented by a tilde because both the left-hand and the right-hand operands reference user data.</p>
<NOT>	<p>The string NOT if the result was negated prior to reporting; otherwise the field has no value. This is the <i>If Not</i> option when creating a condition.</p> <p>In the sample CO trace, this condition result was not negated, therefore the element is represented by a tilde.</p>
<Result>	<p>A numerical result. See “Policy Result Values” on page 480.</p> <p>In the sample CO trace, this is 69 and indicates that the condition evaluated to true.</p>
<ResultOnError>	<p>A string describing the error that occurred. This is an optional field that only appears when the condition evaluation results in an error.</p> <p>The sample CO trace did not result in an error, so it has no string.</p>

Condition Set Evaluation Result

After the initial four fields, a CS trace has the following fields

```
~<JoinType>~<NOT>~<ConditionCount>~<Result>
```

A CS trace looks similar to the following:

```
com.novell.nxpe.NxpePolicyEvaluation$InformationContext@1af33d6~Thu
Jan 25 09:20:30 MST 2007~CS~1~~ANDs~~1~~69
```

Element	Description
<JoinType>	<p>Specifies how the condition results are combined, if there are multiple conditions. Possible values include ANDs and ORs.</p>
<NOT>	<p>The string NOT if the result was negated prior to reporting; otherwise the field has no value. This is the <i>If Not</i> option when creating a condition group.</p> <p>In the sample CS trace, the condition group was not negated, therefore the element is represented by a tilde.</p>
<ConditionCount>	<p>The number of conditions defined in the condition group.</p> <p>In the sample CS trace, this element has the value of 1.</p>

Element	Description
<Result>	A numerical result. See “Policy Result Values” on page 480 . In the sample CS trace, this is 69 and indicates that the condition evaluated to true.

Policy Action Initiation

After the initial four fields, a PA trace has the following fields

```
~<ActionName>~<TraceString1>~<TraceString2>~<TraceString3>~<Result>
```

A PA trace looks similar to the following:

```
com.novell.nxpe.NxpePolicyEvaluation$InformationContext@5af33f5~Fri
Jan 26 13:03:29 EST 2007~PA~1~Deny Access Messasge~Access denied
because you are not a Criminal History User~~~0
```

Element	Description
<ActionName>	The name of the action. In the sample PA trace, this is <code>Deny Access Message</code> .
<TraceString1>	The message specified with the action. In the sample PA trace, this is <code>Access denied because you are not a Criminal History User</code> .
<TraceString2>	The second part of the specified message. In the sample PA trace, this element has no value so only a tilde is added to the line.
<TraceString3>	The third part of the specified message. In the sample PA trace, this element has no value so only a tilde is added to the line.
<Result>	A numerical result. See “Policy Result Values” on page 480 . In the sample PA trace, this is 0 and indicates that the action was successfully assigned to the user.

Policy Action Completion

After the initial four fields, a PC trace has the following fields

```
~<ActionName>~<ActionParameters>~~~<Result>[~<ActionError>]
```

A PC trace looks similar to the following:

```
com.novell.nxpe.NxpePolicyEvaluation$InformationContext@5af33f5~Fri
Jan 26 13:03:29 EST 2007~PC~1~Document=(ou=xpemplPEP,ou=mastercdn,
ou=ContentPublisherContainer,ou=Partition,ou=PartitionsContainer,ou=VC
DN_Root,ou=accessManagerContainer,o=novell:romaContentCollectionXMLDoc
),Policy=(polagauth_DenyNotInCHGroup),Rule=(1::RuleID_1169499645674),A
ction=(Deny::1)~~~0
```

Element	Description
<ActionName>	<p>The fully distinguished name of the action.</p> <p>In the sample PC trace, the action has the following parts in its name:</p> <ul style="list-style-type: none"> ◆ Document=(ou=xpemIPEP,ou=mastercdn,ou=ContentPublisherContainer,ou=Partition,ou=PartitionsContainer,ou=VCDN_Root,ou=accessManagerContainer,o=novell:romaContentCollectionXMLDoc) ◆ Policy =(polagauth_DenyNotInCHGroup) ◆ Rule=(1::RuleID_1169499645674) ◆ Action=(Deny::1)
<ActionParameters>	<p>A list of the parameters passed to the action handler. A Permit action has no parameters. A Deny action has a parameter when the policy specifies a redirect URL.</p> <p>In this sample PC trace, the Deny action has no parameters so the action element is represented by a tilde.</p>
<Result>	<p>A numerical result. See “Policy Result Values” on page 480.</p> <p>In the sample PC trace, this is 0 and indicates success.</p>
<ActionError>	<p>A string describing the error that occurred when invoking the action. This is an optional field that only appears when the Result field contains an error code.</p> <p>The sample PC trace did not result in an error, so it has no string.</p>

Rule Evaluation Result

After the initial four fields, a RU trace has the following fields:

```
~<ParentPolicyName>~<ConditionSetJoinType>~~<ConditionSetCount:ActionCount>~~<Result>
```

A RU trace looks similar to the following:

```
com.novell.nxpe.NxpePolicyEvaluation$InformationContext@5af33f5~Fri
Jan 26 13:03:29 EST 2007~RU~RuleID_1169499645674~polagauth_DenyNotIn
CHGroup~DNF~~1:1~~0
```

Element	Description
<ParentPolicyName>	<p>The name of the parent policy to which the rule is assigned.</p> <p>In this sample RU trace, this element is set to polagauth_DenyNotIn CHGroup.</p>

Element	Description
<ConditionSetJoinType>	<p>The type of joining that occurs between conditions and condition sets. It is set to one of the following:</p> <ul style="list-style-type: none"> ◆ CNF: Indicates that sets are AND'ed and conditions within a condition group are OR'ed. ◆ DNF: Indicates that sets are OR'ed and conditions within a condition group are AND'ed. <p>In the sample RU trace, this element is set to <code>DNF</code>.</p>
<ConditionSetCount:ActionCount>	<p>The number of condition sets and actions defined for this rule.</p> <p>In the sample RU trace, this is 1:1, for one condition set and one action.</p>
<Result>	<p>A numerical result. See "Policy Result Values" on page 480.</p> <p>In the sample RU trace, this is 0, indicating that the rule was successfully evaluated.</p>

Rule List Evaluation Result

After the initial four fields, a RL trace has the following fields:

```
~~~<RuleCount>~~<Result>
```

A RL trace looks similar to the following:

```
com.novell.nxpe.NxpePolicyEvaluation$InformationContext@5af33f5~Fri
Jan 26 13:03:29 EST 2007~RL~0~~~~Rule Count: 1~~0
```

Element	Description
<RuleCount>	<p>The number of rules defined for the policy.</p> <p>In the sample RL trace, this is <code>Rule Count: 1</code>, indicating that there is one rule in the policy.</p>
<Result>	<p>A numerical result. See "Policy Result Values" on page 480.</p> <p>In the sample RL trace, this is 0, indicating success.</p>

41.3.2 Policy Result Values

The last field of all trace string is the `<result>` field. The following values are possible:

Value	Name	Description
0	Success	
1	Error: No memory	The system is out of memory.
2	Error: Bad data	The data sent for evaluation is invalid.

Value	Name	Description
3	Error: Configuration initialization	
4	Error: General failure	
5	Pending	
64	Permit	The rule produced a Permit action.
65	Deny	The rule produced a Deny action.
66	Obligation	The rule triggered an obligation, indicating that additional processing is required. Identity Injection policies trigger obligations.
67	No action	The rule did not initiate any action.
68	Condition false	The condition evaluated to false.
69	Condition true	The condition evaluated to true.
70	Condition unknown	Condition input was not available, thus the results are unknown.
71	Cancel	The current operation has been canceled.
72	Error: Interface unavailable	The current operation is unavailable.
73	Error: Data unavailable	The data required for evaluation was unavailable.
74	Error: Illegal state	Processing error; report to Novell technical support.

41.3.3 Sample Policies

- ◆ [“Simple Trace from a Test Environment” on page 481](#)
- ◆ [“Role Setting Trace” on page 482](#)
- ◆ [“Deny User Trace When the User Is Not Assigned the Role” on page 482](#)

Lines are long and wrap, so line numbers have been added to make the traces more readable. Also the elements before the <Type> element have been removed.

Simple Trace from a Test Environment

The following sample is a trace of a simple authorization policy which has one condition and which results in a permit action.

```

1. ~RL~0~~~~Rule Count: 1~~0
2. ~RU~1~SelectedDateFormat-current-date~DNF~~1:1~~0
3. ~CS~1~~ANDs~~1~~69
4. ~CO~0~1004:no-param~date-equal~1004:no-param~~~69
5. ~PA~1~~doAction()~Permit~~~0
6. ~PC~1~~Document=(), Policy=(SelectedDateFormat-current-date),
Rule=(5000::1), Action=(Permit::1)~~~~0

```

The trace describes the following about this policy:

1. The RL trace indicates that there is one rule in the policy and that the policy evaluated without error.
2. The RU trace indicates that the name of the policy is SelectedDateFormat-current-date, that conditions in a condition group are AND'ed and condition sets are OR'ed (DNF), that there is one condition and one condition set, and that the rule evaluated without error.
3. The CS trace indicates that the condition results will be AND'ed, that there is one condition and that the condition set evaluated to true.
4. The CO trace indicates that the left hand operand references user data, that the condition is for equal dates, that the right-hand operand references user data, and that the condition evaluated to true.
5. The PA trace indicates that the action is a Permit action and that action was successfully applied.
6. The PC trace indicates that action of the policy (specified by the Document, Policy, Rule, and Action parameters) was successfully granted to the requester.

Role Setting Trace

The following trace shows a role evaluation which results in the role not being assigned to the user because the user does not have the required LDAP attribute.

1. ~RL~0~~~~Rule Count: 1~~67
2. ~RU~RuleID_1167918127401~polisrole_InSystemAdministratorsGroup_New~DNF~~1:1~~67
3. ~CS~1~~ANDs~~1~~68
4. ~CO~1~6645:no-param~ldap-is-member-of~66455:null~~~68

This trace describes the following about the policy.

1. The RL trace indicates that the Role policy has one rule, that the policy evaluated without error, but the policy did not result in any actions.
2. The RU trace indicates that the rule (RuleID_1167918127401) whose parent is (polisrole_InSystem AdministratorsGroup_New) has one condition and one condition group and that the rule evaluated without error, but the rule did not result in any actions.
3. The CS trace indicates that the condition set evaluated to false.
4. The CO trace that the condition (ldap-is-member-of) evaluated to false.

Authorization policies verify whether the user matches the conditions. This trace indicates that the user did not have the required value for the ldap-is-member-of condition. If you expected this user to have the appropriate value, you need to evaluate what value you specified for the condition and examine the values stored in this LDAP attribute in your user store.

Deny User Trace When the User Is Not Assigned the Role

The following trace illustrates a trace of a Deny policy that requires the user to be a member of a specific role to be allowed access and denies access to all users who have not been assigned the required role.

1. ~RL~0~~~~Rule Count: 1~~0
2. ~RU~RuleID_1169499645674~polagauth_DenyNotInCHGroup~DNF~~1:1~~0

```

3. ~CS~1~~ANDs~~1~~69
4. ~CO~1~6660:no-param~~6661:null~~NOT~69
5. ~PA~1~~Deny Access Messasge~Access denied because you are not a
Criminal History User~~~0
6. ~PC~1~~Document=(ou=xpemplPEP,ou=mastercdn,ou=ContentPublisher
Container,ou=Partition,ou=PartitionsContainer,ou=VCDN_Root,ou=accessMa
nagerContainer,o=novell:romaContentCollectionXMLDoc),Policy=(polagauth
_DenyNotInCHGroup),Rule=(1::RuleID_1169499645674),Action=(Deny::1)~~~~
0

```

This trace describes the following about the policy.

1. The RL trace indicates that the policy has one rule and that the policy was evaluated without error.
2. The RU trace indicates that the rule (RuleID_1169499645674) has one condition and one action and that the rule evaluated without error.
3. The CS trace indicates that the condition set evaluated to true (the requester is not a member of the group).
4. The CO trace indicates that the condition evaluated to true (the requester is not a member of the group).
5. The PA trace indicates that following message was sent to the user: Access denied because you are not a Criminal History User.
6. The PC trace indicates the policy was successfully applied and that the requester was denied access according to the rules of the policy.

If you expected the requester initiating the request in this trace to have been granted the role, you need to look at the authentication trace. Roles are assigned to users when they authenticate.

41.4 Policy Page Timeout

If your policy page hangs, and you have an LDAP group or LDAP ou being used in the policy, check the health of your user stores (LDAP servers) and ensure that they are communicating.

41.5 Policy Changes Are Not Saved

If you have two or more administrators making changes to policies at the same time and all of these policies reside in the same policy container, some modifications will be lost. If you have multiple administrators managing and configuring policies, we suggest that you create a policy container for each administrator. This allows multiple administrators to modify policies at the same time. See [Section 28.2, “Managing Policy Containers,” on page 312.](#)

41.6 Policy Creation and Storage

For troubleshooting, you can locate the policy XML and send it to Novell® for debugging. Policies are stored as XML documents in the object directory, with one XML document to represent each policy container. The default policy container (Master_Container) resides at:

```

\\novell\accessManagerContainer\VCDN_Root\PartitionsContainer\Partitio
n\ContentPublisherContainer\mastercdn\xpemplPEP\romaContentCollectionXM
LDoc

```

Other policy containers are stored following the same path, with a unique name string representing the policy name that replaces the `ou=mastercdn` portion of the above path.

If you are unsure if the policy is being created correctly or you need to check to see if the policy is enabled, you can view the policy list in the interface. If you think the GUI is not properly displaying the policy, you can also view the XML by navigating to the Policy Conditions on which you edit rules, right click and choose *This Frame > View Frame Source*.

41.7 Policy Distribution

Policy definitions are not replicated but are referenced by the Access Gateways for which the policy is to be evaluated. The policy reference mechanism is a set of XML elements which refer back to the policy definitions stored in the various policy containers. If you have configured a policy for a protected resource and an Access Gateway does not seem to be executing this policy, use the following procedures to verify that the Access Gateway has been configured to use the policy:

- 1 Set the level of Application logging to *config*. See [Section 41.2, “Turning on Logging for Policy Evaluation,” on page 474](#).

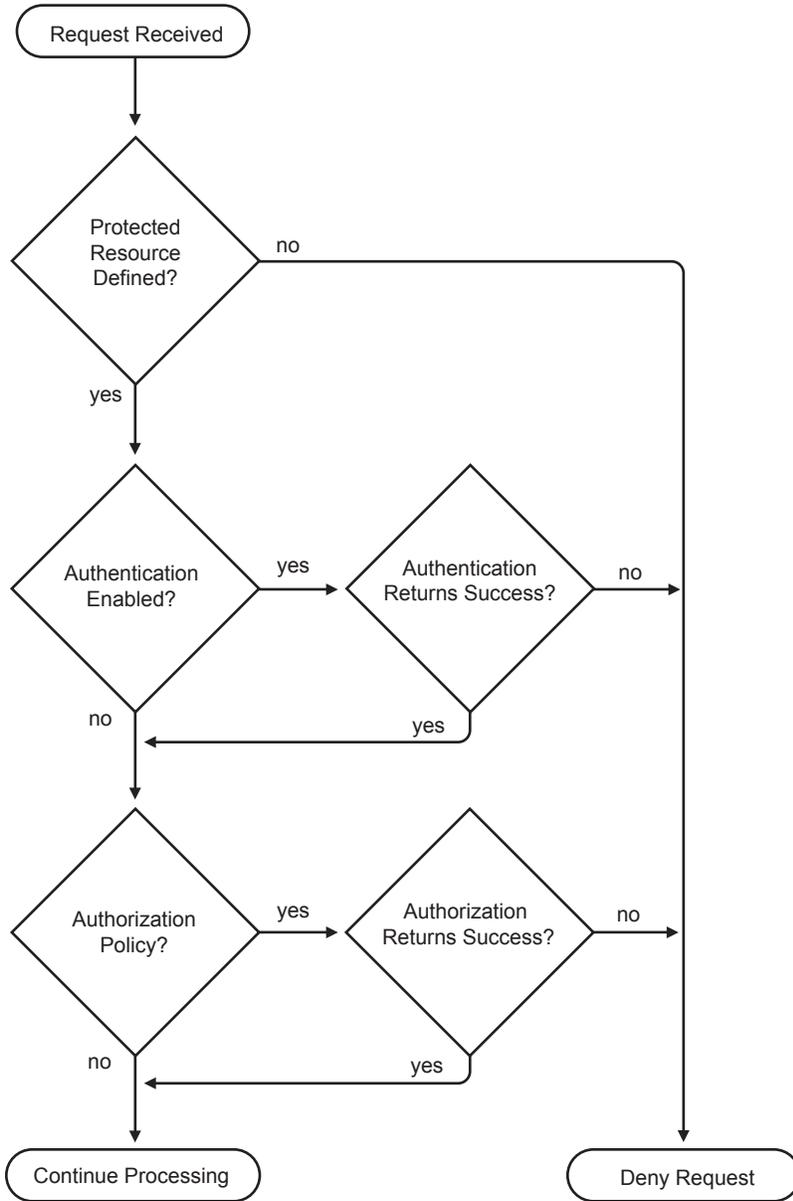
This enables the tracing of the policy enforcement lists.

- 2 Search for name of your policy in a `<PolicyEnforcementList>` element. The `ExternalElementRef` attribute contains a reference to the policy name.
 - ♦ On the Linux Access Gateway, you can find these elements in the `catalina.out` file.
 - ♦ On the NetWare Access Gateway, the trace for these elements goes to the system console.You can also find an XML file named after each protected resource in the `sys:\etc\proxy\pr` directory. These files contain the references to the policy names that have been enabled for the protected resources.
- 3 If you cannot find the policy name, the Access Gateway has not been configured to use the policy. The configuration either needs to be applied or the policy needs to be enabled. For information on how to assign a policy to a protected resource, see [Section 12.4, “Configuring Protected Resources,” on page 150](#).
- 4 If you find the policy name associated with the correct protected resource, you need to evaluate why the policy is not evaluating according to your design. Set the level of Application logging to *info* and examine the policy trace from a user accessing the protected resource. See [Section 41.3, “Understanding Policy Evaluation Traces,” on page 475](#).

41.8 Policy Evaluation: Access Gateway Devices

The following diagram depicts how authorization policies fit into the protected resource processing for the proxy.

Figure 41-1 Policy Evaluation



Policies for the Access Gateway devices are evaluated by the policy engine in Java. A SOAP interface is used to transition from the proxy to Java and back. To see the SOAP messages, you need to set the logging level of the *Application* level to *config*. See [Section 41.2, “Turning on Logging for Policy Evaluation,”](#) on page 474.

For NetWare[®], the SOAP messages are output to the Logger Screen. For Linux, the SOAP messages are output to the `catalina.out` file. Sample SOAP messages are shown in the following examples:

- ♦ [Section 41.8.1, “Successful Policy Configuration Example,” on page 486](#)
- ♦ [Section 41.8.2, “No Policy Defined Configuration Example,” on page 487](#)
- ♦ [Section 41.8.3, “Deny Access Configuration/Evaluation Example,” on page 487](#)

41.8.1 Successful Policy Configuration Example

Note the Policy Enforcement Point (PEP) identifier of `AGIdentityInjection` in the request and the `PolicyID` in the response.

Configuration Request

```
toBufSeg: <?xml version="1.0" encoding="UTF-8"?>
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/
  envelope/">
<SOAP-ENV:Body>
  <NXPPES ID="12">
    <Configure-ag PEPName="AGIdentityInjection">
      <PolicyEnforcementList
        RuleCombiningAlgorithm="DenyOverridesWithPriority"
        schemaVersion="1.32"
        LastModified="1138389868885"
        LastModifiedBy="cn=admin,o=novell">
        <PolicyRef ElementRefType="ExternalWithIDRef"
          ExternalElementRef="PolicyID_xpemplPEP_AGIdentity
            Injection_ii_test"
          ExternalDocRef="ou=xpemplPEP,ou=mastercdn,
            ou=ContentPublisherContainer,ou=Partition,
            ou=PartitionsContainer,ou=VCDN_Root,ou=access
            ManagerContainer,o=novell:romaContentCollection
            XMLDoc"
          UserInterfaceID="PolicyID_xpemplPEP_AGIdentity
            Injection_ii_test"/>
        </PolicyEnforcementList>
      </Configure-ag>
    </NXPPES>
  </SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```

Configuration Response

```
LibertyProcessMsgCB:
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/
  envelope/">
<SOAP-ENV:Body>
  <NXPPES Id="" Status="success">
    <ConfigureResponse PolicyId="755OK8P0-7543-518M-8L8M-N0P2LM2
      N3027">
      <ContextDataElement Enum="2551"/>
    </ConfigureResponse>
  </NXPPES>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```

```

    </NXPES>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>

```

41.8.2 No Policy Defined Configuration Example

The following is a sample of a configuration request where the policy code detects that no policies are in effect for the protected resource and Policy Enforcement Point (PEP).

Configuration Request

```

toBufSeg: <?xml version="1.0" encoding="UTF-8"?>
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/
envelope/">
<SOAP-ENV:Body>
  <NXPES ID="11">
    <Configure-ag PEPName="AGAuthorization">
      <PolicyEnforcementList
        RuleCombiningAlgorithm="DenyOverridesWithPriority"
        schemaVersion="1.32"
        LastModified="1138389868885"
        LastModifiedBy="cn=admin,o=novell">
        <PolicyRef ElementRefType="ExternalWithIDRef"
          ExternalElementRef="PolicyID_xpemplPEP_AGIdentity
            Injection_ii_test"
          ExternalDocRef="ou=xpemplPEP,ou=mastercdn,ou=Content
            PublisherContainer,ou=Partition,ou=Partitions
            Container,ou=VCDN_Root,ou=accessManager
            Container,o=novell:romaContentCollectionXMLDoc"
          UserInterfaceID="PolicyID_xpemplPEP_AGIdentityInjection_
            ii_test"/>
        </PolicyEnforcementList>
      </Configure-ag>
    </NXPES>
  </SOAP-ENV:Body>
</SOAP-ENV:Envelope>

```

Configuration Response

```

LibertyProcessMsgCB:
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/
envelope/">
  <SOAP-ENV:Body>
    <NXPES Id="" Status="emptypolicyset"/>
  </SOAP-ENV:Body>
</SOAP-ENV:Envelope>

```

41.8.3 Deny Access Configuration/Evaluation Example

Configuration Request

```

toBufSeg: <?xml version="1.0" encoding="UTF-8"?>
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/

```

```

envelope/">
<SOAP-ENV:Body>
  <NXPE$ ID="17">
    <Configure-ag PEPName="AGAuthorization">
      <PolicyEnforcementList
        RuleCombiningAlgorithm="DenyOverridesWithPriority"
        schemaVersion="1.32"
        LastModified="1138718667305"
        LastModifiedBy="cn=admin,o=novell">
      <PolicyRef
        ElementRefType="ExternalWithIDRef"
        ExternalElementRef="PolicyID_xpemplPEP_AGIdentityInjection
          _custom_test"
        ExternalDocRef="ou=xpemplPEP,ou=mastercdn,ou=Content
          PublisherContainer,ou=Partition,ou=PartitionsContainer,
          ou=VCDN_Root,ou=accessManagerContainer,o=novell:roma
          ContentCollectionXMLDoc"
        UserInterfaceID="PolicyID_xpemplPEP_AGIdentityInjection
          _custom_test"/>
      <PolicyRef
        ElementRefType="ExternalWithIDRef"
        ExternalElementRef="PolicyID_xpemplPEP_AGAuthorization_
          deny-all"
        ExternalDocRef="ou=xpemplPEP,ou=mastercdn,ou=Content
          PublisherContainer,ou=Partition,ou=PartitionsContainer,
          ou=VCDN_Root,ou=accessManagerContainer,o=novell:roma
          ContentCollectionXMLDoc"
        UserInterfaceID="PolicyID_xpemplPEP_AGAuthorization
          _deny-all"/>
      </PolicyEnforcementList>
    </Configure-ag>
  </NXPE$>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>

```

Configuration Response

```

LibertyProcessMsgCB:
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/
  envelope/">
  <SOAP-ENV:Body>
    <NXPE$ Id="" Status="success">
      <ConfigureResponse
        PolicyId="55N3NL81-L29N-2619-K0M8-2L963M0MM701"/>
    </NXPE$>
  </SOAP-ENV:Body>
</SOAP-ENV:Envelope>

```

Evaluation Request

```

toBufSeg: <?xml version="1.0" encoding="UTF-8"?>
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/
  envelope/">
<SOAP-ENV:Body>

```

```
<NX PES ID="18">
  <Evaluate PolicyId="55N3NL81-L29N-2619-K0M8-2L963M0MM701"
    Verbose="on"/>
</NX PES>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```

Evaluation Response

```
LibertyProcessMsgCB:
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/
  envelope/">
<SOAP-ENV:Body>
  <NX PES Id="" Status="success">
    <EvaluateResponse>
      <DoAction ActionName="Deny" ActionTTL="-1" Enum="2620">
        <Parameter Enum="10" Name="Message" Value=""/>
      </DoAction>
    </EvaluateResponse>
  </NX PES>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```


Troubleshooting the Access Gateway

42

For a solution to an Access Gateway problem, see the following sections:

- ♦ [Section 42.1, “Fixing Problems Common to Both Platforms,” on page 491](#)
- ♦ [Section 42.2, “Troubleshooting the Linux Access Gateway,” on page 492](#)
- ♦ [Section 42.3, “Troubleshooting the NetWare Access Gateway,” on page 507](#)

42.1 Fixing Problems Common to Both Platforms

The following problem occurs on both the Linux Access Gateway and the NetWare Access Gateway.

42.1.1 Mismatch Certificates in a Group of Access Gateways

Sometimes a newly added server to a group does not receive the certificate that the rest of the group is using for SSL. To fix this problem:

- 1 Click *Access Manager > Access Gateways > Edit:[Group Name] > [Name of Reverse Proxy]*.
- 2 For the Server Certificate, click the Select Certificate icon, select a different certificate, for example the test-connector certificate.
- 3 Click *OK* to ignore the warnings that the certificate CN does not match the reverse proxy.
This is what you want
- 4 Click *OK*.
When you click *OK*, the certificate is pushed to the Access Gateway.
- 5 Click *[Name of Reverse Proxy]*.
This needs to be the same reverse proxy that you selected in [Step 1](#).
- 6 For the Server Certificate, click the Select Certificate icon, select the certificate whose CN matches the published DNS name of the parent proxy service, then click *OK*.
- 7 Click *OK*.
When you click *OK*, the correct certificate is pushed to the Access Gateway.
- 8 Repeat [Step 1](#) through [Step 7](#) for each reverse proxy that uses a unique certificate. If all the reverse proxies use the same certificate, continue with [Step 9](#).
- 9 Click *Cancel Changes*.
The certificates have been pushed to the Access Gateway, so no other changes need to be applied.
- 10 Restart the embedded service provider of the newly added server. Click *Access Gateways > [Name of Server] > Actions > Restart Service Provider*.

42.2 Troubleshooting the Linux Access Gateway

This section provides various troubleshooting scenarios and frequently asked questions that you might encounter while using the Linux Access Gateway, and suggests appropriate actions.

- ◆ [Section 42.2.1, “Linux Access Gateway Logs,” on page 492](#)
- ◆ [Section 42.2.2, “Troubleshooting a Linux Access Gateway Crash,” on page 494](#)
- ◆ [Section 42.2.3, “Connection Details,” on page 498](#)
- ◆ [Section 42.2.4, “Network Socket Issues,” on page 498](#)
- ◆ [Section 42.2.5, “Authentication Issues,” on page 498](#)
- ◆ [Section 42.2.6, “Rewriter Issues,” on page 501](#)
- ◆ [Section 42.2.7, “Useful Linux Access Gateway Commands,” on page 502](#)
- ◆ [Section 42.2.8, “Using the Linux Access Gateway Console,” on page 503](#)
- ◆ [Section 42.2.9, “COS Related Issues,” on page 503](#)
- ◆ [Section 42.2.10, “Memory Related Issues,” on page 503](#)
- ◆ [Section 42.2.11, “Authorization and Identity Injection Issues,” on page 504](#)
- ◆ [Section 42.2.12, “Form Fill Issues,” on page 505](#)

42.2.1 Linux Access Gateway Logs

This section contains the following information about the Linux Access Gateway logs:

- ◆ [“Interpreting Log Messages” on page 492](#)
- ◆ [“Configuring Log Levels” on page 493](#)
- ◆ [“Configuring Log Messages” on page 493](#)

Interpreting Log Messages

In Linux Access Gateway, the ICS_dyn logs are logged in the following format:

```
<timestamp> : <hostname> : LINUX_AG : <Component> : <Unique Id> :  
<log message>
```

- ◆ **Timestamp:** Specifies the time when the log was sent.
- ◆ **Hostname:** Specifies the name of the host that sent the log.
- ◆ **LINUX_AG:** Indicates that the logs are sent from the Linux Access Gateway.
- ◆ **Component:** Specifies the Linux Access Gateway component that sent the log.
- ◆ **Unique ID:** Specifies a number that is the unique ID for that request.
- ◆ **Log Message:** Specifies the log message.

A sample log message is given below:

```
Jul 19 02:18:52 proxy140 LINUX_AG:  SERVICEMANAGER :      0 :  
ServiceManager - HTTP_ErrorService service errorservice is running
```

Configuring Log Levels

You can set the following log levels:

1. **LOG_EMERG:** Sends only messages that render the system unusable, if not resolved.
2. **LOG_ALERT:** Sends only messages which require immediate action.
3. **LOG_CRIT:** Sends only messages about critical situations.
4. **LOG_ERROR:** Sends warning messages about recoverable errors.
5. **LOG_NOTICE:** Sends the service configuration logs information about the status of a service.
6. **LOG_INFO:** Sends informational messages such as requests sent to Web servers and the results of authentication requests.
7. **LOG_DEBUG:** Sends debug messages.

When you run the `/etc/init.d/novell-vmc start` command, the default log level is set to `LOG_NOTICE`. You can change the log level to any level from `LOG_EMERG` to `LOG_INFO`.

When you run the `/etc/init.d/novell-vmc debug-start` command, the default log level is set to `LOG_DEBUG`. You can change the log level to any of the available log levels.

Changing Log Levels

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

```
nash
```

- 2 At the `nash` shell prompt, enter the following command:

```
configure .current
```

- 3 To change the log level, enter the following command:

```
log-conf log-level <log level>
```

Replace `<log level>` with the new log level that you want to set.

- 4 To apply changes, enter the following command:

```
apply
```

Configuring Log Messages

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

```
nash
```

- 2 At the `nash` shell prompt, enter the following command:

```
nash
```

- 3 To enter the configuration mode, enter the following command:

```
configure .current
```

- 4 Enter one of the following commands to configure logging:

Command	Purpose
log-conf debug-soap-messages enable	To log all the SOAP messages between the Linux Access Gateway and the Enterprise Server to the <code>/var/log/lagsoapmessages</code> file.
log-conf no debug-soap-messages enable	To disable to logging of SOAP messages between the Linux Access Gateway and the Enterprise Server.
log-conf debug-http-headers enable	To log all the HTTP headers between the browsers and the Linux Access Gateway and between the Linux Access Gateway and the Web servers-lag and lag-webserver to the <code>/var/log/laghttpheaders</code> file.
log-conf no debug-http-headers enable	To disable the logging of HTTP headers to the <code>/var/log/laghttpheaders</code> file.

5 To apply changes, enter the following command:

```
apply
```

42.2.2 Troubleshooting a Linux Access Gateway Crash

The Linux Access Gateway might have crashed due to the following reasons:

- ◆ SIGSEGV
- ◆ ASSERT

In the event of a Linux Access Gateway crash, send the following information to the Novell® Support:

- ◆ Core dump of the proxy. For more information on collecting a core dump, see [“Core Dump” on page 494](#).
- ◆ Event log of the proxy. For more information on collecting an event log, see [“Event Log” on page 495](#).
- ◆ Linux Access Gateway logs. For more information on collecting Linux Access Gateway logs, see [“Linux Access Gateway Logs” on page 495](#)
- ◆ Packet capture. For more information, see [“Packet Capture” on page 497](#).

Core Dump

Before you begin:

- ◆ Make sure there is free space in `root` to hold the core file and that the space is at least equal to the RAM size
- ◆ Make sure the core file size is set in the `/etc/security/limits.conf` file.

To collect a core dump:

- 1 Log in as the root user.

- 2 To disconnect all instances of the Linux Access Gateway, enter the following command:

```
/etc/init.d/novell-vmc stop
```

- 3 At the bash prompt, open the following:

```
/etc/init.d/novell-vmc-chroot
```

Locate the following line and uncomment it:

```
ulimit -c unlimited
```

- 4 Enter the following command to start the Linux Access Gateway:

```
/etc/init.d/novell-vmc start
```

- 5 Repeat the scenarios to reproduce the issue.

The core is dumped to the `/chroot/lag` directory with the filename `core.<pid>`

`<pid>` is the process ID of `ics_dyn` process.

After the core is dumped, the Linux Access Gateway restarts.

- 6 Tar or Zip the core dump and send it to Novell Technical Support.

Linux Access Gateway Logs

- 1 Enter the following command from the bash shell to collect debug log files that are generated:

```
/chroot/lag/opt/novell/bin/getlaglogs.sh
```

- 2 The tar file `laglogs.tgz` will be located in the `/var/log` directory.

Send this tar file to Novell Support.

Event Log

By default the event log size is 15 MB. The size of event log can be controlled by configuring the required event log size in the file `eventlogsize.cfg`, located at the `/chroot/lag/etc/opt/novell` directory. This is a memory mapped file and should contain only the file size information. This file should not contain any other characters or new line. For example, if you specify 350 in the file, you can configure an event log of size 350 MB.

Event Log for Production Build

To get the event log for the production build:

- 1 Log in as the root user.

- 2 To disconnect all instances of Linux Access Gateway, enter the following command:

```
/etc/init.d/novell-vmc stop
```

- 3 Enter the following command to change the root environment:

```
chroot /chroot/lag
```

- 4 To start the process, enter the following command:

```
gdb /opt/novell/bin/ics_dyn 2>/var.log/ics_dyn.log
```

- 5 In the GDB prompt run the following command:

```
run -m <memory>
```

Where, `<memory>` is the percentage of total memory to be used for `ics_dyn` process. It is recommended to set this value in the range of 20-30 per cent.

- 6 Repeat the scenarios to reproduce the issue.
 - 6a If you are trying to reproduce the proxy crash, you will enter the GDB prompt as soon as the crash is reproduced.
 - 6b If you are trying to reproduce a functionality issue, enter the following command to enter the GDB prompt as soon as the issue is reproduced:

```
Ctrl+C
```

NOTE: For a list of commands that can be entered in the debugger, see [“Useful Debugger Commands” on page 497](#).

- 7 To save event logs to a file, enter the following command:

```
d ,save l
```

This stores all the events in the `/chroot/lag/opt/novell/debug/eventlog.txt` file.

- 8 Tar or Zip this file and send it to Novell Support.

If you do not get an event log, send only the core dump, the ICS debug Log, and the packet capture to the Novell Support.

Event Log for Debug Build

To get the event log:

- 1 Log in as the root user.
- 2 To stop all instances of Linux Access Gateway, enter the following command:

```
/etc/init.d/novell-vmc stop
```
- 3 To start the Novell Linux Access Gateway in debugging mode, enter the following command:

```
/etc/init.d/novell-vmc gdb
```
- 4 To run the Linux Access Gateway process, enter the following command at the GDB prompt:

```
run -m <memory> 2>/var/log/ics_dyn.log
```

Where *<memory>* is the percentage of total memory to be used for `ics_dyn` process. It is recommended to set this value in the range of 20-30 per cent.
- 5 Repeat the scenarios to reproduce the issue.
 - 5a If you are trying to reproduce the proxy crash, you will enter the GDB prompt as soon as the crash is reproduced.
 - 5b If you are trying to reproduce a functionality issue, enter the following command to enter the GDB prompt as soon as the issue is reproduced:

```
Ctrl+C
```

NOTE: For a list of commands that can be entered in the debugger, see [“Useful Debugger Commands” on page 497](#).

- 6 To save all event logs to a file, enter the following command:

```
d ,save l
```

This stores all the events in the `/chroot/lag/opt/novell/debug/eventlog.txt` file.

- 7 Tar or Zip this file and send it to Novell Support.

If you do not get an event log, send only the core dump, `laglogs.tgz`, and the packet capture to Novell Support.

Useful Debugger Commands

Table 42-1 GDB Commands

Command	Function
<code>gcore</code>	Generate core file
<code>k</code>	Kill process
<code>q</code>	Quit GDB prompt
<code>bt</code>	Print the back trace

Analyzing Proxy Hang

To analyze the proxy hang:

- 1 Enter the following command to change the root environment:

```
chroot /chroot/lag
```

- 2 Enter the following command to attach the `ics_dyn` process to the debugger:

```
gdb /opt/novell/bin/ics_dyn <pid>
```

Here, `<pid>` refers to the Process ID of the `ics_dyn` process.

- 3 In the GDB prompt, enter the following commands:

```
set logging off
```

```
set logging on <filename>
```

`<filename>` specifies the name of the file where the output of debugger commands executed will be stored.

```
thread apply all bt
```

- 4 Enter the following command to save the core dump in the `/chroot/lag` directory.

```
gcore
```

Core dump will be saved as `core.<pid>`

- 5 Tar or Zip these files and send it to Novell Support.

Packet Capture

The `tcpdump` utility is available allows you to capture a trace of TCP/IP packet activity.

- 1 Log in as the `root` user.

- 2 Enter the following command:

```
tcpdump -s0 -n -t -p -i 'any' -w filename.cap
```

- 3 Tar or Zip this file and send it to Novell Support.

42.2.3 Connection Details

To obtain connection information:

- 1 Log in as the `root` user.
- 2 At the bash prompt, enter one of the following `netstat` commands:

Command	Details
<code>netstat -an</code>	Provides the connection information
<code>netstat -s -t</code>	Provides the connection statistics

42.2.4 Network Socket Issues

This section lists various issues related to network sockets and provides information on how to verify bind and connection issues:

- ♦ [“Socket Listener Bind” on page 498](#)
- ♦ [“Issues with Outgoing Connections” on page 498](#)

Socket Listener Bind

To verify whether the socket listener is bound to the required port:

- 1 Log in as the `root` user.
- 2 At bash prompt, enter the following command:

```
netstat -an | grep LISTEN
```

All ports are displayed.
- 3 Search for the desired port.
If the required port is not visible in the list, a bind failure has occurred.

Issues with Outgoing Connections

To verify that the Access Gateway is able to make outbound connections:

- 1 Log in as the `root` user.
- 2 At bash prompt, view the following log file:

```
/var/log/ics_dyn.log
```
- 3 Search for a connection message. If the service is unavailable, the file contains messages similar to the following:

```
Jul 19 02:18:52 proxy140 LINUX_AG: CONMGR :          0 : ERROR  
Connection FAILED with peer (10.10.10.1), port(35120)
```

42.2.5 Authentication Issues

This section provides information related to authentication:

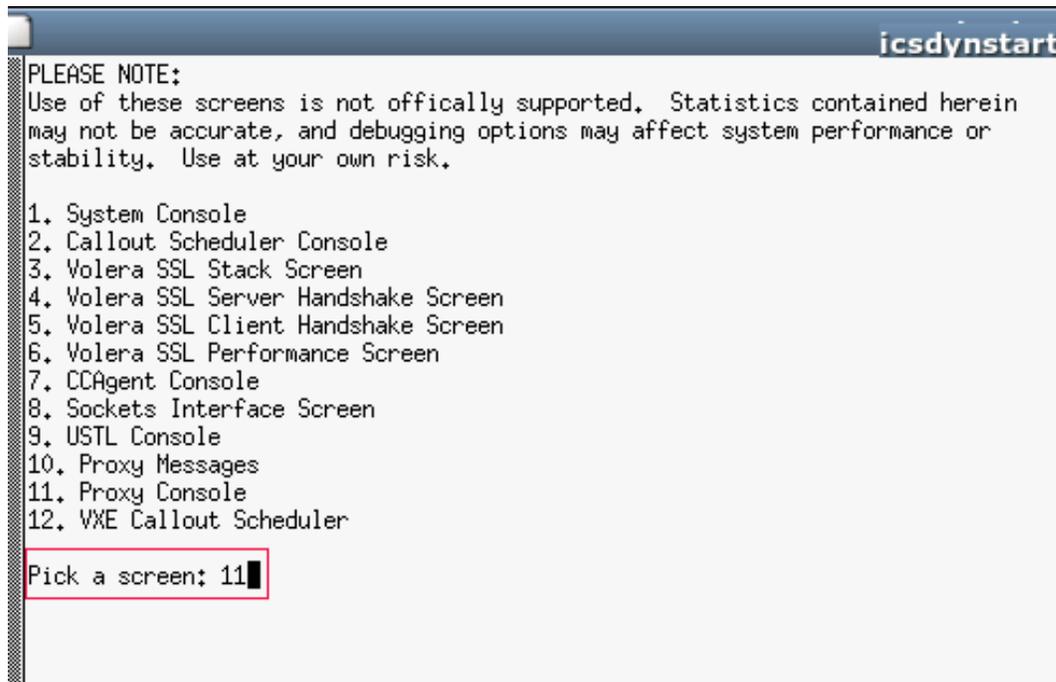
- ♦ [“User Details” on page 499](#)

- ♦ “Error Codes” on page 501

User Details

To check the details about the users logged in to the Linux Access Gateway:

- 1 Access the Linux Access Gateway main screen. For more information on how to access the Linux Access Gateway main screen, see [Section 42.2.8, “Using the Linux Access Gateway Console,”](#) on page 503.



- 2 Select the *Proxy Console* option. Type the option number at *Pick a Screen*.
The Linux Access Gateway Console screen is displayed.

- 3 To select the *Identity Agent Console* option, type the option number at *Enter Option*.

```
icsdynstart
Volera Excelerator Proxy Console

 1. Display current activity
 2. Display memory usage
 3. Display ICP statistics
 4. Display DNS options
 5. Display cache statistics
 6. Display not cached statistics
 7. Display HTTP server statistics
 8. Display HTTP client statistics
 9. Display connection statistics
10. Display FTP client statistics
11. Display GOPHER client statistics
12. Display configured addresses and services
13. Display SOCKS client statistics
14. Application Proxies
15. Transparent Proxy statistics
16. Site download options
17. Debug options
18. Identity Agent Console

Enter option: 18
```

Identity Agent Console screen is displayed.

```
icsdynstart
Total users: 1 (All users currently displayed)
X-Auth, O-UnAuth, R-Rtrd, W-Wrkg, U-Use, I-Idle, TTL-in d:hh:mm:ss format
(5)xw+U0I0: : 10.10.10.1, TTL:0:02:30:00

(1) Previous Page, (2) Next Page, (3) Refresh (4) Exit:
```

- ◆ **X:** Authenticated user.
- ◆ **O:** Unauthenticated user.
- ◆ **R:** Retired user; the user session has timed out. The default time-out is 3 minutes. In this state, the user session is deleted. If the user makes another request from the browser session, the Linux Access Gateway requires the user to authenticate.

- ♦ **W:** User session is functional.
- ♦ **U:** Use count is more than zero.
- ♦ **I:** User session is idle.
- ♦ **TTL:** Time remaining for the user session to go to the retired state if the user session remains idle.

The screen displays 20 users at a time at a time. The screen also displays the browser IP address. The following options are available at the bottom of the screen:

- ♦ **Previous Page:** Lets you go to the previous page.
- ♦ **Next Page:** Lets you go to the next page (to view the next set of users).
- ♦ **Refresh:** Refreshes the page to reflect the latest user status.
- ♦ **Exit:** Exits the console.

Error Codes

The following error codes indicate authentication problems.

500 Internal Server Error

Possible Cause: Authentication failed because of a system error.

Action: Contact Novell Support.

504 Gateway Timed Out

Possible Cause: The authentication back-end channel is not working.

Action: Check if the embedded service provider is listening on the loopback address 127.0.0.1 at port 8080: Use the following command

```
netstat -na | grep 8080
```

42.2.6 Rewriter Issues

The following information can help you troubleshoot problems with the rewriter:

- ♦ [“Reading Configuration Files” on page 501](#)
- ♦ [“Rewriting a URL” on page 502](#)

Reading Configuration Files

Enable log-level to LOG_INFO to see the log messages (see [Section , “Configuring Log Levels,” on page 493](#)) in /var/log/ics_dyn. If the Rewriter is successful in reading the configuration files, the following message is displayed:

```
Reading Config File
```

```
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:Configuration information read successfully
```

If the Rewriter fails to read the configuration files, the following message is displayed:

```
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:Reading configuration failed for ssTypeName=www.mynovell.com
```

If this happens, re-create the corresponding proxy service and restart the Linux Access Gateway service.

Rewriting a URL

Enable log-level to LOG_DEBUG to see the log messages (see [Section , “Configuring Log Levels,” on page 493](#)) in `/var/log/ics_dyn`. If the Rewriter successfully rewrites the URL, the following messages are displayed:

```
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:URL:'http://
www.mynovell.com:9090/common/inc/nav/main.js' Content type match, Will
Rewrite
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:URL:'http://
www.mynovell.com:9090/common/inc/nav/main.js' Unknown Content-Type -
automatic match - Will Rewrite
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0::'http://
www.mynovell.com:9090/common/inc/nav/main.js' NULL Content-Type -
automatic match - Will Rewrite
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:In
RewriterOption::shouldRewriteUrl, returning TRUE.
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:URL:'http://
www.mynovell.com:9090/common/inc/nav/main.js' Unknown extension -
automatic match - Will Rewrite
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:URL:'http://
www.mynovell.com:9090/common/inc/nav/main.js' NULL extension -
automatic match - Will Rewrite
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:URL:'http://
www.mynovell.com:9090/common/inc/nav/main.js' Extension type match -
Will Rewrite
```

If the conditions for rewriting a URL fail, the following messages are displayed:

```
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:URL:'http://
www.mynovell.com:9090/favicon.ico' - Did not match INCLUDE list,
Content-Type and Extension type
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:In
RewriterOption::shouldRewriteUrl, returning FALSE.
```

Check the Rewriter configuration. Ensure that your content type, extension type, and include URL list are valid.

42.2.7 Useful Linux Access Gateway Commands

Command	Description
<code>/etc/init.d/novell-vmc start</code>	Starts the Linux Access Gateway.
<code>/etc/init.d/novell-vmc stop</code>	Stops the Linux Access Gateway.
<code>/etc/init.d/novell-vmc status</code>	Displays the Linux Access Gateway status.
<code>/etc/init.d/novell-vmc debug-start</code>	Starts the Linux Access Gateway in the debug mode.
<code>/etc/init.d/novell-vmc restart</code>	Stops and starts the Linux Access Gateway.

42.2.8 Using the Linux Access Gateway Console

- ◆ To access the console, run the following command:

```
netcat localhost 2300
```

Press Enter at the `Please enter terminal type` prompt. This displays the Linux Access Gateway console screens.

- ◆ To access one specific console, enter the number of the required Linux Access Gateway from the Proxy Console Numbering.
- ◆ To return to the opening page of the console from other console page, press Esc and Enter.
- ◆ To return to the main console screen, press F1 and Enter.

42.2.9 COS Related Issues

You can view COS partition details either through YaST or through the nash prompt.

Using YaST

- 1 Log in as the `root` user.
- 2 At command prompt, enter the following command:

```
fdisk -l
```

The partition details are displayed. Check for COS partition details. Make sure that a partition is created with the Partition ID 68. The file system is created as type `unknown`.

Using nash

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

```
nash
```

- 2 At the nash shell prompt, enter the following command:

```
configure .current
```

- 3 Enter the following command:

```
vm scan
```

If the COS partition is already created, the details are displayed.

42.2.10 Memory Related Issues

This section contains the following memory related issues:

- ◆ [“Checking Memory Details and Related Information” on page 503](#)
- ◆ [“Checking Available Memory” on page 504](#)

Checking Memory Details and Related Information

Most of the information, including the memory details, can be accessed by entering the following command:

```
top
```

Ensure that the Linux Access Gateway does not occupy more than the percentage of the memory requirements you attribute.

Levels	Requirement
Lower Limit	5 Percent
Requirement for Access Gateway	500 MB
Upper Limit	80 percent
Default	20 percent

NOTE: ics_dyn occupies 20 percent of the total memory by default.

Checking Available Memory

As the root user, enter the following command at the bash prompt:

```
cat /proc/meminfo | grep MemTotal
```

42.2.11 Authorization and Identity Injection Issues

- ♦ [“Error: System Is Not Up Properly” on page 504](#)
- ♦ [“Identity Injection Failures” on page 505](#)

Error: System Is Not Up Properly

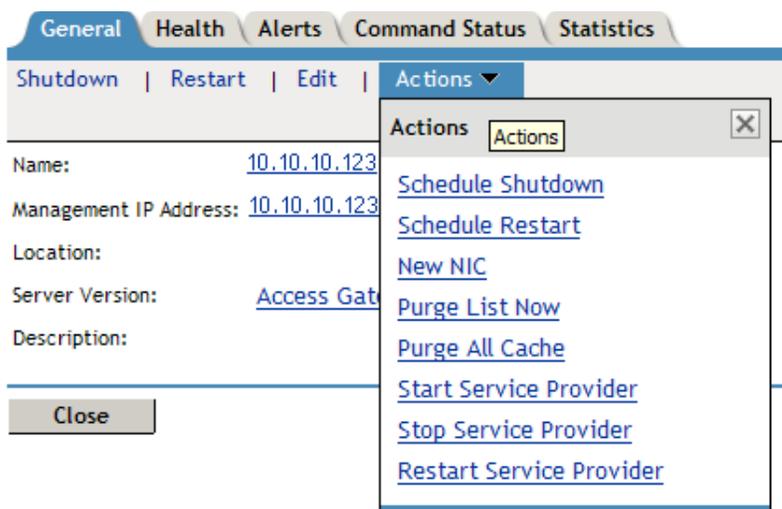
If you have already configured the identity injection policies, you might receive the following error while trying to send a browser request:

```
System is not up properly
```

This error indicates that the embedded service provider is down. Every identity injection policy has a policy ID, which is sent to the Access Gateway by the embedded service provider. If the embedded service provider is down, the Access Gateway does not get the policy ID, and the error is thrown.

To restart the embedded service provider:

- 1 In the Administration Console, select *Access Gateway* > *[Server Name]*.



- 2 Click *Actions* and select *Restart Service Provider*.
- 3 Click *OK*.

Identity Injection Failures

Identity injection might fail while trying to inject authentication headers because of improper policy configuration or because the Identity Server is not sending values to the Access Gateway.

Check the `/var/log/ics_dyn.log` file for the following error messages:

- ◆ Customer Header Injection Failed.
- ◆ Query String Injection Failed.
- ◆ Authentication Header Injection Failed

To receive help resolving identity injection failures, send the following information to Novell Support:

- ◆ Linux Access Gateway logs. For more information on how to get Linux Access Gateway log files, see [Section , “Linux Access Gateway Logs,” on page 495](#).
- ◆ Packet Capture. For more information on how to get packet capture, see [Section , “Packet Capture,” on page 497](#).

42.2.12 Form Fill Issues

Form Fill error messages are logged only if you use the Debug Build. If you want to log the error messages, change the log-level to LOG_DEBUG and reproduce the issue. For more information on changing log-level, see [“Changing Log Levels” on page 493](#).

Send the following information to the Novell Support:

- ◆ The `ics_dyn.log` file
- ◆ Packet capture. For more information, see [“Packet Capture” on page 497](#).

This section contains the following information about form fill issues:

- ◆ [“Form Fill Error Message” on page 506](#)
- ◆ [“Form Fill Failure Due to Incorrect Policy Configuration” on page 506](#)
- ◆ [“Browser Spinning Issues” on page 506](#)

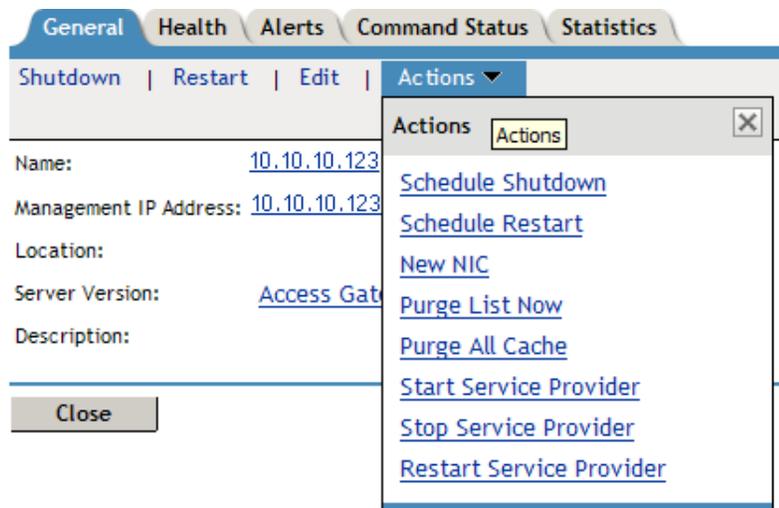
Form Fill Error Message

You might get the following errors when sending browser request:

- ◆ DataStore Error
- ◆ The service provider is not running at the moment. Please retry after a few seconds.

These errors indicate that the Access Gateway is not able to fetch the information such as, the user authentication information and data to fill the form, which are essential to process the browser request or is unable to save the information provided by the user, as the embedded service provider is down. In this case retry after a few seconds. If the error persists, restart the embedded service provider as follows:

- 1 In the Administration Console, select *Access Gateway > [Server Name]*.



- 2 Click *Actions* and select *Restart Service Provider*.
- 3 Click *OK*.

Form Fill Failure Due to Incorrect Policy Configuration

Form fill fails if the policy is not configured correctly. For configuration information, see [Chapter 32, “Creating Form Fill Policies,” on page 385](#).

Browser Spinning Issues

Authentication to the Web server fails if inappropriate data is filled in the form due to one of the following reasons:

- ◆ Shared Secrets are configured and the user provided incorrect data the Linux Access Gateway.

- ◆ Credential Profile - LDAP Attributes are configured and there is a mismatch between the user name used to authenticate to the Linux Access Gateway and the user name used to authenticate to the accelerated Web server.

When the Form Fill succeeds and the authentication to the webserver fails, the Web server redirects the browser to its authentication page again and again, if auto-submit is enabled. In such a situation, if there is no appropriate login-failure action configured in the policy, the browser spins endlessly.

If this happens, do the following:

- ◆ Kill the browser session. If you are unable to do this, run the following commands to restart the Linux Access Gateway:

```
/etc/init.d/novell-vmc stop
/etc/init.d/novell-vmc start
```

- ◆ If the issue is with Credential Profile - LDAP Attributes, then verify and create appropriate users in the Linux Access Gateway.
- ◆ If the issue is with Shared Secrets, delete the corresponding values from the Secret Store. If it is not possible to delete the value, modify the corresponding policy to use a different or a new Custom Attribute/Shared Secret Attribute. For more information on modifying the policy, see [Section 32.3, “Implementing Form Fill Policies,” on page 391](#).

42.3 Troubleshooting the NetWare Access Gateway

- ◆ [Section 42.3.1, “Additional Options during the Boot Process,” on page 507](#)
- ◆ [Section 42.3.2, “Setting the Date and Time at the Console,” on page 508](#)
- ◆ [Section 42.3.3, “Unlocking the NetWare Access Gateway Console,” on page 509](#)
- ◆ [Section 42.3.4, “Command Line Options,” on page 509](#)
- ◆ [Section 42.3.5, “Telnet Fails after Performing an Upgrade,” on page 509](#)
- ◆ [Section 42.3.6, “SSL Certificate Error with X.509 Authentication from NetWare Access Gateway,” on page 510](#)

42.3.1 Additional Options during the Boot Process

You can enter additional commands during the boot process to enable monitoring of the load process and local maintenance.

- 1 Boot the machine and wait for the following screen:

```
=====
Loading Bootstrap ...
Preparing to start NetWare ...

Press any key to Interrupt
=====
```

- 2 Press any key. The following menu appears:

```

=====
Default NetWare configuration file detected (CONFIG.NW) Contents:
-LS 1024 -CON"Booting Novell(R) Access Gateway 3.0" -L

Type:
S to start NetWare
P to specify additional starting parameters
H for help
Enter selection:
=====

```

3 Enter P, then the following parameter:

```
-NetWareOnly
```

4 To start NetWare[®], enter S.

The NetWare Access Gateway boots to the NetWare prompt, so you can do local maintenance.

5 (Optional) During the blue screen where all the modules are counted in the load process, enter one of the following keystrokes:

- ◆ To unlock this screen so you can see the loading process, press
SHIFT+CTRL+ALT+U
- ◆ To boot to the NetWare command line, press
SHIFT+CTRL+ALT+N

As a memory aid for the two key sequences, remember that U indicates unlock, and N indicates NetWare.

42.3.2 Setting the Date and Time at the Console

If you inadvertently set the date and time on the Access Gateway to a time before the certificates are valid, the Administration Console is denied access to the Access Gateway and can no longer interact with it. To correct this problem, you must reset the date and time at the Access Gateway console.

1 Unlock the console. (For instructions, see [Section 42.3.3, “Unlocking the NetWare Access Gateway Console,”](#) on page 509.)

2 Switch to the device manager screen.

3 Enter the following command:

```
set date [year=<yyyy>,< >] [month=<mm>,< >] [day=<dd>,< >] [time=<hh:mm:ss>]
```

Replace the variables with the following values:

<yyyy>	Replace with a four-digit value representing the current year, such as 2006.
<mm>	Replace with a two-digit value representing the current month, with 1 representing January and 12 representing December.
<dd>	Replace with a two-digit value, from 1 to 31, indicating the current day of the month.
<hh:mm:ss>	Replace hh with a two-digit value, from 1 to 24, indicating the current hour. Replace mm with a two-digit value, from 0 to 60, indicating the current minutes. Replace ss with a two-digit value, from 0 to 60, indicating the current seconds.

For example, to set the date to December 1, 2006 and the time to 10:10 am, enter the following:

```
set date year=2006, month=12, day=1, time=10:10:00
```

The `set date` command disables NTP. Use the Administration Console to enable it.

42.3.3 Unlocking the NetWare Access Gateway Console

Before you can enter NetWare commands or view the logger screen, you must unlock the console.

- 1 To unlock the console, enter

```
unlock
```

- 2 When prompted for a password, press Enter.

The console is now unlocked and the active screen is the device manager screen. From this screen you can enter device manager commands.

- 3 To switch to the logger screen or other NetWare screens, enter

```
debug
```

- 4 When prompted for a password, enter

```
proxydebug
```

- 5 To switch from the device manager screen, press Ctrl+Escape and enter the screen number.

42.3.4 Command Line Options

Access Manager has been designed to use the Administration Console for most management and configuration tasks. If you have created a group for your Access Gateways, Novell® highly recommends that you use the Administration Console for these tasks.

The Access Gateway does not push configuration changes to the Administration Console. As soon as you make a change at the Administration Console and save the change, the Administration Console pushes the change to the Access Gateway and wipes any changes that have been made manually with the command line interface. Various troubleshooting tips explain how to use various command line options; other than troubleshooting, you should have very little cause to use them.

The NetWare Access Gateway uses the SET command syntax for its command line options. You must unlock the console to gain access to the command line prompt. (See [Section 42.3.3, “Unlocking the NetWare Access Gateway Console,”](#) on page 509.)

To get a list of possible commands, enter the following command at the command line prompt:

```
help
```

To get help for a particular command, enter

```
help <command_name>
```

Replace `<command_name>` with the name of a command, such as `set`.

42.3.5 Telnet Fails after Performing an Upgrade

After performing an over-the-wire upgrade (OTWUG), Telnet does not allow you to connect. To correct this problem, enter the following command at the NetWare Access Gateway console:

```
clear adminacl serveraddress
```

42.3.6 SSL Certificate Error with X.509 Authentication from NetWare Access Gateway

If you set up an X.509 contract and use it to authenticate from NetWare Access Gateway, you might see an error generated for the SSL certificate, causing possible problems authenticating with certificates. This occurs during SSL re-negotiation between Tomcat and the Internet Explorer browser, and is possibly an IE bug. This error does not occur when using Firefox.

This section provides various troubleshooting scenarios that you might encounter while configuring SSL VPN.

- ◆ [Section 43.1, “Connecting Successfully to the Server,” on page 511](#)
- ◆ [Section 43.2, “SSL VPN Not Reporting,” on page 513](#)
- ◆ [Section 43.3, “Verifying SSL VPN Components,” on page 514](#)
- ◆ [Section 43.4, “Issues With Keep Alive,” on page 514](#)
- ◆ [Section 43.5, “Idle Connection,” on page 515](#)
- ◆ [Section 43.6, “Mozilla Firefox Browser Displays “X” Mark,” on page 515](#)
- ◆ [Section 43.7, “Unable to Contact the SSL VPN Server,” on page 515](#)
- ◆ [Section 43.8, “Unable to Get Authentication Headers,” on page 515](#)
- ◆ [Section 43.9, “The SSL VPN Connection Is Successful But There Is No Data Transfer,” on page 516](#)
- ◆ [Section 43.10, “Applications Are Not Getting SSLized from Terminal After Running the su Command,” on page 516](#)
- ◆ [Section 43.11, “Unable to Connect to SSL VPN Gateway,” on page 516](#)

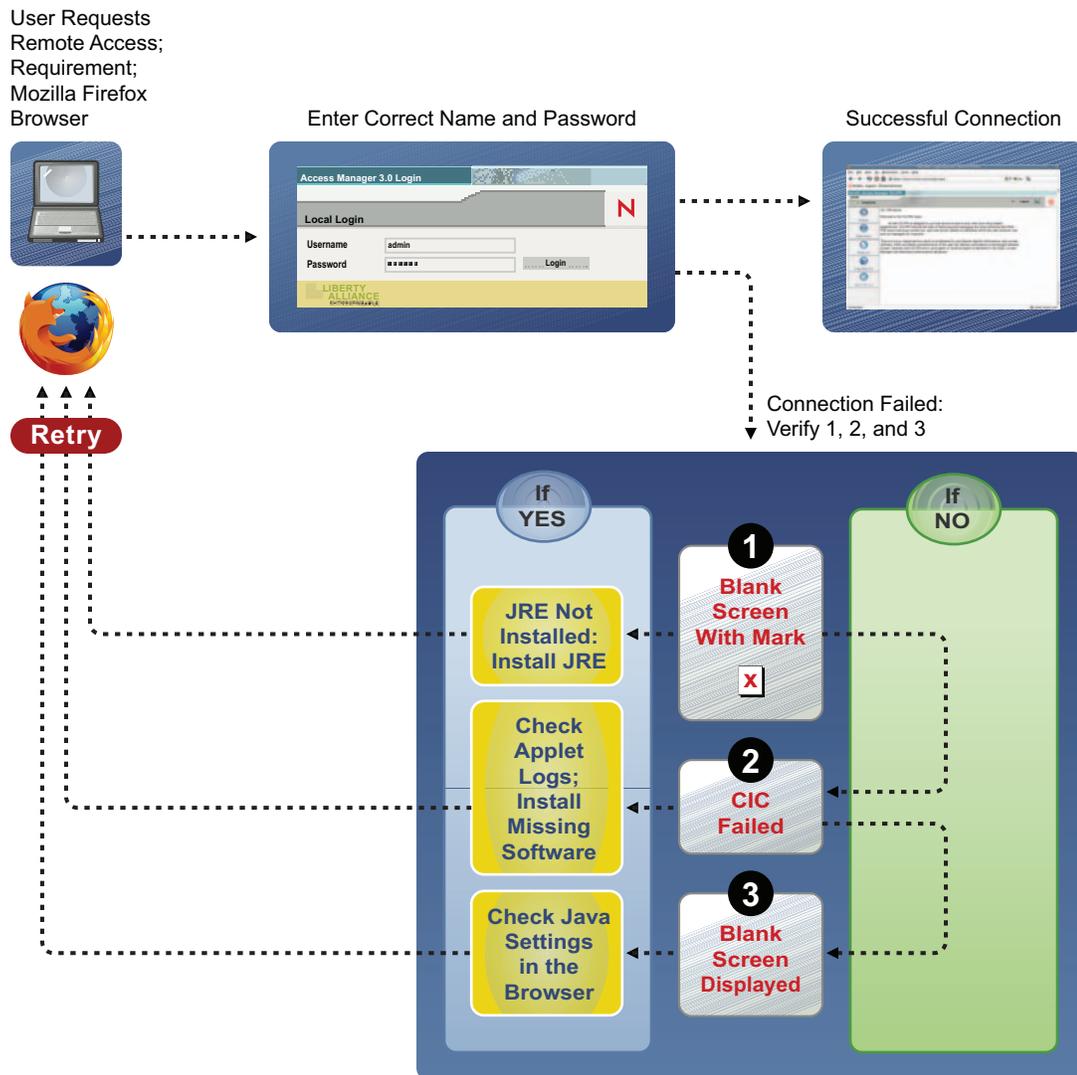
43.1 Connecting Successfully to the Server

You can access the protected resources using SSL VPN by authenticating to the proxy server. The proxy server loads the SSL VPN client on your browser. The following sections describe some of the problems that clients might encounter:

- ◆ [“Connection Problems with Mozilla Firefox” on page 512](#)
- ◆ [“Connection Problems with Internet Explorer” on page 513](#)

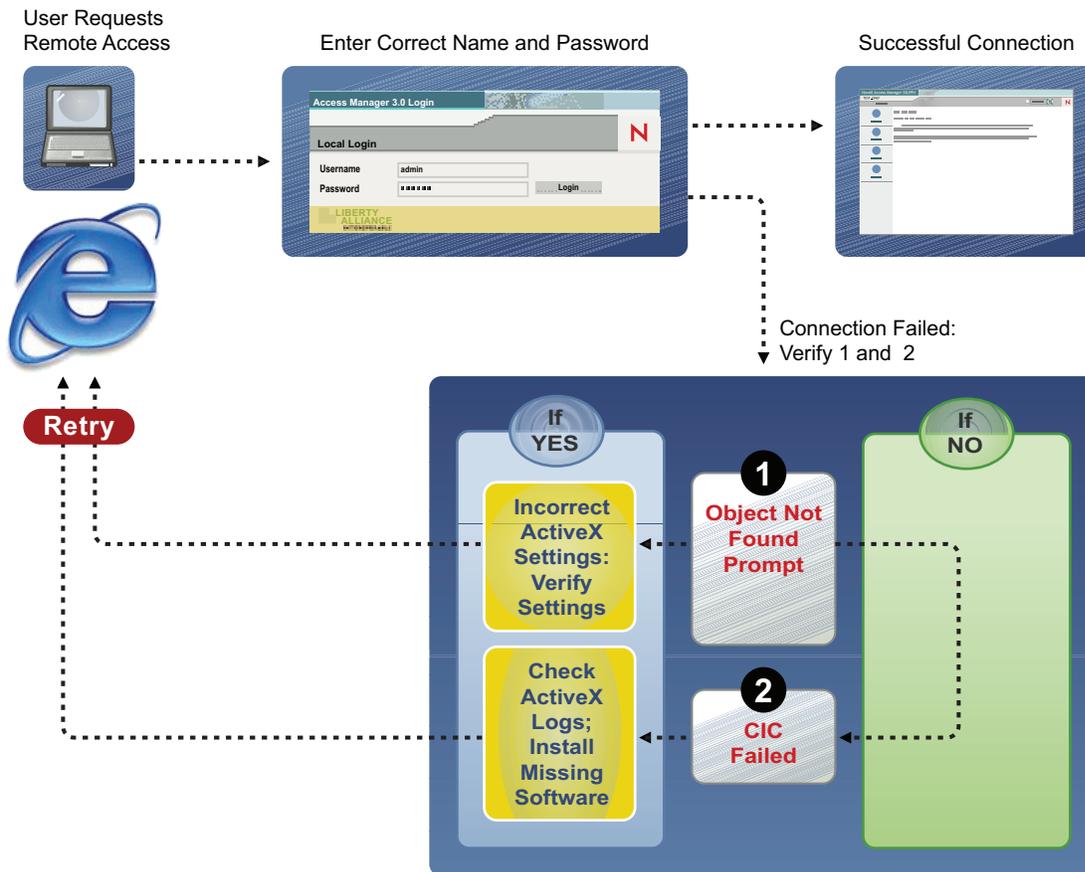
43.1.1 Connection Problems with Mozilla Firefox

Figure 43-1 Connecting to the Server Using Mozilla Firefox



43.1.2 Connection Problems with Internet Explorer

Figure 43-2 Connecting to the Server Using Internet Explorer



43.2 SSL VPN Not Reporting

If the SSL VPN is not reporting and if you have a backup of the `config.xml` and `config.txt` files, do the following:

- 1 In Administration Console, click *SSL VPNs*.
- 2 Select the SSL VPN server which has the problem, then click *Delete*.
- 3 Install the SSL VPN Gateway on a new server.
Specify the IP address of Administration Console and the public and private addresses of SSL VPN Gateway, at the time of installation.
After installation, the SSL VPN Gateway is imported into the Administrator Console. This gateway does not have the configuration of the old SSL VPN Gateway.
- 4 In Administration Console, select *Novell Access Manager > SSL VPN*. Select the newly added SSL VPN server, then click *Delete*.
- 5 Copy `config.xml` from the backup device to the following path:
`/etc/opt/novell/sslvpn/`
- 6 Copy `config.txt` from the backup device to the following path:

```
/var/opt/novell/tomcat4/webapps/sslvpn/WEB-INF/
```

7 As a root user, enter the following command to stop the SSL VPN server:

```
/etc/init.d/novell-sslvpn stop
```

8 Enter the following command to configure SSL VPN:

```
sslvpnc --configure
```

Specify the following information:

- ♦ IP Address of the Administration Console
- ♦ Public IP address of SSL VPN server
- ♦ Private IP Address of the SSL VPN server

9 Enter the following command to start the SSL VPN server:

```
/etc/init.d/novell-sslvpn start
```

10 Enter the following command to restart server communications:

```
/etc/init.d/novell-jcc restart
```

This imports the new SSL VPN server into the Administration Console with the configuration of the old SSL VPN server. If you had configured multiple private IP addresses for the old SSL VPN server, you can change it in the Administration Console.

43.3 Verifying SSL VPN Components

On the SSL VPN Server

To verify the function of all the SSL VPN components, use the commands listed in the table below:

Component	Command
Connection Manager	<code>pgrep connman</code>
Sock Daemon	<code>pgrep sockd</code>
Secure Tunnel	<code>pgrep stunnel</code>

On the SSL VPN Linux Client

Component	Command
Policy Resolver	<code>pgrep polresolver</code>
Secure Tunnel	<code>pgrep stunnel</code>

Windows Users

Check if the stunnel and polresolver processes are up and running.

43.4 Issues With Keep Alive

Possible Cause: Proxy might be down.

Action: Make sure that proxy is up and running.

Possible Cause: Connection to the server is interrupted in the midway. The link might be down.

Action: Retry the connection by getting the link up.

Possible Cause: The SSL VPN service is down.

Action: Make sure that your SSL VPN service is up and running. Verify all the SSL VPN processes. For more information see [Section 43.3, “Verifying SSL VPN Components,” on page 514.](#)

43.5 Idle Connection

If there is no data communication over the SSL VPN channel for more than 30 minutes, the connection becomes inactive.

In this case, reconnect.

NOTE: Do not use the Refresh, Back, or Forward options in the browser.

43.6 Mozilla Firefox Browser Displays “X” Mark

If you see an “X” mark on the top left corner of Mozilla Firefox while trying to access SSL VPN end user portal, it indicates that Java Run-time Environment (JRE) is not installed in the client machine.

Install Sun JRE 1.4 or above from <http://www.java.com/en/download/index.jsp>.

43.7 Unable to Contact the SSL VPN Server

If you encounter the error messages, `SSLVPN Gateway is in bad state` or `SSLVPN Gateway is not available` in the client browser, verify the following:

- ♦ **VPN Status:** Use the following command to check the status:

```
/etc/init.d/novell-sslvpn status
```

- ♦ **Error Status:** Check the status at `/var/log/messages` and `/var/log/stunnel.log`.
- ♦ **Check the Status of SSL VPN:** At the command prompt, enter the following command:

```
/etc/init.d/novell-sslvpn status
```

- ♦ **Message Log:** Check the `/var/log/messages` file for more information.

43.8 Unable to Get Authentication Headers

If the browser displays the `Unable to Get Authentication Headers` error while accessing the SSL VPN URL, check whether the custom HTTP headers required for SSL VPN are configured and enabled in the Access Gateway. In the Administration Console, click *Access Gateways* > *[Configuration Link]* > *[Name of Reverse Proxy]* > *[Name of SSL VPN Proxy Service]* > *[Name of SSL VPN Protected Resource]* > *Identity Injection*.

The `SSLVPN_Default` policy should be enabled. This policy injects an authentication header and two custom headers (`X-SSLVPN-PROXY-SESSION-COOKIE` and `X-SSLVPN-ROLE`).

43.9 The SSL VPN Connection Is Successful But There Is No Data Transfer

Possible Cause: The private address specified during the server configuration might be incorrect.

Action: Click *SSL VPNs > Edit > Gateway Configuration*, then check the private address configured. Make sure that this is the IP address of the private interface of SSL VPN server.

Possible Cause: If the SSLVPN server is behind a NAT, the external IP address specified during server configuration might be incorrect.

Action: Click *SSL VPNs > Edit > Gateway Configuration*. Make sure that the external IP address is configured to be the IP address of a NAT through which the external user on the Internet can access the SSL VPN server.

43.10 Applications Are Not Getting SSLized from Terminal After Running the `su` Command

If you are a Linux or a Macintosh user, do the following to access private network after running the `su` command in a terminal:

- ♦ Run the source `sslize_bashrc` file located in the home directory of the logged-in user, if you are using the `bash` shell.
- ♦ Run the source `sslize_tcshrc` file located in the home directory of the logged-in user, if you are using `tcsh` or `csch` shell.

If you have changed directory after running the `su` command, you must give the complete path to the above files.

43.11 Unable to Connect to SSL VPN Gateway

Possible Cause: Forward proxy is enabled in the Internet Explorer.

Action: Log in as the root user. In the Administration Console, select *Access Manager > Access Gateways > Edit > Reverse Proxy > Proxy List > Path-Based Multi-Homing > HTTP Options*. Select the *Allow Pages to Be Cached by the Browser* check box.

Appendixes

IX

The following section contains additional documentation and information about Novell® Access Manager and the Liberty Alliance.

- ◆ [Appendix A, “About Liberty and SAML 2.0,” on page 519](#)
- ◆ [Appendix B, “Certificates Terminology,” on page 521](#)
- ◆ [Appendix C, “Data Model Extension XML,” on page 523](#)
- ◆ [Appendix D, “Logging: Using the Custom Content Filter,” on page 529](#)
- ◆ [Appendix E, “Authentication Classes and Duplicate Common Names,” on page 535](#)

About Liberty and SAML 2.0



This section provides additional information about the Liberty Alliance and SAML 2.0.

- ♦ [Section A.1, “Liberty Alliance Resources,” on page 519](#)
- ♦ [Section A.2, “What's New in SAML 2.0?,” on page 520](#)

A.1 Liberty Alliance Resources

The Liberty Alliance is a consortium of business leaders with a vision to enable a networked world in which individuals and businesses can more easily conduct transactions while protecting the privacy and security of vital identity information.

To accomplish its vision, the Liberty Alliance established an open standard for federated network identity through open technical specifications. In essence, this open standard is a structured version of the Security Assertions Markup Language, commonly referred to as SAML, with the goal of accelerating the deployment of standards-based single sign-on technology.

For general information about the Liberty Alliance, visit the [Liberty Alliance Project \(http://www.projectliberty.org/index.php\)](http://www.projectliberty.org/index.php) Web site.

Liberty resources, including specifications, white papers, FAQs, and presentations can be found at the [Liberty Alliance Resources \(http://www.projectliberty.org/resources/index.php\)](http://www.projectliberty.org/resources/index.php) Web site.

The following table provides links to specific Liberty Alliance specifications:

Liberty Specification	Location
Liberty Alliance Project Overview	Liberty Alliance Project Overview (http://www.projectliberty.org/)
Liberty Whitepapers	Papers (http://www.projectliberty.org/liberty/resource_center/papers)
Identity Federation Specifications	Liberty ID-FF 1.2 Specification (http://www.projectliberty.org/resources/specifications.php#box1)
Web Service Framework Specifications	Liberty ID-WSF 1.1 Specifications (http://www.projectliberty.org/resources/specifications.php#box2a)
Liberty Profile Service Specifications	Liberty Alliance ID-SIS 1.0 Specifications (http://www.projectliberty.org/resources/specifications.php#box3)
Support Documentation (Glossary, Trust Model, Metadata Description, etc.)	Liberty Alliance Support Documents (http://www.projectliberty.org/resources/specifications.php#box4)
OASIS Standards (SAML)	Oasis Standards (http://www.oasis-open.org/specs/index.php#samlv2.0)

A.2 What's New in SAML 2.0?

SAML 2.0 is an XML-based framework for exchanging authentication and authorization information. Developed and managed by OASIS (Organization for the Advancement of Structured Information Standards), SAML provides a key to simplifying Internet identity management through single sign-on.

- ◆ **Pseudonyms:** An arbitrary name assigned by the identity provider to identify a user to a service provider. The identifier has meaning only in the context of the relationship between the relying parties. They can be a principal's e-mail or account name. Pseudonyms are a key privacy-enabling feature that inhibits collusion between multiple providers.
- ◆ **Metadata:** The SAML metadata specification defines how to express configuration and trust-related data to simplify SAML deployment. Metadata identifies the Identity Servers involved in performing single sign-on between trusted identity providers and service providers.
Metadata includes supported roles, identifiers, supported profiles, URLs, certificates, and keys. System entities must agree upon the data.
- ◆ **Encryption:** SAML permits attribute statements, name identifiers, or entire assertions to be encrypted. Encryption ensures that end-to-end confidentiality of these elements can be supported as needed.
- ◆ **Attribute profiles:** These simplify how you configure and deploy systems that exchange attribute data. Attribute profiles include:
 - ◆ **Basic attribute profile:** Supports string attribute names and attribute values drawn from XML schema primitive type definitions.
 - ◆ **X.500/LDAP:** Supports canonical X.500/LDAP attribute names and values.
 - ◆ **UUID attribute profile:** Supports using UUIDs as attribute names.
 - ◆ **XACML attribute profile:** Defines formats suitable for processing by XACML (Extensible Access Control Markup Language).

You can find detailed information about SAML 2.0 on the [OASIS Security Services \(SAML\) TC Web site \(http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security\)](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security).

Certificates Terminology

B

A public key certificate is a collection of information attached to an electronic message. It is used to verify that the user sending the message is who he or she claims to be. The following is a list of useful certificate terminology used in Access Manager:

CA: A certificate authority that signs a certificate.

Certificate: Public information about the entity identified by the certificate, including the public key. A certificate is signed. The signer of the certificate (a CA), if trusted, verifies the accuracy of the information in the certificate.

Certificate Chain: In addition to identifying a user, server, or computer, certificates can validate the identity and trustworthiness of other certificates. A certificate that asserts an identity is signed by a certificate that trusts the contents of the certificate it is signing. The signing certificate in turn can be signed by another certificate, which can be signed by another certificate, and so forth, thus forming a certificate chain. The last certificate in the certificate chain is referred to as the root certificate and is a self-signed certificate.

When a certificate or certificate chain is sent from one computer to another, the receiving computer examines the certificate chain to determine if it can be trusted. To verify certificate trust in a chain, the receiving computer examines its own configuration store to see if it contains a CA certificate that matches the root certificate of the certificate chain. If so, the receiver compares its copy of the certificate with the chain's root certificate to verify its authenticity.

Certificate Signing Request (CSR): Requesting a signed certificate is accomplished by sending a CSR to the CA. A CSR is created with information about the person or organization that desires the signed certificate. A public key is also generated and included in the CSR. A private key is also generated, but not included in the CSR.

When the CA receives the CSR, the CA uses it in combination with the CA's guidelines and practices to establish that the person or organization represented by the CSR is properly identified and authorized as the owner of the information in CSR. The CA creates and signs a certificate that the requesting person or organization can use. The signature of the CA in the certificate is what identifies to anyone who trusts the CA that the entity is who it claims to be. The signed certificate is delivered to its owner, who adds it to the keystore (usually the same keystore where the private key created with the original CSR resides).

Issuer: The CA that issues a certificate.

Intermediate Certificate: A subordinate certificate issued by the trusted root specifically to issue end-entity server certificates. The result is a certificate chain that begins at the trusted root CA, through the intermediate and ending with the SSL certificate issued to you. Using intermediate certificates adds more levels of security, but does not cause performance, installation, or compatibility issues.

Key: A certificate that also contains a private key.

Key Pair: Public and private keys generated by a cryptography system and used in combination with each other.

Keystore: A storage file containing keys, certificates, and trusted roots. Access Manager agents can access keystores to retrieve certificates, keys, and trusted roots as needed.

Local CA: The CA of the administration console's instance of eDirectory. Also known as Organizational CA.

Private Key: Used for authentication, data encryption/decryption, digital signing, and secure e-mail. One of the most common uses is sending and receiving digitally signed and encrypted e-mail using the S/MIME standard.

Public Key: The publicly distributed key.

Self-Signed Certificate: A certificate whose issuer is itself.

SSL Connections: When two computers connect and need to establish trust and a secure connection, certificates are exchanged and an encryption algorithm is established. Public keys shared in the exchanged certificates, as well as the associated private keys (which are not exchanged) are used as part of the encryption algorithm. After security is established, a secure SSL session is established and the two computers are able to communicate securely.

Trusted Certificate: The certificate of a known CA. These certificates are self-signed and are recognized as representing a CA that is trusted.

Trusted Root: The same as a trusted certificate. A trusted root provides the basis for trust in public key cryptography. Trusted roots enable security for SSL, secure e-mail, and certificate-based authentication. The Identity Server already has a list of trusted certificates installed. These certificates are for root CAs and, hence, are called "trusted roots."

Trust Store: A keystore containing only trusted roots. Intermediate CAs and end entity public certificates can be part of a trust store.

Data Model Extension XML

C

The data model for some Web services is extensible. You can enter XML definitions of data model extensions in this field. Data model extensions hook into the existing Web service data model at predefined locations.

All schema model extensions reside inside of a schema model extension group. The group exists to bind model data items together under a single localized group name and description. Schema model extension groups can reside inside of a schema model extension root or inside of a schema model extension. There can only be one group per root or extension. Each root is hooked into the existing Web service data model. Multiple roots can be hooked into the same location in the existing Web service data model. This conceptual model applies to the structure of the XML that is required to define data model extensions.

The high-level view of the data model extension XML is as follows:

```
<SchemaExtensions>
  <Root>
    <Group>
      <Extension>
        <Group>
          <Extension>...</Extension>
          <Extension>...</Extension>
          ...
        </Group>
      </Extension>
      <Extension>
        <ValueSet>
          <Value/>
          <Value/>
        </ValueSet>
      </Extension>
      ...
    </Group>
  </Root>
</SchemaExtensions>
```

C.1 Elements

The definition of the attributes for each data model extension XML element follows:

Root Element

parent: The unique identifier of the “hook point” in the Web service’s data model. These hook points are defined by the Web service data model schema. These unique identifiers represent the xpaths of each data item within the model schema. Possible values for the Parent attribute are listed in [Table C-1](#):

Table C-1 Root Element

Personal Profile	/pp:PP/pp:Extension /pp:PP/pp:CommonName/pp:Extension /pp:PP/pp:CommonName/pp:AnalyzedName/pp:Extension /pp:PP/pp:LegalIdentity/pp:Extension /pp:PP/pp:LegalIdentity/pp:VAT/pp:Extension /pp:PP/pp:LegalIdentity/pp:AltID/pp:Extension /pp:PP/pp:EmploymentIdentity/pp:Extension /pp:PP/pp:AddressCard/pp:Extension /pp:PP/pp:AddressCard/pp:Address/pp:Extension /pp:PP/pp:MsgContact/pp:Extension /pp:PP/pp:Facade/pp:Extension /pp:PP/pp:Demographics/pp:Extension
Employee Profile	/ep:EP/ep:Extension /ep:EP/ep:CorpCommonName/ep:Extension /ep:EP/ep:CorpLegalIdentity/ep:Extension /ep:EP/ep:CorpLegalIdentity/ep:VAT/ep:Extension /ep:EP/ep:CorpLegalIdentity/ep:AltID/ep:Extension
Open Profile	/op:OP/op:Extension /op:OP/op:CustomizableStringsop:Extension

package (required): The Java package name where all classes for this root are implemented. This includes resource description classes and data model instance classes. For example, com.novell.nids.profile.model.extensions.

resourceClass (required): The Java class name of the resource description class that is used to load all resources associated with this root. Because resource description class files are assumed to reside in the root's package, only the filename is needed. Resource description classes are Java classes that must be created by the person extending the model. You must also extend the com.novell.nidp.resource.NIDPResDesc class.

Group Element

resourceID: The resource ID of the display name of the group. This resource ID is assumed to be a key in the resource bundle supplied by the resource description class file associated with the containing root.

descriptionResourceID: The resource ID of the description of the group. This resource ID is assumed to be a key in the resource bundle supplied by the resource description class file associated with the containing root.

Extension Element

name (required): The name of the data model extension. This name must be the name of the XML element that will be used in the data model.

class (optional): The Java class name of the data model instance class. Because data model instance class files are assumed to reside in the root's package, only the filename is needed. If this attribute is omitted, then the value of the name attribute must be the instance class filename.

syntax: The syntax of this data model extension. Possible values are:

- ◆ String
- ◆ LocalizedString
- ◆ Container

format: (required if the syntax is *String* or *LocalizedString*)

The syntax of this data model extension. Possible values are:

- ◆ CaseIgnore
- ◆ CaseExtract
- ◆ URI
- ◆ URL
- ◆ Date
- ◆ DateNoYear
- ◆ CountryCode
- ◆ LanguageCode
- ◆ KeyInfo
- ◆ Number

upper: The upper bound of a numeric value. Use this attribute only if the format attribute value is Number. The value is a signed integer. If this attribute is omitted, the default value is `java.lang.Integer.MAX_VALUE`.

lower (optional): The lower bound of a numeric value. This attribute is only used if the format attribute value is Number. The value is a signed integer. If this attribute is omitted, the default value is `java.lang.Integer.MIN_VALUE`.

min (required): The cardinality of the XML element represented by this data model extension. It is the minimum number of elements allowed. The value is an unsigned integer. If this attribute is omitted, the default value is 0.

max (required): The cardinality of the XML element represented by this data model extension. It is the maximum number of elements allowed. The value is an unsigned integer. If this attribute is omitted, the default value is 1. The value UNBOUNDED may be used to indicate that there are no bounds.

namingClass: (required if syntax equals Container and max is UNBOUNDED)

The class that is used as the naming attribute for the container. The class must represent one of the immediate children of the container. This class will be used to name each instance of the container.

ValueSet Element

A ValueSet element contains a set of fixed values that a data model entry can contain. If a data model extension has a ValueSet, then the user interface to edit the value of that extension limits the user to these values. The ValueSet element has no attributes.

Value Element

A Value element represents a value in a ValueSet. It contains the actual value to be stored in the data model entry and the display name resource ID associated with the value.

resourceID (required): The resource ID of the display name of the value. This resource ID is assumed to be a key in the resource bundle supplied by the resource description class file associated with the containing root.

value (required): The value stored in the data model entry.

name (required): The name of the data model extension. This name must be the name of the XML element that is used in the data model.

C.2 Writing Data Model Extension XML

Data model extension XML must be defined in the namespace `novell:liberty:wsf:config:1:0:0` and that namespace must be defined on the SchemaExtensions element. Normally, the namespace prefix `wsfc` is used. An example of data model extension XML is:

```
<wsfc:SchemaExtensions xmlns:wsfc="novell:liberty:wsf:config:1:0:0">
  <wsfc:Root parent="/pp:PP/pp:Facade/pp:Extension"
    package="com.novell.nidp.liberty.wsf.idsis.ppservice.extensions"
    resourceClass="PPExtensionsResDesc">
    <wsfc:Group resourceId="PP.EXT.FC.GROUP"
      descriptionResourceId="PP.EXT.FC.GROUP.DESC">
      <wsfc:Extension name="AliasName"
        class="FacadeAliasName"
        syntax="String"
        format="CaseIgnore"
        resourceId="PP.EXT.FC.AliasName"
        min="0" max="1"/>
      <wsfc:Extension name="FavoriteURLs"
        class="FacadeFavoriteURLs"
        syntax="String"
        format="CaseExact"
        resourceId="PP.EXT.FC.FavoriteURLs" min="0" max="UNBOUNDED"/>
    </wsfc:Group> </wsfc:Root>
  <wsfc:Root parent="/pp:PP/pp:Demographics/pp:Extension"
    package="com.novell.nidp.liberty.wsf.idsis.ppservice.extensions"
    resourceClass="PPExtensionsResDesc">
    <wsfc:Group resourceId="PP.EXT.DM.GROUP"
      descriptionResourceId="PP.EXT.DM.GROUP.DESC">
      <wsfc:Extension name="EyeColor"
        class="DemographicsEyeColor"
        syntax="String" format="URI"
        resourceId="PP.EXT.DM.EyeColor"
        min="0"
        max="UNBOUNDED">
      <wsfc:ValueSet>
      <wsfc:Value resourceId="PP.EXT.DM.HC.Blue" value="urn:pp:dm:blue"/>
      <wsfc:Value resourceId="PP.EXT.DM.HC.Brown" value="urn:pp:dm:brown"/>
      <wsfc:Value resourceId="PP.EXT.DM.HC.Green" value="urn:pp:dm:green"/>
      <wsfc:Value resourceId="PP.EXT.DM.HC.Gray" value="urn:pp:dm:gray"/>
      <wsfc:Value resourceId="PP.EXT.DM.HC.Hazel" value="urn:pp:dm:hazel"/>
    </wsfc:ValueSet>
    </wsfc:Group> </wsfc:Root>
  </wsfc:SchemaExtensions>
```

```

</wsfc:ValueSet>
</wsfc:Extension>
</wsfc:Group>
</wsfc:Root>
<wsfc:Root parent="/pp:PP/pp:Extension"
  package="com.novell.nidp.liberty.wsf.idsis.ppservice.extensions"
  resourceClass="PPExtensionsResDesc">
<wsfc:Group resourceId="PP.EXT.AU.GROUP"
  descriptionResourceId="PP.EXT.AU.GROUP.DESC">
<wsfc:Extension name="Automobile"
  class="Automobile"
  syntax="Container"
  resourceId="PP.EXT.Automobile"
  min="0"
  max="UNBOUNDED"
  namingClass="AutomobileLicensePlate">
<wsfc:Group resourceId="PP.EXT.AU.DETAILS.GROUP"
  descriptionResourceId="PP.EXT.AU.DETAILS.GROUP.DESC">
<wsfc:Extension name="AutomobileModel"
  class="AutomobileModel"
  syntax="String"
  resourceId="PP.EXT.AU.Model"
  min="0"
  max="1"/>
<wsfc:Extension name="AutomobileMake"
  class="AutomobileMake"
  syntax="String"
  format="CaseIgnore"
  resourceId="PP.EXT.AU.Make"
  min="0"
  max="1"/>
<wsfc:Extension name="AutomobileLicensePlate"
  class="AutomobileLicensePlate"
  syntax="String"
  format="CaseIgnore"
  resourceId="PP.EXT.AU.LicensePlate"
  min="0" max="1"/>
</wsfc:Group>
</wsfc:Extension>
</wsfc:Group>
</wsfc:Root>
</wsfc:SchemaExtensions>

```


Logging: Using the Custom Content Filter

D

The Custom Content Filter allows you to focus trace content on a specific section of the system where you suspect a problem exists. The filter is an XML document that specifies which trace logging content to send to the trace logger.

You can limit the trace logging to one or more Java class files, or to one or more Java packages, or to one or more Novell-defined thread identifiers. A thread identifier correlates to a group of events that logically should be logged together. In the XML, you can include and exclude content from an entity in the trace log. If trace logging content becomes too verbose, you can exclude Java classes, Java packages, or thread identifiers to reduce the irrelevant logged data.

- ◆ [Section D.1, “Custom Content Filter XML Syntax,” on page 529](#)
- ◆ [Section D.2, “Examples of Custom Content Filter XML,” on page 530](#)
- ◆ [Section D.3, “Custom Content Filter Thread Identifiers,” on page 533](#)

D.1 Custom Content Filter XML Syntax

The following text provides the XML syntax.

```
<Trace flushFrequency="immediate">
  <Classes>
    <Class>...</Class>
    <Class exclude="false">...</Class>
    <Class exclude="true">...</Class>
  </Classes>
  <Packages>
    <Package>...</Package>
    <Package exclude="false">...</Package>
    <Package exclude="true">...</Package>
  </Packages>
  <Threads>
    <ThreadId>...</ThreadId>
    <ThreadId exclude="false">...</ThreadId>
    <ThreadId exclude="true">...</ThreadId>
  </Threads>
</Trace>
```

The `<Trace>` element contains three sub-sections called `<Classes>`, `<Packages>`, and `<Threads>`. Each sub-section is optional and can be omitted. The `<Trace>` element has a single attribute called `flushFrequency` that controls the frequency at which trace log data is flushed out to the file. Keep the value of this attribute set to `immediate` so that data is flushed as soon as possible. When in debugging mode, which is the only recommended use for trace logging, immediate flushing is preferred.

The <Classes> element contains zero or more <Class> elements. Each <Class> element defines a single Java class that is included or excluded in the trace log output. The name of the Java class must include the complete Java package and class name, while omitting the .java extension.

The <Packages> element contains zero or more <Package> elements. Each <Package> element defines a single Java package that is included or excluded in the trace log output. The inclusion or exclusion applied to this Java package also applies to all of this package's child packages.

The <Threads> element contains zero or more <ThreadId> elements. Each <ThreadId> element defines a single Novell-defined thread identifier that is included or excluded in the trace log output.

The elements <Class>, <Package>, and <ThreadId> have a single attribute exclude="true/false". This attribute marks the associated Java class, Java package, or Thread Identifier as being included in the trace log output or as being excluded from the trace log output. If this attribute is not present, the default is "false". Meaning, the default is to include the associated entity in the trace logging output.

The elements <Class>, <Package>, and <ThreadId> accept the single character * as the text value of the element. This wildcard character means "all entities of this type." This wildcard character can be used only as a single character. It cannot be combined with other strings in an attempt to form a wildcard string. For example <Class>*/</Class> causes all Java classes to be included in the trace log output. The following example is invalid: <Class>com.novell.nidp.NIDP*/</Class>.

D.2 Examples of Custom Content Filter XML

This section provides examples of the Custom Content Filter XML.

- ◆ [Section D.2.1, "Example One," on page 530](#)
- ◆ [Section D.2.2, "Example Two," on page 530](#)
- ◆ [Section D.2.3, "Example Three," on page 532](#)

D.2.1 Example One

The following Custom Content Filter causes all Java classes and all thread identifiers to be included in the trace log output. This filter traces everything. Care must be taken when using this filter because large amounts of data is logged, and the performance of the system degrades substantially.

```
<Trace flushFrequency="immediate">
  <Classes>
    <Class>*/</Class>
  </Classes>
  <Threads>
    <ThreadId>*/</ThreadId>
  </Threads>
</Trace>
```

D.2.2 Example Two

The following Custom Content Filter causes all Java classes, except com.novell.nidp.common.authority.ldap.jndi.JNDIUserStoreReplicaConnection, and all thread identifiers, to be included in the trace log output. The <Packages> sub-

section is not needed because this filter already includes all Java classes. Including all Java packages as well would only be redundant.

```
<Trace flushFrequency="immediate">
  <Classes>
    <Class>*</Class>
    <Class
exclude="true">com.novell.nidp.common.authority.ldap.jndi.JNDIUse
StoreReplicaConnection</Class>
  </Classes>
  <Threads>
    <ThreadId>*</ThreadId>
  </Threads>
</Trace>
```

Specific Java classes can be excluded if irrelevant information is bogging down the log file. To determine how to filter out unwanted entries from the trace log content, perform the following steps:

- 1** Locate the header for the entries that you want to exclude from the trace log.

Each trace entry in the log file has a header that names the Java class where the trace entry originated. An example header is:

```
NIDP TRACE LOG Method:
com.novell.nidp.liberty.wsf.WSFFramework.initialize().
```

- 2** Extract the Java class or Java package name from the header.

In the above example, the Java class is

```
com.novell.nidp.liberty.wsf.WSFFramework
```

and the Java package is `com.novell.nidp.liberty.wsf`. The `.initialize()` method is inside the `WSFFramework` class. You do not need to extract the method name. You can ignore it.

- 3** Add a `<Class>` or `<Package>` entry to the Custom Content Filter XML that excludes the Java class or Java package.

Excluding the entire package removes trace log entries from other Java class files in the same package. Using the preceding example, if you want to exclude trace log entries from only the Java class `com.novell.nidp.liberty.wsf.WSFFramework` entry you would add

```
<Class
exclude="true">com.novell.nidp.liberty.wsf.WSFFramework</Class>
```

to the `<Classes>` element sub-section. If you want to exclude the entire package, you add

```
<Package exclude="true">com.novell.nidp.liberty.wsf</Package>
```

to the `<Packages>` element sub-section.

If you follow the preceding steps to exclude a Java class or package, and the trace log entry still gets logged, this is because the log entry is being logged based on a thread identifier. Logs based on thread identifier do not consider the Java class or package when deciding if the trace log should happen. In this case, determine which aspect of the product the trace log entry pertains to, and attempt to match it with a thread identifier. (Thread identifiers are explained below.) Then add a

```
<ThreadId exclude="true">[thread id name]</ThreadId>
```

line to the <Threads> sub-section. Or, if you want to remove all trace logs associated with all thread identifiers, simply remove the <Threads> subsection.

D.2.3 Example Three

The following Custom Content Filter example includes all packages except for the explicitly excluded package: `com.novell.nidp.common.authority.ldap`.

```
<Trace flushFrequency="immediate">
  <Packages>
    <Package>*</Package>
    <Package exclude="true">com.novell.nidp.common.authority.ldap</
Package>
  </Packages>
  <Threads>
    <ThreadId>*</ThreadId>
    <ThreadId exclude="true">tIdPPModelTokenCreate</ThreadId>
    <ThreadId exclude="true">tIdEPModelTokenCreate</ThreadId>
    <ThreadId exclude="true">tIdCUPModelTokenCreate</ThreadId>
    <ThreadId exclude="true">tIdAPModelTokenCreate</ThreadId>
    <ThreadId exclude="true">tIdCPModelTokenCreate</ThreadId>
    <ThreadId exclude="true">tIdWSFSchemaExtensions</ThreadId>
    <ThreadId exclude="true">tIdRequestResponse</ThreadId>
    <ThreadId exclude="true">tIdConfiguration</ThreadId>
    <ThreadId exclude="true">tIdLdapJndiConnShare</ThreadId>
    <ThreadId exclude="true">tIdLdapJndiOperations</ThreadId>
    <ThreadId exclude="true">tIdLdapJndiOperationStats</ThreadId>
    <ThreadId exclude="true">tIdLdapJndiSearch</ThreadId>
    <ThreadId exclude="true">tIdLdapJndiGetObject</ThreadId>
    <ThreadId exclude="true">tIdLdapJndiModifyObject</ThreadId>
    <ThreadId exclude="true">tIdLdapJndiCreateConnection</ThreadId>
    <ThreadId exclude="true">tIdLdapJndiCloseConnection</ThreadId>
    <ThreadId exclude="true">tIdLdapJndiPerUserStoreStats</ThreadId>
    <ThreadId exclude="true">tIdCBPing</ThreadId>
    <ThreadId exclude="true">tIdCBRetirement</ThreadId>
    <ThreadId exclude="true">tIdCBLogouts</ThreadId>
    <ThreadId exclude="true">tIdCBStats</ThreadId>
    <ThreadId exclude="true">tIdHealthCheck</ThreadId>
  </Threads>
</Trace>
```

This example shows the filter for one of the most verbose packages. The example shows that you have chosen to exclude the LDAP package because the issue under investigation was not related to LDAP. This example goes on to include all thread identifiers, and then excludes each thread identifier. Thus, all thread identifiers are excluded by this filter. However, this example shows the complete list of thread identifiers. Therefore, using this filter would require you to only change `exclude="true"` to `exclude="false"` in order to include the relevant thread identifier.

D.3 Custom Content Filter Thread Identifiers

A thread identifier names a sequence of events that can be traced as a group. The events logged for a given thread id can be a sequence of events performed to accomplish a task, or it can be a group of similar events.

The Web Service Framework includes several Web services. Each Web service has a data model associated with it. As the identity provider (IDP) or service provider (SP) initializes, the data model builds the set of data items included in each Web service. Including all of the following five thread identifiers logs the complete data model of the indicated service. This log is written once at startup and once each time the identity server application is restarted.

- ◆ **tIdPPModelTokenCreate:** Logs the data model of the Personal Profile Service.
- ◆ **tIdEPModelTokenCreate:** Logs the data model of the Employee Profile Service.
- ◆ **tIdCUPModelTokenCreate:** Logs the data model of the Custom Profile Service.
- ◆ **tIdAPModelTokenCreate:** Logs the data model of the Role Profile Service.
- ◆ **tIdCPModelTokenCreate:** Logs the data model of the Credential Profile Service.

As each Web Service data model is built, you can configure model extensions (or schema extensions) to add additional data items to the model. You can configure the model extensions for each Web Service by adding XML to the edit box on each Web Service's Details: General Settings page (*Identity Servers > [Configuration] > Liberty > Web Service Provider > [Profile] > Details*).

The following thread identifier logs each new entry that is added to the model. Also, all errors that occur from attempting to add to the model are logged.

- ◆ **tIdWSFSchemaExtensions:** Logs successful and failed additions to all Web service data models.

One of the best ways to debug the IDP/SP is to log the HTTP requests and HTTP responses that are handled by the IDP/SP. The following thread identifier logs the requests and responses for all subsystems. The HealthCheck request is not logged under this thread identifier because it might become verbose and regularly interferes with locating pertinent data. Therefore, the HealthCheck request is only logged if the tIdHealthCheck thread identifier is included.

- ◆ **tIdRequestResponse:** Logs the requests and responses for all subsystems.

As the IDP/SP is initializing after a startup or a reconfigure, the configuration is applied to the IDP/SP. The following thread identifier logs the configuration data that is used to initialize the IDP/SP.

- ◆ **tIdConfiguration:** Logs the versions of various sub-components used in the system. Also logs the details of each Web service.

The IDP/SP includes an LDAP operations sub-system that handles all communications with the LDAP trust/configuration database and LDAP user stores. This subsystem maintains connection pools for general purpose administrative level LDAP operations and for user LDAP operations. A typical administrative LDAP operation is to read a user's identity information from the directory. A typical user LDAP operation is to bind a user to a directory object to prove that a name/password combination is valid.

As the system is pushed to its limits, the LDAP operations sub-system can determine that it needs more connections devoted to administrator operations. Thus, user connections from the user

connection pool will be shared with the administrator connection pool. The same can happen in the opposite direction. The following thread identifiers log data about the current state of the LDAP operations sub-system and the LDAP operations it performs. The LDAP operations sub-system is the most verbose logging section of the IDP/SP. Thus, there is a different thread identifier for each basic LDAP operation. Be careful when including all of these thread identifiers at the same time because large amounts of data will be logged.

- ♦ **tIdLdapJndiConnShare:** Logs details about how the LDAP operations sub-system shares the connection between user and administrator connection pools.
- ♦ **tIdLdapJndiOperations:** Logs details associated with the LDAP operations sub-system.
- ♦ **tIdLdapJndiOperationStats:** Periodically logs the LDAP operations sub-system statistics
- ♦ **tIdLdapJndiSearch:** Logs details about all LDAP Object Search operations.
- ♦ **tIdLdapJndiGetObject:** Logs details about all LDAP Object Get operations.
- ♦ **tIdLdapJndiModifyObject:** Logs details about all LDAP Object Modify operations.
- ♦ **tIdLdapJndiCreateConnection:** Logs details about all LDAP Connection Create operations.
- ♦ **tIdLdapJndiCloseConnection:** Logs details about all LDAP Connection Close operations.
- ♦ **tIdLdapJndiPerUserStoreStats:** Periodically logs generic statistics about each user store.

The Session Broker is a component of the embedded SP (ESP) that works closely with the Access Gateway to monitor user authentications within a clustered environment. As users log in to the system, their login information is registered in the Session Broker of the ESP. The Session Broker communicates with other members of the cluster to share user session information. Therefore, successful communication between cluster members is vital to a properly functioning system.

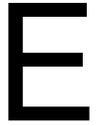
The session broker is also responsible for timing-out or retiring authentication data that has been unused for too long. When an authentication data item times out, or when the user logs out of the system, the session broker is responsible to send a message to each Access Gateway in the cluster to tell the Access Gateway that the logout has taken place, and that user's authentication data must be removed.

- ♦ **tIdCBPing:** Logs a periodic ping that displays all cluster members to which successful communication is available. The computer initiating the ping is not shown in the list.
- ♦ **tIdCBRetirement:** Logs details about user session data that is being retired.
- ♦ **tIdCBLogouts:** Logs details about the messages sent to the Access Gateway indicating that a user session has timed out or was logged out.
- ♦ **tIdCBStats:** Logs generic statistics about the session broker.

A periodic health check of the system can be configured. The following thread identifier logs the details about the system items checked during the health check. If health reports an error and the administrator is not sure why the error is happening, then this health check log detail can provide more information.

- ♦ **tIdHealthCheck:** Logs details about the health check.

Authentication Classes and Duplicate Common Names



If users have the same common name and exist in different containers under the same authentication search base, one or more attributes in addition to the common name must be configured for authentication to uniquely identify the user. You can set up an authentication class to handle duplicate common names.

- 1 Select either the name/password or secure name/password class.
- 2 Add two properties to the class:
 - ♦ **Query:** The value of the Query attribute needs to be a valid LDAP query string. Field names from the JSP login form can be used in the LDAP query string as variables for LDAP attribute values. The variables must be enclosed between two % characters. For example, `(&(objectclass=person)(cn=%Ecom_User_ID%)(mail=%Ecom_Email%))` queries for an object of type person that contained a common name equal to the Ecom_User_ID field from the specified JSP form and mail equal to the Ecom_Email field from the same JSP form.
 - ♦ **JSP:** The JSP property value needs to be the name of a new `jsp` file that includes all the needed fields for the Query property. The value of this attribute does not include the `.jsp` extension of the file. For example, if you create a new `.jsp` file named `login2.jsp`, then the value of the JSP property is `login2`.