NetIQ

Access Gateway Guide
Access Manager 4.0 Appliance SP1

May 2014

Contents

3 Server Configuration Settings 81

4 Access Gateway Maintenance 109

About NetIQ Corporation

We are a global, enterprise software company, with a focus on the three persistent challenges in your environment: Change, complexity and risk—and how we can help you control them.

Our Viewpoint

Adapting to change and managing complexity and risk are nothing new

In fact, of all the challenges you face, these are perhaps the most prominent variables that deny you the control you need to securely measure, monitor, and manage your physical, virtual, and cloud computing environments.

Enabling critical business services, better and faster

We believe that providing as much control as possible to IT organizations is the only way to enable timelier and cost effective delivery of services. Persistent pressures like change and complexity will only continue to increase as organizations continue to change and the technologies needed to manage them become inherently more complex.

Our Philosophy

Selling intelligent solutions, not just software

In order to provide reliable control, we first make sure we understand the real-world scenarios in which IT organizations like yours operate — day in and day out. That's the only way we can develop practical, intelligent IT solutions that successfully yield proven, measurable results. And that's so much more rewarding than simply selling software.

Driving your success is our passion

We place your success at the heart of how we do business. From product inception to deployment, we understand that you need IT solutions that work well and integrate seamlessly with your existing investments; you need ongoing support and training post-deployment; and you need someone that is truly easy to work with — for a change. Ultimately, when you succeed, we all succeed.

Our Solutions

- Identity & Access Governance
- Access Management
- Security Management
- Systems & Application Management
- Workload Management
- Service Management

Contacting Sales Support

For questions about products, pricing, and capabilities, contact your local partner. If you cannot contact your partner, contact our Sales Support team.

Worldwide:	www.netiq.com/about_netiq/officelocations.asp
United States and Canada:	1-888-323-6768
Email:	info@netiq.com
Web Site:	www.netiq.com

Contacting Technical Support

For specific product issues, contact our Technical Support team.

Worldwide:	www.netiq.com/support/contactinfo.asp
North and South America:	1-713-418-5555
Europe, Middle East, and Africa:	+353 (0) 91-782 677
Email:	support@netiq.com
Web Site:	www.netiq.com/support

Contacting Documentation Support

Our goal is to provide documentation that meets your needs. If you have suggestions for improvements, click **Add Comment** at the bottom of any page in the HTML versions of the documentation posted at www.netiq.com/documentation. You can also email Documentation-Feedback@netiq.com. We value your input and look forward to hearing from you.

Contacting the Online User Community

Qmunity, the NetIQ online community, is a collaborative network connecting you to your peers and NetIQ experts. By providing more immediate information, useful links to helpful resources, and access to NetIQ experts, Qmunity helps ensure you are mastering the knowledge you need to realize the full potential of IT investments upon which you rely. For more information, visit http://community.netiq.com.

About this Book and the Library

.This guide is intended to help you understand and configure all the Access Gateway features. The term Access Gateway is used to refer to Access Gateway Service in the Access Manager appliance product.

Before proceeding, you should be familiar with the *NetIQ Access Manager Appliance 4.0 SP1 Installation Guide* and the *NetIQ Access Manager Appliance 4.0 SP1 Setup Guide*, which provide information about setting up the Access Manager system.

The *Installation Guide* provides an introduction to NetIQ Access Manager Appliance and describes the installation and upgrade procedures.

Intended Audience

This book 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/TLS)
- Hypertext Transfer Protocol (HTTP and HTTPS)
- Uniform Resource Identifiers (URIs)
- Domain Name System (DNS)
- Web Services Description Language (WSDL)

Other Information in the Library

For information about other Access Manager devices and features, see the following:

- *NetIQ Access Manager Appliance 4.0 SP1 Administration Console Guide*
- *NetIQ Access Manager Appliance 4.0 Identity Server Guide*
- *NetIQ Access Manager Appliance 4.0 SP1 Policy Guide*
- *NetIQ Access Manager Appliance 4.0 SSL VPN Server Guide*
- *NetIQ Access Manager Appliance 4.0 SP1Event Codes*

NOTE: Contact namsdk@netiq.com for any query related to Access Manager SDK.

1 Configuring the Access Gateway for SSL and Other Security Features

This section describes how the Access Gateway can use SSL in its interactions with other Access Manager components, how you can enable SSL between an Access Gateway and these components, and how you can use other options to increase security:

1.1 Using SSL on the Access Manager Appliance Communication Channels

You can configure the Access Manager Appliance to use SSL in its connections to the Identity Server, to the browsers, and to its Web servers. Figure 1-1 illustrates these communication channels.

Figure 1-1 *Setting Up SSL for the Access Gateway Communication Channels*
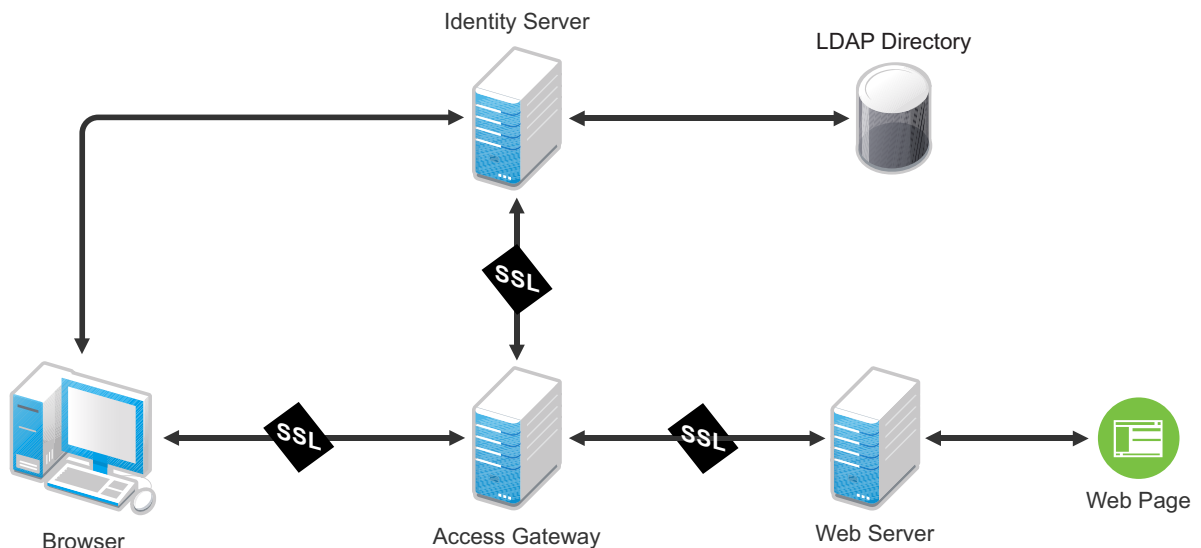
This section only describes how to set up SSL for the Access Gateway communication channels. The Identity Server needs to be configured for SSL before the Access Gateway can be configured for SSL.

When a 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 certificate authority (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.

- You should enable at least SSL if the Access Gateway is injecting authentication credentials into HTTP headers.
- Mutual SSL is probably not needed if you have configured the Web servers so that they can only accept connections with the Access Gateway.

1.2   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 "*Creating a Locally Signed Certificate*" in the *NetIQ Access Manager Appliance 4.0 SP1 Administration Console Guide*.
- You can create a certificate signing request (CSR), send it to an external CA, then import the returned certificates into Access Manager. See "Generating a Certificate Signing Request" and "Importing Public Key Certificates (Trusted Roots)" in the *NetIQ Access Manager Appliance 4.0 SP1 Administration Console Guide*.

1.2.1 Prerequisites for SSL Communication between the Identity Server and Access Manager Appliance

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 have not imported the public certificate of this CA into the trusted root store of the Identity Server, do so now. For more information, see "Importing Public Key Certificates (Trusted Roots)" in the *NetIQ Access Manager Appliance 4.0 SP1 Administration Console Guide*.

2 To add the public certificate to the Access Gateway:

   2a In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Service Provider Certificates** > **Trusted Roots**

   2b In the Trusted Roots section, click **Add**.

   2c Click the **Select trusted root(s)** icon, select the public certificate of the CA that signed the Identity Server certificates, then click **OK**.

   2d Specify an alias, then click **OK** twice.

3 To apply the changes, click **Close**, then click **Update** on the Access Manager Appliance page.

1.2.2 Prerequisites for SSL Communication between the Access Gateway and 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 about 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.

1.3 Configuring SSL Communication with Browsers and the Identity Server

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]**.

Reverse Proxy: doc1 - ag-206

Cluster Member: ag18 ▼

Listening Address(es): ☑ 10.10.15.18

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: jwilson1_provo_novell_com 🔲

Auto-generate Key

Auto-Import Embedded Service Provider Trusted Root

Non-Secure Port: * 80　(Redirected to Secure Port)

Secure Port: * 443　(Used for Trusted IDS Encryption, HTTPS Listening)

2 Configure the reverse proxy for SSL by filling 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 available only for the reverse proxy that 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 this option 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 browsers. Select one of the following methods:

- ◆ To auto-generate a certificate key by 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 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 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.

IMPORTANT: If you select not to redirect HTTP requests (port 80) and your Access Gateway has only one IP address, do not use port 80 to configure another reverse proxy. Although it is not used, it is reserved for this reverse proxy.

Secure Port: Specifies the port on which to listen for HTTPS requests (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.

5 Click **OK**.

6 On the Server Configuration page, click **OK**.

7 On the Access Gateways page, click **Update** > **OK**.

The Embedded Service Provider is restarted during the update.

8 (Conditional) The Identity Server is automatically updated to use the new SSL configuration. If the update is not started and an update is flagged, click Identity Servers > Update.

9 Verify that the trusted relationship between the Identity Server and the Access Gateway has been reestablished.

    9a Enter the URL to a protected resource on the Access Gateway.

    9b Complete one of the following:

        ◆ If you are prompted for login credentials, enter them. The trusted relationship has been reestablished.

        ◆ If you receive a 100101043 or 100101044 error, the trusted relationship has not been established. For information about how to solve this problem, see "Troubleshooting 100101043 and 100101044 Liberty Metadata Load Errors" in the *NetIQ Access Manager Appliance 4.0 Identity Server Guide*.

1.4 Configuring SSL between the Proxy Service and the Web Servers

SSL must be enabled between the Access Gateway and browsers before you can enable it between the Access Gateway and its Web servers.

1 In the Administration Console, click **Devices** > **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 browsers and the Access Gateway. See Section 1.3, "Configuring SSL Communication with Browsers and the Identity Server," on page 14 and select the **Enable SSL between Browser and Access Gateway** field.

3 (Optional) Set up mutual authentication so that the Web server can verify the proxy service certificate:

    3a Click the **Select Certificate** icon,

    3b 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. For instructions, see your Web server documentation.

4 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
WebSphere	9080	9443
JBoss	8080	8443

5 To save your changes to browser cache, click **OK**.

6 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

1.5 Enabling Secure Cookies

The Access Gateway and the Embedded Service Provider of the Access Gateway both use session cookies in their communication with the browser. The following sections explain how to protect these cookies from being intercepted by hackers.

- Section 1.5.1, "Securing the Embedded Service Provider Session Cookie on the Access Gateway," on page 17
- Section 1.5.2, "Securing the Proxy Session Cookie," on page 18

For more information about making cookies secure, see the following documents:

- Secure attribute for cookies in RFC 2965 (http://www.faqs.org/rfcs/rfc2965.html)
- HTTP-only cookies (http://msdn.microsoft.com/en-us/library/ms533046.aspx)

1.5.1 Securing the Embedded Service Provider Session Cookie on the Access Gateway

An attacker can spoof a non-secure browser into sending a JSESSION cookie that contains a valid user session. This might happen because the Access Gateway communicates with its Embedded Service Provider on port 8080, which is a non-secure connection. Because the Embedded Service Provider does not know whether the Access Gateway is using SSL to communicate with the browsers, the Embedded Service Provider does not mark the JSESSION cookie as secure when it creates the cookie. The Access Gateway receives the Set-Cookie header from the Embedded Service Provider and passes it back to the browser, which means that there is a non-secure, clear-text cookie in the browser. If an attacker spoofs the domain of the Access Gateway, the browser sends the non-secure JSESSION cookie over a non-secure channel where the cookie might be sniffed.

To stop this, you must first configure the Access Gateway to use SSL. See Section 1.3, "Configuring SSL Communication with Browsers and the Identity Server," on page 14. After you have SSL configured, you must configure Tomcat to secure the cookie.

1 On the Access Gateway server, log in as an admin user.

2 Change to the Tomcat configuration directory.

`/opt/novell/nam/mag/conf/`

3 In a text editor, open the `server.xml` file.

4 Search for the connector on port 9009.

5 Add the following parameter within the `Connector` element:

`secure="true"`

6 Save the `server.xml` file.

7 Enter one of the following commands to restart Tomcat:

`/etc/init.d/novell-mag restart` OR `rcnovell-mag restart`

1.5.2   Securing the Proxy Session Cookie

The proxy session cookies store authentication information and other information in temporary memory that is transferred between the browser and the proxy. These cookies are deleted when the browser is closed. However if these cookies are sent through a non-secure channel, there is a threat of hackers intercepting the cookies and impersonating a user on Web sites. To stop this, you can use the following configuration options:

- ◆ "Setting an Authentication Cookie with a Secure Keyword for HTTP" on page 18
- ◆ "Preventing Cross-Site Scripting Vulnerabilities" on page 18

Setting an Authentication Cookie with a Secure Keyword for HTTP

You can configure the Access Gateway to force the HTTP services to have the authentication cookie set with the keyword secure.

To enable this option:

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Reverse Proxy / Authentication**.

2 Select the **Enable Secure Cookies** option, then click **OK** twice.

3 Update the Access Gateway.

This option is used to secure the cookie when the Access Gateway is placed behind an SSL accelerator, such as the Cisco SSL accelerator, and the Access Gateway is configured to communicate by using only HTTP.

Preventing Cross-Site Scripting Vulnerabilities

Cross-site scripting vulnerabilities in Web browsers allow malicious sites to grab cookies from a vulnerable site. The goal of such attacks might be to perform session fixation or to impersonate the valid user. You can configure the Access Gateway to set its authentication cookie with the `HttpOnly` keyword, to prevent scripts from accessing the cookie.

To enable this option:

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Reverse Proxy / Authentication**.

2 Enable the **Force HTTP-Only Cookies** option, then click **OK** twice.

3 Update the Access Gateway.

1.6 Managing Access Gateway Certificates

1.6.1 Managing Reverse Proxy and Web Server Certificates

You select Access Gateway certificates on two pages in the Administration Console:

* **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]**
* **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **Web Servers**

When you configure certificates on these pages, you need to be aware that two phases are used to push the certificates into active use.

Phase 1: When you select a certificate on one of these pages, then click **OK**, the certificate is placed in the keystore on the Administration Console and it is pushed to the Access Gateway. The certificate is available for use, but it is not used until you update the Access Gateway.

Phase 2: When you select to update the Access Gateway, the configuration for the Access Gateway is modified to contain references to the new certificate and the configuration change is sent to the Access Gateway. The Access Gateway loads and uses the new certificate.

IMPORTANT: The certificate associated with the NAM-RP Reverse Proxy is available for use for Identity Server. When you select the certificate, it is pushed to Access Gateway and the Identity Server. Along with updating the Access Gateway, Identity Server should also be updated.

---

2 Configuring the Access Gateway to Protect Web Resources

The 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.

This section describes the following tasks:

2.1 Configuration Options

A typical Access Manager Appliance configuration includes an Identity Server with LDAP directories and an Access Gateway with a protected Web server. Figure 2-1 illustrates the process flow that allows an authorized user to access the protected resource on the Web server.

Figure 2-1   Accessing a Web Resource
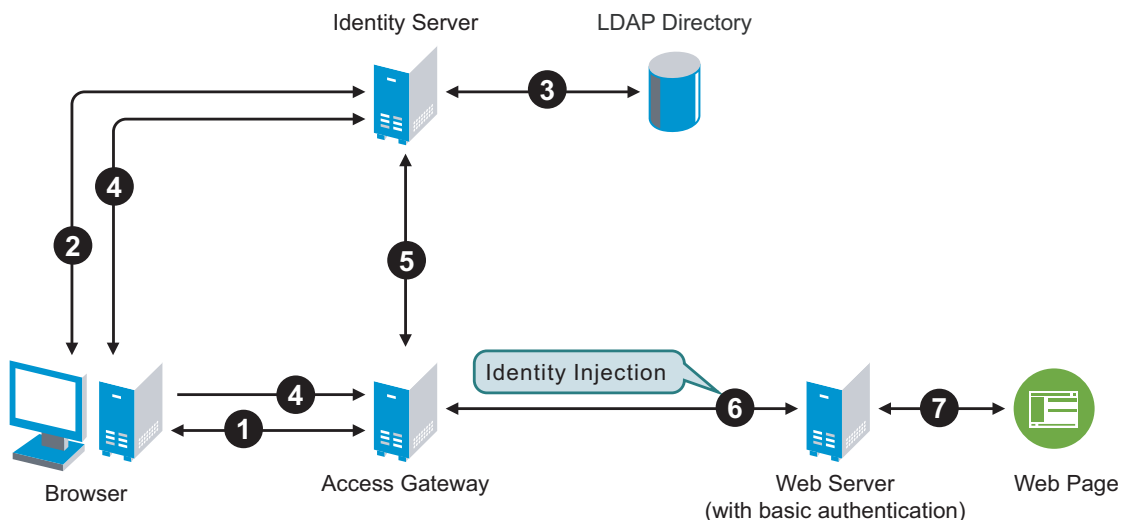
1. The user requests access to a resource protected by the Access Gateway.

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 2-2 *Access Gateway Modules and Their Configuration Options*

Module Hierarchy **Configuration Options**

◇ Reverse Proxy
{ Listening Address & Port
  SSL Requirements

◇ Proxy Service
{ Web Servers
  Caching
  HTML Rewriting
  Logging

◇ Protected Resource
{ URLs
  Authentication Contracts and Procedures
  Authorization
  Identity Injection
  Form Fill
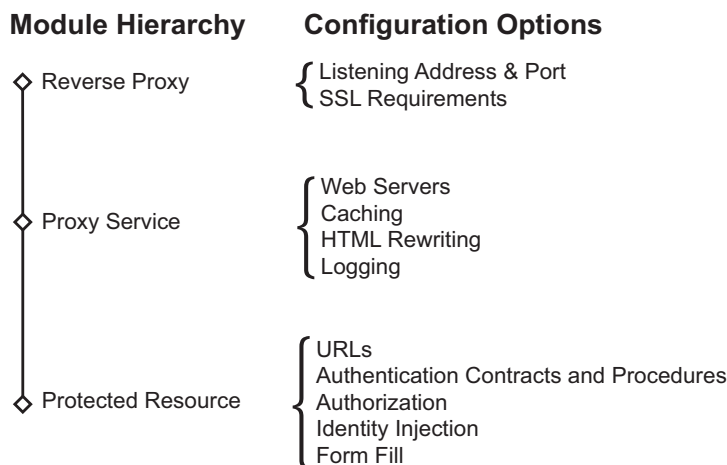
This hierarchy allows you to have precise control over what is required to access a particular resource, and also allows you to provide a single sign-on solution for all the resources protected by the Access Gateway. The authentication contract, authentication procedure, Authorization policy, Identity Injection policy, and Form Fill policy are configured at the resource level so that you can enable exactly what the resource requires. This allows you to decide where access decisions are made:

 ◆ You can configure the Access Gateway to control access to the resource.

 ◆ You can configure the Web server for access control and configure the Access Gateway to supply the required information.

 ◆ You can use the first method for some resources and the second method for other resources or use both methods on the same resource.

2.2 Managing Reverse Proxies and Authentication

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`.

The first reverse proxy and proxy service you create are automatically assigned to be the authenticating proxy.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit**

The **Edit** link is either for a single Access Gateway or for a cluster of Access Gateways.

2 Click **Reverse Proxy / Authentication**.

Reverse Proxies / Authentication: ag18

Authentication Settings

Identity Server Cluster: IDP-Cluster

Embedded Service Provider

Reverse Proxy: NAM-RP

Metadata URL: https://idp201.com:443/nesp/idff/metadata

Health-Check URL: https://idp201.com:443/nesp/app/heartbeat

Logout URL: https://idp201.com:443/AGLogout

Proxy Settings

☐ Behind Third Party SSL Terminator

☑ Enable Via Header

Cookies Settings

☐ Enable Secure Cookies

☐ Force HTTP-Only Cookie

Reverse Proxy List

New.. | Delete | Rename... | Enable | Disable   1 item(s)

☐ **Name Enabled Listening Address Port**

☐ NAM-RP   ✔   Multiple   443

3 (Conditional) If you have already created at least one reverse proxy, you can view the Embedded Service Provider options and configure some of them:

Reverse Proxy: Specifies which proxy service is used for authentication. If you have configured only one proxy service, only one appears in the list and it is selected. If you change the reverse proxy that is used for authentication, certificates must be updated to match this new configuration.

Specifies which proxy service is used for authentication. If you have configured only one proxy service, only one appears in the list and it is selected. If you change the reverse proxy that is used for authentication, certificates must be updated to match this new configuration.

Metadata URL: Displays the location of the metadata.

Health-Check URL: Displays the location of the health check.

Logout URL: Displays the URL that you need to use for logging users out of protected resources. This value is empty until you have created at least one reverse proxy and it has been assigned to be used for authentication. 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.

If any of your protected resources have a logout page or button, you need to redirect the user's logout request to the page specified by this URL. 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 reauthenticates the user to the resource, and it appears that the logout did nothing.

Auto-Import Identity Server Configuration Trusted Root: Allows you to import the public key from the Identity Server cluster 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. This option is not available until you have selected an **Identity Server Cluster** and have configured the use of SSL on the Embedded Service Provider of the reverse proxy that is performing authentication (see the **Enable SSL with Embedded Service Provider** option on the Reverse Proxy page).

If the Identity Server cluster is using a certificate created by the Access Manager certificate authority (CA), the public key is automatically added to this trust store, so you do not need to use this option. If the Identity Server cluster is using a certificate created by an external CA, you need to use this option to import the public key into the trust store.

4 (Optional) Configure the proxy settings:

Behind Third Party SSL Terminator: Enable this option if you have installed an SSL terminator between the users and the Access Gateway. This allows the terminator to handle the SSL traffic between the browsers and the terminator. The terminator and the Access Gateway can use HTTP for their communication.

Enable Via Header: Enables the sending of the Via header to the Web server. The Via header contains the DNS name of the Access Gateway and a device ID. It has the following format:

```
Via: 1.1 www.mymag.com (Access Gateway-ag-BFBA9849520DB63B-5)
```

Deselect this option when your Web server does not need this information or does not know what to do with it.

5 (Optional) Configure the cookie settings:

For more information and other options for securing Access Manager cookies, see Section 1.5, "Enabling Secure Cookies," on page 17.

Enable Secure Cookies: Enabling this option sets secure keyword on HTTPS request. If you have enabled the **Behind Third Party SSL Terminator** option and also enabled the **Enable Secure Cookies** option, the secure keyword on HTTP and HTTPS requests are set.

WARNING: Do not enable the **Enable Secure Cookies** option if you have both HTTP and HTTPS reverse proxies. The HTTP services become unavailable because authentication requests to the non-HTTP services fail.

Force HTTP-Only Cookie: Forces the Access Gateway to set the HttpOnly keyword, which prevent scripts from accessing the cookie. This helps protect browsers from cross-site scripting vulnerabilities that allow malicious sites to grab cookies from a vulnerable site. The goal of such attacks might be to perform session fixation or to impersonate the valid user.

IMPORTANT: The HttpOnly keyword can prevent applets from loading and can interfere with JavaScript. Do not enable this option if you are using the traditional SSL VPN server (which is configured as a protected resource of the Access Gateway) or if you have the Access Gateway protecting applications that download applets or use JavaScript.

---

6 To create a proxy service, continue with Section 2.2.1, "Creating a Proxy Service," on page 26.

2.2.1 Creating a Proxy Service

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Reverse Proxy /
Authentication**.

2 In the **Reverse Proxy List**, click **New**, specify a display name for the reverse proxy, then click
OK.

3 Enable a listening address. Fill in the following fields:

Cluster Member: (Available only if the Access Gateway is a member of a cluster.) 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 cluster.

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 cluster, you must select a listening address for each cluster
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 2.7.1,
"Configuring TCP Listen Options for Clients," on page 77.
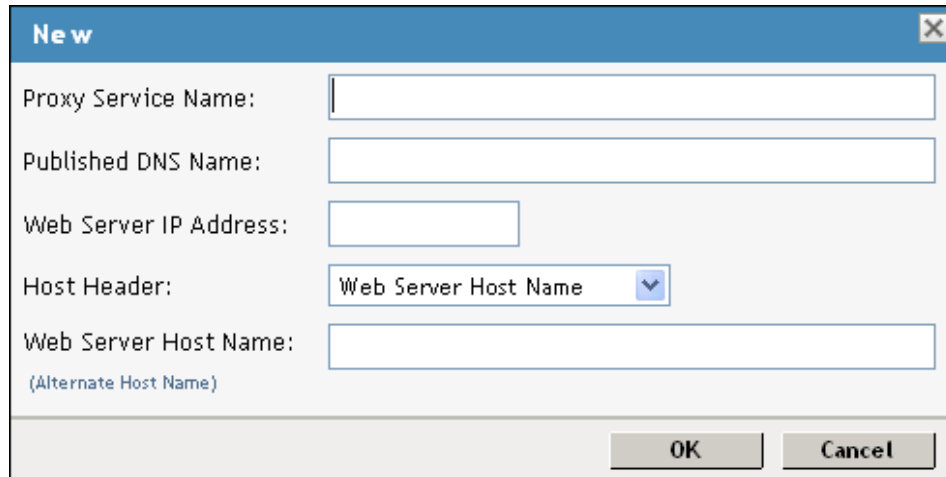
4 Configure the listening ports:

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 1, "Configuring the Access Gateway for SSL
and Other Security Features," on page 11.

5 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 7.2, "Using Multi-Homing to Access Multiple Resources," on page 170.

6 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 when 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 hostname when configuring a Web Server Accelerator.

---

7 Click **OK**.

8 Continue with Section 2.2.2, "Configuring a Proxy Service," on page 28 or select one of the following tasks:

 ◆ For information about how to create multiple reverse proxies, see Section 7.3, "Managing Multiple Reverse Proxies," on page 179.

 ◆ For information about how to create multiple proxy services for a reverse proxy, see Section 7.2, "Using Multi-Homing to Access Multiple Resources," on page 170.

2.2.2   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]**.

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.

If one proxy service has a DNS name of www.support.novell.com and the second proxy service has a DNS name of www.developernet.novell.com, the cookie domains are support.novell.com for the first proxy service and developernet.novell.com for the second proxy service. You can configure them to share the same cookie domain by selecting novell.com for each proxy service. Single sign-on between the proxy services is simplified when the proxy services share the same cookie domain.

HTTP Options: Allows you to set up custom caching options for this proxy service. See the following:

- Section 6.2, "Controlling Browser Caching," on page 160
- Section 6.3, "Configuring Custom Cache Control Headers," on page 161

Advanced Options: (Access Gateway Service) Specifies how the proxy service handles specific conditions, such as Web server error pages. If similar options are configured globally, the proxy service configuration overwrites the global setting. For configuration information on the proxy service options, see Section 5.2, "Configuring the Advanced Options for a Domain-Based and Path-Based Multi-Homing Proxy Service," on page 157.

3 Click **OK** to save your changes to browser cache.

4 Click **Devices** > **Access Gateways**.

5 To apply your changes, click **Update** > **OK**.

Until this step, nothing has been permanently saved or applied. The **Update** status pushes the configuration to the server and writes the configuration to the configuration data store. When the update has completed successfully, the server returns the status of **Current**.

To save the changes to the configuration store without applying them, do not click **Update**. Instead, click **Edit**. On the Configuration page, click **OK**. The **OK** button on this pages saves the cached changes to the configuration store. The changes are not applied until you click **Update** on the Access Gateways page.

6 Update the Identity Server to accept the new trusted relationship. Click **Identity Servers > Update**.

7 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 2.3, "Configuring Web Servers of a Proxy Service," on page 30.
- 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 2.4, "Configuring Protected Resources," on page 32.

2.2.3   Modifying the DNS Setting for a Proxy Service

1 Get the SSL certificate for the new DNS name.

For more information, see "Creating Certificates"in the *NetIQ Access Manager Appliance 4.0 SP1 Administration Console Guide*.

2 In the Administration Console, click **Devices** > **Access Gateways**.

3 Edit AG-Cluster and click on any reverse proxy listed under **Reverse Proxy/Authentication**.

4 Change the **Server Certificate** to the new one for your new DNS name.

Ignore any warning displayed about CN name mismatch because the proxy service is not yet updated.

5 Under the **Proxy Service List** tab, click the proxy which DNS name you want to modify.

6 Change the **Published DNS Name** for the proxy service.

7 Click **OK** > **OK**.

8 The Cluster Configuration page is displayed.

9 Click **Network Settings** > **Hosts** > *IP address of your system*.

10 Add the new DNS name in the list of host names.

11 Click **OK**.

12 Go to **Devices** > **Access Gateway**.

13 Click **Update All**.

14 When the Access Gateway **Health** turns green, check the Identity Server **Health** and ensure that it is green as well.

2.3 Configuring 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.

IMPORTANT: For caching to work correctly, the Web servers must be configured to maintain a valid time. They should be configured to use an NTP server.

---

1 Click **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **Web Servers**.

| Path-Based Multi-Homing | **Web Servers** | HTML Rewriting | Logging |

Host Header: Forward Received Host Name ▾
    Web Server Host Name: [   ]
    (Alternate Host Name)
    ☐ Error on DNS Mismatch (singlebox.server.com)

☐ Enable Force HTTP 1.0 to Origin
☑ Connect Using SSL
    Web Server Trusted Root: Do not verify ▾
    SSL Mutual Certificate: [   ] 🗐
Connect Port: * [8443]
TCP Connect Options

2 Specify the hostname 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 name of the Web server. Use the **Web Server Host Name** field to specify this name. You can also append the port number to the **Web Server Host Name** field. For example, *<web server hostname>*:*<web server port number>*.

3 Select **Error on DNS Mismatch** to have the proxy determine whether the proxy service should compare the hostname in the DNS header that came from the browser with the DNS name specified in the **Web Server Host Name** option. The value in the parentheses is the value that comes in the header from the browser.

If you enable this option and the names don't match, the request is not forwarded to the Web server. Instead, the proxy service returns an error to the requesting browser. This option is only available when you select to send the **Web Server Host Name** in the HTTP header.

> **NOTE:** The **Error on DNS Mismatch** option does not work in the following scenarios:
> - If the option is enabled in a protected resource.
> - If the option is enabled in a master host based service, and disabled in a path-based child services, then the Access Gateway does a strict check of DNS match for path-based child.

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 Session Stickiness: Selecting this option makes the proxy server to use the same Web server for all fills during a session.

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 1.4, "Configuring SSL between the Proxy Service and the Web Servers," on page 16.

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 2.7.2, "Configuring TCP Connect Options for Web Servers," on page 78.

8 To add a Web server, click **New** in the **Web Server List** and specify the IP address or the fully qualified DNS name 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 2.7.4, "Configuring Web Servers," on page 79.

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.

NOTE: Do not remove all configured Web servers to the cluster if any of the cluster member does not have at least one Web server configured.

10 To save your changes to browser cache, click **OK**.

11 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

2.4 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 procedures and the policies that should be used to enforce protection. The authentication procedures and the policies (Authorization, Identity Injection, and Form Fill) enable the single sign-on environment for the user. The type of protection 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 Procedures: 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 procedure, which allows the resource to be a public resource, with no login required.

In addition to selecting the contract, you can also configure how the authentication procedure handles subsequent authentication requests from an application.

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 prepopulate 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 reenter the information only when something changes, such as a password.

These policies allow you to design a custom access 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.

To configure a protected resource:

1 In the Administration Console, click **Devices** > **Access Gateways > Edit > [Name of Reverse Proxy] > [Name of Domain-Based Proxy Service or Primary Proxy Service] > Protected Resources**.

The **Resource View** of the **Protected Resource List** is used to create new protected resources or manage existing protected resources. The **Policy View** is used to see which policies are being used by multiple protected resources. For more information about the **Policy View**, see Section 2.4.7, "Assigning a Policy to Multiple Protected Resources," on page 44.

2 Select one of the following actions:

New: To create a new protected resource, click this option and specify a display name for the resource. For configuration information, see Section 2.4.1, "Setting Up a Protected Resource," on page 33.

Delete: To delete a protected resource, select a protected resource, then click **Delete**.

Enable: To enable a resource so that the Access Gateway protects it, select a protected resource, then click **Enable**.

Disable: To disable protection for a resource, select a protected resource, then click **Disable**. After a resource is disabled, its path no longer has special protection. For example, you can set up a resource that allows access to all pages (for example /*) and another resource with special protection for a subpath. If you disable the subpath, make sure the security requirements of the /* resource are sufficient for the subpath.

Also, when a protected resource is disabled, the resource no longer shows up in the Path List for a path-based multi-homing proxy.

3 Select the name of a protected resource to perform the following tasks:

- ◆ Section 2.4.2, "Configuring an Authentication Procedure for Non-Redirected Login," on page 36
- ◆ Section 2.4.3, "Assigning an Authorization Policy to a Protected Resource," on page 38
- ◆ Section 2.4.4, "Assigning an Identity Injection Policy to a Protected Resource," on page 39
- ◆ Section 2.4.5, "Assigning a Form Fill Policy to a Protected Resource," on page 40
- ◆ Section 2.4.6, "Assigning a Timeout Per Protected Resource," on page 41

2.4.1   Setting Up a Protected Resource

1 Click **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service] > Protected Resources**.

2 Either click the name of an existing resource or click **New**, then specify a display name for the resource.

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 to use for the authentication procedure. The contract 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, this option allows any contract defined for the Identity Server to be valid, or if the user has not authenticated, it 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 provided by the Web browser.

Name/Password - Form: Specifies a form-based authentication over HTTP or HTTPS, using the Access Manager login form.

Secure Name/Password - Basic: Specifies basic authentication over HTTPS, using a standard login pop-up provided by the Web browser.

Secure Name/Password - Form: Specifies a form-based authentication over HTTPS, using the Access Manager login form.

The contract also determines the session timeout for inactive connections. If you have some resources that need to time out quickly to protect sensitive data and other resources that don't need this kind of protection, you need to configure contracts for these resources. For more information about this feature, see Section 2.4.6, "Assigning a Timeout Per Protected Resource," on page 41.

If no contracts are available, you have not configured a relationship between the Access Gateway and the Identity Server. See Section 2.2, "Managing Reverse Proxies and Authentication," on page 23.

5 (Conditional) To modify how the authentication procedures are handled for a specific resource and contract, click the **Edit Authentication Procedures** icon.

For configuration information, see Section 2.4.2, "Configuring an Authentication Procedure for Non-Redirected Login," on page 36.

6 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 allow access to all the files in a directory, but not to the subdirectories and their files, specify the following:

```
/?
```

```
/public/?
```

The /? allows access to the root directory, but not the subdirectories. The /public/? allows access to the files in the public directory, but not the subdirectories.

To allow access to files of a specific type, specify the following:

```
/public/*.pdf
```

This allows access to all the files in the public directory that have a PDF extension. Access to other file types and subdirectories is denied.

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.

For more information about how a user's request is matched to a protected resource, see "Understanding URL Path Matching" on page 35.

For more information about using a query string, see "Using a Query String in the URL Path" on page 36.

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**.

7 Click **OK**.

8 In the **Protected Resource List**, ensure that the protected resource you created is enabled.

9 (Optional) To add policies for protecting this resource, continue with one of the following:

* "Assigning an Authorization Policy to a Protected Resource" on page 38
* "Assigning an Identity Injection Policy to a Protected Resource" on page 39
* "Assigning a Form Fill Policy to a Protected Resource" on page 40
* "Assigning a Policy to Multiple Protected Resources" on page 44

10 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

Understanding URL Path Matching

The URL path determines which protected resource is used for a user request. Suppose you create one protected resource with the following URL paths:

```
/*
/test/*
/test/
```

You create a second protected resource with the following path:

```
/test/*.php
```

Users then send the following paths in their access requests:

```
/test/
/test/1/2/3/file.php
/file.php
/test/file.php
/test/file.php?param1=1234
```

The first three requests (`/test/`, `/test/1/2/3/file.php`, and `/file.php`) match the first protected resource, and the last two requests (`/test/file.php` and `/test/file.php?param1=1234`) match the second protected resource.

You then add the following URL path to the first protected resource:

```
/test/?
```

This URL path in the first protected resource causes all the requests to match the first protected resource, and the second protected resource is ignored. The ? wildcard, which matches all content in the current directory, takes precedence over the more specific wildcard (*.php).

Using a Query String in the URL Path

You can specify a query string in the URL path of a protected resource. For example:

URL path: `/test/index.html?test=test`

When the requested URL has a query string, the Access Gateway searches for a protected resource with a matching URL path and query string. If it can't find a match, the request returns a `resource not found` error.

 The Access Gateway Service does not have this option. If you want the query string ignored, you must remove it from the URL path of the protected resource.

2.4.2 Configuring an Authentication Procedure for Non-Redirected Login

When a contract is created, it is assigned an authentication procedure that allows the user to be redirected to the Identity Server for authentication. Some applications, such as AJAX and WebDAV applications, do not support redirection for authentication. You can change the authentication behavior of a contract so that redirection does not occur.

When non-redirected login is enabled, the Access Gateway prompts the user to supply basic authentication credentials. The SOAP back channel between the Access Gateway and the Identity Server is used to complete the authentication on the user's behalf rather than a redirect. The SOAP back channel is also used for the session renewals.

Non-redirected login has the following restrictions:

* **Password Expiration Services:** When you modify the authentication procedures to use non-redirected login, you cannot also use a password expiration service. Even when the **Password expiration servlet** and **Allow user interaction** options are configured, users are not redirected when their passwords are expiring and they are not prompted to change their passwords.
* **Locked Shared Secrets:** When non-redirected login is enabled, users are not prompted for their passphrase for locked shared secrets.
* **Session Limits:** Non-redirected login can cause the user to create more than one session with the Identity Server because the SOAP back channel uses a different process than authentication requests that are directed to the Identity Server. Therefore, do not limit your users to one session. Session limits are set by clicking **Devices** > **Identity Servers** > **Edit**.

If the contract you are going to use for non-redirected login is also assigned to protected resources that do not require non-redirected login, you should create a new authentication procedure for the resource requiring non-redirected login. Multiple authentication procedures can be configured to use the same contract.

To configure an authentication procedure:

1 Click **Access Gateways** > **Edit > [Name of Reverse Proxy] > [Name of Proxy Service] > Protected Resources** > **[Name of Protected Resource]**.

2 On the Authentication Procedure line, click the **Edit Authentication Procedure** icon.

The Authentication Procedure List displays all available contracts, the name of the authentication procedure they are assigned to, the protected resources that the authentication procedure has been assigned to, and whether the procedure has been enabled for non-redirected login.

3 Select one of the following actions:

  ◆ To create an new authentication procedure, click **New**, specify a name, then click **OK**. Continue with Step 4.

  ◆ To modify an existing authentication procedure, click the name of the procedure. Continue with Step 4.

  ◆ To delete an existing authentication procedure, select the procedure, then click **Delete**. Continue with Step 7.

    If the procedure is used by a resource, it cannot be deleted until it is not being used to protect resources. An authentication procedure must exist for each contract. If you delete an authentication procedure for a contract without also deleting the contract, the system automatically re-creates an authentication procedure for the contract.

4 To specify the method for obtaining the credentials, fill in the following fields:

Contract: Select the contract that you want to use for this protected resource. This needs to be a contract that supports basic authentication credentials such as Name/Password- Basic or Secure Name/Password-Basic. You can also configure Non-Redirected Login with a Kerberos contract.

Non-Redirected Login: Select this option to use the SOAP back channel to verify the user's credentials rather than a redirected request to the Identity Server.

Realm: Specify a name that your users can use to identify the site that they are authenticating to. This could be your company name or the name of the application. The realm is displayed as a heading when the application requests a basic authentication.

Redirect to Identity Server When No Authentication Header Is Provided: The response should provide an authentication header. If the first request does not contain the authentication header, you can select this option to allow the first request to be redirected to the Identity Server.

5 Click **OK**.

6 For the Authentication Procedure, select the authentication procedure you created or modified in Step 4.

7 Click **OK**.

8 Click **Devices** > **Access Gateways**, then update the Access Gateway.

9 (Optional) For some configuration scenarios that use this feature, see

  ◆ "Configuring Protected Resource for a SharePoint Server" on page 45

  ◆ "Configuring a Protected Resource for a SharePoint Server with an ADFS Server" on page 46

  ◆ "Configuring a Protected Resource for Outlook Web Access" on page 49

  ◆ "Configuring a Protected Resource for a Novell Vibe 3.3 Server" on page 52

2.4.3 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**.

The **Authorization Policy List** contains all the Access Gateway Authorization policies that have been created on this Administration Console for the selected policy container.

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. For configuration information, see "Creating Access Gateway Authorization Policies" in the *NetIQ Access Manager Appliance 4.0 SP1 Policy Guide*.

  When you have completed your policy modifications, 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 "Creating Access Gateway Authorization Policies" in the *NetIQ Access Manager Appliance 4.0 SP1 Policy Guide*.

  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 to the browser cache, click **OK**.

5 To apply the changes, click the **Access Gateways** link, then click **Update** > **OK**.

2.4.4 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**.

The **Identity Injection Policy List** contains all the Identity Injection policies that have been created on this Administration Console for the selected policy container.

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. For configuration information, see "Creating Identity Injection Policies" in the *NetIQ Access Manager Appliance 4.0 SP1 Policy Guide*.

    When you have finished your policy modifications, 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 "Creating Identity Injection Policies" in the *NetIQ Access Manager Appliance 4.0 SP1 Policy Guide*.

    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. If you use the same policy for multiple protected resources, use the policy description field to indicate this.

4 To save your changes to the browser cache, click **OK**.

5 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

IMPORTANT: If you enable an Identity Injection policy for a protected resource that has been assigned to use a contract that does not prompt the user for a password and the Identity Injection policy injects the user's password, single sign-on cannot be enabled because the password is not available. However, you can create a contract that retrieves the user's password when the user is not prompted for a password when authenticating. See "Configuring Password Retrieval" in the *NetIQ Access Manager Appliance 4.0 Identity Server Guide*.

2.4.5 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 prompts 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, whether it is 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 "Creating Form Fill Policies" in the *NetIQ Access Manager Appliance 4.0 SP1 Policy Guide*). To assign a Form Fill policy to a protected resource:

1 In the Administration Console, click **Devices** > **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 possible, it should not be a URL that ends in a wildcard (for example, an asterisk) and therefore matches many pages.

   IMPORTANT: When the URL ends in a wildcard, the Access Gateway must 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. For more information about the performance problems this can cause, see "Creating a Form Matching Rule" in the *NetIQ Access Manager Appliance 4.0 SP1 Policy Guide*.

3 (Conditional) If the URL is not specific, click the name of the path and modify it.

4 Click **Form Fill**.

The **Form Fill Policy List** contains all the Form Fill policies that have been created on this Administration Console for the selected policy container.

5 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 7.

- To disable an existing policy, select the policy, then click **Disable**. Continue with Step 7.

- 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. For configuration information, see "Creating Form Fill Policies" in the *NetIQ Access Manager Appliance 4.0 SP1 Policy Guide*.

  When you have finished the policy modifications, continue with Step 7.

- 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 "Creating Form Fill Policies" in the *NetIQ Access Manager Appliance 4.0 SP1 Policy Guide*.

  When you have created your new policy, continue with Step 6.

6 To enable the policy you just created, select the policy, then click **Enable**.

Only the policies that are enabled are applied to this resource. If you use the same policy for multiple protected resources, use the policy description field to indicate this.

7 To save your changes to the browser cache, click **OK**.

8 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

IMPORTANT: If you enable a Form Fill policy for a protected resource that has been assigned to use a contract that does not prompt the user for a password and the Form Fill policy contains a field for the user's password, single sign-on cannot be enabled because the password is not available. To enable single sign-on, you need to use an Authentication class that retrieves the user's password and injects it into the user's credentials when the user authenticates using a non-password method such as X.509, RADIUS, smart card, or Kerberos. For information about such a class, see "Configuring Password Retrieval" in the *NetIQ Access Manager Appliance 4.0 Identity Server Guide*.

2.4.6 Assigning a Timeout Per Protected Resource

If all your resources are using the same contract and you want them all to have the same timeout for inactivity, you set the **Authentication Timeout** option on the contract to the required limit and leave the **Activity Realm** option blank. The user logs in, and activity by the user on any resource keeps the user's session active. The user is prompted to reauthenticate only when the user has no activity on any resources for longer than the authentication timeout value.

If you have some resources that require a shorter timeout than other resources, you need to balance the need for single sign-on with the timeout requirements:

- To strictly enforce a timeout, the resource needs to be assigned to a custom contract.
- To preserve single sign-on, resources need to be assigned to the same contract.

The protected resource is assigned to use a contract, and the timeout is assigned to the contract. For information about how to configure the contract, see "Configuring Authentication Contracts" in the *NetIQ Access Manager Appliance 4.0 Identity Server Guide*.

The following sections describe four configuration scenarios and the user experience that they create.

Scenario 1: If strictly adhering to the timeout value is more important than preserving the session or single sign-on, configure your resources as follows:

- Protected resource 1 (PR1) is configured to use contract 1 (C1), which has been created from method 1 (M1) and placed in its own activity realm (AR1). For this scenario you set the authentication timeout to 30 minutes.
- Protected resource 2 (PR2) is configured to use contract 2 (C2), which has been created from method 2 (M2) and placed in its own activity realm (AR2). For this scenario, you set the authentication timeout to 15 minutes.

With this scenario, the user is prompted to log in when accessing PR1 and when accessing PR2. Each resource has its own time line, because each resource belongs to its own activity realm. Figure 2-3 The figure below illustrates this scenario.

Figure 2-3 *Login Requirements with Separate Methods and Separate Activity Realms*
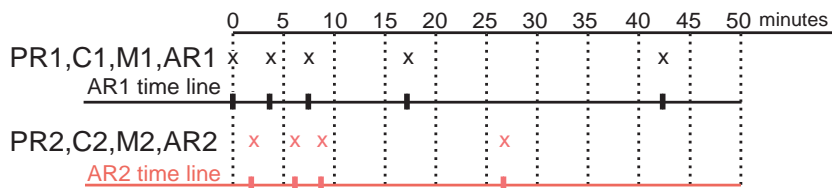
After authenticating to both resources and remaining active on both resources for the first 10 minutes, the sessions remain active. The user then stays active on PR1 without accessing PR2 for over 15 minutes. The AR1 time line is updated with this activity. The AR2 time line is not updated. When the user accesses PR2 after more than 15 minutes of inactivity on the AR2 time line, the user is prompted to authenticate. The user then returns to PR1 after over 20 minutes of inactivity, but AR1 time line shows activity within the 30-minute timeout. The user is granted access and does not need to log in again to access PR1.

In this scenario, the resources are independent of each other and do not influence each other's timeout limits.

Scenario 2: If you are willing to allow a resource to influence the timeout of another resource, configure your resources as follows:

◆ Protected resource 1 (PR1) is configured to use contract 1 (C1), which has been created from method 1 (M1) and placed in a shared activity realm (shared1). For this scenario you set the authentication timeout to 30 minutes.

◆ Protected resource 2 (PR2) is configured to use contract 2 (C2), which has been created from method 2 (M2) and placed in a shared activity realm (shared1). For this scenario, you set the authentication timeout to 15 minutes.

With this scenario, the user is prompted to log in when accessing PR1 and when accessing PR2. Activity at either resource updates the shared1 time line. Figure 2-4 illustrates this scenario.

Figure 2-4   Login Requirements for Separate Methods with a Shared Activity Realm
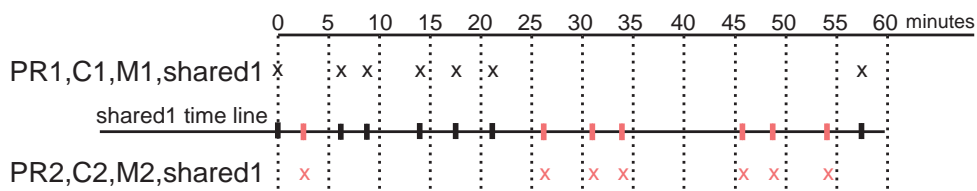
As long as the user is active on PR1, the user's session to PR2 remains active. After 20 minutes of activity on PR1, the user returns to PR2. The user is allowed access and does not need to log in because the shared1 time line shows activity within the last 5 minutes. The user remains active on PR2 for over 30 minutes, then accesses PR1. Again, the shared1 time line shows activity within the last 5 minutes, so the user is granted access to PR1 without logging in again.

With this configuration, activity at other resources influences the time limits so that they are not strictly enforced.

Scenario 3: If single sign-on is more important than strictly enforcing a timeout value, NetIQ recommends that you configure all contracts to have the same authentication timeout value.

If you configure your resources as follows, you might not get the behavior you require:

◆ Protected resource 1 (PR1) is configured to use contract 1 (C1), which has been created from method 1 (M1) and placed in a shared activity realm (shared1). For this scenario you set the authentication timeout to 30 minutes.

◆ Protected resource 2 (PR2) is configured to use contract 2 (C2), which has been created from method 1 (M1) and placed in a shared activity realm (shared1). For this scenario, you set the authentication timeout to 15 minutes.

Because C1 and C2 are created from the same method (M1), the user does not need to log in twice to access both resources. Logging in to one resource allows them access to the other resource. Figure 2-5 illustrates this scenario.

Figure 2-5   Login Requirements for Shared Methods and Shared Realms
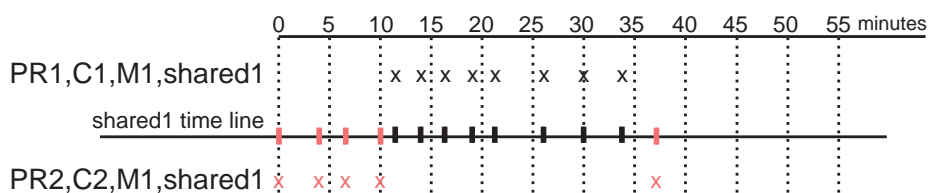
The user first logs in to PR2 and is active for 10 minutes. The shared1 time line gets updated with this activity. When the user requests access to PR1, the user is granted access without being prompted for credentials. The user is then active on PR1 for over 20 minutes. When the user requests access to PR2, even though the user has been inactive on this resource for over 20 minutes, the user is granted access because the time line shows activity within the last five minutes.

With this configuration, PR2 does not time out as long as the user remains active on PR1. However, when the user goes inactive on both PR2 and PR1 for over 15 minutes and the user requests access to PR1, the time line shows no activity within the time limit specified for PR2 and the user is prompted to log in.

Scenario 4: NetIQ does not recommend that you set different authentication timeouts on contracts and then use the Any contract option for protected resources. If you want to use the Any contract, then you should set the authentication timeout to the same value on all contracts. If the timeouts are not the same, you cannot consistently predict what timeouts are being applied to the various protected resources. For example, the user requests access to a resource that is protected with a contract with a short timeout. The user logs in, then accesses resources that use the Any contract option. All of these resources are assigned a short timeout.The user then goes inactive and the session times out. The user then requests access to a resource with a contract with a long timeout. The user logs in, and after a few minutes, accesses same resources protected with the Any contract option. These resources are now assigned the long timeout value.

2.4.7 Assigning a Policy to Multiple Protected Resources

If you have created multiple protected resources that need to be protected by the same policy or policies, you can use the policy view to assign a policy to multiple protected resources. However, the protected resources must belong to the same proxy service.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit > [Reverse Proxy Name] > [Name of Proxy Service] > Protected Resources**.

2 Select the **Policy View**.

Policy View ∨

Policy List

Name	Type	Policy Container	Used By ℹ
Innerweb_Identity_Injection	Access Gateway: Identity Injection	Innerweb	Third Party, ... (4)
Innerweb_Login	Access Gateway: Form Fill	Innerweb	[None]
Partners_Auth	Access Gateway: Authorization	Innerweb	Partners
Third_Party_Auth	Access Gateway: Authorization	Innerweb	Third Party

3 Select the **Used By** link of the policy you want to assign to multiple resources.

Policy: Innerweb_Identity_Injection
Policy Container: Innerweb

Enable/Disable this Policy on the Protected Resources defined for this Proxy Service.

Protected Resource Policy Usage List

Enable | Disable

	Name	**Enabled**	**Description**
☐	Human_Resources		
☐	Innerweb_General		
☐	Partners		
☐	Third_Party		

The **Policy** and **Policy Container** fields identify the policy. The **Protected Resource Policy Usage List** displays the protected resources defined for this proxy service and indicates which resources the policy has been enabled on.

4 To enable the policy for multiple resources, either select them one by one or click **Name** to select all of them, then click **Enable**. To disable a policy for a resource, select the resource, then click **Disable**.

5 To save your changes to browser cache, click **OK**.

6 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

2.5 Configuring Protected Resources for Specific Applications

2.5.1 Configuring Protected Resource for a SharePoint Server

You can protect a SharePoint server as a domain-based or a path-based multi-homing resource on the Access Gateway. When you protect a SharePoint server on Access Gateway, you might see issues with rewriting if the published DNS name is not the same as the DNS name of the original server. Also, if you access SharePoint folder by using non-browser clients such as Microsoft Network Place, Nautilus in SUSE Linux Enterprise Server (SLES), or the MAC finder, you might see issues

because these WebDAV clients do not support 302 redirection for authentication. You must modify the authentication procedure to prevent redirection on initial authentication or redirection to Identity Server when the user session expires.

For more information about how to configure a protected resource for a SharePoint server, see "Protecting SharePoint 2010" in the *NetIQ Access Manager 4.0 SP1 Best Practices Guide*.

2.5.2 Configuring a Protected Resource for a SharePoint Server with an ADFS Server

If your SharePoint server is configured to use an ADFS server and you want to create a protected resource for the SharePoint server, you need to configure the following Access Manager features. The instructions assume that you have a functioning SharePoint server and a functioning Access Manager system:

- "Configuring a Custom Contract" on page 46
- "Creating a Reverse Proxy Service" on page 47
- "Configuring Multiple Protected Resources" on page 47

Configuring a Custom Contract

ADFS requires a different format for a contract URI than the format used in the default contracts. It expects the URI to conform to the format of a URL. You need to create a custom contract.

1 In the Administration Console, click **Devices > Identity Servers > Servers > Edit > Local > Contracts**

2 Click **New**, then fill in the following fields:

Display name: Specifies the name of the authentication contract.

URI: Specifies a value that uniquely identifies the contract from all other contracts. No spaces can exist in the URI field. For SharePoint, specify the following format for the URI:

```
https://<baseurl>/name/password/uri
```

Replace *<baseurl>* with the base URL of your Identity Server. If the DNS name of your Identity Server is `idp-50.amlab.net`, the URI would have the following format:

```
https://idp-50.amlab.net:8443/nidp/name/password/uri
```

Methods and Available Methods: Move a name/password method to the **Methods** list. We recommend **Secure Name/Password - Basic**, but you can use **Name/Password - Basic**.

Do not configure a password expiration servlet. This contract is going to be used with non-redirected login, which prevents all redirection, including redirection to a password expiration service.

For more information about other options, see "Configuring Authentication Contracts" in the *NetIQ Access Manager Appliance 4.0 Identity Server Guide*.

3 Click **Next**.

4 Configure a card for the contract by filling in the following:

Text: Specify the text that is displayed on the card to the user.

Image: Specify the image to be displayed on the card. To use an existing image, select an image from the drop-down list. To add an image to the list, click **Select local image**.

Show Card: Determine whether the card is shown to the user, which allows the user to select and use the card for authentication. If this option is not selected, the card is only used when a service provider makes a request for the card.

5 Click **Finish**, then **OK**.

6 Update the Identity Server and the Access Gateway.

7 Continue with .

Creating a Reverse Proxy Service

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy].**

2 In the **Proxy Service List** section, click **New**.

3 Fill in the following fields:

Proxy Service Name: Specify a display name for the proxy service that the Administration Console uses for its interfaces.

Multi-Homing Type: Select **Domain-Based** as the multi-homing method that the Access Gateway should use to identify this proxy service.

Published DNS Name: Specify the DNS name you want the public to use to access the SharePoint server. This DNS name must resolve to the IP address you set up as the listening address.

If the DNS name of the reverse proxy is the same as the DNS name of the SharePoint server, no rewriting configuration is required. If they are different, there is a high probability that the application will respond incorrectly to user requests.

Web Server IP Address: Specify the IP address of the IIS Web server with the SharePoint server.

Host Header: Select the **Web Server Host Name** option.

Web Server Host Name: Specify the DNS name of the SharePoint server that the Access Gateway should forward to the Web server.

For more information about how to create a reverse proxy, see

4 Click **OK**.

5 Continue with .

Configuring Multiple Protected Resources

If your SharePoint server has been configured for multiple domains, you need to create three protected resources to enable single sign-on. The server has two ways to access the home page. You need to create a protected resource for each of these paths, and then a protected resource for the other pages. These protected resources should have a configuration similar to the following:

SharePoint Page	URL Path	Contract	Authentication Procedure
home page	default.aspx	custom	Normal
root	/	custom	Normal
all others	/*	custom	Non-redirected login

For single sign-on, all the protected resources need to specify the same contract. When assigning the contract for the /* resource, the contract needs to be configured to use non-redirected login for its authentication procedure. When a user first accesses the SharePoint server, the users are directed either to the home page or the root of the server. From either of these locations, the users can be redirected to the Identity Server for authentication. After the users have authenticated and the SharePoint server requests authentication for access to any of the other pages, these pages need to be configured to use non-redirected login.

1 In the **Proxy Service List**, click the name of the Proxy Service you created, then click **Protected Resources**.

2 To create a protected resource for the home page:

   2a In the **Protected Resource List**, click **New**, specify a name such as `homepage`, then click **OK**.

   2b For the home page of the SharePoint server, specify the following values:

      Authentication Procedure: Select the custom contract you created.

      URL Path: Click `/*` and replace it with `default.aspx`, then click **OK** twice.

3 To create a protected resource for the root page:

   3a In the **Protected Resource List**, click **New**, specify a name such as `root`, then click **OK**.

   3b For the root of the SharePoint server, specify the following values:

      Authentication Procedure: Select the custom contract you created.

      URL Path: Click `/*` and remove the asterisk, then click **OK** twice.

4 To create a protected resource for all other pages:

   4a In the **Protected Resource List**, click **New**, specify a name such as `allothers`, then click **OK**.

   4b For all other pages of the SharePoint server, specify the following values:

      Authentication Procedure: Select the custom contract you created.

      URL Path: Leave the default value.

   4c Click the **Edit Authentication Procedures** icon on the **Authentication Procedure** line.

   4d Click the name of your custom contract, then fill in the following:

      Non-Redirected Login: Select this option.

      Realm: Specify a name that your users associate with the SharePoint server. This name is displayed when the user needs to reauthenticate.

      For more information about this feature, see Section 2.4.2, "Configuring an Authentication Procedure for Non-Redirected Login," on page 36.

5 Click **OK** three times.

   In the **Protected Resource List**, you should have three protected resources that use the same Authentication Procedure.

   For information about configuring protected resources, see Section 2.4.1, "Setting Up a Protected Resource," on page 33.

6 Click **Access Gateways**, then update the Access Gateway.

7 (Conditional) If you have limited your users to one session, modify this limitation:

   7a Click **Devices** > **Identity Servers** > **Edit.**

   7b Increase the value of the **Limit user sessions** option.

   7c Click **OK**, then update the Identity Server.

2.5.3 Configuring a Protected Resource for Outlook Web Access

If you want to protect your Outlook Web Access server with the Access Gateway, you need to configure the following Access Manager features. The instructions assume that you have a functioning Outlook Web Access server and a functioning Access Manager system:

Configuring a Protected Resource for Outlook Web Access

The following instructions assume that you have a basic setup with at least one reverse proxy and proxy service. If you don't have this basic setup, see Section 2.2, "Managing Reverse Proxies and Authentication," on page 23 and complete a basic setup before continuing.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit > [Name of Reverse Proxy]**.

2 In the **Proxy Service List** section, click **New**.

3 Specify a name for the proxy service, then click **OK**.

4 Click the newly added proxy service. 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.

   Multi-Homing Type: Select the multi-homing method that the Access Gateway should use to identify this proxy service.

   Web Server IP Address: Specify the IP address of the IIS Web server.

   Host Header: Select the **Web Server Host Name** option.

   Web Server Host Name: Specify the DNS name of the Outlook Web Access server that the Access Gateway should forward to the Web server.

5 Click **OK**.

6 Continue with "Configuring an Authentication Procedure" on page 49.

Configuring an Authentication Procedure

1 Click **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service] > Protected Resources**.

2 Click **New**, then specify a display name for the resource.

3 (Optional) Specify a description for the protected resource. You can use it to briefly describe the purpose for protecting this resource.

4 Select an authentication contract. If you want to enable non-redirected login, select **Name/Password - Basic** as the authentication contract.

5 (Optional) If you want to enable non-redirected login, click the **Edit Authentication Procedure** icon, then click the contract that you have added to specify the following information:

Non-Redirected Login: Select the option to enable non-redirected login.

Realm: Specify the security realm configured for the IIS server running the Outlook Web Access server.

To check the security realm configured for the IIS server, open the IIS Administration Console, right-click the Outlook Web Access Server the Access Gateway is protecting, then select **Properties**. The **Directory Security** tab contains the **Security realm** field.

6 Create protected resource:

   6a In the **Protected Resource List**, click **New**, specify a name such as `root`, then click **OK**.

   6b Specify the following values:

      Authentication Procedure: Select the contract you created.

      URL Path: Make sure that `/*` is selected. If you have configured Outlook Web Access as a path-based service, then click the URL path and add the path name of the service. For example, `/owa/*`, where owa is the path name.

      Click **OK** twice.

7 Create a second protected resource:

   7a In the **Protected Resource List**, click **New**, specify a unique name, then click **OK**.

   7b Specify the following values:

      Authentication Procedure: Do not select any authentication procedure because the URL path is a public resource.

      URL Path: Specify `/exchweb/*as the URL path`. If you have configured Outlook Web Access as a path-based service, click the URL path and add the path name of the service. For example, `/owa/exchweb/*`, where `owa` is the path name.

      Click **OK** twice.

8 Click **OK**.

9 In the **Protected Resource List**, ensure that the protected resource you created is enabled.

10 If you want to enable single sign-on, then configure Identity Injection or Form Fill policy, depending on the Outlook Web Access configuration. For more information, see "Configuring Identity Injection" on page 51.

11 Continue with "Configuring a Rewriter Profile" on page 50.

Configuring a Rewriter Profile

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **HTML Rewriting**.

2 Click **New** in the **HTML Rewriter Profile List**.

3 Configure a Word profile:

   3a Specify a name for the profile, select **Word** as the search boundary, then click **OK**.

   3b Click **New** in the **Variable or Attribute Name to Search for Is** section, then specify `value`.

   3c Click **OK**.

   3d Select **Rewrite Inbound Query String Data**.

   3e Select **Rewrite Inbound Post Data.**

   3f Select **Rewrite Inbound Headers**.

   3g Ensure that **Enable Rewrite Actions** remains selected.

4 (Optional) If you have configured the path-based multi-homing service, do the following:

4a Add the following content types for the **And Document Content-Type Header Is** option in the Word profile:

- text/x-component

- extension/htc

4b Configure the following options for **Strings to Search for Is**:

- Specify **Search** as `/exchange` and **Replace With** as `$path/exchange`

- Specify **Search** as `/exchweb` and **Replace With** as `$path/exchweb`

5 To save your changes to browser cache, click **OK**.

6 Use the up-arrow button to move your profile to the top of the **HTML Rewriter Profile List**.

7 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

Configuring Identity Injection

You must configure an Identity Injection policy to enable single sign-on with the Outlook Web Access server that has basic authentication configured. This Identity Injection policy should be configured to inject an authentication header. For information about creating this policy, see "Configuring an Authentication Header Policy" in the *NetIQ Access Manager Appliance 4.0 SP1 Policy Guide*.

Configuring Form Fill

You can configure a Form Fill policy to prepopulate fields in the form when you log into the Outlook Web Access first time and then save the information in the completed form to the config store for subsequent logins. For information about how to create this policy, see "Creating Form Fill Policies" in the *NetIQ Access Manager Appliance 4.0 SP1 Policy Guide*.

Enabling the **Auto Submit** option requires additional entries apart from the username and password fields. To enable the **Auto Submit** option:

1 In the Administration Console, click **Policies** > **Policies** > *<Policy Name>*.

2 In the Edit Policy page, add the following details under **Fill Options**:

Input Field Name	Input Field Type	Input Field Value	Data Conversion
destination	Hidden	String Constant : http:// <webserver DNS Name/ owa> (when Web server is configured for http.) String Constant : https:// <webserver DNS Name/ owa> (when Web server is configured for https.)	None
flags	Hidden	String Constant : 0	None
forcedownlevel	Hidden	String Constant : 0	None
isUt8	Hidden	String Constant : 1	None
trusted	Radio Button	String Constant : 0	None

3 Under the **Submit Options** section, select the **Enable JavaScript Handling** check box.

4 Enter `document.cookie="PBack=0; path=/"` in the **Statements to Execute on Submit** field.

5 Click **OK** and apply the changes.

2.5.4 Configuring a Protected Resource for a Novell Vibe 3.3 Server

The following sections explain how to configure the Access Gateway with a domain-base multi-homing service. The instructions assume that you have a functioning Novell Vibe 3.3 server on Linux and a functioning Access Manager system with a reverse proxy configured for SSL communication between the browsers and the Access Gateway.

The Novell Vibe server needs to be configured to trust the Access Gateway to allow single sign-on with Identity Injection and to provide simultaneous logout. You also need to create an Access Gateway proxy service and configure it.

- "Configuring the Novell Vibe Server to Trust the Access Gateway" on page 52
- "Configuring a Domain-Based Multi-Homing Service for Novell Vibe" on page 53
- "Creating a Pin List" on page 56

For information about other possible Access Gateway configurations, see "Teaming 2.0: Integrating with Linux Access Gateway" (http://www.novell.com/communities/node/9580/teaming-20-integration-linux-access-gateway).

Configuring the Novell Vibe Server to Trust the Access Gateway

To use Novell Vibe as a protected resource of an Access Gateway and to use Identity Injection for single sign-on, the Teaming server needs a trusted relationship with the Access Gateway. With a trusted relationship, the Teaming server can process the authorization header credentials. The Teaming server accepts only a simple username (such as user1) and password in the authorization header.

This section explains how to set up the trusted relationship and how to enable simultaneous logout, so that when the user logs out of Teaming, the user is also logged out of the Access Gateway.

To configure the trusted relationship:

1 Log in to the Novell Vibe server.

2 Stop the Teaming server with the following command:

`/etc/init.d/teaming stop`

3 Run the `installer-teaming.linux` script.

4 Follow the prompts, then select **Reconfigure settings**.

5 Follow the prompts, then select **Advanced installation**.

6 Follow the prompts, selecting the defaults until the **Enable Access Gateway** option appears, then type `Yes`.

7 In the **Access Gateway address(es)** section, include the IP address of the Access Gateway that is used for the connection to the Teaming server.

If the Access Gateway is part of a cluster, add the IP address for each cluster member. Wildcards such as 164.99.*.* are allowed.

When you specify IP addresses in this option, Novell Vibe logins are allowed only from the specified addresses. Also, if authorization header credentials are not present or are incorrect, the user is prompted for login by using Basic Authentication.

8 When prompted for the Logout URL, specify the URL of the published DNS name of the proxy service plus `/AGLogout`.

For example, if the published DNS name of the proxy service is vibe.`doc.provo.novell.com`, specify the following URL:

`https://Vibe.doc.provo.novell.com/AGLogout`

9 When you are prompted to use the Access Gateway for WebDAV connections, specify `No`.

10 Follow the prompts to complete the reconfiguration process.

11 Start the Vibe server with the following command:

`/etc/init.d/teaming start`

12 Continue with "Configuring a Domain-Based Multi-Homing Service for Novell Vibe" on page 53.

Configuring a Domain-Based Multi-Homing Service for Novell Vibe

The following instructions describe how to set up a domain-based service to protect the Novell Vibe server. In this example, the published DNS name of the service is `Vibe.doc.provo.novell.com`. Users would access the Vibe server with a URL similar to `http://Vibe.doc.provo.novell.com`.

To configure a domain-based service for Vibe, complete the following tasks:

- "Configuring the Domain-Based Proxy Service" on page 53
- "Configuring Protected Resources" on page 54
- "Configuring a Rewriter Profile" on page 55

Configuring the Domain-Based Proxy Service

You must create a new reverse proxy before you configure the domain-based proxy service. Configure the Vibe domain as the primary proxy service and enable SSL between browser and the Access Gateway. For more information about how to create a new reverse proxy, see Section 2.2.1, "Creating a Proxy Service," on page 26.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]**.

2 In the **Reverse Proxy List**, click **New**, then specify the following details:

Proxy Service Name: Specify a display name for the proxy service that the Administration Console uses for its interfaces.

Multi-Homing Type: Select **Domain-Based**.

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. For example, `vibe.doc.provo.novell.com`.

Web Server IP Address: Specify the IP address of the Vibe server.

Host Header: Select the **Forward Received Host Name** option.

Web Server Host Name: Specify the DNS name of the Vibe server.

3 Click **OK**.

4 Click the newly added proxy service, then select the **Web Servers** tab.

5 Change the **Connect Port** to 8080.

If the Novell Vibe server has port forwarding enabled, you do not need to change from the default port 80.

6 Click **TCP Connect Options**.

7 Change the value of **Data Read Timeout** option to 300 seconds.

This longer timeout is needed for file uploads.

8 Click **OK**.

9 Continue with .

Configuring Protected Resources

You must configure an Identity Injection policy to enable single sign-on with the Novell Vibe server. This Identity Injection policy should be configured to inject the authentication credentials into the authorization headers.

1 In the Administration Console, click **Policies > Policies**.

2 Select the policy container, then click **New**.

3 Specify a name for the policy, select **Access Gateway: Identity Injection** for the type, then click **OK**.

4 (Optional) Specify a description for the injection policy. This is useful if you plan to create multiple policies to be used by multiple resources.

5 In the **Actions** section, click **New**, then select **Inject into Authentication Header**.

6 Fill in the following fields:

User Name: Select **Credential Profile** > **LDAP User Name**.

Password: Select **Credential Profile** > **LDAP Password**.

7 Click **OK** twice.

8 Click **Apply Changes**.

For more information about how to create such a policy, see "Configuring an Authentication Header Policy" in the *NetIQ Access Manager Appliance 4.0 SP1 Policy Guide*.

Assign this policy to the protected resources. You need to create two protected resources, one for HTML content and one for WebDAV and AJAX content.

9 Click **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service] > Protected Resources**.

10 Create a protected resource for HTML content:

10a In the **Protected Resource List**, click **New**, specify a name, then click **OK**.

10b (Optional) Specify a description for the protected resource. You can use it to briefly describe the purpose for protecting this resource.

10c Specify a value for **Authentication Procedure.** For example, select the **Secure Name/ Password - Form** contract.

10d In the URL Path List, remove the /* path and add the following two paths:

```
/teaming/*

/ssf/*
```

10e Click **OK**.

11 Create a protected resource for WebDAV and AJAX content:

11a In the **Protected Resource List**, click **New**, specify a unique name, then click **OK**.

11b (Optional) Specify a description for the protected resource. You can use it to briefly describe the purpose for protecting this resource.

11c Click the **Edit Authentication Procedure** icon.

11d In **Authentication Procedure List**, click **New**, specify a name, then click **OK**.

11e Specify details in the following fields:

Contract: Select the **Secure Name/Password - Form** contract, which is same contract that you selected for the HTML content protected resource.

Non-Redirected Login: Select this option.

Realm: Specify a name that you want to use for the Teaming server. This name does not correspond to a Vibe configuration option. It appears when the user is prompted for credentials.

Redirect to Identity Server When No Authentication Header is Provided: Deselect this option.

11f Click **OK** twice.

11g For the Authentication Procedure, select the procedure you just created.

11h In the **URL Path List**, remove the `/*` path and add the following paths:

```
/ssfs/*
/ssf/rss/*
/ssf/atom/*
/ssf/ical/*
/ssf/ws/*
/ssr/*
/rest/*
```

The `/ssfs/*` path is for WebDAV content and the `/ssf/rss/*` path enables non-redirected login for RSS reader connections.

11i Click **OK**.

12 In the **Protected Resource List**, ensure that the protected resources you created are enabled.

13 To apply your changes, click **Devices** > **Access Gateways**, then click **Update**.

14 Continue with .

Configuring a Rewriter Profile

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **HTML Rewriting**.

2 In **HTML Rewriter Profile List**, click **New**.

3 Specify a name for the profile, select **Word** as the search boundary, then click **OK**.

4 In the **And Document Content-Type Header Is** section, click **New**, then specify the following type:

```
application/rss+xml
```

5 In the **Variable or Attribute Name to Search for Is** section, click **New**, then specify the following as the variable to search for:

```
value
```

6 Click **OK**.

7 Ensure that **Enable Rewrite Actions** remains selected.

8 Click **OK**.

9 In **HTML Rewriter Profile List**, move the Word profile you created to be the first profile in the list, and move the default profile to be the second profile in the list.

10 Click **OK**.

11 To apply your changes, click **Devices** > **Access Gateways**, **Update**.

12 Continue with "Creating a Pin List" on page 56.

NOTE: If Vibe is configured to send the binary content in the JSON format, you must disable the HTML Rewriter to prevent errors.

Creating a Pin List

Configure the Access Gateway to bypass the published URL of the proxy service:

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit.**

2 Click **Pin List** in the configuration page.

3 Click **New**, then specify the published DNS name of the proxy service. For example, `vibe.doc.provo.novell.com`.

4 Select **Bypass** as the Pin type.

5 Click **OK**.

6 To save the configuration changes, click **Devices** > **Access Gateways**, then click **Update**.

NOTE: If you do not want Access Manager to cache site information, do not create a Pin List. Instead, you should configure Access Manager to forward cache control headers to the browser. This is the recommended configuration for Novell Vibe. For information about how to forward cache control headers to the browser, see Section 6.2, "Controlling Browser Caching," on page 160.

2.5.5 Configuring Access to the Filr Site through Access Manager

For information about configuring Access Manager to configure a protected resource for a Novell Filr server, see Allowing Access Manager to configure a protected resource for a Novell Filr server (http://www.novell.com/documentation/novell-filr1/filr1_admin/data/btk7698.html)

2.6 Configuring HTML Rewriting

Access Gateway configurations generally require HTML rewriting because the Web servers are not aware that the Access Gateway machine is obfuscating their DNS names. URLs contained in their pages must be checked to ensure that these references contain the DNS names that the client browser understands. 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 2-6 illustrates these processes.

Figure 2-6   *HTML Rewriting*
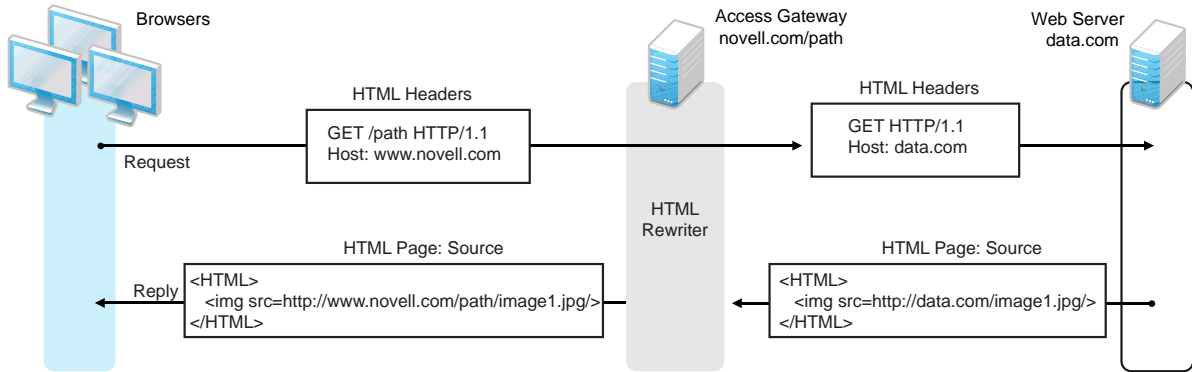
The following sections describe the HTML rewriting process:

- Section 2.6.1, "Understanding the Rewriting Process," on page 57
- Section 2.6.2, "Specifying DNS Names to Rewrite," on page 59
- Section 2.6.3, "Defining the Requirements for the Rewriter Profile," on page 62
- Section 2.6.4, "Configuring the HTML Rewriter and Profile," on page 69
- Section 2.6.5, "Creating or Modifying a Rewriter Profile," on page 72
- Section 2.6.6, "Disabling the Rewriter," on page 74

2.6.1 Understanding the Rewriting Process

The Access Gateway needs to rewrite URL references under the following conditions:

- To ensure that 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 sessions 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 `novell.com` might have a URL reference to `http://novell.com/path/image1.jpg`. If the reverse proxy for `novell.com/path` is using SSL sessions between the browser and Access Gateway, the URL reference `http://novell.com/path/image1.jpg` must be rewritten to `https://novell.com/path/image1.jpg`. Otherwise, when the user clicks the HTTP link, the browser must change from HTTP to HTTPS and establish a new SSL session.

- To ensure that URL references containing 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 named `data.com`, and wants to expose this site to Internet users through the Access Gateway by using a published DNS name of `novell.com`. Many of the HTML pages on this Web site have URL references that contain the private DNS name, such as `http://data.com/image1.jpg`. Because Internet users are unable to resolve `data.com/image1.jpg`, 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 `novell.com`, which users can resolve through a public DNS server to the Access Gateway. The rewriter parses the Web page, 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 rewritten to the published DNS name `novell.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 understood by the internal Web server.

Using the example in Figure 2-6 on page 57, suppose that the internal Web server expects all HTTP or HTTPS requests to have the **Host** field set to `data.com`. When users send requests using the published DNS name `novell.com/path`, the **Host** field of the packets in those requests received by the Access Gateway is set to `novell.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 `data.com`. Before the Access Gateway forwards packets to the Web server, the **Host** field is changed (rewritten) from `novell.com` to `data.com`. For information about configuring this option, see Section 2.3, "Configuring Web Servers of a Proxy Service," on page 30.

The rewriter searches for URLs in the following HTML contexts. They must meet the following criteria to be rewritten:

Context	Criteria
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 `index.html`) are not attempted. Absolute paths (such as `/docs/file.html`) are evaluated if the page is read from a path-based multi-homing Web server and the reference follows an HTML tag. For example, the string `href='/docs/file.html'` is rewritten if `/docs` is a multi-homing path that has been configured to be removed.
HTML Tags	URL references occurring within the following HTML tag attributes are evaluated for rewriting:

action	archive	background
cite	code	codebase
data	dynscr	filterLink
href	longdesc	lowsrc
o:WebQuerySourceHref	onclick	onmenuclick
pluginspage	src	usemap
usermapborderimage		

Context	Criteria
References	An absolute reference is a reference that has all the information needed to locate a resource, including the hostname, such as `http://internal.web.site.com/index.html`. The rewriter always attempts to rewrite absolute references.
	The rewriter attempts to rewrite an absolute path when it is the multi-homing path of a path-based multi-homing service. For example, `/docs/file1.html` is rewritten if `/docs` is a multi-homing path that has been configured to be removed.
	Relative references are not rewritten.
Query Strings	URL references contained within query strings can be configured for rewriting by enabling the **Rewrite Inbound Query String Data** option.
Post Data	URL references specified in Post Data can be configured for rewriting by enabling the **Rewrite Inbound Post Data** option.

2.6.2 Specifying DNS Names to Rewrite

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 when they 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.
- The DNS names in the **Exclude DNS Name List** specify the names that the rewriter should skip and not rewrite.

NOTE: Excludes in the **Exclude DNS Name List** are processed first, then the includes in the **Additional DNS Name List**. If you put the same DNS name in both lists, the DNS name is rewritten.

The following sections describe the conditions to consider when adding DNS names to the lists:

Determining Whether You Need to Specify Additional DNS Names

Sometimes Web pages contain URL references to a hostname 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 2-7 illustrates a scenario that requires an entry in the **Additional DNS Name List**.

Figure 2-7 Rewriting a URLs for Web Servers
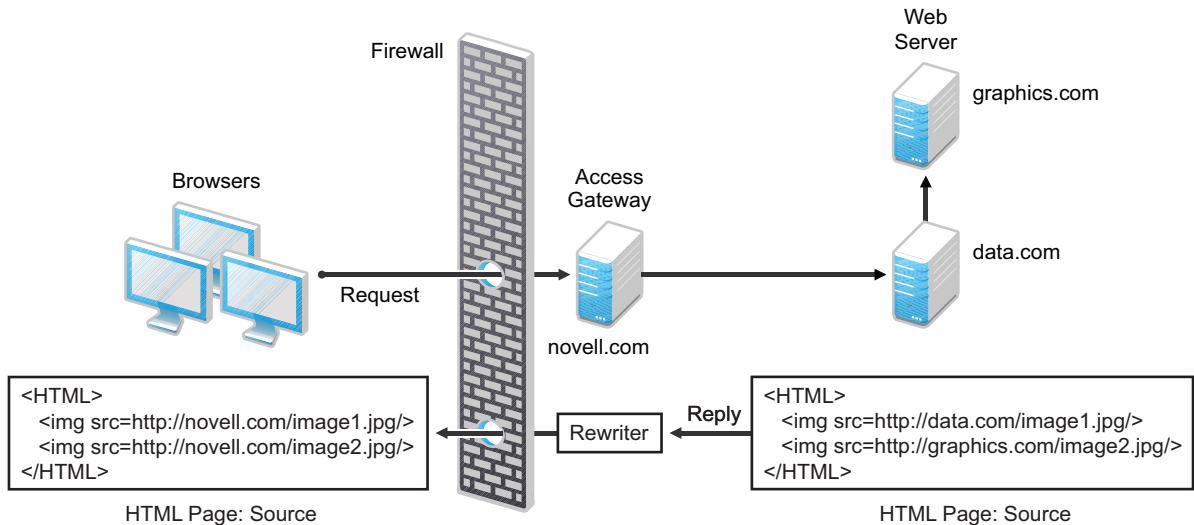
The page on the `data.com` Web server contains two links, one to an image on the `data.com` server and one to an image on the `graphics.com` server. The link to the `data.com` server is automatically rewritten to `novell.com`, when rewriting is enabled. The link to the image on `graphics.com` is not 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 for this additional Web server needs to be added to the list.

- If the Web server listens on one port (for example, 80), and redirects the request to a secure port (for example, 443), the DNS name needs to be added to the list. The response to the user comes back on `https://<DNS_name>:443`. This does not match the request that 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.

- If an application is written to use a private hostname, you need to add the private hostname to the list. For example, assume that an application URL reference contains the hostname of home (`http://home/index.html`). This hostname needs to be added to the **Additional DNS Name List**.

- If you enable the **Forward Received Host Name** option on your path-based multi-homing service and your Web server is configured to use a different port, you need to add the DNS name with the port to the **Additional DNS Name List**.

For example, if the public DNS name of the proxy service is `www.myag.com`, the path for the path-based multi-homing service is `/sales`, and the Web server port is 801, the following DNS name needs to be added to the **Additional DNS Name List** of the `/sales` service:

```
http://www.myag.com:801
```

When you enter a name in the list, it can use any of the following formats:

```
DNS_name
host_name
IP_address
scheme://DNS_name
scheme://IP_address
scheme://DNS_name:port
scheme://IP_address:port
```

For example:

```
HOME
https://www.backend.com
https://10.10.15.206:444
```

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 `novell.com.uk`, the response is returned with references to `novell.com.uk` and not `novell.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 2-8.

Figure 2-8   *Excluding URLs*
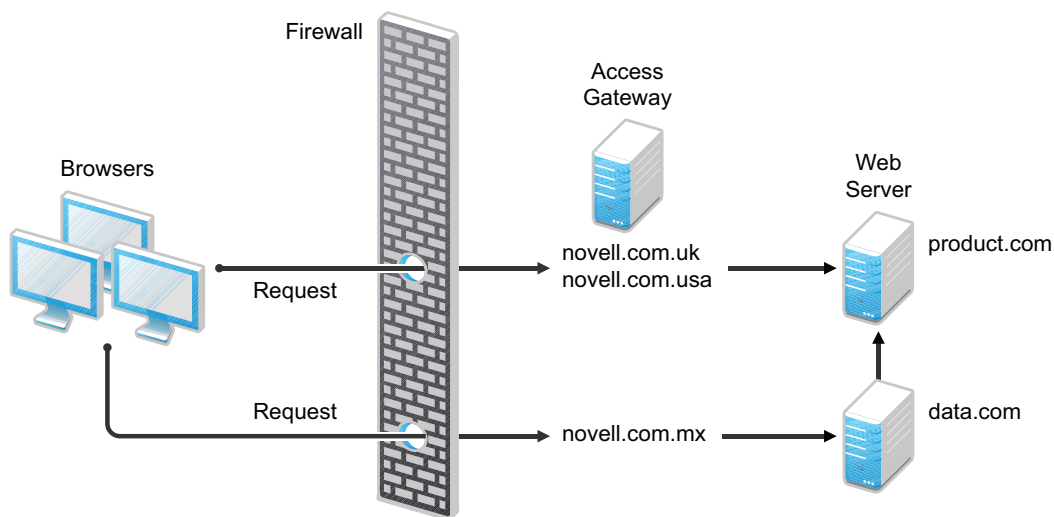
A user accesses `data.com` through the published DNS name of `novell.com.mx`. The `data.com` server has references to product.com. The `novell.com.mx proxy` has two ways to get to the `product.com` server because this Web server has two published DNS names (`novell.com.uk` and `novell.com.usa`). The rewriter could use either of these names to rewrite references to `product.com`.

- If you want all users coming through `novell.com.mx` to use the `novell.com.usa` proxy, you need to block the rewriting of `product.com` to `novell.com.uk`. On the HTML Rewriting page of the reverse proxy for `novell.com.uk`, add `product.com` and any aliases to the **Exclude DNS Name List**.
- 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**.

2.6.3 Defining the Requirements for the Rewriter Profile

An HTML rewriter profile allows you to customize the rewriting process and specify the profile that is selected to rewrite content on a page. This section describes the following features of the rewriter profile:

Types of Rewriter Profiles

The Access Gateway has the following types of profiles:

Default Word Profile

The default Word profile, named `default`, is not specific to a reverse proxy or its proxy services.

If you enable HTML rewriting, but you do not define a custom Word profile for the proxy service, the `default` Word profile is used. This profile is preconfigured to rewrite the **Web Server Host Name** and any other names listed in the **Additional DNS Name List**. The preconfigured profile matches all URLs with the following content-types:

text/html	text/javascript
text/xml	application/javascript
text/css	application/x-javascript

When you modify the behavior of the default profile, remember its scope. If the default profile does not match your requirements, you should usually create your own custom Word profile or custom Character profile.

Custom Word Profile

A Word profile searches for matches on words. For example, "get" matches the word "get" and any word that begins with "get" such as "getaway" but it does not match the "get" in "together" or "beget."

For information about how strings are replaced in a Word profile, see the following:

- "String Replacement Rules for Word Profiles" on page 66
- "Using $path to Rewrite Paths in JavaScript Methods or Variables" on page 68

You should create a custom Word profile when an application requires rewrites of paths in JavaScript. If the application needs strings replaced or new content-types, these can also be added to the custom profile. In a custom Word profile, you can also configure the match criteria so that the profile matches specific URLs. For more information, see "Page Matching Criteria for Rewriter Profiles" on page 63.

When you create a custom Word profile, you need to position it before the `default` profile in the list of profiles. Only one Word profile is applied per page. The first Word profile that matches the page is applied. Profiles lower in the list are ignored.

Custom Character Profile

A custom Character profile searches for matches on a specified set of characters. For example, "top" matches the word "top" and the "top" in "tabletop," "stopwatch," and "topic." If you need to replace strings that require this type of search, you should create a custom Character profile.

For information about how strings are replaced in a Character profile, see "String Replacement Rules for Character Profiles" on page 67.

In a custom Character profile, you can also configure the match criteria so that the profile matches specific URLs. For more information, see "Page Matching Criteria for Rewriter Profiles" on page 63.

After the rewriter finds and applies the Word profile that matches the page, it finds and applies one Character profile. The first Character profile that matches the page is applied. Character profiles lower in the list are ignored.

Page Matching Criteria for Rewriter Profiles

You specify the following matching criteria for selecting the profile:

- The URLs to match
- The URLs that cannot match
- The content types to match

You use the **Requested URLs to Search** section of the profile to set up the matching policy. The first Word profile and the first Character profile that matches the page is applied. Profiles lower in the list are ignored.

URLs: The URLs specified in the policy should use the following formats:

Sample URL	Description
`http://www.a.com/content`	Matches pages only if the requested URL does not contain a trailing slash.
`http://www.a.com/content/`	Matches pages only if the requested URL does contain a trailing slash.
`http://www.a.com/content/index.html`	Matches only this specific file.
`http://www.a.com/content/*`	Matches the requested URL whether or not it has a trailing slash and matches all files in the directory.
`http://www.a.com/*`	Matches the proxy service and everything it is protecting.

You can specify two types of URLs. In the **If Requested URL Is** list, you specify the URLs of the pages you want this profile to match. In the **And Requested URL Is Not** list, you specify the URLs you don't want this profile to match. You can use the asterisk wildcard for a URL in the **If Requested URL Is** list to match pages you really don't want this profile to match, then use a URL in the **And Requested URL Is Not** list to exclude them from matching. If a page matches both a URL in the **If Requested URL Is** list and in the **And Requested URL Is Not** list, the profile does not match the page.

For example, you could specify the following URL in the **If Requested URL Is** list:

`http://www.a.com/*`

You could then specify the following URL in the **And Requested URL Is Not** list:

`http://www.a.com/content/*`

These two entries cause the profile to match all pages on the `www.a.com` Web server except for the pages in the `/content` directory and its subdirectories.

IMPORTANT: If nothing is specified in either of the two lists, the profile skips the URL matching requirements and uses the content-type to determine if a page matches.

Content-Type: In the **And Document Content-Type Is** section, you specify the content-types you want this profile to match. To add a new content-type, click **New** and specify the name, such as `text/dns`. Search your Web pages for content-types to determine if you need to add new types. To add multiple values, enter each value on a separate line.

Regardless of content-types you specify, the page matches the profile if the file extension is `html`, `htm`, `shtml`, `jhtml`, `asp`, or `jsp` and you have not specified any URL matching criteria.

Possible Actions for Rewriter Profiles

The rewriter action section of the profile determines the actions the rewriter performs when a page matches the profile. Select from the following:

- Inbound Actions
- Enabling or Disabling Rewriting
- Additional Names to Search for URL Strings to Rewrite with Host Name
- String Replacement

Inbound Actions: A profile might require these options if the proxy service has the following characteristics:

- URLs appear in query strings, Post Data, or headers.
- The Web server uses WebDAV methods.

If your profile needs to match pages from this type of proxy service, you might need to enable the options listed below. They control the rewriting of query strings, Post Data, and headers from the Access Gateway to the Web server.

- **Rewrite Inbound Query String Data:** Select this option to rewrite the domain and URL in the query string to match the Web server configuration or to remove the path from the query string on a path-based multi-homing proxy with the **Remove Path on Fill** option enabled.
- **Rewrite Inbound Post Data:** Select this option to rewrite the domain and URL in the Post Data to match the Web server configuration or to remove the path from the Post Data on a path-based multi-homing proxy with the **Remove Path on Fill** option enabled.
- **Rewrite Inbound Headers:** Select this option to rewrite the following headers:

  Call-Back
  Destination
  If
  Notification-Type
  Referer

The inbound options are not available for a Character profile.

Enabling or Disabling Rewriting: The **Enable Rewriter Actions** option determines whether the rewriter performs any actions:

- Select the option to have the rewriter rewrite the references and data on the page.
- Leave the option deselected to disable rewriting. This allows you to create a profile for the pages you do not want rewritten.

Additional Names to Search for URL Strings to Rewrite with Host Name: Use this section to specify the name of the variable, attribute, or method in which the hostname might appear. These options are not available for a Character profile.

- **Variable and Attribute Name to Search for Is:** Use this section to specify the HTML attributes or JavaScript variables that you want searched for DNS names that might need to be rewritten. For the list of HTML attribute names that are automatically searched, see "HTML Tags" on page 58. You might want to add the following attributes:
  - **value:** This attribute enables the rewriter to search the `<param>` elements on the HTML page for value attributes and rewrite the value attributes that are URL strings.

    If you need more granular control (some need to be rewritten but others do not) and you can modify the page, see "Disabling with Page Modifications" on page 75.
  - **formvalue:** This attribute enables the rewriter to search the `<form>` element on the HTML page for `<input>`, `<button>`, and `<option>` elements and rewrite the value attributes that are URL strings. For example, if your multi-homing path is `/test` and the form line is `<input name="navUrl" type="hidden" value="/IDM/portal/cn/ GuestContainerPage/656gwmail">`, this line would be rewritten to the following value before sending the response to the client:

```
<input name="navUrl" type="hidden" value="/test/IDM/portal/cn/
GuestContainerPage/656gwmail">
```

The formvalue attribute enables the rewriting of all URLs in the `<input>`, `<button>`, and `<option>` elements in the form.If you need more granular control (some need to be rewritten but others do not) and you can modify the form page, see "Disabling with Page Modifications" on page 75.

 ◆ **Replacing URLs in Java Methods:** The **JavaScript Method to Search for Is** list allows you to specify the Java methods to search to see if their parameters contain a URL string.

String Replacement: The **Additional Strings to Replace** list allows you to search for a string and replace it. The search boundary (word or character) that you specified when creating the profile is used when searching for the string.

Word profile search and replace actions take precedence over character profile actions.

For the rules and tokens that can be used in the search strings, see the following:

 ◆ "String Replacement Rules for Word Profiles" on page 66
 ◆ "String Tokens" on page 66
 ◆ "String Replacement Rules for Character Profiles" on page 67

For information about how the **Additional Strings to Replace** list can be used to reduce the number of Java methods you need to list, see "Using $path to Rewrite Paths in JavaScript Methods or Variables" on page 68.

String Replacement Rules for Word Profiles

In a Word profile, a string matches all paths that start with the characters in the specified string. For example:

Search String	Matches This String	Doesn't Match This String
/path	/path	/mypath
	/pathother	
	/path/other	
	/path.html	

String Tokens

On the Access Gateway Service, you can use the following special tokens to modify the default matching rules.

 ◆ [w] to match one white space character
 ◆ [ow] to match 0 or more white space characters
 ◆ [ep] to match a path element in a URL path, excluding words that end in a period
 ◆ [ew] to match a word element in a URL path, including words that end in a period
 ◆ [oa] to match one or more alphanumeric characters

White Space Tokens: You use the [w] and the [ow] tokens 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.

Path Tokens: You use the [ep] and [ew] tokens to match path strings. The [ep] token can be used to match the following types of paths:

Search String	Matches This String	Doesn't Match This String
/path[ep]	/path	/path.html
	/home/path/other	/home/pathother

The [ew] token can be used to match the following types of paths:

Search String	Matches This String	Does not Match This String
/path[ew]	/path.html	/paths
	/home/path	

Name Tokens: You use the [oa] token to match function or parameter names that have a set string to start the name and end the name, but the middle part of the name is a computer-generated alphanumeric string. For example, the [oa] token can be used to match the following types of names:

Search String	Matches This String	Doesn't Match This String
javaFunction-[oa](	javaFunction-1234a56()	javaFunction()
	javaFunction-a()	

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. Only one replacement pass 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.

Using $path to Rewrite Paths in JavaScript Methods or Variables

You can use the $path token to rewrite paths on a path-based multi-homing service that has the **Remove Path on Fill** option enabled. This token is useful for Web applications that require a dedicated Web server and are therefore installed in the root directory of the Web server. If you protect this type of application with Access Manager using a path-based multi-homing service, your clients access the application with a URL that contains a /path value. The proxy service uses the path to determine which Web server a request is sent to, and the path must be removed from the URL before sending the request to the Web server.

The application responds to the requests. If it uses JavaScript methods or variables to generate paths to resources, these paths are sent to client without prepending the path for the proxy service. When the client tries to access the resource specified by the Web server path, the proxy service cannot locate the resource because the multi-homing path is missing. The figure below illustrates this flow with the rewriter adding the multi-homing path in the reply.

Figure 2-9   *Rewriting with a Multi-homing Path*
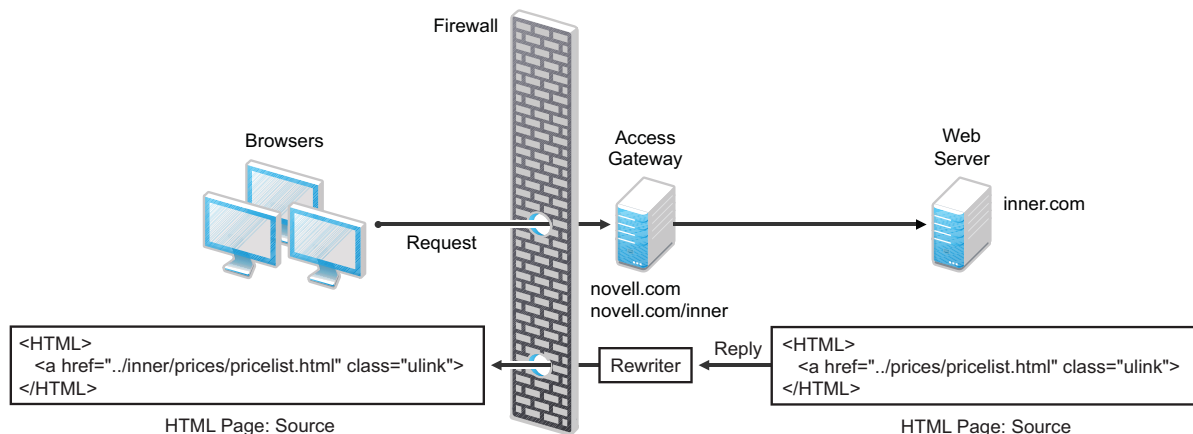
To make sure all the paths generated by JavaScript are rewritten, you must search the Web pages of the application. You can then either list all the JavaScript methods and variables in the **Additional Names to Search for URL Strings to Rewrite with Host Name** section of the rewriter profile, or you can use the $path token in the **Additional Strings to Replace** section. The $path token reduces the number of JavaScript methods and variables that you otherwise need to list individually.

To use the $path token, you add a search string and a replace string that uses the token. For example, if the /prices/pricelist.html page is generated by JavaScript and the multi-homing path for the proxy service is /inner, you would specify the following stings:

Search String	Replacement String
/prices	$path/prices

This configuration allows the following paths to be rewritten before the Web server sends the information to the browser.

Web Server String	Rewritten String for the Browser
/prices/pricelist.html	/inner/prices/pricelist.html
/prices	/inner/prices

This token can cause strings that shouldn't be changed to be rewritten. If you enable the **Rewrite Inbound Query String Data**, **Rewrite Inbound Post Data,** and **Rewrite Inbound Header** actions, the rewriter checks these strings and ensures that they contain the information the Web server expects.

For example, when these options are enabled, the following paths and domain names are rewritten when found in query strings, in Post Data, or in the Call-Back, Destination, If, Notification-Type, or Referer headers.

Browser String	Rewritten String for the Web Server
/inner/prices/pricelist.html	/prices/pricelist.html
/inner/prices	/prices
novell.com/inner/prices	inner.com/prices

2.6.4 Configuring the HTML Rewriter and Profile

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 **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **HTML Rewriting**.

The HTML Rewriting page specifies which DNS names are 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 .

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 .

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 .
  - **Word:** A Word profile searches for matches on words. For example, "get" matches the word "get" and any word that begins with "get" such as "getaway" but it does not match the "get" in "together" or "beget."

If you create multiple Word profiles, order is important. The first Word profile that matches the page is applied. Word profiles lower in the list are ignored.

- **Character:** A Character profile searches for matches on a specified set of characters. For example, "top" matches the word "top" and the "top" in "tabletop," "stopwatch," and "topic."

  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 applied. Character profiles lower in the list are ignored.

- **Delete:** To delete a profile, select the profile, then click **Delete**.

- **Enable:** To enable a profile, select the profile, then click **Enable**.

- **Disable:** To disable a profile, select the profile, then click **Disable**.

- **Modify:** To view or modify the current configuration for a profile, click the name of the profile. Continue with Section 2.6.5, "Creating or Modifying a Rewriter Profile," on page 72.

  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 custom Word profile and enable it.

6 If you have more than one profile in the **HTML Rewriter Profile List**, use the up-arrow and down-arrow buttons 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.

Even if multiple Word or Character profiles are enabled, a maximum of one Word profile and one Character profile is executed per page. The first Word profile and Character profile in the list that matches a page are executed, and the others are ignored.

7 Enable the profiles you want to use for this protected resource. Select the profile, then click **Enable**.

The `default` profile cannot be disabled. However, it is not executed if you have enabled another Word profile that matches your pages, and this profile comes before the `default` profile in the list.

8 To save your changes to browser cache, click **OK**.

9 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

10 The cached pages affected by the rewriter changes must be updated on the Access Gateway. Do one of the following:

- If the changes affect numerous pages, click **Access Gateways**, select the name of the server, then click **Actions** > **Purge All Cache**.

- If the changes affect only a few pages, you can refresh or reload the pages within the browser.

2.6.5 Creating or Modifying a Rewriter Profile

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **HTML Rewriting**.

2 Select one of the following:

- ◆ To create a new profile, click **New**, specify a name, select a profile type, then click **OK**.
- ◆ To modify a profile, click the name of the profile.

3 Use the **Requested URLs to Search** section to set up a policy for specifying the URLs you want this profile to match.

Fill in the following fields:

If Requested URL Is: Specify the URLs of the pages you want this profile to match. Click **New** to add a URL to the text box. To add multiple values, enter each value on a separate line.

And Requested URL Is Not: Specify the URLs of pages that this profile should not match. If a page matches the URL in both the **If Requested URL Is** list and **And Requested URL Is Not** list, the profile does not match the page. Click **New** to add a URL to the text box. To add multiple values, enter each value on a separate line.

And Document Content-Type Is: Select the content-types you want this profile to match. To add a new content-type, click **New** and specify the name such as `text/dns`. Search your Web pages for content-types to determine if you need to add new types. To add multiple values, enter each value on a separate line.

For more information about how to use these options, see "Page Matching Criteria for Rewriter Profiles" on page 63.

4 Use the **Actions** section to specify the actions the rewriter should perform if the page matches the criteria in the **Requested URLs to Search** section.

Configure the following actions:

Rewrite Inbound Query String Data: (Not available for Character profiles) Select this option to rewrite the domain and URL in the query string to match the Web server. To use this option, your proxy service must meet the conditions listed in "Possible Actions for Rewriter Profiles" on page 64.

Rewrite Inbound Post Data: (Not available for Character profiles) Select this option to rewrite the domain and URL in the Post Data to match the Web server. To use this option, your proxy service must meet the conditions listed in "Possible Actions for Rewriter Profiles" on page 64.

Rewrite Inbound Headers: Select this option to rewrite the following headers:

Call-Back

Destination

If

Notification-Type

Referer

Enable Rewriter Actions: Select this action to enable the rewriter to perform any actions:

- Select it to have the rewriter use the profile to rewrite references and data on the page. If this option is not selected, you cannot configure the action options.

- Leave it unselected to disable rewriting. This allows you to create a profile for the pages you do not want rewritten.

5 (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 58.)

Fill in the following fields:

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 text box. To add multiple values, enter each value on a separate line.

JavaScript Method to Search for Is: Lists the names of Java methods to search to see if their parameters contain a URL string. Click **New** to add a method to the text box. To add multiple values, enter each value on a separate line.

6 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.

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 67
- "String Replacement Rules for Word Profiles" on page 66
- "Using $path to Rewrite Paths in JavaScript Methods or Variables" on page 68

Replace With: Specify the string you want to use in place of the search string.

7 Click **OK**.

8 If you have more than one profile in the **HTML Rewriter Profile List**, use the up-arrow and down-arrow buttons 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.

Even if multiple Word or Character profiles are enabled, a maximum of one Word profile and one Character profile is executed per page. The first Word profile and Character profile in the list that matches a page are executed, and the others are ignored.

9 Enable the profiles you want to use for this protected resource. Select the profile, then click **Enable**.

The default profile cannot be disabled. However, it is not executed if you have enabled another Word profile that matches your pages, and this profile comes before the default profile in the list.

10 To save your changes to browser cache, click **OK**.

11 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

12 The cached pages affected by the rewriter changes must be updated on the Access Gateway. Do one of the following:

- If the changes affect numerous pages, click **Access Gateways**, select the name of the server, then click **Actions** > **Purge All Cache**.
- If the changes affect only a few pages, refresh or reload the page within the browser.

2.6.6 Disabling the Rewriter

There are three methods you can use to disable the internal rewriter:

- "Disabling per Proxy Service" on page 75
- "Disabling per URL" on page 75
- "Disabling with Page Modifications" on page 75

Disabling per Proxy Service

By default, the rewriter is enabled for all proxy services. The rewriter can slow performance because of the parsing overhead. In some cases, a Web site might not have content with URL references that need to be rewritten. The rewriter can be disabled on the proxy service that protects that Web site.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **HTML Rewriting**.

2 Deselect the **Enable HTML Rewriting** option, then click **OK**.

3 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

4 Select the Access Gateway, then click **Actions** > **Purge All Cache** > **OK**.

Disabling per URL

You can also specify a list of URLs that are to be excluded from being rewritten for the selected proxy service.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **HTML Rewriting**.

2 Click the name of the Word profile defined for this proxy service.

If you have not defined a custom Word profile for the proxy service, you might want to create one. If you modify the `default` profile, those changes are applied to all proxy services.

3 In the **And Requested URL Is Not** section, click **New**, then specify the names of the URLs you do not want rewritten.

Specify each URL on a separate line.

4 Click **OK** twice.

5 In the **HTML Rewriter Profile List**, make sure the profile you have modified is enabled and at the top of the list, then click **OK**.

6 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

7 Select the Access Gateway, then click **Actions** > **Purge All Cache** > **OK**.

Disabling with Page Modifications

There are cases when the URLs in only part of a page or in some of the JavaScript or form can be rewritten and the rest should not be rewritten. When this is the case, you might need to modify the content on the Web server. Although this deviates from the design behind Access Manager, you might encounter circumstances where it cannot be avoided.

You can add the following types of tags to the pages on the Web server:

- Page Tags
- Param Tags
- Form Tags

These tags are seen by browsers as a comment mark, and do not show up on the screen (except possibly on older browser versions).

NOTE: If the pages you modify are cached on the Access Gateway, you need to purge the cache before the changes become effective. Click **Access Gateways**, select the name of the server, then click **Actions** > **Purge All Cache**

Page Tags: If you want only portions of a page rewritten, you can add the following tags to the page.

```
<!--NOVELL_REWRITER_OFF-->
.
.
HTML data not to be rewritten
.
.
<!--NOVELL_REWRITER_ON-->
```

The last tag is optional, and if omitted, it prevents the rest of the page from being rewritten after the `<!--NOVELL_REWRITER_OFF-->` tag is encountered.

Param Tags: Sometimes the JavaScript on the page contains `<param>` elements that contain a value attribute with a URL. You can enable global rewriting of this attribute by adding `value` to the list of variable and attribute names to search for. If you need more control because some URLs need to be rewritten but others cannot be rewritten, you can turn on and turn off the `value` rewriting by adding the following tags before and after the `<param>` element in the JavaScript.

```
<!--NOVELL_REWRITE_ATTRIBUTE_ON='value'-->
.
.
<param> elements to be rewritten
.
.
<!--NOVELL_REWRITE_ATTRIBUTE_OFF='value'-->
.
.
<param> elements that shouldn't be rewritten
```

Form Tags: Some applications have forms in which the `<input>`, `<button>`, and `<option>` elements contain a value attribute with a URL. You can enable global rewriting of these attributes by adding `formvalue` to the list of variable and attribute names to search for. If you need more control because some URLs need to be rewritten but others cannot be rewritten, you can turn on and turn off the `formvalue` rewriting by adding the following tags before and after the `<input>`, `<button>`, and `<option>` elements in the form.

```
<!--NOVELL_REWRITE_ATTRIBUTE_ON='formvalue'-->
.
.
<input>, <button>, and <option> elements to be rewritten
.
.
<!--NOVELL_REWRITE_ATTRIBUTE_OFF='formvalue'-->
.
.
<input>, <button>, and <option> elements that shouldn't be rewritten
```

2.7 Configuring Connection and Session Limits

The Access Gateway establishes connections with clients and with Web servers. For most networks, the default values for unresponsive connections and sessions provide adequate performance, but you can fine-tune the options for your network, its performance requirements, and your users:

Authentication time limits for inactivity sessions are configured on the contract and enforced by the Identity Server. For information about how to configure this limit, see Section 2.4.6, "Assigning a Timeout Per Protected Resource," on page 41.

2.7.1 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. If your network is congested and slow, you might want to increase some of the limits.

1 In the Administration Console, click **Devices** > **Access Gateways > Edit > [Name of Reverse Proxy] > TCP Listen Options**.

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 **TCP Listen Options**:

Keep Alive Interval: 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 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).

4 To configure the encryption key, select one or more of the following:

Enforce 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.

Enforce 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.

NOTE: These SSL listening options appear disabled if you are configuring the tunneling services.

5 To save your changes to browser cache, click **OK**.

6 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

2.7.2 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. If your network is congested and slow, you might want to increase some of the limits.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **Web Servers** > **TCP Connect Options**.

2 Configure the IP address to use when establishing connections with Web servers:

Cluster Member: (Available only if the Access Gateway is a member of a cluster.) Select the server you want to configure from the list of servers. Only the value of the **Make Outbound Connection Using** option applies to the selected server.

3 Select how the Web servers should be contacted when multiple Web servers are available. Select one of the following for the **Policy for Multiple Destination IP Addresses** option:

 ◆ **Simple Failover:** Allows the next available Web server in the group to be contacted when the first server in the list is no longer available.

 ◆ **Round Robin:** Moves in order through the list of Web servers, allowing each to service requests before starting at the beginning of the list for a second group of requests.

4 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.

5 To modify the connection timeouts between the Access Gateway and the Web servers, configure the following fields:

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 (Keep Alive Interval): 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).

6 To save your changes to browser cache, click **OK**.

7 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

2.7.3 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 connections to the browser and the Access Gateway connections to the Web server 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 **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **TCP Listen Options** and configure the **Enable Persistent Connections** option.

- **Access Gateway to Web Server:** Click **Devices** > **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.

2.7.4 Configuring Web Servers

The Web server configuration determines how the Access Gateway handles connections and packets between itself and the Web servers. For more information about Web Server configuration, see Section 2.3, "Configuring Web Servers of a Proxy Service," on page 30

1 Click **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **Web Servers**.

2 If your browsers are capable of sending HTTP 1.1 requests, configure the following field 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.

3 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 1.4, "Configuring SSL between the Proxy Service and the Web Servers," on page 16.

4 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

5 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 2.7.2, "Configuring TCP Connect Options for Web Servers," on page 78.

6 To add a Web server, click **New** in the **Web Server List** and specify the IP address or the fully qualified DNS name of the Web server.

   ◆ **New:** To create a new Web server, click **New**. Specify the Web Server IP Address or DNS. Click OK to add the new Web server to the list or Cancel to discard the changes.

   After creating the Web server in the list, you can configure it as primary server and prioritize the list of Web servers based on your requirement.

   ◆ **Delete:** To delete a Web server, select the Web server from the list, then click **Delete**.

   If you delete the selected Web server, then all the Web servers which are corresponding to the device in the cluster gets deleted.

7 To save your changes to browser cache, click **OK**.

8 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

3 Server Configuration Settings

This section describes the configuration settings that affect the Access Gateway as a server, such as changing its name or setting the time.

For logging and audit options, see the following:

For cache management, see Chapter 5, "Advanced Access Gateway Options," on page 145.

3.1 Configuration Overview

The Configuration page allows you to view the configuration status and to configure the features of the cluster or the Access Gateway. After an Access Gateway has been made a member of a cluster, you can only configure it from the cluster configuration. Some options are specific to an Access Gateway. For these options, you must select the Access Gateway and then configure the options.

1 In the Administration Console, **Devices** > **Access Gateways** > **Edit**.

   To edit an Access Gateway that is not a member of a cluster, click the **Edit** button on the Access Gateway row.

   To edit an Access Gateway cluster, click the **Edit** button on the Access Gateway cluster row.

2 Select one of the following options:

   Reverse Proxy / Authentication: Allows you to configure a reverse proxy so that it hides the IP address of a Web server and accelerates access by caching the most frequently used pages. This option displays the list of configured proxies and allows you to add new proxies and modify existing proxies. To add a new reverse proxy or manage the existing proxies, click **Reverse Proxy / Authentication** (see Section 2.2, "Managing Reverse Proxies and Authentication," on page 23). To manage a specific reverse proxy, click its name (see Section 2.2.1, "Creating a Proxy Service," on page 26).

Tunneling: Allows you to tunnel non-HTTP traffic through the Access Gateway to a Web server. For more information, see Section 3.5, "Setting Up a Tunnel," on page 92.

Date & Time: Allows you to configure the server's time source. For more information, see Section 3.6, "Setting the Date and Time," on page 93.

Alerts: Allows you to select the alerts and then configure whether they are sent to a server, a log file, or to selected individuals via e-mail. For more information, see Section 4.6.3, "Managing Access Gateway Alert Profiles," on page 134.

Auditing: Allows you to select the events to send to a NetIQ Sentinel or Audit server. For more information, see Section 4.7, "Enabling Access Gateway Audit Events," on page 137.

Adapter List: Displays the list of configured network cards and allows you to edit an existing configuration or to add a new one. For more information, see Section 3.8.1, "Viewing and Modifying Adapter Settings," on page 98. To manage a specific adapter, click the name of the adapter.

Gateways: Displays the list of configured gateways and allows you to edit an existing configuration or to add a new one. For more information, see Section 3.8.2, "Viewing and Modifying Gateway Settings," on page 99.

DNS: Displays the current DNS configuration that the Access Gateway is using to resolve names and allows you to modify it. For more information, see Section 3.8.3, "Viewing and Modifying DNS Settings," on page 101.

Hosts: Allows you to create a static mapping between the host IP addresses and host names. For more information, see Section 3.8.4, "Configuring Hosts," on page 102.

Purge List: Allows you to prevent Web objects from being cached. For more information, see Section 6.5, "Configuring a Purge List," on page 166.

Pin List: Allows you to prepopulate the cache with the Web objects that you want cached, before a user has requested the object. For more information, see Section 6.4, "Configuring a Pin List," on page 164.

Cache Options: Allows you to globally disable caching or configure which objects are cached and how frequently they are refreshed. For more information, see Configuring Caching Options.

Advanced Options: Allows you to configure how all reverse proxies handle specific items in cache. For more information, see Section 5.1, "Configuring the Global Advanced Options," on page 145.

3 For information about using the **OK**, **Cancel**, and **Revert** buttons, see Section 3.2, "Saving, Applying, or Canceling Configuration Changes," on page 82.

3.2 Saving, Applying, or Canceling Configuration Changes

When you make configuration changes on a page accessed from **Devices** > **Access Gateways** > **Edit** and click **OK** on that page, the changes are saved to the browser cache. If your session expires or you close the browser session before you update the Access Gateway with the changes, the changes are lost.

The Configuration page allows you to control how your changes are saved so they can be applied with the update options (see Section 3.3.2, "Configuration Options," on page 86).

If you have any configuration changes saved to the browser cache, use the following options to control what happens to the changes:

OK: To save the configuration changes to the configuration store, click **OK**. This allows you to return at a later time to review or modify the changes before they are applied. If your Access Gateways are clustered and you prefer to update them one at a time, you need to save the configuration change. This ensures that the changes aren't lost before the last cluster member is updated. When your session times out or you log out, the configuration changes are flushed from the browser cache. If this happens before the changes have been applied to some servers in the cluster, the changes cannot be applied to those servers.

If you decide to cancel the saved changes, click the **Revert** button and the saved configuration is overwritten by the last successfully applied configuration.

Cancel: To cancel changes that are pending in the browser cache, click the **Cancel** button. To cancel modifications to specific services, click the **Cancel** link by the service. The **Cancel** button does not affect the changes that have been saved to the configuration store.

Revert: To cancel any saved changes, click **Revert**, then confirm the cancellation. The saved configuration is overwritten by the last successfully applied configuration.

If you have applied the changes to one member of the cluster, you cannot use the **Revert** button to revert to the configuration you had before applying the changes. If you decide you do not want to apply these changes to other members of the cluster, remove the server that you updated with the changes from the cluster. Then click **Revert** to cancel the saved changes. The members of the cluster return to the last successfully applied configuration. To apply this configuration to the removed server, add this server to the cluster.

The **Revert** button and the **Cancel** button cannot cancel the following configuration changes:

- ◆ **Identity Server Cluster:** If you change the **Identity Server Cluster** option on the Reverse Proxy/Authentication page, then click **OK**, the **Revert** button cannot cancel this change. It is saved, and the next time you apply a configuration change, the Identity Server cluster configuration is applied. To cancel the change, you need to return to the Reverse Proxy/ Authentication page, set the **Identity Server Cluster** option to the original selection, then click **OK on the Configuration page**.

- ◆ **Reverse Proxy for the Embedded Service Provider:** If you change the **Reverse Proxy** option on the Reverse Proxy/Authentication page, then click **OK**, the **Revert** button cannot cancel this change. It is saved, and the next time you apply a configuration change, the **Reverse Proxy** option change is applied. To cancel the change, return to the Reverse Proxy/Authentication page, set the **Reverse Proxy** option to the original selection, then click **OK** on the Configuration page.

- ◆ **Port of the Reverse Proxy for the Embedded Service Provider:** If you change the port of the reverse proxy that is used by the Embedded Service Provider (click **Edit** > **[Name of Reverse Proxy]**), then click **OK**, the **Revert** button cannot cancel this change. It is saved, and the next time you apply a configuration change, the port change is applied. To cancel the change, return to the Reverse Proxy page, set the port to the original value, then click **OK** on the Configuration page.

- ◆ **Published DNS Name of the Proxy Service for the Embedded Service Provider:** If you change the Published DNS Name of the proxy service that is used by the Embedded Service Provider (click **Edit** > **[Name of Reverse Proxy] > [Name of Proxy Service]**), then click **OK**, the **Revert** button cannot cancel this change. It is saved, and the next time you apply a configuration change, the Published DNS Name is changed. To undo the change, return to the Proxy Service page, set the Published DNS Name to its original value, then click **OK** on the Configuration page.

- ◆ **Certificates:** Certificates are pushed as soon as they are selected. If you change the server certificate for the reverse proxy (click **Edit** > **[Name of Reverse Proxy]**) or change the Web server certificates (click **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **Web Servers)**, the **Revert** button cannot cancel these changes. To undo the change, return to the page, select the original certificate, then click **OK**.

◆ **Renaming a Reverse Proxy:** If you change the name of a reverse proxy (click **Edit** > **Reverse Proxies / Authentication**), then click **OK**, you cannot cancel this change. To undo the change, return to the Reverse Proxies / Authentication page, rename the reverse proxy to its original name, then click **OK** and update the Access Gateway.

3.3 Managing Access Gateways

The following sections contain information about settings available with Access Gateways, changing the settings, and their impact on users:

3.3.1 Viewing and Modifying Gateway Settings

Use the Servers page to view the status of Access Gateways, to modify their configuration, and to perform other actions such as creating a new cluster or stopping and starting an Access Gateway or its Embedded Service Provider.

1 In the Administration Console, click **Devices** > **Access Gateways**.

2 Select one of the following:

Stop: To stop an Access Gateway, select the service, then click **Stop**. You can use the **Restart** option to start the Access Gateway.

Restart: To stop and start an Access Gateway, select it, then click **Restart**. If the Access Gateway is already stopped, use **Restart** to start it.

Refresh: To update the list of Access Gateways and the status columns (**Status**, **Health**, **Alerts**, **Commands**), click **Refresh**.

3 To perform an action available in the **Actions** drop-down menu, select an Access Gateway, then select one of the following:

Schedule Restart: To schedule when the selected Access Gateway should be stopped and then started, select **Schedule Restart**. On an Access Gateway Service, a restart stops the Access Gateway Service, then starts it. For information about how to schedule this command, see Section 3.3.3, "Scheduling a Command," on page 89.

Schedule Stop: To schedule when the selected Access Gateway or cluster should be stopped, select **Schedule Stop**.

You can use the **Restart** option to start it again.

For more information about how to schedule this command, see Section 3.3.3, "Scheduling a Command," on page 89

Purge List Now: Click **Purge List Now** to cause all objects in the current purge list to be purged from the cache of the selected server or cluster.

Purge All Cache: Click **Purge All Cache** to purge the server cache for the selected server or cluster. All cached content is lost.

When you make certain configuration changes such as updating or changing certificates, changing the IP addresses of Web servers, or modifying the rewriter configuration, you are prompted to purge the cache. The cached objects must be updated for users to see the effects of such configuration changes. If your Access Gateways are in a cluster, you need to manage the

purge process so your site remains accessible to your users. You should apply the configuration changes to one member of a cluster. When its status returns to healthy and current, issue the command to purge its cache. Then apply the changes to the next cluster member.

IMPORTANT: Do not issue a purge cache command when an Access Gateway has a pending configuration change. Wait until the configuration change is complete.

Update Health from Server: Click this action to send a request to the server for updated health information. If you have selected multiple servers, a request is sent to each one. The health status changes to an animated circle until the reply returns.

Service Provider: Select one of the following actions:

- **Start Service Provider:** To start the Embedded Service Provider associated with the selected Access Gateway, click **Start Service Provider**. The Embedded 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.

- **Stop Service Provider:** To stop the Embedded Service Provider associated with the selected Access Gateway, click **Stop Service Provider**. The Embedded Service Provider is the module within the Access Gateway that communicates with the Identity Server.

  When an Access Gateway is not functioning correctly, you should always try stopping and starting the service provider before stopping and starting the Access Gateway.

- **Restart Service Provider:** To restart the Embedded Service Provider associated with the selected Access Gateway, click **Restart Service Provider**. This command stops the Embedded Service Provider and then starts it. The Embedded Service Provider is the module within the Access Gateway that communicates with the Identity Server.

  When an Access Gateway is not functioning correctly, you should always try restarting the service provider before stopping and starting the Access Gateway.

4 Use the following links to manage a cluster or an Access Gateway.

Name: Displays a list of the Access Gateway servers and the clusters that can be managed from this Administration Console.

- To view or modify the general details of a particular server, click the name of the server.
- To view or modify general details of a cluster, click the name of the cluster.

Status: Indicates the configuration status of the clusters and the Access Gateways. Possible states are pending, update, current, and update all. For more information, see Section 3.3.2, "Configuration Options," on page 86.

Health: Indicates whether a cluster or an Access Gateway is functional. Click the icon to view additional information about the operational status of an Access Gateway.

- For information about the health of a specific Access Gateway, click the health icon on the Access Gateway row. For more information, see Section 4.8.2, "Monitoring the Health of an Access Gateway," on page 139.
- For information about the health of a Access Gateway cluster, click the health icon on the cluster row. For more information, see Section 4.8.3, "Viewing the Health of an Access Gateway Cluster," on page 141.

Alerts: Indicates whether any alerts have been sent. If the alert count is non-zero, click the count to view more information.

- For information about the alerts of a specific Access Gateway, click the link on the Access Gateway row. For more information, see Section 4.6.1, "Viewing Access Gateway Alerts," on page 133.

- For information about the alerts sent to the cluster, click the link on the cluster row. For more information, see Section 4.6.2, "Viewing Access Gateway Cluster Alerts," on page 133.

Commands: Indicates the status of the last executed command and whether any commands are pending. Click the link to view more information. For more information, see Section 4.9, "Viewing the Command Status of the Access Gateway," on page 142.

Statistics: Provides a link to the statistic pages.

- For information about the statistics of a specific Access Gateway, click the **View** link on the Access Gateway row and see Section 4.3, "Viewing Access Gateway Statistics," on page 118.

- For information about statistics sent to the cluster, click the **View** link on the cluster row and see Section 4.4, "Viewing Cluster Statistics," on page 128.

Edit: Provides a link to the configuration page. If the server belongs to a cluster, the **Edit** link appears on the cluster row. Otherwise, the link is on the server row. See Section 3.1, "Configuration Overview," on page 81.

3.3.2 Configuration Options

Use the information in this section to modify the Status options described in Step 4 on page 85.

1 In the Administration Console, click **Devices** > **Access Gateways**.

Access Manager	Devices	Policies	Auditing	Security

Access Gateways

Access Gateway Servers

Restart | Stop | Refresh | Actions ▼

☐ Name	Status	Health	Alerts	Commands	Statistics	Type	Configuration
AG-Cluster	Current	🟢	50		View		Edit
☐ ↳ 164.99.184.149 ‡	Current	🟢	50	Succeeded	View	Gateway Service	

2 View the **Status** column and make changes as necessary.

Status	Description
Current	Indicates that all configuration changes have been applied.

Status	Description
Update	Indicates that a configuration change has been made, but not applied. To apply the changes, click the Update link, then select one of the following:.

- **All Configuration:** The **All Configuration** option causes the Access Gateway to read its complete configuration file and restarts the Embedded Service Provider.

  The configuration update causes logged-in users to lose their connections unless the server is a member of a cluster. When the server is a member of a cluster, the users are sent to another Access Gateway and they experience no interruption of service.

- **Logging Settings:** When the ESP logging settings have been modified on the Identity Server, the update option for **Logging Settings** is available. The **Logging Settings** option causes no interruption in services. When you modify Access Gateway logging settings, this option is not available because they are considered configuration settings.

- **Policy Settings:** If a policy is modified for a protected resource of the Access Gateway and the policy change is the only modification that has occurred, the update option for **Policy Settings** is available. This option causes no interruption in services.

- **Rewriter Profile Changes:** When the administrator changes the rewriter profile, a purge cache command is issued to a Gateway from the administration console, the connection is lost and the service is interrupted for a few seconds. Similar experience is observed during the rewriter profile configuration change, as this internally triggers the purge cache command.

- **Changing Certificates:** When a certificate configuration is changed from the administration console, the service is interrupted due to the Tomcat restart.

Status	Description

Update All	This link is available when a server belongs to a cluster. You can select to update all the servers at the same time, or you can select to update them one at a time. If the modification is a policy or a logging change, then use **Update All**. If the modification is a configuration change, we recommend that you update the servers one at a time.
	
	  ◆ When you select **Update All** for a configuration change, users experience an interruption of service.
	
	  ◆ When you update servers one at a time for a configuration change, users experience no interruption of service.
	
	When you make the following configuration changes, the **Update All** option is the only option available and your site will be unavailable while the update occurs:
	
	  ◆ The Identity Server configuration that is used for authentication is changed (**Access Gateways** > **Edit** > **Reverse Proxy/Authentication**, then select a different value for the **Identity Server Cluster** option).
	
	  ◆ A different reverse proxy is selected to be used for authentication (**Access Gateways** > **Edit** > **Reverse Proxy/Authentication**, then select a different value for the **Reverse Proxy** option).
	
	  ◆ The protocol or port of the authenticating reverse proxy is modified (**Access Gateways** > **Edit** > **Reverse Proxy/Authentication** > **[Name of Reverse Proxy]**, then change the SSL options or the port options).
	
	  ◆ The published DNS name of the authentication proxy service is modified (**Access Gateways** > **Edit** > **Reverse Proxy/Authentication** > **[Name of Reverse Proxy]** > [**Name of First Proxy Service**], then modify the **Published DNS Name** option).
	
	For more information, see Section 7.4.3, "Applying Changes to Cluster Members," on page 182.

Update ⚠	If the configuration update contains a configuration error, the **Update** link is disabled and the **Configuration Error** icon is displayed. Click the icon to discover which objects have been misconfigured. You need to fix the error by either canceling or modifying the changes before you can perform an update.
Update All ⚠	If the configuration update contains a configuration error, the **Update All** and the member **Update** links are disabled and the **Configuration Error** icon is displayed. Click the icon to discover which objects have been misconfigured. You need to fix the error by either canceling or modifying the changes before you can perform an update.
Pending	Indicates that the server is processing a configuration change, but has not completed the process.
Locked	Indicates that another administrator is making configuration changes. Before you proceed with any configuration changes, you need to coordinate with this administrator and wait until the Access Gateway has been updated with the other administrator's changes.

3.3.3 Scheduling a Command

Use the Schedule New Command page to schedule a command, such as a shutdown, restart, or upgrade.

1 In the Administration Console, click **Devices** > **Access Gateways**.

2 (Conditional) To schedule a shutdown or restart, select a server, then click **Actions** > **Schedule Restart** or **Schedule Stop**. Continue with Step 3.

3 Fill in the following fields:

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

Description: (Optional) Specify a reason for the command.

Date & Time: Select the day, month, year, hour, and minute when the command should execute.

The following fields display information about the command you are scheduling:

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

Server: Displays the name of the server that the command is being scheduled for.

4 Click **OK** to schedule the command.

3.4 Managing General Details of the Access Gateway

The Server Details page allows you to perform general maintenance actions on the selected Access Gateway.

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

2 Select one of the following options:

Edit: Click this option to edit the general details of the Access Gateway. See Section 3.4.1, "Changing the Name of an Access Gateway and Modifying Other Server Details," on page 90.

New IP: Click this action to trigger a scan to detect new IP addresses. This might take some time. If you have used a system utility to add an IP address after you have installed the Access Gateway Service, use this option to update the Access Gateway Service to display the new IP address as a configuration option. For more information about this option, see Section 3.8.5, "Adding a New IP Address to the Access Gateway," on page 103.

Configuration: Click this option to export the configuration of this Access Gateway or to import the configuration of a saved configuration file. See Section 3.4.2, "Exporting and Importing an Access Gateway Configuration," on page 90.

3 Click **Close**.

3.4.1 Changing the Name of an Access Gateway and Modifying Other Server Details

The default name of an Access Gateway is its IP address. You can change this to a more descriptive name as well as modifying other details that can help you identify one Access Gateway from another.

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

Server Details Edit: ag18

Name: ag18

Management IP Address: 10.10.14.18   Port: 1443

Location:

Description:

OK   Cancel

2 Modify the values in the following fields:

Name: Specify the Administration Console display name for the Access Gateway. This is a required field. The default name is the IP address of the Access Gateway. If you modify the name, the name must use alphanumeric characters and can include spaces, hyphens, and underscores.

Management IP Address: Specify the IP address used to manage the Access Gateway. Select an IP address from the list.

Port: Specify the port to use for communication with the Administration Console.

Location: Specify the location of the Access Gateway server. This is optional, but useful if your network has multiple Access Gateway servers.

Description: Describe the purpose of this Access Gateway. This is optional, but useful if your network has multiple Access Gateways.

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

When you click **OK**, any changes are immediately applied to the Access Gateway.

3.4.2 Exporting and Importing an Access Gateway Configuration

You can export an existing Access Gateway configuration and its dependent policies, and then import this configuration to a new server. This feature is especially useful for deployments that set up configurations in a staging environment, test and validate the configuration, then want to deploy the configuration on new hardware that exists in the production environment.

IMPORTANT: The export feature is not a backup tool. The export feature is designed to handle configuration information applicable to all members of a cluster, and network IP addresses and DNS names are filtered out during the import. (The server-specific information that is filtered out is the

information you set specifically for each member in a cluster.) If you want a copy of all configuration information, including server-specific information, you need to perform a backup. See "Backing Up and Restoring" in the *NetIQ Access Manager Appliance 4.0 SP1 Administration Console Guide*

The export feature is not an upgrade tool. You cannot export a configuration from one version of Access Manager and import it into a newer version of Access Manager.

If your Access Gateway is not a member of a cluster and you have configured it to use multiple IP addresses, be aware that the export feature filters out multiple IP addresses and uses only eth0. You need to use the backup utility to save this type of information. If you need to reinstall the machine, leave the Access Gateway configuration in the Administration Console and reinstall the Access Gateway. If you use the same IP address for the Access Gateway, it imports into the Administration Console and inherits the configuration.

When exporting the file, you can select to password-protect the file, which encrypts the file. If you are using the exported file to move an Access Gateway from a staging area to a production area and you need to change the names of the proxy services and DNS names from a staging name to a to a production area and you need to change the names of the proxy services and DNS names from a staging name to a production name, do not select to encrypt the file. You need a simple text file so you can search and replace these names. If you select not to encrypt the file, remember that the file contains sensitive information and protect it accordingly production name, do not select to encrypt the file. You need a simple text file so you can search and replace these names. If you select not to encrypt the file, remember that the file contains sensitive information and protect it accordingly.

Exporting the Configuration

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

2 Click **Configuration** > **Export**.

3 (Conditional) If you want to encrypt the file, fill in the following fields:

Password protect: Select this option to encrypt the file.

Password: Specify a password to use for encrypting the file. When you import the configuration onto another device, you are prompted for this password.

4 Click **OK**, then select to save the configuration to a file.

The filename is the name of the Access Gateway with an `xml` extension.

5 Export the policies used by the Access Gateway. In the Administration Console, click **Policies** > **Policies**, then either select **Name** to include all policies or individually select the policies to export.

You need to export all Access Gateway policies and any Role policies used by the Access Gateway policies.

6 Click **Export** and modify the proposed filename if needed.

7 Click **OK**, then select to save the policy configurations to a file.

8 (Conditional) If you have created multiple policy containers, select the next policy container in the list, and repeat Step 5 through Step 7.

The policies for each container must be saved to a separate export file.

3.5   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 a 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 **Devices** > **Access Gateways** > **Edit** > **Tunneling**.

2 Click **New**, enter a display name for the tunnel, then click **OK**.

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 or the virtual IP address assigned to the Access Gateway cluster by the L4 switch. If you specify a DNS name, the 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:

Cluster Member: (Available only if the Access Gateway is a member of a cluster.) 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 cluster.

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 check box.

TCP Listen Options: Provides additional options for configuring how requests are handled. See Section 2.7.1, "Configuring TCP Listen Options for Clients," on page 77. 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 2.7.2, "Configuring TCP Connect Options for Web Servers," on page 78.

6 Specify a Web server to receive the traffic. In the Web Server List section, click **New**, specify the IP address or DNS name of the Web server, then click **OK**.

At least one Web server must be specified in the list before you can save a tunnel configuration.

7 To save your changes to browser cache, click **OK**.

8 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

3.6 Setting the 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 **Devices** > **Access Gateways** > **Edit** > **Date & Time**.

Group Date and Time: doc2

Cluster Member: ag18 ▼

Server Date and Time

June 12, 2009 11:12 AM   Set Date & Time Manually

Network Time Protocol

Set Up NTP

Time Zone

Name: US/Mountain   ▼

2 (Conditional) If the Access Gateway belongs to a cluster of Access Gateways, select the Access Gateway from the list displayed in the **Cluster Member** field. The modifications you make on this page apply only to the selected Access Gateway.

If the Access Gateway does not belong to a cluster, 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.

IMPORTANT: 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.

Use the `yast` command and select **System** > **Date and Time**.

Set Up NTP: Click this option to specify the DNS name or IP address of a Network Time Protocol server. The installation program enters the name of `pool.ntp.org`, the DNS name of a public NTP server. To disable this feature, you must remove all servers from the NTP Server List. This is not recommended.

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 Click **OK**.

5 On the Server Configuration page, click **OK**.

6 To apply your changes, click **Update** > **OK**.

3.7 Customizing Error Messages and Error Pages on Access Gateway

Access Gateway uses the custom error page template to rebrand and localize the language of error pages that are published to the browser.

By default, Access Gateway contains the following files to help customize and localize the error messages:

- The error page configuration file, `ErrorPagesConfig.xml`
- The error messages file, `ErrorMessages.xml.en`

NOTE: If you are modifying any of the above files, ensure that you retain the original filenames.

Access Gateway maintains `/opt/novell/nam/mag/webapps/agm/WEB-INF/config/current/` directory to save files that are used for error page configuration.

You can customize and localize the error template and the error messages:

- Section 3.7.1, "Customizing and Localizing Error Messages," on page 95
- Section 3.7.2, "Customizing the Error Pages," on page 96

3.7.1 Customizing and Localizing Error Messages

When Access Gateway serves an error message to the browser by using the `Accept-Language` header value received from the browser, it selects a suitable error template and an error message file. To localize the error messages, you must to do the following:

Localize or customize the error messages in the `ErrorPagesConfig.xml` file and save it with the language extension. For more information, see "Localizing and Customizing the Error Messages" on page 95.

Localizing and Customizing the Error Messages

The error messages contained in the `ErrorMessages.xml.en` file can be localized in various languages and stored as `ErrorMessages.xml.<lang>`, where *<lang>* is the `fileXn` attribute value. You can also customize the English error messages present in the `ErrorMessages.xml.en` file.

NOTE: You cannot customize an error message that is not present in the `ErrorMessages.xml.en` file.

To localize the error messages:

1. Log in as `root`.
2. Open the `ErrorMessages.xml.<lang>` file.
3. Copy the error messages that you have localized or customized to within the `<TranslatedMessage></TranslatedMessage>` tags. For example:

```
</Message>
  <Message id="<ID No>" name="<ERROR_MESSAGE_NAME>" enable="yes">
    <EnglishMessage>English Message goes here</EnglishMessage>
<TranslatedMessage>
Localized message goes here
</TranslatedMessage>
</Message>
```

Do not delete the contents within the `<TranslatedMessage></TranslatedMessage>` tags from an English file because, the `ErrorPagesConfig.xml` file selects the error message within these tags for display.

4 Save the file.

5 If the Access Gateway belongs to a cluster, copy the modified file to each member of the cluster, then restart that member.

6 Edit the configuration and make dummy changes and push the configuration.

3.7.2 Customizing the Error Pages

Access Gateway uses the Apache method for localizing error messages. You can modify these messages or customize the page they are displayed on.

1 To change a message:

   1a Change to the Apache message configuration directory:

     `/etc/opt/novell/apache2/conf/extra`

   1b Open the `http-multilang-errordoc.conf` file.

     The first few lines of this file contains comments on how Apache recommends modifying the error messages. You can select to use their method or continue with the following steps.

   1c Locate the `ErrorDocument` section and determine the error code message you want to modify. Make note of the `*.var` filename.

   1d Change to the Apache error directory:

     `/opt/novell/apache2/share/apache2/error`

   1e Open the `*.var` file that you want to modify.

     The message is listed alphabetically by language code.

   1f Save the changes.

2 To change the header of the error page:

   2a Change to the Apache error include directory:

     `/opt/novell/apache2/share/apache2/error/include`

   2b Open the `top.html` page.

   2c To change the title of the page, locate the following line:

     `<title>Access Manager 4.0<\title>`

   2d Replace the `Access Manager 4.0` string with the content you require.

   2e To replace the image in the header, locate the following line:

     `<img src="NAGErrors/images/Odyssey_LoginHead.gif" alt="" height="80" width="550" border="0">`

   2f Replace `Odyssey_LoginHead.gif` with the filename of the image you want to display.

   2g Adjust the height and width values to match your image.

2h Save the file.

2i Copy your image to the `images` directory:

`/opt/novell/apache2/share/apache2/error/images`

3 To change the footer of the error page:

3a Change to the Apache error include directory:

`/opt/novell/apache2/share/apache2/error/include`

3b Open the `bottom.html` page.

3c To change the image, find the following line:

```
<td style="background-color: #E6D88C; padding-left: 10px"><img
style="padding-right: 200px" src="/NAGErrors/images/
LAP_interoperable_logo_100.gif" align="absmiddle" border="0">
```

3d Change `LAP_interoperable_logo_100.gif` to the filename of the image you want to display.

3e Save the file.

3f Copy your image to the `images` directory:

`/opt/novell/apache2/share/apache2/error/images`

4 Copy all modified files and image files to all Access Gateways in the cluster.

The `err.jsp` file will also log the ESP error messages. For more information on customizing the `err.jsp` page, see Customizing Identity Server Messages. The procedure for customizing is the same but the paths referred to will change for the Access Gateway. Following are the path changes:

♦ In Customizing Messages (https://www.netiq.com/documentation/netiqaccessmanager32/identityserverhelp/data/bhz5kn8.html), the paths for Access Gateway are as follows:

Step 3, path on Linux will be `/opt/novell/nam/mag/webapps/nesp/WEB-INF/lib` and on Windows `/Program Files/Novell/Tomcat/webapps/nesp/WEB-INF/lib/`.

Step 10, path on Linux will be `/opt/novell/nam/mag/webapps/nesp/WEB-INF/classes` and on Windows `/Program Files/Novell/Tomcat/webapps/nesp/WEB-INF/classes`.

Step 12, restart Access Gateway `/etc/init.d/novell-mag restart`.

♦ In Customizing the Branding of the Error Page (https://www.netiq.com/documentation/netiqaccessmanager32/identityserverhelp/data/bhz5kn8.html#bixm5ek), the path for err.jsp in the ESP on Linux will be `/opt/novell/nam/mag/webapps/nesp/jsp` and on Windows `/Program Files/Novell/Tomcat/webapps/nesp/jsp/`.

3.8 Configuring Network Settings

After initial setup, you seldom need to change the network settings unless something in your network changes, such as adding a new gateway or DNS server. These options are for the Access Gateway Appliance. For the Linux or Windows Access Gateway Service, use the utilities supplied by the operating system. However, if you add an new network interface card to the Access Gateway Service

machine and use system utilities to configure it and assign it an IP address, you need to update the Access Gateway Service with this information. See Section 3.8.5, "Adding a New IP Address to the Access Gateway," on page 103.

This section describes the following tasks:

- Section 3.8.1, "Viewing and Modifying Adapter Settings," on page 98
- Section 3.8.2, "Viewing and Modifying Gateway Settings," on page 99
- Section 3.8.3, "Viewing and Modifying DNS Settings," on page 101
- Section 3.8.4, "Configuring Hosts," on page 102
- Section 3.8.5, "Adding a New IP Address to the Access Gateway," on page 103

3.8.1 Viewing and Modifying Adapter Settings

The adapter settings allow you to view the current configuration for the network adapters installed in the Access Gateway Appliance 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 these settings to add them.
- If you have multiple adapters installed on an Access Gateway Appliance machine, you can only configure eth0 during installation. Use the procedure described in this section to configure the others.

To view or modify your current adapter settings:

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Adapter List**.

2 (Conditional) If the Access Gateway is a member of a cluster, select the server you want to configure from the list of servers in the **Cluster 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 subnet must be configured for you to delete one.
- 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:

Subnet: Displays the address of the subnet that you are modifying. This is empty if you are creating a new subnet.

Subnet Mask: (Required) Specifies the subnet mask address for this subnet. The address can be specified in standard dotted format or in CIDR format.

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, select the address, then click **Change IP Address**, specify the new IP address, then click **OK**.

5 Click **OK**.

6 Click **OK**.

7 On the Server Configuration page, click **OK**, then click **Update** > **OK**.

3.8.2 Viewing and Modifying Gateway Settings

The gateway settings display the current gateway configuration that the Access Gateway Appliance is using to route packets. On this page, you can also 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.

Gateways fall within the following three basic groups:

- ◆ Host gateways for specific destination addresses.
- ◆ Network gateways for destination addresses that fall within specific subnets.
- ◆ The default gateway for destination addresses that aren't covered by host or network gateways.

The Access Gateway uses additional gateways only when the **Act As Router** option is selected. When this option is selected, you can add Host Gateways and Network Gateways. When configuring a Host Gateway or Network Gateway, you specify the IP address of the host or network gateway in the **Next Hop** field. This address must be on the same subnetwork as the IP address for the Access Gateway.

IMPORTANT: If you enter an IP address that is on a different subnetwork, the Access Gateway reports this error on the Health page, after the configuration has been applied.

To modify your current gateway configuration:

1 In the Administration Console, click **Devices** > **Access Gateways > Edit** > **Gateways**.

2 Configure your default gateway, which specifies the gateway to use when no other routes apply. Configure the following:

Next Hop: 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.

3 Configure your host gateways, which are 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: The address of the host gateway that is to be used.

Host: 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.

Click **OK** when the fields are configured.

4 Configure your network gateways, which are 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: The address of the gateway that is to be used.

Network Address: The subnet address for the destination IP address range. You should enter the valid subnet address.

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 a Class A mask is 255.0.0.0, a Class B mask is 255.255.0.0, and Class C, D, and 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.

Click **OK** when the fields are configured.

5 Click **OK**.

6 On the Server Configuration page, click **OK**, then click **Update** > **OK**.

3.8.3 Viewing and Modifying DNS Settings

The DNS page displays the current configuration for domain name services for the Access Gateway Appliance and allows you to modify it.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **DNS**.

2 (Conditional) If the Access Gateway is a member of a cluster, select the server you want to configure from the list of servers in the **Cluster 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.
- **Order:** To modify the order in which the DNS servers are listed, select the server, then click either the up-arrow or the down-arrow buttons. The first server in the list is the first server contacted when a DNS name needs to be resolved.

4 Click **OK**.

5 On the Server Configuration page, click **OK**, then click **Update** > **OK**.

3.8.4 Configuring Hosts

You can configure the Access Gateway Appliance to have multiple hostnames or to resolve DNS names to IP addresses. If you manually edit the `/etc/hosts` file, your modifications are lost when the Access Gateway Appliance is updated. However, if you use the Hosts page to specify the entries, the entries are written to the `/etc/hosts` file whenever the configuration of the Access Gateway Appliance is updated.

1 (Access Gateway Appliance) In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Hosts**.

2 (Conditional) If the Access Gateway is a member of a cluster, select the server you want to configure from the list of servers in the **Cluster Member** field. All changes made to this page apply to the selected server.

3 To add a new hostname to an existing IP address, click the name of a **Host IP Address**.

4 In the **Host Name(s)** text box, specify a name for the host. Place each hostname on a separate line, then click **OK**.

5 To add a new IP address and hostname, click **New** in the **Host IP Address List** section, then specify the IP address. In the **Host Name(s)** text box, specify a hostname, then click **OK**.

6 To delete a host, select the check box next to the host you want to delete, then click **Delete**.

7 Click **OK**.

8 On the Server Configuration page, click **OK**, then update the Access Gateway.

3.8.5 Adding a New IP Address to the Access Gateway

Before you can configure Access Gateway to use a new IP address, you must first use an operating system utility to add the IP address.

Start YaST, click **Network Devices** > **Network Card**, then select the **Traditional Method**.

After you have used a system utility to add an IP address, you need to update the Access Gateway Service to display the new IP address as a configuration option.

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

2 On the Server Details page, click **New IP**, then click **OK**.

   Access Gateway scans the operating system for its configured IP addresses and adds any new addresses. Any new address is then available for assignment on the Access Gateway configuration pages.

3 (Optional) To verify that the scan has completed, click the **Command Status** tab.

3.9 Customizing Logout Requests

- Section 3.9.1, "Customizing Applications to Use the Access Gateway Logout Page," on page 103
- Section 3.9.2, "Customizing the Access Gateway Logout Page," on page 104
- Section 3.9.3, "Configuring the Logout Disconnect Interval," on page 105

3.9.1 Customizing Applications to Use the Access Gateway Logout Page

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 reauthenticates the user to the resource, and it appears that the logout did nothing.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Reverse Proxy / Authentication**.

2 In the **Embedded Service Provider** section, view the path to the AGLogout page 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.

3 Redirect application logout requests to the AGLogout page.

4 Click **OK**.

The Access Gateway does not support the following logout pages that were used in previous version of Access Manager and iChain:

- /cmd/BM-Logout
- /cmd/ICSLogout

3.9.2 Customizing the Access Gateway Logout Page

You can create your own logout page and configure the Access Gateway to use it. To do this, you need to modify the `logoutSuccess.jsp` file on the Access Gateway. It is located in the following directory:

`/opt/novell/nesp/lib/webapp/jsp`

You can modify the file to display what you want or you can modify it to redirect the user to your custom page. The following sections provide some tips for accomplishing this task:

- "Modifying the Header" on page 104
- "Redirecting to Your Custom Page" on page 104
- "Calling Different Logout Pages" on page 105

Modifying the Header

The `logoutSuccess.jsp` file is called in a frame from the `nidp.jsp` file. The branding in the header of the logout page is controlled by the branding of the `nidp.jsp` file. For information about how to modify `nidp.jsp` for logos, titles, and colors, see "Rebranding the Header" in the *NetIQ Access Manager Appliance 4.0 Identity Server Guide*.

IMPORTANT: Take a backup of `nidp.jsp` file before modifications. Every time you upgrade your Access Gateway, upgrade process overrides any custom changes made to JSP files that use the same filename as those included with the product. If you want the modified file, you need to restore the `nidp.jsp` file. During an upgrade, you can select to restore custom login pages, but NetIQ still recommends that you have your own backup of any customized files.

Redirecting to Your Custom Page

One way to provide redirection is to replace the information in the `<body>` element of the `logoutSuccess.jsp` file with something similar to the following:

```
<body>
      <script language="JavaScript">
        top.location.href='http://<hostname/path>';
      </script>
</body>
```

Replace the *<hostname/path>* string with the location of your customized logout page.

IMPORTANT: Take a backup of `logoutSuccess.jsp` file before modifications. Every time you upgrade your Access Gateway, upgrade process overrides any custom changes made to JSP files that use the same filename as those included with the product. If you want the modified file, you need to restore the `nidp.jsp` file. During an upgrade, you can select to restore custom login pages, but NetIQ still recommends that you have your own backup of any customized files

---

Calling Different Logout Pages

If you need to use a different logout page for specific protected resources, you need to modify the logout button of the applications to use the plogout URL rather that the AGLogout URL (see). The AGLogout page redirects to the plogout page, which calls the `logoutSuccess.jsp`. Any parameter added to the AGLogout or plogout URL is saved and passed to the `logoutSuccess.jsp` file.

The parameter passed to the `logoutSuccess.jsp` file can be used with if/else logic in the body of the page to load different custom logout pages based on the parameter value.

To use the plogout URL, you need to modify the application's logout button to call the following URL:

`<ESP Domain>/nesp/app/plogout`

Replace *<ESP Domain>* with the same value as the AGLogout value. For example, suppose your AGLogout value is the following:

`https://jwilson1.provo.novell.com:443/AGLogout`

You would replace it with the following value:

`https://jwilson1.provo.novell.com:443/nesp/app/plogout`

If you add a parameter to the URL, it would look similar to the following:

`https://jwilson1.provo.novell.com:443/nesp/app/plogout?app=email`

Logging Out of Sessions to the Access Gateway and SAML Connectors when Branding Exists in the Customized Logout Page

When you have both Liberty and SAML 2.0 sessions running on the Identity Server and you log out of the Access Gateway, the `logoutSuccess.jsp` page is not executed with the customizations you have made to the logout page. You will be able to log out of the Access Gateway but the customizations you made are lost.

If the `logutSuccess.jsp` file is not loaded in a frame, the banner will not be displayed, and the Access Gateway will comment out the content in the `logoutSuccess.jsp` file. Add the below line after the <body> tag in the `logoutSuccess.jsp` file.

`<!-- BANNER LOADS IF THIS PAGE IS NOT LOADED IN REGULAR FRAME -->`

`<%@include file="logoutHeader.jsp"%>`

3.9.3 Configuring the Logout Disconnect Interval

When a user clicks the logout button and the user is logging out of an Access Gateway that is a member of a cluster, the user is not immediately disconnected from the resource. The logout message must be sent to each member of the cluster. The default interval for checking the pending logout message queue is 30 seconds. If this interval is too long, you can configure a shorter interval in the `web.xml` file of the Embedded Service Provider. This must be set on each Access Gateway in the cluster.

1 Log in to the Access Gateway as the root or administrator user.

2 Open the `web.xml`.

   `/opt/novell/nesp/lib/webapps/WEB-INF/web.xml`

3 Find the `<context-param>` section in the file.

4 Add the following parameter to the `<context-param>` section.

```
<context-param>
    <param-name>logoutRetirementFrequency</param-name>
    <param-value>15000</param-value>
</context-param>
```

5 Set the `<param-value>` element to a value between 5000 and 30000 milliseconds (5 seconds and 30 seconds).

6 Restart the Embedded Service Provider.

For information about how to restart the Embedded Service Provider from the Administration Console, see Section 3.3, "Managing Access Gateways," on page 84.

3.10 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.

Access Gateway Service: Apache is configured to always send the X-Forwarded-For, X-Forwarded-Host, and X-Forwarded-Server headers. There are no options in the Administration Console to change this behavior.

To enable the X-Forwarded-For header on the Access Gateway:

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **HTTP Options** > **Header Options**.

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 Click **OK**.

4 To apply your changes, click the **Access Gateways** link, then click **Update** > **OK**.

3.11 Enabling the Access Gateway to Display Post-Authentication Message

When the Identity Server authentication process is completed, the user-agents are redirected to their originally requested URL. The originally requested URL is then retrieved by the proxy. This process requires SSO and authentication process of its own. As a result, retrieving the requested URL may take a long time. It is not clear how much time the authentication process takes and how much time the origin server request and authentication processes take.

To remove this ambiguity, you can enable the Access Gateway to display a message before redirecting the user-agent to the originally requested URL.

To enable this enhancement, complete the following steps:

1 Open `/opt/novell/nam/mag/webapps/nesp/WEB-INF/classes/nidpconfig.properties`.

2 Set the `IS_DISPLAY_AUTH_DONE_PAGE` parameter to true.

When this option is enabled, the following message is displayed before the final redirect to the requested URL:

`Authentication successful, please wait while your requested page loads.`

The Web page that display this message is a JSP page. Location of this page is `/opt/novell/nam/mag/webapps/nesp/jsp/waitredir.jsp`. You can perform further customization on this page.

4 Access Gateway Maintenance

4.1 Managing Access Gateway Logs

In the Access Gateway, logging can be configured by using Advanced Options.

4.1.1 Configuring the Log Level

1 In the Administration Console, select **Devices > Access Gateways > Edit > Advanced Options.**

2 Add the following line with appropriate log level:

```
LogLevel <loglevel>
```

Replace *loglevel* option with `emerg`, `alert`, `crit`, `error`, `warn`, `notice`, `info` or `debug`. The default log level is warn.

Option	Description
emerg	Sends only messages that render the system unusable, if they are not resolved.
alert	Sends only messages that require immediate action.
crit	Sends only messages about critical situations
error	Sends warning messages about recoverable errors.
warn	Sends warning messages.
Notice	Sends information about the status of a service to the service configuration logs.

Option	Description
Info	Sends informational messages such as requests sent to Web servers and the results of authentication requests.
Debug	Sends debug messages

3 Click **OK**.

4 Click the Access Gateways link, then click Update > OK.

The `error_log` file is available at `/var/opt/novell/nam/logs/mag/apache2/`.

4.1.2 Configuring the Log File

1 In the Admin console, click **Devices > Access Gateways > Edit > Advanced Options**

2 Add the following line:

`ErrorLog <path to the file where logs should be recorded>`

3 Click **OK**.

4 Click the Access Gateways link, then click **Update > OK**.

4.2 Configuring Logging for a Proxy Service

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. For more information about custom log formats, see Apache Log Configuration Module.

4.2.1 Determining Logging Requirements

Because logging requirements and transaction volume vary widely, NetIQ 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 transactions 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 4.2.5, "Configuring Extended Log Options," on page 115.

4 Design a rollover 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 rollover. To calculate the best procedure, see Section 4.2.2, "Calculating Rollover Requirements," on page 111.

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 rollover files by number or age. To calculate the best procedure, see Section 4.2.2, "Calculating Rollover Requirements," on page 111.

4.2.2 Calculating Rollover 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 4.2.6, "Configuring the Size of the Log Partition," on page 118.

◆ **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 4.3, "Viewing Access Gateway Statistics," on page 118.

◆ **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 more.

The following formulas can help you estimate when the system would run out of resources:

- "Calculating diskfull_time" on page 112
- "Calculating max_roll_time" on page 112
- "Calculating max_log_roll_size" on page 113

Calculating diskfull_time

Use the following formula to calculate how long it takes the Access Gateway to fill your logging disk space:

```
diskfull_time in seconds = logpartition_size / (request_rate *
   logentry_size * 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 fills 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 Use the diskfull_time formula to 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 rollover 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

Use the following formula to calculate the maximum rollover time value you should specify in the **Roll over every** field

```
max_roll_time = diskfull_time / 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 Use the max_roll_time formula and the diskfull_time value obtained in "Calculating diskfull_time" on page 112 to 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

Use the following formula to 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:

1 Determine the values of the three variables listed above.

2 Use the max_log_roll_size formula to calculate the maximum size a log file should reach before the cache device rolls it over.

4.2.3 Enabling Logging

Do not enable logging until you have designed a logging strategy. See Section 4.2.1, "Determining Logging Requirements," on page 110.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **Logging.**

2 Fill in the following fields:

Enable Logging: Select this field to enable logging.

Log Directory: Default location for log files of proxy service is `/var/log/novell/reverse/<reverse_proxy_name>`.

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.

A logging profile determines the type of information that is written to the log file; it also manages rollover and old file options.

4 Continue with one of the following:

- Section 4.2.4, "Configuring Common Log Options," on page 114
- Section 4.2.5, "Configuring Extended Log Options," on page 115

4.2.4 Configuring Common Log Options

Use the common log options page to control log rollover 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
- IP address 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 rollover 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 **Devices > Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **Logging** > **[Name of Common Log Profile]**.

2 Select one of the following rollover options:

Rollover When File Size Reaches: 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 hours or days.

- **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:

Limit Number of Files to: 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.

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.

Do Not Delete: Prevents the system from automatically deleting the log files.

4 Click **OK**.

5 Click the **Access Gateways** link, then click **Update** > **OK**.

4.2.5 Configuring Extended Log Options

Use the extended log options page to control log entry content, log rollover, 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 rollover 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 **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **Logging** > **[Name of Extended Log Profile]**.

2 Select one or more of the log data options:

Name	Description	Entry in Configuration File	Sample Entry in the Log file
User Name	The name of the user sending the request.	%u	"public","cn=admin,o=novell"
Server IP	The IP address of the Access Gateway.	%a	123.1.2.3
Site Name	The name of the reverse proxy.	%v	www.lagssl.com
Method	The HTTP method the browser sent to the Access Gateway.	%m	GET,POST
URI	The HTTP URL the browser sent to the Access Gateway.	%U	nam/acme_ss_js7.html
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 **URI Stem** field is the same as the **URI** field. **URI Stem** is redundant if **URI** is selected.		/path/to/resource

Name	Description	Entry in Configuration File	Sample Entry in the Log file
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. **URI Query** is redundant if URI is selected.	%q	?page=catalog&x=100 &y=0
Version	The HTTP version specified in the URL the browser sent to the Access Gateway.		HTTP/1.1
Status	The HTTP status code the Access Gateway sent to the browser.	%s	200, 304, 404
Bytes Sent	The number of bytes of HTTP response data the Access Gateway sent to the browser.	%I	14378
Bytes Received	The number of bytes of HTTP request data the proxy service received from the browser.	%O	14378
Time Taken	The time it took the Access Gateway resources to deal with the request in microseconds.	%D	0.062, 0.392, 2, 802.1
User Agent	The User-Agent HTTP request header value the browser sent to the Access Gateway.	%{user-agent}i	Mozilla/5.0 (X11; Linux x86_64; rv:19.0) Gecko/20100101 Firefox/19.0
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.	%{cookie}	IPCZQX0355730a2b=0 1001300a463874a93ef 23e89e9acc94468beb4 b; ZNPCQ003- 37323400=c2e51552
Referer	The Referer HTTP request header value the browser sent to the Access Gateway.	%{Referer}	https://www.lagssl.com/ netiq/nam/ acme_ss_js7.html
Cached Status	The value indicates whether the request was filled from cache. 1 = filled from cache 0 = not filled from cache		0,1
Origin Server	The IP address of the Web server. This assumes the Access Gateway retrieved the requested information directly from the Web server.	%{BALANCER_WORK ER_IP}e	125.1.2.5

Name	Description	Entry in Configuration File	Sample Entry in the Log file
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, which causes the Access Gateway to generate or forward headers to upstream proxies or Web servers.	%{x-forward-for}i	10.0.0.1,10.0.02,10.0.03,10.0.0.4
Bytes Filled	The total bytes filled in response to the request.	%l	184
Content Range	The byte ranges sent from the Access Gateway to a requesting browser.	%{Content-Range}o	
E Tag	The tag sent from the Access Gateway to a requesting browser.	%{ETag}	604888-1077-466372c0
Completio n Status	The completion status for the transaction, indicating that it completed successfully or that it failed. Possible values: success, timeout, reset (the client terminated the connection), administrative (the Access Gateway terminated the connection).	%X	success, timeout, reset
Reply Header Size	The size in bytes of the HTTP header associated with a response to a client.	%L	361
X Cache Info	Brief status statement for cached objects; brief reasons why an object was not cached.	%{Cache-Control}o	no-store
Range	The Range header value.	%{Range}o	
If Range	The If Range header value, which indicates whether the browser request was a conditional range request.	%{If-Range}	bytes 0-200/736
Content Length	The size in bytes of the entire object delivered to a requesting browser.	%O	741
Request Pragma	The pragma value associated with a browser request.	%{Pragma}o	No-cache , no-store
Reply Pragma	The pragma value associated with a server response to a requesting browser.	%{Pragma}i	no-cache

3 Select one of the following rollover options:

Rollover When File Size Reaches: 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 hours or days.

- ◆ **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:

Limit Number of Files to: 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.

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.

Do Not Delete: Prevents the system from automatically deleting the log files.

5 Click **OK**.

6 Click the **Access Gateways** link, then click **Update** > **OK**.

4.2.6 Configuring the Size of the Log Partition

The size of the log partition should be configured as part of the installation process. The Access Gateway logs are stored in the `/root` partition by default. You can create a `/var` partition to store the logs. The size of this partition depends on your requirements.

4.3 Viewing Access Gateway Statistics

The Statistics page allows you to monitor the amount of data and the type of data the Access Gateway is processing.

1 In the Administration Console, click **Devices** > **Access Gateways** > **[Name of Server]** > **Statistics**.

General \ Health \ Alerts \ Command Status \ **Statistics**	
Server Activity | Server Benefits | Service Provider Activity	
[Statistics | Live Statistics Monitoring]	
Server Activity	**Last Reported Time: July 3, 2007 8:12 AM**
CPU Utilization	60.0 % ☑ Graphs
Cache Hit	93.0 % ☑ Graphs
Mounted Partitions Disk Space	73.82 GB
Mounted Partitions Disk Space Used	32.62 GB
Mounted Partitions Disk Space Free	41.20 GB
Swap Partition Disk Space	4.006 GB
Swap Partition Disk Space Used	2.921 MB
Swap Partition Disk Space Free	4.003 GB
Cache Disk Space	73433088 KB
Cache Disk Space Utilization	0.0 %
Total Installed Memory	1993 MB
Start Up Time	Tuesday, July 3, 2007 8:06:55 AM GMT
Up Time	0 Days, 6 Hours, 7 Minutes, 8 Seconds
Number of Objects Cached	179

2 Select from the following types:

3 Click **Close**.

NOTE: The statistics graphs of the Identity Server and Access Gateway are available in only the primary Administration Console. The periodic stats are sent to the secondary Administration Console only when the primary console is not available. Hence, the statistics graphs of the Identity Server and Access Gateway do not display any statistics values in the secondary Administration Console.

4.3.1 Server Activity Statistics

Select whether to monitor live or static statistics:

Statistics: Select this option to view the statistics as currently gathered. The page is static and the statistics are not updated until you click **Live Statistics Monitoring**.

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

The Server Activity section displays general server utilization statistics.

Statistic	Description
CPU Utilization	Displays the current CPU utilization rate. Use the available graph for capacity planning.
	Click **Graphs** to view the CPU usage for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the percentage of use.
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.
	Click **Graphs** to view the number of cache hits for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the number of hits.
Mounted Partitions Disk Space	Displays the total disk space configured on mounted partitions.

Statistic	Description
Mounted Partitions Disk Space Used	Displays the disk space in use on mounted partitions.
Mounted Partitions Disk Space Free	Displays the disk space available on mounted partitions.
Swap Partition Disk Space	Displays the total disk space configured for the swap partition. The Linux Gateway Service displays the available swap space reported by the Linux kernel (see sysinfo for details).
Swap Partition Disk Space Used	Displays the disk space in use on the swap partition.
Swap Partition Disk Space Free	Displays the disk space available on the swap partition.
Cache Disk Space	Displays the total disk space available for caching.
Cache Disk Space Utilization	Reserved. Not currently used.
Total Installed Memory	Displays the amount of memory that is installed on the Access Gateway.
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 objects that have been cached since the Access Gateway was last started.

Connections

The connection statistics show the current and peak levels of usage in terms of TCP connections.

Statistic	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. Click **Graphs** to view the number of connections for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the number of connections.
Peak Connections from Origin Server	Displays the peak number of connections that the Access Gateway has established with Web servers.

Statistic	Description
Connections to Browsers	Displays the total number of connections that the Access Gateway has established with browsers since it was last started. Click **Graphs** to view the number of connections for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the number of connections.
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 Web 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

Statistic	Description
Throughput of the Origin Server	Displays the average number of bytes of data being sent each second from the Web servers to the Access Gateway. Average number of bytes = total number of bytes sent from origin server to the Access Gateway per system uptime in seconds. Click **Graphs** to view the number of bytes for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the number of bytes.
Throughput of the Browser	Displays the average number of bytes of data being sent each second from the Access Gateway to the browsers. Average number of bytes = total number of bytes sent from the Access Gateway to browsers per system uptime in seconds. Click **Graphs** to view the number of bytes for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the number of bytes.
Total Bytes per Second	Displays the total number of bytes of data being sent each second from the Access Gateway and from the Web servers. Click **Graphs** to view the number of bytes for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the number of bytes.
Bytes Sent to Origin Server	Displays the total number of bytes sent to the origin server after the server is started.
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.

Statistic	Description
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.
Bytes Received from Browser	The total number of bytes received from the browser after the server is started.
Total Bytes	Displays the total number of bytes of data sent from the Access Gateway and from the Web servers since the Access Gateway was 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.

Statistic	Description
Current Requests to Origin Server	Displays the current number of requests that the Access Gateway has made to the Web servers. Click **Graphs** to view the number of requests for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the number of requests.
Current Requests from Browsers	Displays the current number of requests that the browsers have made to the Access Gateway. Click **Graphs** to view the number of requests for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the number of requests.
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.
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 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 to Origin Server	Displays the number of requests that are being sent each second from the Access Gateway to the Web servers. Click **Graphs** to view the number of requests for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the number of requests.
Requests per Second from Browsers	Displays the number of requests that are being sent each second from the browsers to the Access Gateway. Click **Graphs** to view the number of requests for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the number of requests.

Statistic	Description
Total Requests per Second	Displays the total number of requests that are being sent each second from the Access Gateway and from the browsers.
Peak Requests per Second to Origin Server	Displays the peak number of requests that have been sent in one second from the Access Gateway to the Web servers.
Peak Requests per Second from 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.

Statistic	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 are considered fresh.
	Click **Graphs** to view the percentage of fresh objects for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the percentage of fresh objects.
Oldest Object in Memory	Displays how long the oldest cache object has been cached.

4.3.2 Server Benefits Statistics

Select whether to monitor live or static statistics:

Statistics: Select this option to view the statistics as currently gathered. The page is static and the statistics are not updated until you click **Live Statistics Monitoring**.

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 caching:

Statistic	Description
Total 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.
Bytes Saved per Second	Displays how many bytes of data the Access Gateway was able to send from cache rather than requesting it from the Web servers.

4.3.3 Service Provider Activity Statistics

Select whether to monitor live or static statistics:

Statistics: Select this option to view the statistics as currently gathered. The page is static and the statistics are not updated until you click **Live Statistics Monitoring**.

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 ESP Activity page displays information about the communication process between the Access Gateway (ESP) and the Identity Server. These statistics are grouped into the following categories:

- Application
- Authentications
- Incoming HTTP Requests
- Outgoing HTTP Requests
- Liberty
- Clustering
- SP Brokering

Click **Graphs** to review historical statistics.

Application

Statistic	Description
Free Memory	The percentage of free memory available to the JVM (Java Virtual Machine).
	Click **Graphs** to view the free memory for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the percentage of free memory.

Authentications

Statistic	Description
Provided Authentications	The number, since the Identity Server was started, of successful provided authentications given out to external entities.
Consumed Authentications	The number, since the Identity Server was started, of successful consumed authentications.
Provided Authentication Failures	The number, since the Identity Server was started, of failed provided authentications given out to external entities.
Consumed Authentication Failures	The number, since the Identity Server was started, of failed consumed authentications.
	NOTE: The consumed authentication failures does not show the number of invalid password attempt failures of the Identity Provider in the statistics page.
Historical Maximum Logins Served	The maximum number of logins served during an interval and displayed after completion of the interval.
Logins in Last Interval	The number of active user sessions during the last interval.
Logouts	The number of explicit logouts performed by users. This does not include logouts where an inactive session was destroyed.

Statistic	Description
Cached Sessions	The number of currently active cached user sessions. This represents the number of users currently logged into the system with the following caveat: If a single person has two browser windows open on the same client and if that person performed two distinct authentications, then that person has two user sessions.

Click **Graphs** to view the number of cached sessions for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the number of cached sessions. If no sessions have been cached, the value axis is not meaningful. |
Cached Ancestral Sessions	The number of cached ancestral session IDs. An ancestral session ID is created during the failover process. When failover occurs, a new session is created to represent the previous session. The ID of the previous session is termed an "ancestral session ID," and it is persisted for subsequent failover operations.
Cached Subjects	The number of current cached subject objects. Conceptually, the cached subjects are identical to the cached principals.
Cached Principals	The number of current cached principal objects. A principal can be thought of as a single directory user object. Multiple users can log in using a single directory user object, in which case multiple cached sessions would exist sharing a single cached principal.
Cached Artifacts	The number of current cached artifact objects. During authentication, an artifact is generated that maps to an assertion. This cache holds the artifact to assertion mapping until the artifact resolution request is received. Under normal operations, artifacts are resolved within milliseconds of being placed in this cache.

Incoming HTTP Requests

Incoming HTTP requests are divided into three categories: active, interval, and historical. As soon as a request is complete, it is placed into the interval category. The interval represents the last 60 seconds of processed requests. At the completion of the 60-second interval, all requests in the interval category are merged into the historical category.

Statistic	Description
Total Requests	The total number of incoming HTTP requests that have been processed since the Identity Server was started.

Click **Graphs** to view the number of requests for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the number of requests for the selected time period. |
Currently Active Requests	The number of currently active incoming HTTP requests.
Oldest Active Request (Milliseconds)	The age of the oldest currently active incoming HTTP request.
Last Interval Maximum Request Duration (Milliseconds)	The age of the longest incoming HTTP request that was processed during the last 60-second interval.
Last Interval Mean Request Duration (Milliseconds)	The mean age of all incoming HTTP requests that were processed during the last 60-second interval.

Statistic	Description
Historical Maximum Request Duration (Milliseconds)	The age of the longest incoming HTTP request that was processed since the Identity Server was started.
Historical Mean Request Duration (Milliseconds)	The mean age of all incoming HTTP requests that were processed since the Identity Server was started.

Outgoing HTTP Requests

Outgoing HTTP requests are divided into three categories: active, interval, and historical. As soon as a request is complete, it is placed into the interval category. The interval represents the last 60 seconds of processed requests. At the completion of the 60-second interval, all requests in the interval category are merged into the historical category.

Statistic	Description
Total Requests	The total number of outgoing HTTP requests that have been processed since the Identity Server was started.
	Click **Graphs** to view the number of requests for a specific unit of time (1 hour, 1 day, 1 week, 1 month, 6 months, or 12 months). The Value axis displays the number of requests for the selected time period.
Currently Active Requests	The number of currently active outgoing HTTP requests.
Oldest Active Request (Milliseconds)	The age of the oldest currently active outgoing HTTP request.
Last Interval Maximum Request Duration (Milliseconds)	The age of the longest outgoing HTTP request that was processed during the last 60-second interval.
Last Interval Mean Request Duration (Milliseconds)	The mean age of all outgoing HTTP requests that were processed during the last 60-second interval.
Historical Maximum Request Duration (Milliseconds)	The age of the longest outgoing HTTP request that was processed, since the Identity Server was started.
Historical Mean Request Duration (Milliseconds)	The mean age of all outgoing HTTP requests that were processed, since the Identity Server was started.

Liberty

Statistic	Description
Liberty Federation	The number of Liberty protocol federations performed, since the Identity Server was started.
Liberty De-Federations	The number of Liberty protocol de-federations performed, since the Identity Server was started.
Liberty Register-Names	The number of Liberty protocol register names performed, since the Identity Server was started.

Clustering

An authoritative server is the cluster member that holds the authentication information for a given user session. For a request associated with a given session to be processed, it must be routed ("proxied") to the authoritative cluster member. If an L4 switch causes a request to go to a non-authoritative cluster member, then that cluster member proxies that request to the authoritative cluster member.

When a request is received, a cluster member uses multiple means to determine which cluster member is the authoritative server for the request. It looks for a parameter on the query string of the URL indicating the authoritative server. It looks for an HTTP cookie indicating the authoritative server. If these do not exist, the cluster member examines the payload of the HTTP request to determine the authoritative server. Payload examinations result in immediate identification of the authoritative server or a user session ID or user identity ID that can be used to locate the authoritative server.

If a user session ID or user identity ID is found, the ID is broadcast to all cluster members asking which member is the authoritative server for the given ID. The authoritative server receives the broadcast message, determines that it indeed holds the given session or user, and responds accordingly.

The higher the number of proxied requests, the lower the performance of the entire system. Furthermore, the higher the number of payload examinations and ID broadcasts, the lower the performance of the entire system.

Statistic	Description
Currently Active Proxied Requests	The number of currently active proxied HTTP requests.
Total Proxied Requests	The total number of proxied requests that have been processed, since the Identity Server was started. These requests were sent to a non-authoritative (wrong) box.
Total Non-Proxied Requests	The total number of non-proxied requests that have been processed, since the Identity Server was started. These requests were sent to the authoritative (correct) box.
Authoritative Server Obtained from URL Parameter	The total number of authoritative servers identified by using the parameter from the URL query string, since the Identity Server was started.
Authoritative Server Obtained from Cookie	The total number of authoritative servers identified by using the HTTP cookie, since the Identity Server was started.
Payload Examinations	The total number of attempted payload examinations to identify the authoritative server, since the Identity Server was started.
Successful Payload Examinations	The total number of successful payload examinations to identify the authoritative server, since the Identity Server was started.
Identity ID Broadcasts	The total number of attempted Identity ID Broadcasts to identify the authoritative server, since the Identity Server was started.
Successful Identity ID Broadcasts	The total number of successful Identity ID Broadcasts to identify the authoritative server, since the Identity Server was started.
Session ID Broadcasts	The total number of attempted Session ID Broadcasts to identify the authoritative server, since the Identity Server was started.
Successful Session ID Broadcasts	The total number of successful Session ID Broadcasts to identify the authoritative server, since the Identity Server was started.

SP Brokering

Statistic	Description
Total Brokering Requests	The total number of brokering requests created after the Identity Server was started. This count is a sum of all connections created to all replicas of the configuration datastore and all user stores.
Total Brokering Requests Denied Due to Group Check	The total number of brokering authentication requests denied in a target service provider. The brokering group can either be the identity provider or target service provider but both does not belong to the same group.
Total Brokering Requests Denied Due to Role Deny	The total number of brokering authentication requests to a target service provider denied due to broker policy evaluation denying the role.
Total Brokering Requests Passed	The total number of brokering requests passed after the Identity Server was started.

4.4 Viewing Cluster Statistics

You can view and configure general performance statistics for the servers and service providers assigned to the selected cluster.

Server Statistics

On the Cluster Statistics page, you can configure the list of server statistics to show the desired statistics for an Access Gateway cluster. See Section 4.3.1, "Server Activity Statistics," on page 119 and Section 4.3.2, "Server Benefits Statistics," on page 123 for the complete list of statistics for each server in an Access Gateway cluster.

Perform the following steps:

1 In the Administration Console, click **Devices** > **Access Gateways** > **[Name of Cluster]** > **Statistics** > **Server Statistics**.

2 Click **Configure**.

3 Select **Set default statistics under Selected Statistics** if you want to replace the selected statistics with the default statistics.

The default statistics include CPU Utilization, Cache Hit, Current Total Connections, Throughput of the Origin Server, and Throughput of the Browser.

4 Select statistics from **Available Statistics** and move to **Selected Statistics**.

5 Click **OK**.

The Cluster Statistics page displays the summary of configured statistics for each individual member of the cluster.

To view additional statistical information about a specific Access Gateway, click the name of the Access Gateway in the **Server Name** column of the summary.

You can also view all the statistics for an individual server of the cluster. Click **View** in the **Statistics** column of the summary to see these additional statistics.

For more information, see Section 4.3, "Viewing Access Gateway Statistics," on page 118.

6 Click **Close**.

NOTE: In the cluster level statistics, you can view the list of servers, which are associated with that cluster. If a statistics is not applicable for a particular Access Gateway server, the value of the statistics is displayed as `Not Supported` for that server.

Service Provider Statistics

On the Cluster Statistics page, you can configure the list of service provider statistics to show the desired statistics for an Access Gateway cluster. See Section 4.3.3, "Service Provider Activity Statistics," on page 123 for the complete list of statistics for each server in an Access Gateway cluster.

Perform the following steps:

1 In the Administration Console, click **Devices** > **Access Gateways** > **[Name of Cluster]** > **Statistics** > **Service Provider Statistics**.

2 Click **Configure**.

3 Select **Set default statistics under Selected Statistics** if you want to replace the selected statistics with the default statistics.

   The default statistics include Free Memory, Provided Authentications Failures, Consumed Authentications Failures, Cached Sessions Currently Active Requests - Incoming HTTP Requests, and Currently Active Requests - Outgoing HTTP Requests Statistics.

4 Select statistics from **Available Statistics** and move to **Selected Statistics**.

5 Click **OK**.

   The Cluster Statistics page displays the summary of configured statistics for each individual member of the cluster.

   To view additional statistical information about a specific Access Gateway, click the name of the Access Gateway in the **Server Name** column of the summary.

   You can also view all the statistics for an individual server of the cluster. Click **View** in the **Statistics** column of the summary to see these additional statistics.

   For more information, see Section 4.3.3, "Service Provider Activity Statistics," on page 123.

6 Click **Close**.

4.5   Monitoring API for the Access Gateway Statistics

For programmatic access to the Access Gateway statistics, you require to enable the global advanced option NAGStatsClientIPWhitelist. This option takes a list of IP addresses of servers that can access the Access Gateway statistics.

To access the statistics, run the HTTP GET command on the resource: https://<mag-host-name>/mag-stats.

NOTE: Frequent requests to get the statistics impact the system' performance. It is recommended to keep a five minutes interval between every probe for the statistics.

To enable this option:

1 In the Administration Console, select **Devices** > **Access Gateway Servers** > **Edit** > **Advanced Options**.

2 Add this line: `NAGStatsClientIPWhitelist <ip1> <ip2>`.

3 Replace <ip1> and <ip2> with the IP addresses of the servers from which you want to access the statistics.

4 Click **OK**.

This request displays the following:

- Https related statistics
  - Requests received
  - Active requests
- Server related statistics
  - Product start time
  - Product up time
  - Product CPU utilization
  - Disk swap (KB)
  - Disk swap used (KB)
  - Memory total (KB)
- Cache statistics

NOTE: Cache statistics are 0 because they are not implemented currently in the server side.

  - Cache stats (KB)
  - Cache stats utilization percentage
  - Cache hit ratio since last reset
  - Cache stats object count
- Summary Statistics Byte
  - Total bytes sent to the origin server
  - Total bytes read from the Web server
  - Total bytes sent to the browsers
  - Total bytes received from the browsers
  - Bytes per sec read from the Web server
  - Bytes per sec sent to the browsers
- Summary Statistics Benefits
  - Total bytes saved
  - Total bytes saved per second

Example output:

```
<?xml version="1.0" encoding="UTF-8"?>

<MAGStatistics>

<httpStats>

  <NamedValues>

    <NamedValue name="RequestsReceived" value="0" />

    <NamedValue name="ActiveRequests" value="1" />

  </NamedValues>
```

```
        </httpStats>


<boxStats>
  <NamedValues>
    <NamedValue name="ProductStartTime" value="Fri, 27 Jul 2012 11:01:11 GMT"
/>
    <NamedValue name="ProductUpTime" value="0:0:0:26" />
    <NamedValue name="ProductCPUUtilization" value="-294" />
    <NamedValue name="DiskSwapKb" value="4088532" />
    <NamedValue name="DiskSwapUsedKb" value="0" />
    <NamedValue name="MemoryTotalKb" value="7835" />
  </NamedValues>
</boxStats>


<cacheStats>
  <NamedValues>
    <NamedValue name="cacheStatsKb" value="0" />
    <NamedValue name="cacheStatsUtilPercentage" value="0" />
    <NamedValue name="cacheHitRatioSinceReset" value="0" />
    <NamedValue name="cacheStatsObjectCount" value="0" />
  </NamedValues>
</cacheStats
<summaryStatsByte>
  <NamedValues>
    <NamedValue name="TotalBytesSentToOriginServer" value="0" />
    <NamedValue name="TotalBytesReadFromWS" value="0" />
    <NamedValue name="TotalBytesSentToBrowsers" value="0" />
    <NamedValue name="TotalBytesReceivedFromBrowsers" value="0" />
    <NamedValue name="BytesPsecReadFromWS" value="0" />
    <NamedValue name="BytesPsecSentToBrowsers" value="0" />
  </NamedValues>
</summaryStatsByte>


<summaryStatsBenefits>
  <NamedValues>
    <NamedValue name="TotalBytesSaved" value="0" />
```

```
      <NamedValue name="TotalBytesSavedPerSecond" value="0" />

   </NamedValues>

</summaryStatsBenefits>


<summaryStatsRequests>

  <NamedValues>

    <NamedValue name="TotalRequestsPSecBrowsers" value="0" />

    <NamedValue name="PeakTotalRequestsPSecBrowsers" value="1" />

    <NamedValue name="TotalRequestsPSecOriginServer" value="0" />

    <NamedValue name="PeakTotalRequestsPSecOriginServer" value="0" />

    <NamedValue name="CurrentTotalRequestsToOriginServer" value="0" />

    <NamedValue name="CurrentTotalRequestsReceivedFromBrowser" value="1" />

    <NamedValue name="FailedRequestsToWS" value="0" />

    <NamedValue name="CumulativeRequestsToWS" value="0" />

  </NamedValues>

</summaryStatsRequests>


<summaryStatsConnections>

  <NamedValues>

    <NamedValue name="CurrentConnectionsBrowser" value="10" />

    <NamedValue name="CurrentConnectionsBackend" value="0" />

    <NamedValue name="TotalConnectionsBrowser" value="28" />

    <NamedValue name="TotalConnectionsBackend" value="0" />

    <NamedValue name="PeakConnectionsBrowser" value="6" />

    <NamedValue name="PeakConnectionsBackend" value="0" />

    <NamedValue name="FailedConnectionsBackend" value="0" />

  </NamedValues>

</summaryStatsConnections>

</MAGStatistics>
```

4.6 Monitoring Access Gateway Alerts

The Access Gateway has been programmed to issue events to various types of systems (such as a Novell Audit server, a NetIQ Sentinel 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 . Alerts can also be configured so that the administrator is informed when significant changes occur.

4.6.1 Viewing Access Gateway 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 **Devices** > **Access Gateways** > **[Name of Server]** > **Alerts**.

2 To delete an alert from the list, select the check box for the alert, then click **Acknowledge Alert(s)**. To remove all alerts from the list, click the **Severity** check box, then click **Acknowledge Alert(s)**.

3 Click **Close**.

4 (Optional) To verify that the problem has been solved, click **Access Gateways** > **[Server Name]** > **Health** > **Update from Server**.

4.6.2 Viewing Access Gateway Cluster Alerts

To view information about current alerts for all members of a cluster:

1 In the Administration Console, click **Devices** > **Access Gateways** > **[Name of Cluster]** > **Alerts**.

2 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.

Column	Description
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 acknowledge all alerts for an Access Gateway, select the check box for the Access Gateway, then click **Acknowledge Alert(s)**. When you acknowledge an alert, you clear the alert from the list.

4 To view information about a particular alert, click the server name.

4.6.3   Managing Access Gateway Alert Profiles

For an Access Gateway, this option allows you to send notification of generated system alerts to the Administration Console, to a Syslog server, to a log file, or to a list of e-mail recipients.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Alerts**.

2 Select one of the following actions:

New: To add a new profile, click **New**. Specify a name for the profile, then click **OK**. For configuration information, see Section 4.6.4, "Configuring an Alert Profile," on page 134.

Enable: To enable a profile, select the check box next to the profile, then click **Enable**.

Disable: To disable a profile, select the check box next to the profile, then click **Disable**.

Delete: To delete a profile, select the check box next to the profile, then click **Delete**.

3 Click **OK** twice.

4 On the **Access Gateways** page, click **Update**.

4.6.4   Configuring an Alert Profile

The alert profile determines which alerts are sent and where the alerts are sent.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Alerts** > **[Name of Profile]**.

2 Select one or more of the following:

Connection Refused: Generated when a 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.

Failure in Audit, Stopping Services: Generated when the audit server has failed, and the Access Gateway has been configured to stop services.

To configure the Access Gateway to continue when auditing services are not available, click **Auditing** > **Novell Auditing**, deselect the **Stop Services on Audit Server Failure** option, then click **Apply**.

Failure in Audit, Will lose events, but continuing services: Generated when the audit agent has failed. The Access Gateway continues to run, but no audit events are generated.

As a workaround while solving this problem, you can enable proxy service logging (see Section 4.2, "Configuring Logging for a Proxy Service," on page 110). The common and extended log files provide some details on the HTTP traffic.

If you do not want the Access Gateway to run without generating events, you need to manually shut down the Access Gateway.

Failure in Audit, Server is offline: Generated when the audit agent is unable to contact the audit server. When this condition occurs, the audit agent uses local caching for the audit events.

Do not allow this condition to continue indefinitely. The Access Gateway soon reaches the limits of its local cache. If this happens, events can be lost and the Access Gateway might need to stop services.

For troubleshooting information, see "Troubleshooting Novell Audit" (http://www.novell.com/documentation/novellaudit20/novellaudit20/data/al0lh30.html) in the *Novell Audit Administration Guide* (http://www.novell.com/documentation/novellaudit20/novellaudit20/data/bookinfo.html).

3 Select where you want the alerts sent:

Send to Device Manager: Select this option to send alerts to the Administration Console.

Send to SNMP: (Access Gateway Service) Select this option to send alerts to an SNMP server. To configure the SNMP server, click the **Send to SNMP** link. For configuration information, see "SNMP Profile" on page 136.

Send to Log File: Select this option to send alerts to a log file. To send alerts to a log file, click **New**, specify a name for the log profile, then click **OK**. For configuration information, see "Configuring a Log Profile" on page 136.

To enable a log profile, select the profile, then click **Enable**.

To disable a log profile, select the profile, then click **Disable**.

To delete a log profile, select the profile, then click **Delete**. Click **OK** in the confirmation dialog box.

Send E-mail Notifications: Select this option to send alerts through e-mail notifications. To enable e-mail notification click **New**, specify a name for the e-mail profile, then click **OK**. For configuration information, see "Configuring an E-Mail Profile" on page 136.

To enable an e-mail profile, select the profile, then click **Enable**.

To disable an e-mail profile, select the profile, then click **Disable**.

To delete an e-mail profile, select the profile, then click **Delete**. Click **OK** in the confirmation dialog box.

Send to Syslog: Select this option to enable syslog alerts. Click **New**, specify a name for the syslog profile, then click **OK**. For configuration information, see "Configuring a Syslog Profile" on page 137.

To enable a syslog profile, select the profile, then click **Enable**.

To disable a syslog profile, select the profile, then click **Disable**.

To delete a syslog profile, select the profile, then click **Delete**. Click **OK** in the confirmation dialog box.

4 To enable an alert action profile, select the action profile, click **Enable**, then click **OK**.

The action to send the alerts to a log file, to e-mail addresses, or to a syslog file is not performed until the action profile is enabled.

5 On the Alert Profiles page, verify that the alert profile you have created is enabled, then click **OK** twice.

6 Update the Access Gateway.

4.6.5 SNMP Profile

1 (Access Gateway Service) To add the IP address of a SNMP server, click **New**, specify the IP addresses, then click **OK**.

2 (Optional) To delete an IP address, select the IP address, then click **Delete**.

3 Click **OK**.

4 Select one of the following:

- ◆ To add another profile, continue with .
- ◆ To save your modifications, continue with .

4.6.6 Configuring a Log Profile

The **Send to Log File** field displays the name of the log profile you are configuring.

1 Fill in the following fields:

Log File Name: Specify a name for the log file and a path where the file should be stored.

You must specify the full path.

```
/var/opt/novell/amlogging/logs/
```

Max File Size: Specify a maximum size for the log file in KB. The size can be from 50 to 100000 KB. Specify 0 to indicate that there is no maximum file size.

2 Click **OK**.

3 Select one of the following:

- ◆ To add another profile, continue with .
- ◆ To save your modifications, continue with .

4.6.7 Configuring an E-Mail Profile

The **Send E-Mail Notifications** field displays the name of the e-mail profile you are configuring.

1 Fill in the following fields:

E-mail Recipients: To add a recipient to the list, click **New**, specify the e-mail address of the recipient, then click **OK**. You can add multiple e-mail addresses. To delete a recipient, select the user's email address, click **Delete**, then click **OK**.

Mail Exchange Servers: To add a mail server, click **New**, specify the IP address or the DNS name of the mail exchange server, then click **OK**. You can add multiple mail exchange servers. To delete a server, select the server, click **Delete**, then click **OK**.

2 Click **OK**.

3 Select one of the following:

- To add another profile, continue with .

- To save your modifications, continue with .

4.6.8 Configuring a Syslog Profile

The **Send to Syslog** field displays the name of the syslog profile you are configuring.

1 Fill in the following field:

Facility Name: Specify a facility name for the Syslog server. It can be any name from local0 to local7. If you specify local0 - local7 as your facility name, the alerts are stored at `/var/log/ localmessages`.

2 Click **OK**.

3 Select one of the following:

- To add another profile, continue with .

- To save your modifications, continue with .

To configure the syslog profile for Access Gateway Service on RedHat Enterprise Linux, use the following procedure:

1 Go to `/etc/rsyslog.conf` file.

2 Add the following under `# Provides UDP syslog reception`

$ModLoad imudp.so

$UDPServerRun 514

3 Restart the syslog service using one of the following commands:

`/etc/init.d/rsyslog restart` OR `rcsyslog start`

4.7 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 auditing on your network. For more information, see "Configuring Access Manager Appliance for Auditing" in the *NetIQ Access Manager Appliance 4.0 SP1 Administration Console Guide*.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Auditing**.

Events

☐ Select All

☐ Access Denied	☐ Access Allowed	☐ Identity Injection Failed	☐ Identity Injection Parameters
☐ System Started	☐ System Shutdown	☐ Form Fill Success	☐ Form Fill Failed
☐ URL Accessed	☐ URL Not Found	☐ IP Access Attempted	

Changes made on this panel must be applied or scheduled from the Configuration Panel.

OK   Cancel

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.
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 **Identity Injection Failed** event.
Form Fill Failed	Generated when a Form Fill policy fails to successfully fill in a form.

IMPORTANT: Enabling **URL Accessed** and **Access Allowed** events may generate high volume of audit events on a system with heavy load. This may degrade the performance of the Access Gateway.

3 To save your modifications, click **OK** twice.

4 On the Access Gateways page, click **Update**.

NOTE: For the Access Gateway Service, caching of audit events is disabled when the audit server is not reachable. To enable caching of audit events on Apache Gateway Service, see .

4.8 Managing Server Health

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 > Dashboard**

  The Dashboard page shows the heath status at the component-level.

- **Auditing > Device Health**

  The Device Health page shows the health status for all devices in one list.

◆ **Devices > [Component]**

The Servers page for each component provides a health status for each device.

This section discusses the following topics:

◆
◆
◆

4.8.1 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 green status with a yellow diamond indicates that the server has not detected any problems but the configuration isn't completely up-to-date because commands are pending.
	A green status with a red x indicates that the server has not detected any problems but that the configuration might not be what you want because one or more commands have failed.
	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 in not running or functional, or that the server is having a runtime problem.

4.8.2 Monitoring the Health of an Access Gateway

To view detailed health status information of an Access Gateway:

1 In the Administration Console, click **Devices > Access Gateways** > **[Name of Server]** > **Health**.

The status icon is followed by a description that explains the significance of the current state. For more information about these icons, see Section 4.8.1, "Health States," on page 139.

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. If you have made changes that affect the health of the Access Gateway, select this option. Otherwise, it can take up to five minutes for the health status to change.

3 Examine the **Services Detail** section that displays the status of each service. For an Access Gateway, this includes information such as the following:

   ◆ "Service Categories of the Access Gateway Service" on page 140

4 Click **Close**.

Service Categories of the Access Gateway Service

Service Category	If Not Healthy
Reverse Proxy - <*Proxy Service Name*>: Indicates the general health of all configured proxy services. A separate row is created for each proxy service.	Check the health of the Web server.
AGM - Configuration: Indicates whether all configuration changes have been applied.	Do the following: ◆ To re-push the current configuration, click **Auditing** > **Troubleshooting**, select the gateway from the list of the **Current Access Gateway Configurations**, then click **Re-push Current Configuration**. ◆ To revert to last applied configuration, click **Devices** > **Access Gateways** > **Edit**, then click **Revert**. If these options do not fix the problem, view the Apache error.log file to discover the cause. The file is located in the following directory:
TCP Listener - <*IP Address:Port*>: Indicates whether the Access Gateway Service is listening on the specified port. A separate row is created for each port the Gateway Service is configured to listen on.	Restart the Apache service.
ApacheGateway.log: Appears when the Access Gateway Service is not healthy. It displays the latest error from the Apache error.log file.	For more information about the problem, view the error.log file in the following directory:
Embedded Service Provider Configuration: Indicates whether the Access Gateway has been configured to trust an Identity Server and whether that configuration has been applied. At least one Identity Server must be configured and set up as a trusted authentication source for the Access Gateway. A green status indicates that a configuration has been applied; it does not indicate that it is a functioning configuration.	See Section 2.2, "Managing Reverse Proxies and Authentication," on page 23 for information about assigning an Identity Server configuration to the Access Gateway.
Configuration Datastore: Indicates whether the configuration datastore is functioning correctly.	Restore the configuration datastore. See "Repairing the Configuration Datastore" in the *NetIQ Access Manager Appliance 4.0 SP1 Administration Console Guide*.
Clustering: Indicates whether all the cluster members are active and processing requests.	Restart the cluster members that are not active or remove them from the cluster.

Service Category	If Not Healthy
Signing, Encryption and SSL Connector Keys: Indicates whether these keystores contain a valid key.	Click **Access Gateways** > **Edit** > **Service Provider Certificates** and replace any missing or expired keys.
System Incoming and Outgoing HTTP Requests: Appears when throughput is slow. This health check monitors incoming HTTP requests, outgoing HTTP requests on the SOAP back channel, and HTTP proxy requests to cluster members. If one or more requests remain in the queue for over 2 minutes, this health check appears.	Verify that all members of the cluster have sufficient bandwidth to handle requests. If a cluster member is going down, the problem resolves itself as other members of the cluster are informed that the member is down. If a cluster member is slow because it doesn't have enough physical resources (speed or memory) to handle the load, upgrade the hardware.
TCP Listener(s): Indicates whether the listening port for the Embedded Service Provider is healthy.	Restart the Access Gateway.
Embedded Service Provider's Trusted Identity Provider: Indicates whether the configuration that the Access Gateway trusts has been configured to contain at least one Identity Server.	Modify the Identity Server configuration and add an Identity Server. Configure the Access Gateway to trust an Identity Server configuration. See Section 2.2.1, "Creating a Proxy Service," on page 26.
Audit Logging Server: Indicates whether the audit agent is functioning and able to log events to the auditing server. Auditing must be enabled on the Identity Server to activate this health check (click **Devices** > **Identity Servers** > **Edit** > **Logging**).	Check the network connection between the Identity Server and the auditing server. See "Troubleshooting Novell Audit" (http://www.novell.com/documentation/novellaudit20/novellaudit20/data/al0lh30.html).

4.8.3 Viewing the Health of an Access Gateway Cluster

The **Health** icon on the cluster row displays the status of the least healthy member of the cluster. For information about the meaning of health icons, see Section 4.8.1, "Health States," on page 139.

To view details about the status of the cluster:

1 In the Administration Console, click **Devices** > **Access Gateways**.

2 On the cluster row, click the **Health** icon.

3 To ensure that the information is current, click **Refresh**.

4 To view specific information about the status of an Access Gateway, click the **Health** icon in the Access Gateway row.

4.9 Viewing the Command Status of the Access Gateway

Commands are issued to an Access Gateway when you make configuration changes and when you select an action such as stopping or starting the gateway.

Certain commands, such as start and stop commands, retry up to 10 times before they fail. The first few retries are spaced a few minutes apart, then they move to 10-minute intervals. These commands can take over an hour to result in a failure. As long as the command is in the retry cycle, the command has a status of pending.

- If you do not want to wait for the cycle to complete, you need to manually delete the command.
- If you enter the same command and it succeeds before the first command has completed its retry cycle, the first command always stays in the pending state. You need to manually delete the command.

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 4.9.1, "Viewing the Status of Current Commands," on page 142
- Section 4.9.2, "Viewing Detailed Command Information," on page 143

4.9.1 Viewing the Status of Current Commands

1 In the Administration Console, click **Devices** > **Access Gateways** > **[Name of Server]** > **Command Status**.

	Name	Status	Type	Admin	Date & Time (Note)
☐	10.10.15.206 Start	EXECUTING	Service Provider Start	cn=admin,o=novell	Feb 27, 2007 3:12 PM
☐	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 4.9.2, "Viewing Detailed Command Information," on page 143.
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**.

4.9.2 Viewing Detailed Command Information

To view information about an individual command:

1 In Administration Console, click **Devices > 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

The following command information is listed:

Name: Specifies the display name that has been given to the command.

Type: Specifies the type of command.

Admin: Specifies whether the system or a user issued the command. If a user issued the command, the field contains the DN of the user.

Status: Specifies the status of the command, and includes such states as **Pending**, **Incomplete**, **Executing**, and **Succeeded**.

Last Executed On: Specifies when the command was issued. The date and time are displayed in local time. If the command failed, additional information is available.

For a command that the Administration Console can successfully send to the Access Gateway, the page displays a **Command Execution Details** section with the name of the command and the command results.

For a command that fails because the Administration Console cannot communicate with the Access Gateway, the page displays the following additional fields:

Number of Tries: Specifies the number of times the command was executed.

Command Try Log: Lists each try and the results.

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.

5 Advanced Access Gateway Options

The following sections describe the advanced options along with examples:

- Section 5.1, "Configuring the Global Advanced Options," on page 145
- Section 5.2, "Configuring the Advanced Options for a Domain-Based and Path-Based Multi-Homing Proxy Service," on page 157

5.1 Configuring the Global Advanced Options

The following settings apply to all reverse proxies, unless the option is overwritten by an advance proxy service setting (see Section 5.2, "Configuring the Advanced Options for a Domain-Based and Path-Based Multi-Homing Proxy Service," on page 157). The advanced options are disabled by default and will be enabled when they are added.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Advanced Options.**

2 To activate these options, configure the value, save your changes, then update the Access Gateway. To deactivate these options, add the pound (#) symbol.

5.1.1 Advanced Access Gateway Options

The following table lists the advanced options along with their descriptions, default values, and examples. Most of the global advanced options are prefixed with `NAGGlobalOptions` and the domain-based and path-based multi-homing proxy service options have been identified and mentioned below the option name.

Table 5-1 *Advanced Access Gateway Options*

Advanced Option	Description
`NAGGlobalOptions FlushUserCache=on` This is a global advanced option.	Specifies whether cached credential data of the user is updated when the session expires or the user changes an expiring password. This option is equivalent to `PasswordMgmt` in the 3.1 SP4 Access Gateway Appliance. ◆ When this option is on, which is the default setting, the credentials and the Identity Injection data are refreshed. ◆ When this option is turned off, the cached user data can become stale. For example, if your password management service is a protected resource of the Access Gateway and this option is turned off, every time a user changes an expiring password, the user's data is not flushed and the Access Gateway continues to use stale data for that user.

Advanced Option	Description
NAGGlobalOptions DebugHeaders=on	When this option is enabled, an X-Mag header is added with debug information. The information can be seen in sniffer traces and with plug-ins such as ieHTTPHeaders, Live HTTP Headers, and FireBug. This option should only be enabled when you are working with NetIQ Support and they instruct you to enable the option.
NAGGlobalOptions DebugFormFill=on This is a global advanced option.	When this option is enabled, additional debug information related to the processing of a Form Fill policy is added to the Apache error log files (error_log file under /var/log/novell-apache2 and to the X-Mag header in the response to browser. The Form Fill entries generated by this option begin with a FF: marker. For example, Oct 23 12:38:29 mag326 httpd[29345]: [warn] AMEVENTID#36: FF:fillSilent: kfh5ummigbq6uGeneral_SS_non_SS_autosumit_Page_13310, referer: https://www.idp.com:8443/nidp/idff/sso?sid=0 Oct 23 12:38:29 mag326 httpd[29345]: [warn] AMEVENTID#36: FF:fillInplaceSilent: kfh5ummigbq6uGeneral_SS_non_SS_autosumit_Page_13310, referer: https://www.idp.com:8443/nidp/idff/sso?sid=0
NAGGlobalOptions ESP_Busy_Threshold=<value> This is a global advanced option.	Proxy starts sending errors to the browser if ESP's average response time in the last one minute is more than the specified value (time in milliseconds).
NAGGlobalOptions noTOPR This is a global advanced option.	Disables the activity based time-out in proxy. The proxy redirects browser requests after soft timeout of configured timeout value. This option is equivalent to disabletoppr in the 3.1 SP4 Access Gateway Appliance.
NAGGlobalOptions InPlaceSilent=on This is a global advanced option.	This enables single sign on to certain Web sites that require the login page to remain as is without any modifications to its structure. If you are using this advanced option for a Form Fill on a page with multiple forms, by default, the first form is posted. If you want to post forms other than the first form, use NAGGlobalOptions InPlaceSilentPolicyDoesSubmit=on. For more information, refer to TID 7011817. This option is equivalent to .enableInPlaceSilentFill in the 3.1 SP4 Access Gateway Appliance.
NAGGlobalOptions ForceUTF8 This is a global advanced option.	When this file is enabled, the Access Gateway uses the UTF-8 character set to serve the Form Fill page to the browser. This option is equivalent to forceUTF8Charset in the 3.1 SP4 Access Gateway Appliance.

Advanced Option	Description
NAGGlobalOptions AllowMSWebDavMiniRedir This is a global advanced option.	This file helps the user to disable the following functionality, which is enabled by default. If a Microsoft Network Places client sends an OPTIONS request with MS-WebDAV-MiniRedir useragent to the Access Gateway, then it receives 409 conflict response. The client uses this response to change the user agent to MS Data Access Internet Publishing Provider DAV. This option is equivalent to AllowMSWebDavRedir in the 3.1 SP4 Access Gateway Appliance.
NAGGlobalOptions noURLNormalize=on This is a global advanced option.	When this option is enabled, it disables the URL normalization protection for back-end Web servers. This option resolves issues in serving Web content from Web servers that have double-byte characters such as Japanese language characters. By default, this option is set to off and URL is normalized before sending it to back end Web server.
NAGAdditionalRewriterScheme webcal:// This is a global advanced option.	When this option is enabled, the rewriter rewrites URLs that have a scheme of webcal://. The default rewriter configuration only rewrites URLs with a scheme of http:// or https://.
NAGGlobalOptions AppendProviderID=on This is a global advanced option.	When this option is enabled, it displays the ESP Provider ID in the Access Gateway authorization audit logs. This option helps to know the issues related to ESP provider ID in the audit log file.
NAGGlobalOptions InPlaceSilentPolicyDoesSubmit=on This is a global advanced option.	This option should be used to fill forms with complex JavaScript or VBScripts. This option is equivalent to .enableInPlaceSilentFillNew in the 3.1 SP4 Access Gateway Appliance.

Advanced Option	Description

`NAGGlobalOptions`
`NAGErrorOnIPMismatch=off`
 	
This is a global advanced option.
	If the value for this option is set to off, the Access Gateway does not perform the IP address check on incoming session cookies. Use this in a setup where two L4 switches are configured in parallel and the browser requests are bounced between these L4 switches.
	
	This option is equivalent to `.lagdisableAuthIPCheck` in the 3.1 SP4 Access Gateway Appliance.
	
	For example, if multiple back-end Web servers are accelerated by the Access Gateway, some users complain that they are not able to complete logging in. When they access the protected resources, they are redirected to the Identity Server for authentication, but they are not redirected to the original URL.
	
	If multiple paths (at the network level) exist between a browser and the Access Gateway and proxies or NAT devices exist on these paths, it is possible that the source IP address of the incoming requests into the Access Gateway might change. For example, assume that user A connects to an ISP. This ISP has multiple transparent proxies in parallel for performance reasons.
	
	User A accesses the Access Gateway for the first time. The request from User A goes through a local transparent proxy TP1, so the incoming IP address of the initial request has that transparent proxy's (TP1) IP address. The Access Gateway session cookie is set and the user is redirected back to the page the user was going to originally.
	
	User A then sends the next request for this original page, but it goes through a different proxy, TP2. The incoming IP address of the request into the Access Gateway is now different than the one that the user used for authentication (TP1 IP address) and the validation fails. The Access Gateway loops as it continues to request the user to send a valid session cookie.
	
	NOTE: On receiving IPC cookie from browser, the Access Gateway checks for the client IP address in the cookie. If the IP address in the cookie and the client IP address from which the request came do not match, Access Gateway displays an error page.

`NAGGlobalOptions`
`NAGDisableExternalRewrite=on`
 	
This is a global advanced option.
	Access Gateway does not insert the path for the links with external published DNS when you enable this option.
	
	This option is equivalent to `.disableExternalDNSRewrite` in the 3.1 SP4 Access Gateway Appliance.

`DisableGWSHealth on` This is a global advanced option.	When this option is enabled, Access Gateway does not check health of the Web server with the back-end server. This option is equivalent to `.disableWSHealth` in the 3.1 SP4 Access Gateway Appliance.

Advanced Option	Description
`NAGIchainCookieVersion on` This is a global advanced option.	When this option is enabled, Access Gateway sends the proxy session cookie to the back-end server as IPCZQX01<clusterid>.
`IgnoreDNSServerHealth on` This is a global advanced option.	When this option is used, the Access Gateway does not send the DNS server health status when the Access Gateway health is reported to the Administration Console. When you set the option to `IgnoreDNSServerHealth off <lookupname>`, the Access Gateway sends a DNS query with the specified *<lookupname>*. The Access Gateway sends a successful message to the Administration Console if it connects to the DNS server, else it will send an unable to connect message. By default if you have not specified any option, the Access Gateway sets the option as `IgnoreDNSServerHealth off www.novell.com`. This option is equivalent to `ignoreDnsServerHealth` in the 3.1 SP4 Access Gateway Appliance.
`NAGHostOptions primaryWebdav=<path of pbmh service>` This is a global advanced option.	This option enables users who use the Microsoft Network Places client to connect to the WebDAV folders of a SharePoint server when the SharePoint server has been configured as a path?based multi-homing service on the Access Gateway. This should be added to master proxy service **Advanced Options** whose path based child services accelerates webdav resources with remove path on fill option enabled. This option is equivalent to `.modifyRequestURI` in the 3.1 SP4 Access Gateway Appliance.
`NAGHostOptions mangleCookies=on` This is a proxy service option.	This option invalidates the cookies set by the Web server when the user logs out of Access Manager. By default, the Access Gateway does not mangle the cookies that are sent by the Web server. Proxy mangles the cookies that are sent by the Web server using the user information and sets these mangled cookies at the browser. When a browser sends the mangled cookies to proxy, it de-mangles them using the user information and sends the de-mangled cookies to the Web server. For more information about this option, see "Enabling Cookie Mangling" on page 156.
`NAGWSMangleCookiePrefix` This is a global advanced option.	Use the `NAGWSMangleCookiePrefix <AnyString>` option to specify the string added to the application cookie after manipulation. For more information about this option, see "Enabling Cookie Mangling" on page 156.
`NAGHostOptions webdavPath=/_vti_bin` This is a global advanced option.	This can be added to master proxy service **Advanced Options** which path based child services with **remove path on fill** option enabled accelerating webdav resources.

Advanced Option	Description
`NAGChildOptions WebDav=<path of pbmh service>` This is a global advanced option.	This option can be added to any path based service that accelerates webdav resources with **remove path on fill** enabled. This option is equivalent to `.modifyRequestURI` in the 3.1 SP4 Access Gateway Appliance.
`EnableWSHandshake on` This is a global advanced option.	Setup a firewall between the Access Gateway and the back-end Web server. When the Access Gateway performs heartbeat check with a simple TCP connect to the Web server, the Web server may throw a TLS handshake error. This may cause the firewall, after a certain threshold, to block the connection. This option enables the Access Gateway to perform a SSL handshake while performing a heartbeat check on the back-end SSL-enabled Web server so that the Web server does not respond with a TLS handshake error. By default, the Access Gateway performs a simple TCP connect while performing a heartbeat check on the back-end Web server.
`NAGGlobalOptions IIRemoveEmptyHeaderValue` This is a global advanced option.	This option enables the Identity Injection policy not to send an empty header with null value when a value is not available. By default, the Access Gateway sends an empty header with a null value if a value is not available. For example, applications may have a public and a protected resource configured. Both resources may use an identity injection policy such as to inject an USERID. The public resource uses the user name if authenticated. If the user accesses the public resource (before authentication), the Access Gateway sends an empty header variable USERID. Web servers may not handle an empty header and may respond with an error. In such a scenario use the advanced option to stop the Access Gateway from sending an empty header with null value.

`DumpHeaders on`
 	
`DumpHeadersFacility user`
 	
This is a global advanced option.
	These options ensure that the proxy, logs the user headers to `/var/opt/novell/nam/logs/mag/apache2/error_log` file.
	
	For information about the equivalent touch files in the 3.1 SP4 Access Gateway Appliance, see list of touch files.

`NAGFilteroutUrlForAudit` This option is available for both domain-based and path-based multi-homing proxy services.	You can add this option to proxy service that filters out specific URLs from auditing (URL Accessed). For example, `NAGFilteroutUrlForAudit ".*.jpg"`, and `NAGFilteroutUrlForAudit ".*.gif"`.

Advanced Option	Description

`FlushUserCache=on`
 	
This is a global advanced option.
	Specifies whether cached credential data of the user is updated when the session expires or the user changes an expiring password.
	
	  ◆ When this option is on, which is the default setting, the credentials and the Identity Injection data are refreshed.
	
	  ◆ When this option is turned off, the cached user data can become stale.
	
	      For example, if your password management service is a protected resource of the Access Gateway and this option is turned off, every time a user changes an expiring password, the user's data is not flushed and the Access Gateway continues to use stale data for that user.

`SSLProxyVerifyDepth=3` This is a global advanced option.	Specifies how many certificates are in a Web server certificate chain. When you activate the verification of the Web server certificate with the **Any in Reverse Proxy Trust Store** and the public certificate is part of a chain, you need to specify the number of certificates that are in the certificate chain. For more information about configuring Web servers for SSL, see Section 1.4, "Configuring SSL between the Proxy Service and the Web Servers," on page 16. ◆ The default search level that is when the attribute **SSLProxyVerifyDepth** is commented to1, if the number of certificates in the Web server certificate chain is greater than 1, then the **SSLProxyVerifyDepth** option should be enabled and should be assigned to the respective value (equal to the number of certificates in the chain).

Advanced Option	Description

`ProxyErrorOverride`
 	
This is a global advanced option.
	Allows you to specify which errors you want returned to the browser unchanged by the Gateway Service. The default behavior of the Gateway Service is to replace Web server errors with Gateway Service errors.
	
	However, some applications put more information, such as keys and JavaScript in the message. If this information is critical, specify an override and allow the error message to be returned to the browser without any modifications.
	
	For example, NetStorage requires an override for the 401 error because it includes a key in the 401 error. The portal page for the Novell Open Enterprise Server requires an override for error 403 because it includes JavaScript.
	
	You can use the following syntax to set this option:
	
	◆ `ProxyErrorOverride on -401 -403:` Allows all errors to be changed to Gateway Service errors except errors 401 and 403, which are sent unchanged.
	
	  This syntax allows you to list the few errors you want to forward without change while allowing all the others to be changed to Gateway Service errors.
	
	◆ `ProxyErrorOverride off +401 +403:` Disables the changing of Web server errors to Gateway Service errors except for errors 401 and 403, which are changed to Gateway Service errors.
	
	  Use this option when you have only a few errors that you want changed to Gateway Service errors.
	
	NOTE: Enable the error codes 401 and 403 for override if you are using Identity Manager 4.0 with Role Mapping Administrator.

`CacheIgnoreHeaders` This option is available only for domain-based proxy service.	Prevents the Access Gateway from writing any Authorization headers to disk. This option is enabled by default, because writing Authorization headers to disk is a potential security risk. You can allow Authorization headers to be written to disk by placing a pound (#) symbol in front of the option or by setting it to `None`. For more information about this Apache option, see "CacheIgnoreHeaders Directive" (http://httpd.apache.org/docs/2.2/mod/mod_cache.html#cacheignoreheaders). **NOTE:** All the path-based services under the domain-based service will inherit the new value.

Advanced Option	Description
`CacheMaxFileSize` This option is available only for domain-based proxy service.	Configuring this value in the **Advanced Options** of a proxy service allows you to set the size of the file that can be stored in the cache. By default the size is set to 5 MB. Add the line CacheMaxFileSize *<bytes>*, for example, CacheMaxFileSize 99900000. **NOTE:** All the path-based services under the domain-based service will inherit the new value.
`NAGErrorOnDNSMismatch` This is a global advanced option.	If SSL is not enabled in reverse proxy, an error message stating **Host Name Does Not Match** is displayed.
`NAGChildOptions WebDav=/Path` This option is valid only for path-based multi-homing proxy service.	Allows the proxy service to handle the specified path. Remove the pound (#) symbol and replace */Path* with the path you want the proxy service to handle.

`SSLHonorCipherOrder`
 	
This is a global advanced option.
	This option enables you to customize the `SSLCipherSuite` used by the Access Gateway. This helps you in taking preventive measures when new vulnerabilities are published.
	
	To avoid Browser Exploit Against SSL/TLS (BEAST) attacks, use the advanced option as follows:
	
	`SSLHonorCipherOrder on`
	
	`SSLCipherSuite !aNULL:!eNULL:!EXPORT:!DSS:!DES:RC4-SHA:RC4-MD5:ALL`
	
	For more information about the format and set of options you can specify in the value, see OpenSSL documentation.

`NAGGlobalOptions onFormFillPolicyRedirUseHttp=on` This is a global advanced option.	This option enables Access Gateway to redirect based on HTTP status code 302 along with the location header when Form Fill policy requires redirect. By default, Access Gateway uses JavaScript to trigger redirect in Form Fill policy. You can use this advanced option when there are issues with JavaScript redirects.
`NAGGlobalOptions OverWriteAuthHeaderWithIIData=on`	The global advanced option ensures that when a browser sends an authentication header, the Access Gateway overwrites it with the authentication header configured in the Identity Injection policy. Setting it to off ensures that when the browser sends an authentication header, it is not overwritten with the authentication header configured in the Identity Injection policy. Setting it to on ensures that when a browser sends an authentication header, it is overwritten with the authentication header configured in the Identity Injection policy.

Advanced Option	Description
NAGLAGCompatiability on	This option enables sharing of session information between the 3.1 SP4 Access Gateway Appliance and the 4.0 Access Gateway Appliance during the process of migration. This option is added by default during the process of migration to ensure communication between the two appliances. You can disable or remove this option after the migration is complete.
ProxyPassIgnorePathCase on	Use this option to make the path-based multi-homing path URL case-insensitive. For example, if you have set up a path based proxy /profile in Administration Console and the end user wants to access the URL https://www.lagssl.com/Profile/Security/login.aspx and not https://www.lagssl.com/profile/Security/login.aspx. By default url path is case sensitive.
NAGPostParkingSizeInKiloBytes	This option allows you to change the post data parking size limit if an error occurs after you post large data (more than 56 KiloBytes in size) after a session timeout.
NAGSendURLinErrorResponses on	This option will not include a href when you access a protected resource and a 302 redirect occurs.
SSLProtocol	This option is supported by the Access Gateway when listening as a server to clients (typically browsers). This directive specifies SSL protocols for mod_ssl to use when establishing the server environment. Clients can only connect with one of the specified protocols. The accepted values are SSLv3, TLSv1, TLSv1.1, TLSv1.2 and all of these. The syntax for this is SSLProtocol [+-]protocol. For example, SSLProtocol +SSLv3. For more information about configuring the SSL versions, see Apache documentation.
SSLProxyProtocol	This option is supported by the Access Gateway when listening as a server to clients (typically browsers). This directive specifies SSL protocols for mod_ssl to use when establishing a proxy connection in the server environment. Proxies can only connect with one of the specified protocols. The accepted values are SSLv3, TLSv1, TLSv1.1, TLSv1.2 and all of these. The syntax for this is SSLProxyProtocol [+-]protocol. For example, SSLProxyProtocol +SSLv3. For more information about configuring the SSL versions, see Apache documentation.

Advanced Option	Description
`NAGSessionKey` This is a global advanced option.	For additional security in case of cross-domain authentication, the Access Gateway session Cookie is encrypted before sending it as a URL query parameter. For example: In earlier releases of Access Manager, the URL is: `https://novell.blr.com:9443/ - CIPCZQX03218a425f=01000300a463892f582b517 22510f334a4223149` In Access Manager 4.1, the URL is: `https:// novell.blr.com:9443/%20- CECCjdOOBPIqZZNtF+dRlAyDfTFvOPwnO0xzOQTcn rubNzJ6GFe6FF8dWRzzg7RY9iZJYxNLaU80KnJOoq tqf6u2g==` This advanced option `NAGSessionKey` can be used to specify the password as per the administrator's needs. Passwords with more characters increase the strength of the password and therefore leads to better security. For example: `NAGSessionKey NAM-CROSS-DOMAINSESSION-KEY-ENCRYPTION-PASSWORD.` By default, the password is set to "default"

Options to Clean Up Thick Client State At Browser

Currently, when the idle timeout is detected by the Access Gateway, the user is redirected to the Identity server for authentication. If the content type and url pattern used by the client (as defined in the advanced options `NAGUrlPattern` and `NAGContentType`), the user should be redirected to a pre-defined timeout as defined in the `NAGAuthFrontChannel` advanced option.

The redirected URL will also have additional information like the ESP login URL, the contract name as well as the landing page URL as defined in the advanced options. The following advanced options must be used together to clean up the thick client.

Advanced Option	Description
`NAGLauncher`	URL that launches the client.
`NAGUrlPattern /messagebroker/*`	URL pattern that identifies if a specific request came from a client.
`NAGContentType application/x-amf`	Content type in the Request header that is used to identify if the request is a client.
`NAGAuthBackChannel /namtimeout/ timeoutamf`	Timeout handler on the server.
`NAGAuthFrontChannel`	Timeout handler on the server which includes the published DNS name of the server.

Enabling Cookie Mangling

When you log out of Access Manager, the Access Manager session cookies will be invalidated on all Identity and Access Gateway Servers. However, the application session cookie is left unchanged on both the browser and the origin Web server. If a different user authenticates to Access Manager again on the same browser and accesses the proxied Web server, the browser may resume the previously established HTTP session with the Web server so that the new user inherits the old logged out users session. The Cookie Mangling feature in Access Gateway prevents this scenario from occurring by manipulating the application cookies set by the Web servers, and invalidating these cookies when the user logs out of Access Manager.

The two advanced Access Gateway options required to enable this functionality are the `NAGHostOptions mangleCookies` and `NAGWSMangleCookiePrefix`. By default, the option `NAGHostOptions mangleCookies` is set to Off.

To enable this feature, add the options, `NAGHostOptions mangleCookies=on` and `NAGWSMangleCookiePrefix <AnyString>` in the Global Advanced Option.

Use the `NAGWSMangleCookiePrefix <AnyString>` option to specify the string added to the application cookie after manipulation. You can replace *<AnyString>* with a string of your choice. For example, adding the `NAGWSMangleCookiePrefix AGMANGLE` results in the `Set-Cookie: AGMANGLEa50b_DzkN=5a8G0` application level cookie set in the browser.

URL Attribute Filter

This feature lets you define filtering options for each proxy service. It helps in filtering out specified URLs from the ones audited as part of the URL Accessed audit event. These filtered out URLs will not be displayed in the Audit Server. This is helpful where auditing every URL is not required and may increase the load on the Audit Server. Unnecessary URLs for example, public images, public javascripts, css and favicons can be safely ignored from auditing.

 The option to set this feature is `NAGFilteroutUrlForAudit <regular expression>`. This option should be added to the Advanced options section of each service. The regular expression is standard perl based regular expressions. For more information about "Regular Expressions", see http://perldoc.perl.org/perlre.html

Each URL (path?querystring) is matched against this expression. If the match is successful, the URL will not be audited for URL access. For example, `NAGFilteroutUrlForAudit ".*.jpg"` and `NAGFilteroutUrlForAudit ".*.gif"`. If these options are added to a service, all the `*.jpg` and `*.gif` files accessed will not be audited under the 'URL Accessed' audit event.

NOTE: If you enable 'URL Accessed' audit events in the Access Gateway, it can overload the Audit subsystem if the requests sent to the Gateway per second is high. There maybe a delay in Web pages getting loaded. NetIQ recommends to use the `http common/extended logging` option for this purpose.

5.2 Configuring the Advanced Options for a Domain-Based and Path-Based Multi-Homing Proxy Service

The following procedure helps you configure the advanced options for domain-based and path-based multi-homing proxy service of an Access Gateway.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **Advanced Options**.

2 Configure the advanced option by removing the pound(#) symbol. To disable an option, add the # symbol in front of the option, save your changes, then update the Access Gateway.

6 Configuring the Content Settings

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.

IMPORTANT: For caching to work correctly, the Web servers must be configured to maintain a valid time. If possible, they should be configured to use an NTP server.

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 Webmasters in such a way that it cannot become stale unless a caching system ignores the Webmaster's settings. The Access Gateway honors all RFC 2616 directives that affect cache freshness such as Cache-Control, If-Modified-Since, and Expires.

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 Access Gateway follow RFC directives. In addition, the Access Gateway Service uses the "Apache Module mod_file_cache" (http://httpd.apache.org/docs/2.2/mod/mod_file_cache.html).

The following sections describe the features available to fine-tune this process for your network:

- Section 6.1, "Configuring Caching Options," on page 160
- Section 6.2, "Controlling Browser Caching," on page 160
- Section 6.3, "Configuring Custom Cache Control Headers," on page 161
- Section 6.4, "Configuring a Pin List," on page 164
- Section 6.5, "Configuring a Purge List," on page 166
- Section 6.6, "Purging Cached Content," on page 167
- Section 6.7, "Apache htcacheclean Tool," on page 168

6.1 Configuring Caching Options

The Cache Options allow you to control how the Access Gateway caches objects.

1 Click **Access Gateways** > **Edit** > **Cache Options**.

2 To disable caching of all Web server content, select **Disable Caching**.

When this option is selected, all other caching options are disabled.

3 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.

WARNING: Enter whole number values. Decimal values (2.5) are not supported and generate an XML validation error.

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 Webmasters if they 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 hours.

HTTP Default: Specifies the maximum time the proxy service serves HTTP data for which Webmasters have not specified a freshness or Time to Expire directive. The default is 2 hours.

4 To save your changes to browser cache, click **OK**.

5 To apply the changes, click the **Access Gateways** link, then click **Update** > **OK**.

6.2 Controlling Browser Caching

Webmasters 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 how the proxy service responds to these directives in the HTTP header.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **HTTP Options**.

☐ Allow Pages to Be Cached by the Browser
☑ Enable X-Forwarded-For

Server(s) must be updated before changes made on this panel will be used. See Configuration Panel for summary of changes.

[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**.

When this option is enabled, the no-cache and no-store headers are not injected into the HTTP header.

You need to select this option if you have a back-end application that updates the data in the Last-Modified or ETag HTTP headers. These changes are forwarded from the Web server to the browser only when this option is enabled.

You need to select this option if you want the Expires HTTP header forwarded from the Web server to 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.

3 For the Access Gateway Service, it is always enabled. For information about this option, see Section 3.10, "Configuring X-Forwarded-For Headers," on page 106.

4 Click **OK**.

5 To apply the changes, click the **Access Gateways** link, then click **Update** > **OK**.

6.3 Configuring Custom Cache Control Headers

In addition to fine-tuning cache freshness by using the HTTP timers, as explained in Section 6.1, "Configuring Caching Options," on page 160, 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 6.3.1, "Understanding How Custom Cache Control Headers Work," on page 161
- Section 6.3.2, "Enabling Custom Cache Control Headers," on page 162

6.3.1 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 such as transparent caches and client accelerators 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 external caches cannot cache the objects, but the proxy service can.

  This lets you off-load 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 6.3.2, "Enabling Custom Cache Control Headers," on page 162).

- 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 proxy service 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
```

6.3.2   Enabling Custom Cache Control Headers

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Proxy Service]** > **HTTP Options**.

☐ Allow Pages to Be Cached by the Browser

☑ Enable X-Forwarded-For

Server(s) must be updated before changes made on this panel will be used. See Configuration Panel for summary of changes.

[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 To save your changes to browser cache, click **OK**.

6 To apply the changes, click the **Access Gateways** link, then click **Update** > **OK**.

7 Modify the pages on the Web server that you want to the 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 6.3.1, "Understanding How Custom Cache Control Headers Work," on page 161).

8 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 7). 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 because of their relatively brief time-to-live value.

6.4   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 to cache
 * Which objects you never want cached

The pin list is global to the Access Gateway and affects all protected resources. The objects remain in cache until their normal cache limits are reached or they are bumped out by more recently requested objects.

To configure a pin list:

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Pin List**.

2 Select the **Enable Pin List** option to enable the use of pinned objects. If this option is not selected, the pinned objects in the pin list are not used.

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 "URL Mask" on page 164.

   Pin Type: Specifies how the URL is to be used to cache objects. Select from **Normal** and **Bypass**. For more information, see "Pin Type" on page 166.

5 To save the list item, click **OK**.

6 To save your changes to browser cache, click **OK**.

7 To apply the changes, click the **Access Gateways** link, then click **Update** > **OK**.

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 hostname 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	`http://www.foo.gov/documents/picture.gif`
	`http://www.foo.gov/documents/*`
	`http://www.foo.gov`
	`foo.gov/documents/*`
	`foo.gov/*`
	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 is processed first. The last item is the most general and is processed last.
path	`/documents/picture.gif`
	`/documents/pictures.gif/*`
	`/documents/*`
	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, `documents` is the root directory.
	The `/*` at the end of the path indicates that the entry is a directory. Its absence indicates that the entry is a file. In these examples, `picture.gif` is a file and `pictures.gif/*` and `documents/*` are directories.
	If you enter a path without the trailing *, the path matches only the directory. With the trailing *, the path matches everything in the directory and its subdirectories.
	These path entry examples are ordered by specificity. The objects in the `/documents/picture.gif`directory are processed before the objects in the `/documents` directory.
filename	`/picture.gif`
	`/widget.js`
	`/widget.jp*g`
	`/picture*group.gif`
	`/DailyTask`
	`/DailyTask*`
	Filenames are processed after paths. A leading forward slash must always be used when specifying a filename.
	You can add asterisks in the file names.
file extension	`/*.gif`
	`/*.js`
	`/*.htm`
	File extensions are processed last. They consist of a leading forward slash, an asterisk, a period, and a file extension.

NOTE: More than one wildcard is not allowed in a URL mask. For example, `/*picture.g*f` is not correct.

Also, the wildcard must be only in the last part of the path. For example:

Correct: `/picture/*.gif`

Incorrect: `/documents/*/picture.gif`

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 are in the pin list:

URL Mask	Pin Type	Pin Links
`http://www.foo.gov/documents/*`	normal	1
`www.foo*`	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 such as `.com`,`.net`, or `.org` whose DNS names begin with `www.foo` from being cached.

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/*` with a pin type of normal. This causes all files, including `.jpg` files, in the graphics directory on the `foo.gov` Web site to be cached as requested. Assuming there are no other URL masks in the pin list, all other JPG graphics are not cached because of the `/*.jpg` mask.

◆ **Bypass:** The Access Gateway does not cache the objects. In other words, you can use this option to prevent objects from being cached.

6.5 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 you specify the URL mask, do not specify a port. Ports are not stored in the cache file that is used to match the URLs that should be purged.

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
/*.pdf	Causes all PDF files to be purged from cache.
www.foo.gov/contracts/*	Causes all objects in the `contracts` 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. Make sure you don't configure a pin list that adds objects to the cache and a purge list that removes the same objects.

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Purge List**.

2 Click **New,** enter a URL pattern, then click **OK**.

3 (Optional) Repeat Step 2 to add additional URL patterns.

4 To save your changes to browser cache, click **OK**.

5 To apply the changes, click the **Access Gateways** link, then click **Update** > **OK**.

6.6 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 **Devices** > **Access Gateways**.

2 Select the name of the server, then click **Actions**.

3 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.

4 Click either **OK** or **Cancel**.

When you make certain configuration changes such as updating or changing certificates, changing the IP addresses of Web servers, or modifying the rewriter configuration, you are prompted to purge the cache. The cached objects must be updated for users to see the effects of such configuration

changes. If your Access Gateways are in a cluster, you need to manage the purge process so your site remains accessible to your users. You should apply the configuration changes to one member of a cluster. When its status returns to healthy and current, issue the command to purge its cache. Then apply the changes to the next cluster member.

IMPORTANT: Do not issue a purge cache command when an Access Gateway has a pending configuration change. Wait until the configuration change completes.

6.7 Apache htcacheclean Tool

If you have caching issues with inodes, disk space, and cache corruption in the Access Gateway, use Apache htcacheclean tool which is used to keep the size of mod_disk_cache's storage within a certain limit. This tool can run either manually or in daemon mode. When running in daemon mode, it sleeps in the background and checks the cache directories at regular intervals for cached content to be removed. For more information on Apache htcacheclean tool, see (https://httpd.apache.org/docs/2.2/programs/htcacheclean.html).

7 Protecting Multiple Resources

This section describes how to create multiple resources for Access Gateways.

7.1 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 7-1 *Adding Redundant Web Servers*
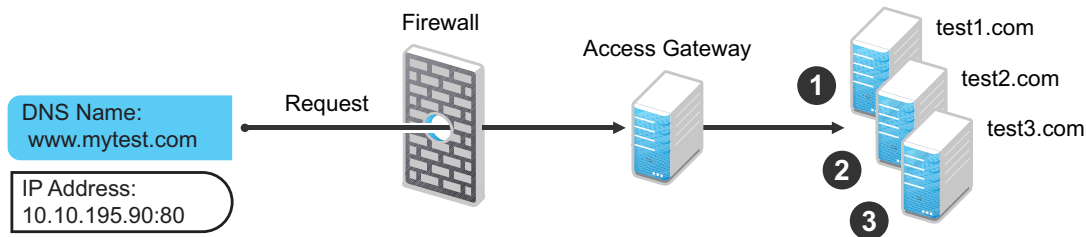
For this configuration, you use a single reverse proxy and proxy service. To add multiple Web servers to a host:

1. In the Administration Console, click **Devices** > **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 or the fully qualified DNS name of another Web server for the "virtual" group, then click **OK**.

4. Repeat Step2 and Step3 to add additional Web servers to the group.

5. Click **OK**.

6. To apply the changes, click the **Access Gateways** link, then click **Update** > **OK**.

The Access Gateway uses the round robin algorithm by default. You can configure it to use the simple failover algorithm.

Simple failover sends all the traffic to the first Web server as long as it is available. Traffic is sent to another Web server in the list only when the first Web server is no longer available. To configure this option, see Section 2.7.2, "Configuring TCP Connect Options for Web Servers," on page 78.

Connection persistence is enabled by default. This allows the Access Gateway to send multiple HTTP requests to the Web server to be serviced before the connection is closed. To configure this option, see Section 2.7.2, "Configuring TCP Connect Options for Web Servers," on page 78.

Session stickiness option is used if multiple Web Servers are configured for a service. Selecting this option makes the proxy server to use the same Web server for all fills during a session. This option is enabled by default. For more information about persistent connections, see Section 2.7.3, "Configuring Connection and Session Persistence," on page 78.

7.2 Using Multi-Homing to Access Multiple Resources

You can configure an Access Gateway to use one public IP address to protect multiple types of Web resources. This is one of the major benefits of the Access Gateway, because it conserves valuable resources such as IP addresses. This feature also makes an 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 Access Gateway. Subsequent proxy services can be configured to use one of the following methods:

- Section 7.2.1, "Domain-Based Multi-Homing," on page 170
- Section 7.2.2, "Path-Based Multi-Homing," on page 172
- Section 7.2.3, "Virtual Multi-Homing," on page 175

This section describes these multi-homing methods, then explains the following:

- Section 7.2.4, "Creating a Second Proxy Service," on page 176
- Section 7.2.5, "Configuring a Path-Based Multi-Homing Proxy Service," on page 177

7.2.1 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 prefix hostnames to a cookie domain name. For a test resource, you can prefix 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 7-2 illustrates this type of configuration for three proxy services.

Figure 7-2 *Using a Base Domain Name with Host Names*
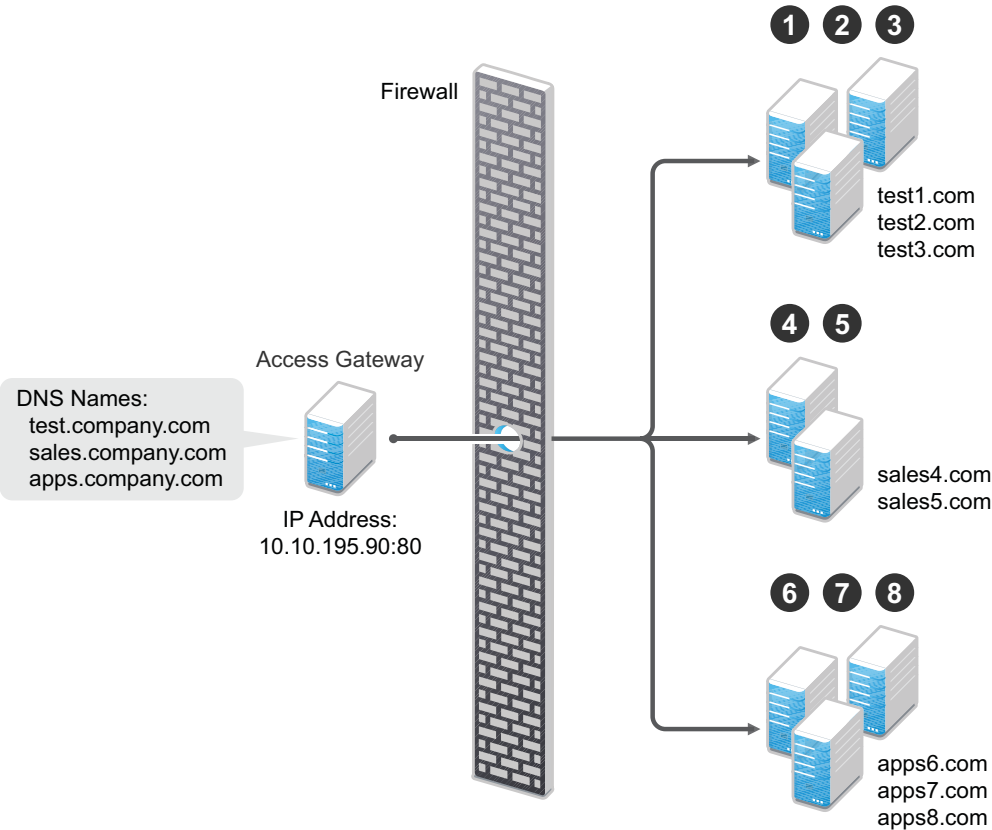
Domain-based multi-homing has the following characteristics:

- If you are using SSL, the back-end servers can all listen on the same SSL port (the 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 hostname 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 names that use `company.com` as a common domain name, lists the IP address that these DNS names resolve to, and the Web servers they 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	10.10.15.10
sales.company.com	10.10.195.90:80	sales.internal.com	10.10.15.20
apps.company.com	10.10.195.90:80	apps.internal.com	10.10.15.30

- 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.

◆ Create three proxy services for these published DNS names.

To create a domain-based multi-homing proxy service, see Section 7.2.4, "Creating a Second Proxy Service," on page 176, and select domain-based for the multi-homing type.

7.2.2 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 7-3 illustrates this type of configuration.

Figure 7-3 *Using a Domain Name with Path Elements*
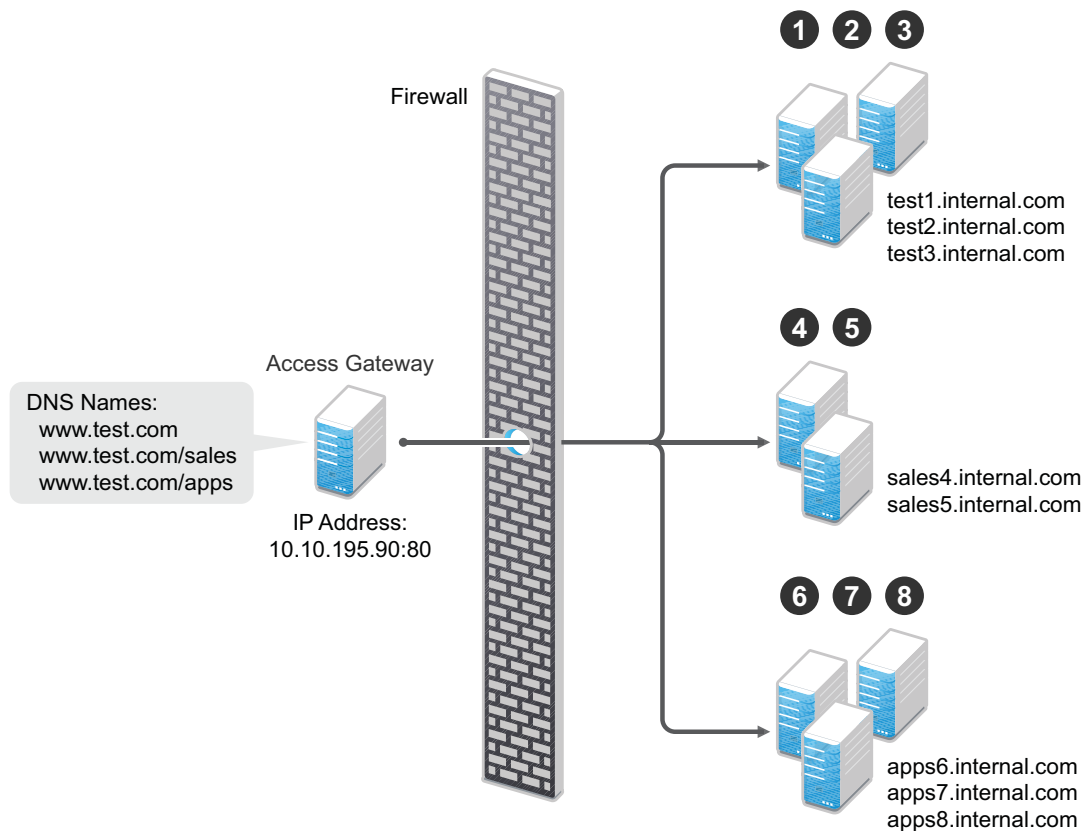
Path-based multi-homing has the following characteristics:

◆ It is considered to be more secure than domain-based 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 following sections explain how to configure path-based proxy services and your network so that the Access Gateway can find the correct protected resources:

- "Configuring the Remove the Path on Fill Option" on page 173
- "Configuring the Host Header Option" on page 173
- "Preparing for Path-Based Multi-Homing" on page 174

Configuring the Remove the Path on Fill Option

If the path that is part of the published DNS name (`/sales` or `/apps`) is used to identify a resource but is not part of directory configuration on the Web server, the path needs to be removed from the URL before the request is sent to 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. The Access Gateway does not allow you to set up multiple paths to this type of Web server, so all pages must have the same authentication requirements.

If the path in the published DNS name is a path on the Web server, the path needs to 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

Because the path component specifies 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

Such a configuration also allows you to set up different authentication and authorization requirements for each path.

Configuring the Host Header Option

When you create path-based proxy services and also enable the **Remove Path on Fill** option, you need to know what types of links exist on the Web servers. For example, you need to know if the sales Web servers in Figure 7-3 on page 172 have links to the app Web servers or to the test Web servers. If they don't, you can set the **Host Header** option to either **Forward Received Host Name** or to **Web Server Host Name**. However, if they do contain links to each other, you need to set the **Host Header** option to **Web Server Host Name** and specify a DNS name for the Web server in the **Web**

Server Host Name option. The Access Gateway needs a method to distinguish between the Web servers other than the path, because after the path is removed, all the Web servers in Figure 7-3 on page 172 have the same name: `www.test.com`.

If you select to use the **Forward Received Host Name** option for a path-based service, you might also need to add entries to the **Additional DNS Name List** for the rewriter. For more information, see "Determining Whether You Need to Specify Additional DNS Names" on page 60.

Preparing for 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	10.10.15.10
test.com/sales	10.10.195.90:80	sales.internal.com	10.10.15.20
test.com/apps	10.10.195.90:80	apps.internal.com	10.10.15.30

 ◆ 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. If they have links to each other, set up DNS names for the Web servers.

 ◆ Create one proxy service that uses `test.com` as its published DNS name and two path-based proxy services.

 To create a path-based multi-homing proxy service, see Section 7.2.4, "Creating a Second Proxy Service," on page 176, and select path-based for the multi-homing type.

7.2.3 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 7-4 illustrates this type of configuration.
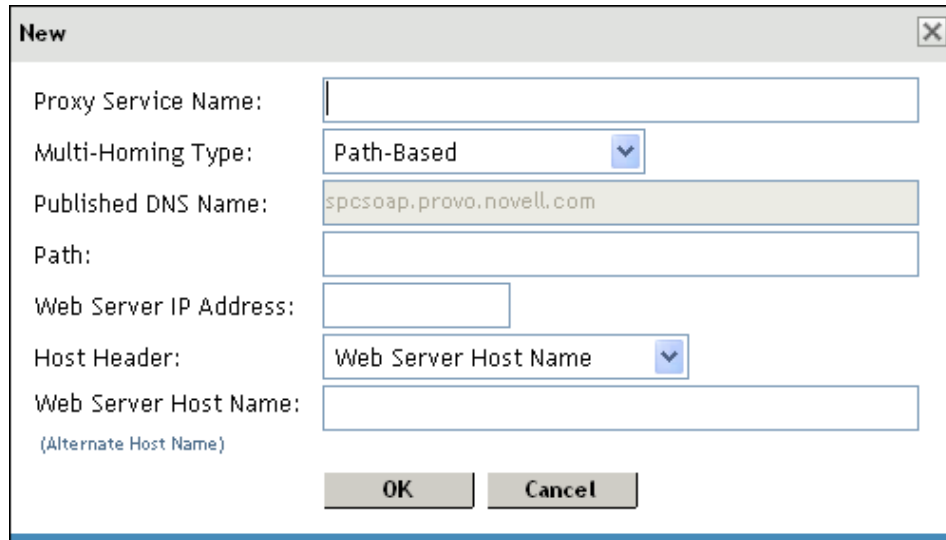
Figure 7-4 *Using Multiple DNS Names*
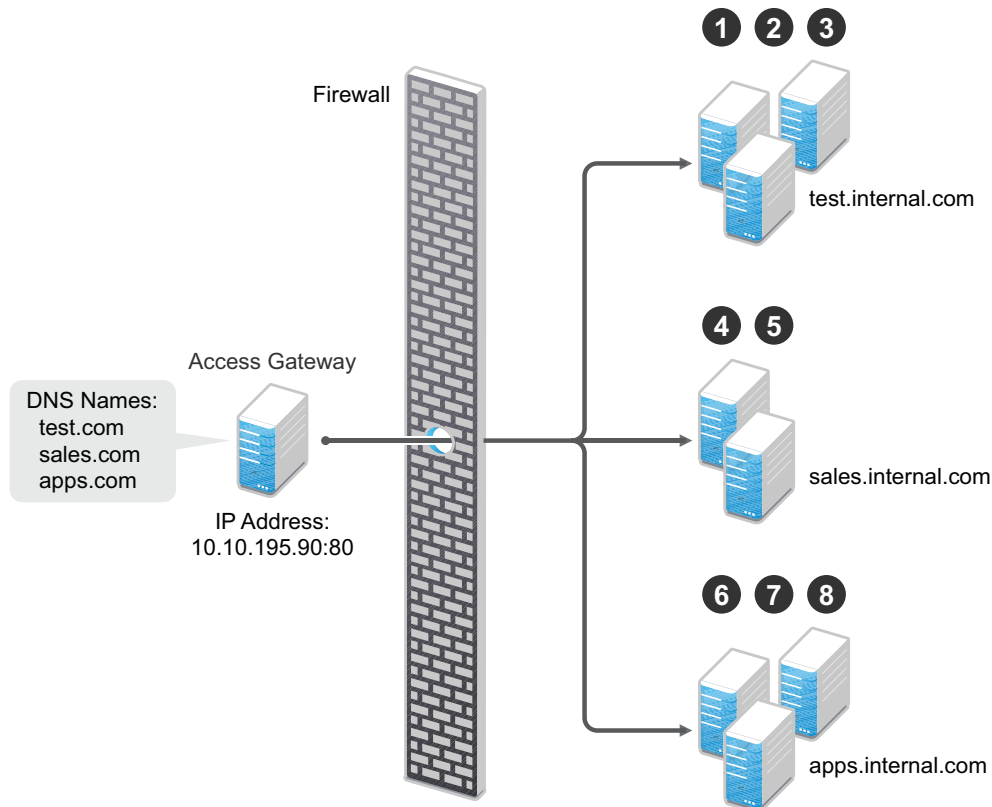
Virtual multi-homing cannot be used with SSL. You should use this configuration with resources that need to be protected, but the information exchanged should be public information that does not need to be secure. For example, you could use this configuration to protect your Web servers that contain 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 configuration of your DNS server. After you have configured your DNS server to allow multiple names to resolve to the same IP address, you are ready to configure the Access Gateway.

To create a virtual multi-homing proxy service, see Section 7.2.4, "Creating a Second Proxy Service," on page 176, and select **Virtual** for the multi-homing type.

7.2.4 Creating a Second Proxy Service

1 In the Administration Console, click **Devices** > **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 hostname (`www.newsite.test.com`). For more information, see Section 7.2.1, "Domain-Based Multi-Homing," on page 170.

 ♦ **Path-Based:** Uses the published DNS name (`www.test.com`) with a path (`www.test.com/path`). For more information, see Section 7.2.2, "Path-Based Multi-Homing," on page 172.

 ♦ **Virtual:** Uses a unique DNS name (`www.newsite.newcompany.com`). Virtual multi-homing cannot be used with SSL. For more information, see Section 7.2.3, "Virtual Multi-Homing," on page 175. If you need a unique DNS name and SSL, you need to create a reverse proxy rather than a proxy service. For information about creating a second reverse proxy, see Section 7.3, "Managing Multiple Reverse Proxies," on page 179.

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).

For a path-based multi-homing service, it is usually best to select the **Web Server Host Name** option. For more information, see "Configuring the Host Header Option" on page 173.

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.

For iChain administrators, the **Web Server Host Name** is the alternate hostname 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 2.2.2, "Configuring a Proxy Service," on page 28.

  ◆ To configure a path-based proxy service, see Section 7.2.5, "Configuring a Path-Based Multi-Homing Proxy Service," on page 177.

7.2.5 Configuring a Path-Based Multi-Homing Proxy Service

To configure a path-based proxy service:

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **[Name of Reverse Proxy]** > **[Name of Path-Based Multi-Homing Proxy Service]**.

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) Provide a description of 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 6.3, "Configuring Custom Cache Control Headers," on page 161 and Section 6.2, "Controlling Browser Caching," on page 160.

Advanced Options: (Access Gateway Service) See Section 5.2, "Configuring the Advanced Options for a Domain-Based and Path-Based Multi-Homing Proxy Service," on page 157.

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.

The Access Gateway automatically selects the protected resource that is used with the specified path. It selects the current protected resource whose URL path most closely matches the specified path.

- If you have a protected resource with a URL path of /*, the Access Gateway selects that resource unless you have configured a protected resource that has a URL path that more closely matches the path specified on this page.

- If you add a protected resource at a future time and its URL path more closely matches the path specified on this page, the Access Gateway automatically reconfigures to use this new protected resource.

- If you disable a protected resource that the Access Gateway has assigned to a path-based service, the Access Gateway automatically reconfigures and selects the next protected resource that most closely matches the path specified on this page.

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 to ensure that they meet the requirements for your path-based service.

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 the protected resource does not meet your needs, you must create a protected resource for the path-based proxy service. Continue with Step 4d.

4d Click **OK**, select the name of the parent proxy service, then click **Protected Resources**.

4e In the **Protected Resource List**, click **New**, specify a name, then click **OK**.

4f Select an Authentication Procedure.

4g In the **URL Path List**, specify the path you used when creating the path-based proxy service. For example, if your path was /apps, specify /apps/* or /apps in the URL Path List.

> **IMPORTANT:** If you create multiple protected resources that exactly match the path-based multi-homing service, there is no guarantee that a specific protected resource will be used. For example, if you create protected resources for both of the paths specified above (`/apps` and `/apps/*`) and you have a path-based service with a path of `/apps`, either of these protected resources could be assigned to this path-based service in the Administration Console or used when access is requested.

    4h Make sure the protected resource you created is enabled. If the resource is disabled, it does not appear in the Path List for the path-based proxy service.

    4i (Optional) Enable the policies the path-based proxy service requires. Click **Authorization**, **Identity Injection**, or **Form Fill** and enable the appropriate policies.

    4j Click **OK**.

5 Click **OK**.

6 To apply the changes, click the **Access Gateways** link, then click **Update** > **OK**.

7.3 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.

You need to use system utilities to configure network interface cards and new IP addresses. After they are configured, you can use the **New IP** option to make them available for Gateway Service configuration. See Section 3.8.5, "Adding a New IP Address to the Access Gateway," on page 103.

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 7.3.1, "Managing Entries in the Reverse Proxy List," on page 180

7.3.1 Managing Entries in the Reverse Proxy List

1 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Reverse Proxy / Authentication**.

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 2.2, "Managing Reverse Proxies and Authentication," on page 23.

  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 next to a specific reverse proxy, then click **Delete**. To delete all reverse proxies, select the check box next to the **Name** column, then click **Delete**.

- ◆ **Enable:** To enable a reverse proxy, select the check box next to a specific reverse proxy, then click **Enable**. To enable all reverse proxies, select the check box next to the **Name** column, then click **Enable**.

◆ **Disable:** To disable a reverse proxy, select the check box next to a specific reverse proxy, then click **Disable**. To enable all reverse proxies, select the check box next to the **Name** column, then click **Disable**.

3 Click **OK**.

4 To apply the changes, click the **Access Gateways** link, then click **Update** > **OK**.

7.4 Managing a Cluster of Access Gateways

Most of the configuration tasks are the same for a single Access Gateway and a cluster of Access Gateways. This section describes the tasks that are specific to managing the servers in a cluster:

For information about monitoring the health or statistics of a cluster, see Section 4.8.3, "Viewing the Health of an Access Gateway Cluster," on page 141 and Section 4.4, "Viewing Cluster Statistics," on page 128.

7.4.1 Managing Cluster Details

Use the Cluster Details page to perform general maintenance actions on the selected cluster and to display server information about the selected cluster.

1 In the Administration Console, click **Devices** > **Access Gateways** > **[Cluster Name]**.

2 View the following fields:

Name: Specifies the name of the cluster.

Description: Specifies the purpose of the cluster. This is optional, but useful if your network has multiple Access Gateway clusters. If the field is empty, click **Edit** to add a description.

Primary Server: Indicates which server in the cluster has been assigned to be the primary server.

3 To modify the information, click **Edit**. For more information, see Section 7.4.2, "Editing Cluster Details," on page 181.

4 To select a different Access Gateway to be the primary cluster member, click **Edit**.

5 To modify details about a cluster member, click the server name in the **Cluster member** list.

6 Click **Close**.

7.4.2 Editing Cluster Details

Use the Cluster Detail Edit to change the name of the cluster and assign a different server to be the primary cluster member.

1 In the Administration Console, click **Devices** > **Access Gateways** > **[Cluster Name]** > **Edit**.

2 Modify the following fields:

Name: Specify a name for the cluster.

Description: Specify the purpose of the cluster. This is optional, but useful if your network has multiple Access Gateway clusters.

Primary Server: Indicates which server in the cluster has been assigned to be the primary server. To change this assignment, select the server from the drop-down list.

3 Click **OK**.

7.4.3 Applying Changes to Cluster Members

When you are configuring services of the Access Gateway, the **OK** button saves the change to browser cache except on the Configuration page. The Configuration page (**Devices** > **Access Gateways** > **Edit**) provides a summary of the changes you have made. The **Cancel Change** column allows you to cancel changes to individual services. When you click **OK**, the changes are saved to the configuration datastore, and you no longer have the option to cancel changes to individual services.

If you don't save the changes to the configuration datastore and your session times out or you log out, any configuration changes that are saved to browser cache are flushed. These changes cannot be applied to other members of the cluster because they are no longer available. To prevent this from happening, save the changes to the configuration datastore.

It is especially important to save the changes to the configuration datastore when you select to update individual members one at a time rather than update all members of the cluster at the same time. Updating members one at a time has the following benefits:

 ◆ When you update all servers at the same time, the site goes down until one server has finished updating its configuration. If you update the cluster members one at a time, only the member that is updating its configuration becomes unavailable.

 ◆ If you update the servers one at time, you can verify that the changes are behaving as expected. After testing the configuration on one server, you can then apply the saved changes to the other servers in the cluster. If you decide that the configuration changes are not behaving as expected, you can revert to the previously applied configuration. See "Reverting to a Previous Configuration" on page 182

Some configuration changes cannot be applied to individual cluster members. For a list of these changes, see "Modifications Requiring an Update All" on page 182.

Reverting to a Previous Configuration

If you have updated only one server in the cluster, you can use the following procedure to revert back to the previous configuration.

1 Remove the server that you have applied the configuration changes from the cluster.

2 Access the Configuration page for the cluster, then click **Revert**.

The servers in the cluster revert to the last applied configuration.

3 Add the removed server to the cluster.

The server is configured to use the same configuration as the other cluster members.

Modifications Requiring an Update All

When you make the following configuration changes, the **Update All** option is the only option available and your site is unavailable while the update occurs:

 ◆ If you change the Identity Server configuration that is used for authentication (**Access Gateways** > **Edit** > **Reverse Proxy/Authentication,** then select a different value for the **Identity Server Cluster** option).

- If you select a different reverse proxy to use for authentication (**Access Gateways** > **Edit** > **Reverse Proxy/Authentication**, then select a different value for the **Reverse Proxy** option).

- If you modify the protocol or port of the authenticating reverse proxy (**Access Gateways** > **Edit** > **Reverse Proxy/Authentication** > **[Name of Reverse Proxy]**, then change the SSL options or the port options).

- If you modify the published DNS name of the authentication proxy service (**Access Gateways** > **Edit** > **Reverse Proxy/Authentication** > **[Name of Reverse Proxy]** > **[Name of First Proxy Service]**, then modify the **Published DNS Name** option).

8 Troubleshooting the Access Gateway

8.1 Useful Troubleshooting Files

The Access Gateway Service consists of two main modules, a Gateway Manager module that runs on top of Tomcat and a Proxy Service module that runs on top of Apache. Figure 8-1 illustrates these modules and the communication paths that the Access Gateway Service has with other devices.

Figure 8-1 *Access Gateway Service Modules*
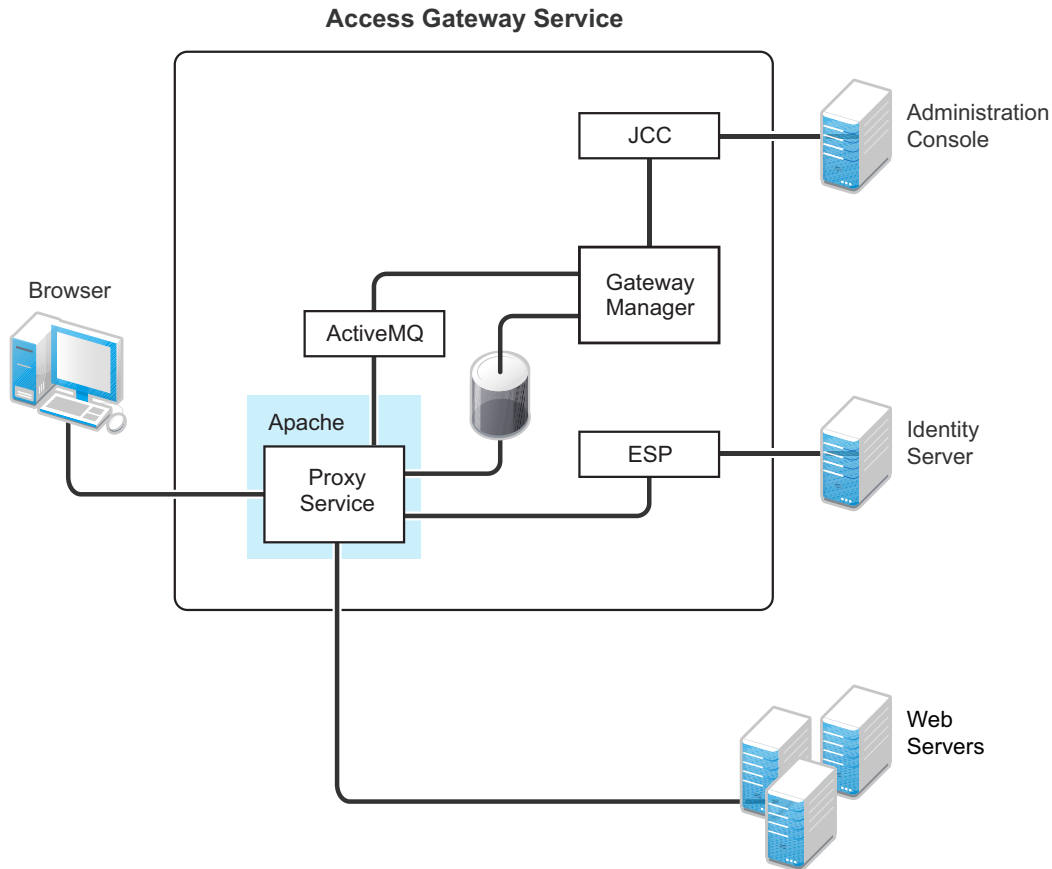
Proxy Service: This component runs as an instance of Apache and is responsible for controlling access to the configured protected resources on the Web servers. Low-level errors are reported in the Apache logs. Some higher-level errors are also reported to the files in the `amlogging/logs` directory.

ESP: The Embedded Service Provider is responsible for handling all communications with the Identity Server and is responsible for the communication that verifies the authentication credentials of users. Log entries for this communication process, including errors, are logged in the `catalina.out` file and the `stdout.log` file.

ActiveMQ: This module is used for real-time communication between the Administration Console and the Proxy Service. Errors generated from the Gateway Manager to the ActiveMQ module are logged to the Tomcat logs. Errors generated from the Proxy Service to the ActiveMQ module are logged to the Apache error logs.

JCC: The Java Communication Controller is the interface to the Administration Console. It handles health, statistics, configuration updates, and purge cache requests from the Administration Console. It is also responsible for certificate management. Errors generated between the JCC module and the Gateway Manager are logged to the `ags_error.log` file. Errors generated between the Administration Console and the JCC module are logged to the `jcc-0.log.`*x* file

Gateway Manager: This module is responsible for handling communication from JCC to the Proxy Service. It also writes the configuration commands to the Apache configuration files and the Proxy Service configuration file on disk. Errors generated while performing these tasks are logged to the `ags_error.log` file.

For more information about these various log files, see the following:

8.1.1 Apache Logging Options for the Gateway Service

The Proxy Service module of the Access Gateway Service is built on top of Apache as an Apache application. This module handles the browser requests for access to resources and is responsible for sending authorized requests to the Web servers. Entries for these events are logged to the Apache log files.

```
/var/log/novell-apache2/
```

For more information, see sections "Ignoring Some Standard Messages" on page 187 and Section 4.1, "Managing Access Gateway Logs," on page 109.

Ignoring Some Standard Messages

Apache cannot detect the proper use of domain-based multi-homing with wildcard certificates, which allows multiple proxy services to share the same SSL port. If you create reverse proxy services that are configured for domain-based multi-homing with SSL, Apache considers this a possible port conflict and logs it as a warning in the `error.log` file.

The error messages look similar to the following:

```
[<time and date stamp>] [warn] Init: SSL server IP/port conflict:
dbmhnsnetid.dsm.cit.novell.com:443 (C:/Program
Files/Novell/apache/conf/vhosts.d/dbmhNS-NetID.conf:18) vs.
magwin1430external.dsm.cit.novell.com:443 (C:/Program
Files/Novell/apache/conf/vhosts.d/magMaster.conf:18)

[<time and date stamp>] [warn] Init: SSL server IP/port conflict:
magdbmheguide.dsm.cit.novell.com:443 (C:/Program
Files/Novell/apache/conf/vhosts.d/dbmhMagEguide.conf:18) vs.
magwin1430external.dsm.cit.novell.com:443 (C:/Program
Files/Novell/apache/conf/vhosts.d/magMaster.conf:18)
```

You can ignore these errors because the Access Gateway Service knows how to handle the traffic and send the packets to the correct proxy service.

For more information about Apache log files, see "Log Files" (http://httpd.apache.org/docs/2.2/logs.html).

Modifying the Logging Level for the Apache Logs

If the Apache error log file does not contain enough information, you can modify the log level and the types of messages written to the file.

WARNING: If you set the log level to debug, the size of the file can grow quickly, consume all available disk space, and crash the system. If you change the log level, you need to carefully monitor available disk space and the size of the error log file.

To modify what is written to the Apache error log file:

1 Change to the Apache configuration directory.

```
/etc/opt/novell/apache2/conf
```

2 Open the `httpd.conf` file.

3 Find the `LogLevel` directive and set it to one of the following:

debug, info, notice, warn, error, crit, alert, emerg

4 Save the file.

5 Restart Apache:

`/etc/init.d/novell-apache2 restart` OR `rcnovell-apache2 restart`

6 (Optional) If you set the level to debug and the log file still does not supply enough information, see Section 8.4, "Enabling Debug Mode and Core Dumps," on page 191.

8.1.2 The Access Gateway Service Log Files

Filename	Description
`error_log`	Contains the Apache messages generated for get and response traffic from the browsers.
	For information about setting the logging level for this file, see "Modifying the Logging Level for the Apache Logs" on page 187.
	To add X-Mag header or Form Fill messages, enable the advanced options. See Section 5.1, "Configuring the Global Advanced Options," on page 145.
	The log file is located in the `/var/opt/novell/nam/logs/mag/apache2` directory.
`rcnovell-apache2.out`	Contains the Apache messages generated at start up.
	For information about setting the logging level for this file, see "Modifying the Logging Level for the Apache Logs" on page 187.
	The log file is located in the `/var/opt/novell/nam/logs/mag/apache2` directory.
`ags_error.log`	Contains the messages generated for configuration, device imports, health, and statistics. It also contains entries for the policy evaluation processes done by the Gateway Service Manager module. When you apply changes on the Access Gateway, all enabled global advanced options or service level advanced options will be dumped to this log file. For additional information, see TID 7012563.
	The log file is located in the `/var/opt/novell/nam/logs/mag/amlogging` directory.
`jcc-0.log.x`	Contains the messages generated between the Administration Console and the JCC module.
	The log file is located in the `/var/opt/novell/nam/logs/jcc` directory.
`catalina.out`	Contains the log messages generated by the Embedded Service Provider. Logging to this file occurs only if you have selected the **Echo to Console** option from the **Identity Servers** > **Servers** > **Edit** > **Logging** page.
	Check this file for entries that trace the evaluation of Authorization, Identity Injection, and Form Fill policies.
	The file is located in the `/var/opt/novell/nam/logs/mag/tomcat` directory.

Filename	Description
<log_profile>	If you enable log profiles, you specify the name for the log file.
	You can select to log entries for the messages sent between the Gateway Service and the Embedded Service Provider (ESP), between the Gateway Service and the Web servers, and between the Gateway Service and the Administration Console.
	The log profile files are located in the `/var/opt/novell/amlogging/logs` directory.
<reverse_proxy>	If you enable logging on one or more reverse proxies, you specify the name for log file.
	To enable this type of logging, see Section 4.2, "Configuring Logging for a Proxy Service," on page 110.
	The reverse proxy log files are located in the `/var/log/novell-apache2` directory.

8.2 Verifying That All Services Are Running

1. Log in to the server as the `root` user.

2. Verify that the ActiveMQ service is running by entering the following command:

   ```
   ps -fea | grep novell-activemq | grep -v "grep"
   ```

   Lines similar to the following are displayed:

   ```
   107      18941      1  0 Apr01 ?         03:11:11 /opt/novell/java/bin/java -
   Xmx512M -Dorg.apache.activemq.UseDedicatedTaskRunner=true -
   Dcom.sun.management.jmxremote -Djavax.net.ssl.keyStorePassword=xxxxx -
   Djavax.net.ssl.trustStorePassword=xxxxx -Djavax.net.ssl.keyStore=/opt/novell/
   activemq/conf/broker.ks -Djavax.net.ssl.trustStore=/opt/novell/activemq/conf/
   broker.ts -Dactivemq.classpath=/opt/novell/activemq/conf; -Dactivemq.home=/
   opt/novell/activemq -Dactivemq.base=/opt/novell/activemq -jar /opt/novell/
   activemq/bin/run.jar start
   ```

3. Verify that one or more Apache proxy services are running by entering the following command:

   ```
   ps -ef | grep httpd
   ```

   Lines similar to the following are displayed:

   ```
   root    2983 30290  0 12:53 pts/0   00:00:00 egrep httpd
   root    3163     1  0 May12 ?      00:00:29 /opt/novell/apache2/sbin/httpd
   wwwrun  3165  3163  0 May12 ?      00:01:00 /opt/novell/apache2/sbin/httpd
   wwwrun  3184  3163  0 May12 ?      00:00:01 /opt/novell/apache2/sbin/httpd
   wwwrun  3188  3163  0 May12 ?      00:00:01 /opt/novell/apache2/sbin/httpd
   ```

4. Verify that the user session cache service is running by entering the following command:

   ```
   ps -ef | grep novell-agscd
   ```

   Lines similar to the following are displayed:

   ```
   root   3259 30290  0 12:56 pts/0    00:00:00 egrep novell-agscd
   108    5525     1  0 May11 ?    00:00:00 /opt/novell/ag/bin/novell-agscd -d
   108    5526  5525  0 May11 ?    00:00:09 /opt/novell/ag/bin/novell-agscd -d
   ```

5. Verify that the Tomcat service is running by entering the following command:

```
ps -ef | grep catalina.base
```

Lines similar to the following are displayed:

```
ps -eaf | grep catalina.base
novlwww  28764    1  0 Jul05 pts/0   00:02:05 /opt/novell/java/bin/java -Dnop
-server -Xmx2048m -Xms512m -Xss128k -Djava.library.path=/usr/lib64:/opt/
novell/eDirectory/lib64:/opt/novell/lib64 -
Dcom.novell.nam.common.util.DeploymentMode=MAGAppliance -
Dsun.net.client.defaultConnectTimeout=29000 -
Dsun.net.client.defaultReadTimeout=28000 -Dnids.freemem.threshold=10 -
Djavax.net.ssl.sessionCacheSize=10000 -
Dsun.net.http.allowRestrictedHeaders=true -Djava.awt.headless=true -
Djava.util.logging.manager=org.apache.juli.ClassLoaderLogManager -
Djava.endorsed.dirs=/var/opt/novell/tomcat7/endorsed -classpath /lib/
tools.jar:/var/opt/novell/tomcat7/bin/bootstrap.jar:/var/opt/novell/tomcat7/
bin/tomcat-juli.jar -Dcatalina.base=/opt/novell/nam/mag -Dcatalina.home=/var/
opt/novell/tomcat7 -Djava.io.tmpdir=/opt/novell/nam/mag/temp
org.apache.catalina.startup.Bootstrap -config /opt/novell/nam/mag/conf/
server.xml start
```

6 Verify that the JCC service is running by entering the following command:

```
ps -ef | grep /opt/novell/devman/jcc/conf/run.sh
```

Lines similar to the following are displayed:

```
root  3777 30290  0 13:03 pts/0 00:00:00 egrep /opt/novell/devman/jcc/
                                    conf/run.sh
root  5506     1  0 May11 ?       00:00:00 /bin/bash /opt/novell/devman/jcc/
                                    conf/run.sh
```

When you are familiar with the services, you can use the following command to display information about all the services:

```
ps -ef | egrep "novell-activemq|novell-agscd|/opt/novell/devman/jcc/conf/
run.sh|catalina.base|httpd"
```

7 If one or more services are not running, use the following commands to start the services:

`/etc/init.d/novell-jcc start` OR `rcnovell-jcc start`

`/etc/init.d/novell-apache2 start` OR `rcnovell-apache2 start`

`/etc/init.d/novell-agcsd start`

`/etc/init.d/novell-activemq start` OR `rcnovell-activemq start`

`/etc/init.d/novell-mag start` OR `rcnovell-mag start`

8 If a service does not start, view the log files to determine the cause. See the following:

- Section 8.7, "Solving Apache Restart Issues," on page 194
- Section 8.1.2, "The Access Gateway Service Log Files," on page 188

8.3 Troubleshooting SSL Connection Issues

SSL handshakes fail when there is a discrepancy between the cipher suites and cipher strengths used by the clients and the servers. If you enable SSL connections between the Access Gateway and the browser or between the Access Gateway and the Web servers, you need to make sure that both

sides are configured to support the same cipher suites and cipher strengths. This is especially important if you enable the options to enforce 128-bit encryption (see Section 2.7.1, "Configuring TCP Listen Options for Clients," on page 77).

The Access Gateway Service relies upon Apache to perform the SSL handshake, and Apache does not log the cause of SSL handshake failures, even when the log level is set to debug. To determine whether cipher strengths are the source of your problem, disable the options to enforce 128-bit encryption (see Section 2.7.1, "Configuring TCP Listen Options for Clients," on page 77). If users are then able to authenticate, verify the cipher strengths, which are configured for the browsers and for the Web servers, are compatible with the Access Gateway.

8.4 Enabling Debug Mode and Core Dumps

If the log files are not generating enough information to identify the cause of a problem, you can run the Access Gateway Service in debug mode. You should not be running in debug mode except when you are trying to isolate a problem because of the following side effects:

- Debug mode causes the size of the log files to grow quickly. They can grow large enough to consume all available disk space and crash the system. When running in debug mode, you need to carefully monitor available disk space and the size of the log files.
- Debug mode opens additional ports. Anyone who has local access to the Access Gateway machine can see the information displayed in the following local URLs:

  ```
  http://127.0.0.1:8181/server-status
  http://127.0.0.1:8181/server-info
  ```

- Debug mode causes load and response times to slow.

Debug mode enables core dumps, X-Mag headers in LAN traces, and increases log levels by enabling advanced option in the Access Gateway configuration. For example LogLevel debug. This will set apache log level to debug in the `error_log` file.

You can generate core dumps in the following two ways:

1 Start `/etc/init.d/novell-apache2` in debug mode. When there is a crash, core file will be created as `/var/cache/novell-apache2/core`.

2 Without starting `novell-apache2` in debug mode, perform the following:

   2a Set `ulimit -c unlimited` in `/etc/init.d/novell-apache2` startup script.

   2b You can create the core directory under `/tmp`. Choose the file path based on the availability of disk space. Give the following command to create a directory in the Access Gateway component:

   ```
   # mkdir -p /tmp/apache2-gdb-dump
   ```

   2c Set permission as follows:

   ```
   # chown novlwww:www /tmp/apache2-gdb-dump
   # chmod 0777 /tmp/apache2-gdb-dump
   ```

   2d Add the following advanced option in the Access Gateway configuration as follows:

   ```
   CoreDumpDirectory /tmp/apache2-gdb-dump
   ```

   2e Apply changes to the Access Gateway.

Whenever there is a crash, core file will be created as `/tmp/apache2-gdb-dump/core`.

For some crashes, the `/tmp/debug000.log` file is created. For more information about the log, see TID 7011804.

This section describes the following tasks:

8.4.1 Starting Apache in Debug Mode

Use the following commands to start debug mode:

`/etc/init.d/novell-apache2 stop` OR `rcnovell-apache2 stop`

`/etc/init.d/novell-apache2 start debug` OR `rcnovell-apache2 start debug`

8.4.2 Examining the Debug Information

1 Examine the Apache error log file or copy it so you can send it to NetIQ Technical Support:

   `/var/log/novell-apache2`

2 View the information at the local URLs or copy the pages to send to NetIQ Support:

   - http://127.0.0.1:8181/server-status

      This page displays debug information about caching, SSL, workers, and proxy information.

   - http://127.0.0.1:8181/server-info

      This page displays module and configuration information.

3 If a crash occurred, examine the core dump file or copy it so you can send it to NetIQ Technical Support.

   `/var/cache/novell-apache2`

8.4.3 Disabling Debug Mode

Use the following commands to disable debug mode:

`/etc/init.d/novell-apache2 stop` OR `rcnovell-apache2 stop`

`/etc/init.d/novell-apache2 start nodebug` OR `rcnovell-apache2 start nodebug`

8.5 Useful Troubleshooting Tools for the Access Gateway Service

Table 8-1 describes some of the tools available in the Administration Console for solving potential problems:

Table 8-1 *Useful Tools*

Tool	Description
Re-push Current Configuration	If you have an Access Gateway that does not seem to be using the current configuration, you can use the Administration Console to push the current configuration to the Access Gateway. Click **Auditing** > **Troubleshooting**. In the **Current Access Gateway Configuration** section, select an Access Gateway, then click **Re-push Current Configuration**.
Health icon	In the Administration Console, click the **Health** icon to view details about the health of the Access Gateway. For more information, see Section 4.8.2, "Monitoring the Health of an Access Gateway," on page 139.
`curl`	Use this command to view identity provider metadata from the Access Gateway. See "Testing Whether the Provider Can Access the Metadata" in the *NetIQ Access Manager Appliance 4.0 Identity Server Guide*.
`netstat -a`	Use this command to view statistics about the listeners on the Access Gateway.
`tail -f`	Use this command to view real time activity in key log files. For information about useful files to tail, see Section 8.1.2, "The Access Gateway Service Log Files," on page 188.
`top`	Use this command to check resources available on the system.
`tcpdump`	Use this command to capture data on standard and loopback interfaces and to view SSL data with imported keys.

8.6 A Few Performance Tips

- **Caching:** Use a high performance disk system for the cache directory, such as tempfs on Linux.

  You can improve the speed of adding files to cache and retrieving them from cache if you turn off gathering cache statistics. Click **Devices > Access Gateways > Edit > Advanced Options** and add the following command:

  ```
  DiskCacheMonitorStats off
  ```

- **SSL Terminator:** Install an SSL terminator between the browsers and the Access Gateway. This reduces the amount of rewriting required when the browsers are using SSL and the Web servers protected by the Access Gateway aren't configured for SSL.

  Click **Devices > Access Gateways > Edit > Reverse Proxy / Authentication**. Enable the **Behind Third Party SSL Terminator** option.

- **SSL Cipher Suites:** Use the advanced options from Apache to set the cipher suites that you want to allow. Some cipher suites take longer than others to process.

  For more information, see "SSLCipherSuite Directive" (http://httpd.apache.org/docs/2.2/mod/mod_ssl.html#sslciphersuite).

- **Statistics:** If additional performance is desired and statistics are not important, you can unload the `mod_status` module. If you unload the `mod_status` module, the extended information is also not gathered.

To unload the module, open the `httpd.conf` file in the `apache` directory, and add a comment symbol (#) to the line that loads the `mod_status` module in the Load Module section.

If you turn on debug mode, the mod_status module is automatically loaded in order to gather as much information as possible.

8.7 Solving Apache Restart Issues

When you make configuration changes and update the Access Gateway, the Administration Console uses the JCC channel to send the configuration changes to the Access Gateway. Figure 8-2 illustrates this flow.

Figure 8-2   *Sending Configuration Updates to the Access Gateway*
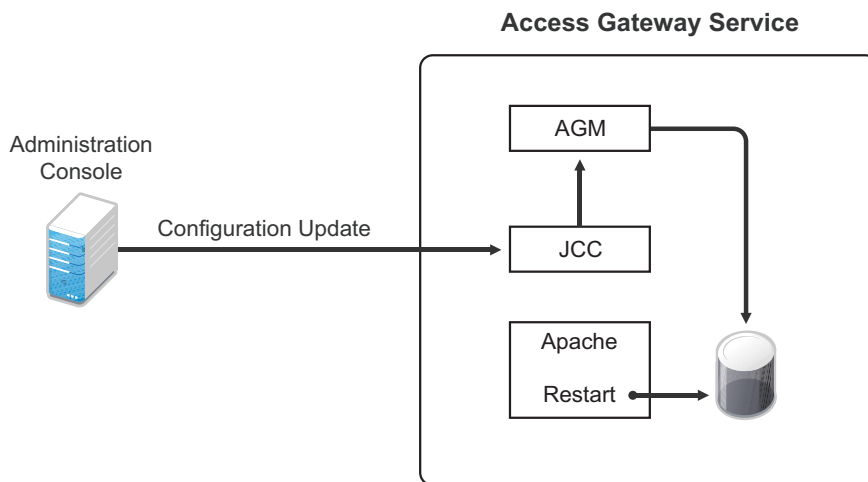
JCC sends the configuration changes to the Access Gateway Manager (AGM), which writes the Apache configuration to disk. Apache is sent a restart command, which causes Apache to read the new configuration, then Apache validates the configuration.

- If the configuration is valid, Apache starts.
- If the configuration is invalid, Apache fails to start.

If Apache fails to start after a configuration change, roll back to the previous configuration. Restore a backup if you have one, or use the Administration Console to manually remove the modifications that have caused the problem. If this does not solve the problem, try the following:

- Section 8.7.1, "Removing Any Advanced Configuration Settings," on page 194
- Section 8.7.2, "Viewing the Logged Apache Errors," on page 195
- Section 8.7.3, "Viewing the Errors as Apache Generates Them," on page 195
- Section 8.7.4, "The ActiveMQ Module Fails to Start," on page 196

8.7.1 Removing Any Advanced Configuration Settings

Apache fails to start when it discovers a syntax error in any of the advanced options.

1 Click **Devices > Edit > Advanced Options**.
2 To reset all options to their default values, delete all options from the text box.
3 Click **OK**.

When you return to the Advanced Options page, all options are set to their default values.

4 Click **[Name of Reverse Proxy] > [Name of Proxy Service] > Advanced Options**.

5 To reset all options to their default value, delete all options from the text box.

6 Click **OK**.

When you return to the Advanced Options page, all options are set to their default values.

7 Repeat these steps for each proxy service that has advanced options configured.

8 Update the Access Gateway.

8.7.2 Viewing the Logged Apache Errors

Apache generates and logs errors when it fails to start. A summary is displayed on the health page.

1 In the Administration Console, click **Devices > Access Gateways > Health**.

The page displays a summary of the problem from the Apache error log file. For the Access Gateway Service, information from the `rcnovell-apache2.out` file might also be displayed.

2 To view the entire contents of the Apache error log file, open a terminal window to the Access Gateway.

3 Change to the following directory and open the Apache error log file.

`/var/log/novell-apache2`

4 View the contents of the `rcnovell-apache2.out` file too.

5 If you still do not have enough information to solve the configuration problem, continue with Section 8.7.3, "Viewing the Errors as Apache Generates Them," on page 195.

8.7.3 Viewing the Errors as Apache Generates Them

Apache allows only a few errors to be sent to log files. To view all the errors, use the following procedure to display the errors in a terminal window.

1 Copy the `config.xml` file in the `current` directory to a temporary location. The Access Gateway allows only one XML file to reside in the `current` directory.

`/opt/novell/nam/mag/webapps/agm/WEB-INF/config/current`

2 Copy the XML file from the `pending` directory to the current directory and rename it `config.xml`.

The file in the `pending` directory has a long numeric name.

3 Change the ownership of the file from `root` to `novlwww:novlwww`.

4 Use one of the following commands to restart Tomcat:

`/etc/init.d/novell-mag restart` OR `rcnovell-mag restart`

5 Restart Apache by using the following command:

`/etc/init.d/novell-apache2 restart` OR `rcnovell-apache2 restart`

Apache uses the terminal window to write the errors it discovers as it tries to process the `config.xml` file.

6 At the Administration Console, fix the configuration problems, then update the Access Gateway.

8.7.4 The ActiveMQ Module Fails to Start

The Active MQ module is used for real-time communication between the Administration Console and the Access Gateway Service. Real-time communication is needed for commands such as purging cache, gathering statistics, and updating health. Figure 8-3 illustrates this communication flow.

Figure 8-3 *Real-Time Communication*
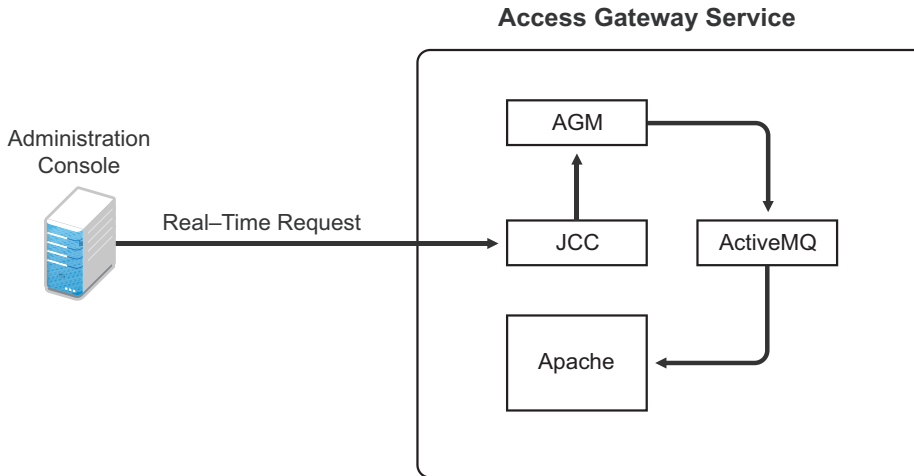
When the ActiveMQ module fails to start, you cannot apply any configuration changes, and the Access Gateway does not set a listener for the configured port.

To start the module, it must be able to resolve the listening IP address to a DNS name. To install an Access Gateway Service, the machine must have a DNS name and the IP address must resolve to this name.

8.8 Understanding the Authentication Process of the Access Gateway Service

When a user requests access to a protected resource, the request can be in one of the following states:

- No session or cookie is established, because this is the user's first request.
- The user's session is a public session because only public resources have been accessed.
- A session is established, the user is authenticated, and the requested resource is from the same cookie domain and uses the same contract.
- A session is established, the user is authenticated, and the requested resource is from the same cookie domain but uses a different contract or the contract has expired.
- A session is established, the user is authenticated, but the request doesn't have a session cookie because the resource is on a different cookie domain.
- A session no longer exists or doesn't exist on the proxy servicing the request.

The Access Gateway Service must handle these conditions and others as it determines whether it needs to forward a plogin request to the Embedded Service Provider or use the user's existing authentication credentials. The following flow charts take you through this process:

-

Figure 8-4   *Identifying the Requester*
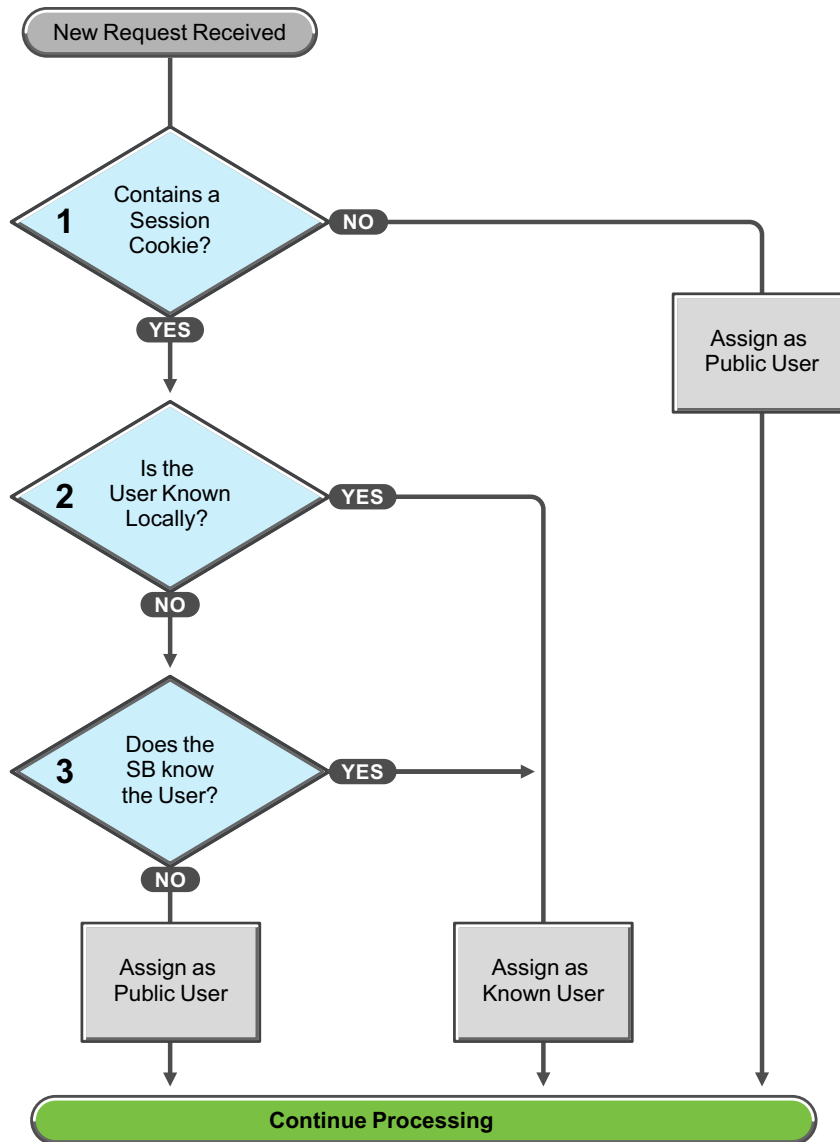
These first steps determine whether the Access Gateway knows the user that has submitted the request. In decision point 1, the Access Gateway checks for a session cookie in the request.

- If the request contains a session cookie, the session cookie needs to be validated. Processing continues with the task in decision point 2.
- If the request does not contain a session cookie, the user is unknown and is assigned as a public user. The Access Gateway continues processing with the tasks outlined in Figure 8-5 on page 198.

When the request contains a session cookie, the Access Gateway checks its local user store for a user that matches the session cookie. Each Access Gateway in the cluster maintains its own list of known users.

- If the session cookie matches one of the locally known users, the user is assigned that identity. The Access Gateway continues with the tasks outlined in Figure 8-5 on page 198.
- If the session cookie doesn't match one of the locally known users, the Access Gateway needs to know if one of the other Access Gateways in the cluster knows the user. Processing continues with the task in decision point 3.

The Access Gateway queries the session broker to see if one of the other Access Gateways in the cluster knows this user.

- If a match is found, the user is assigned that identity. The Access Gateway continues with tasks outlined in Figure 8-5 on page 198.
- If a match is not found, the user is unknown and is assigned as a public user. The Access Gateway continues with the tasks outlined in Figure 8-5 on page 198.

Figure 8-5 *Determining the Type of Request*
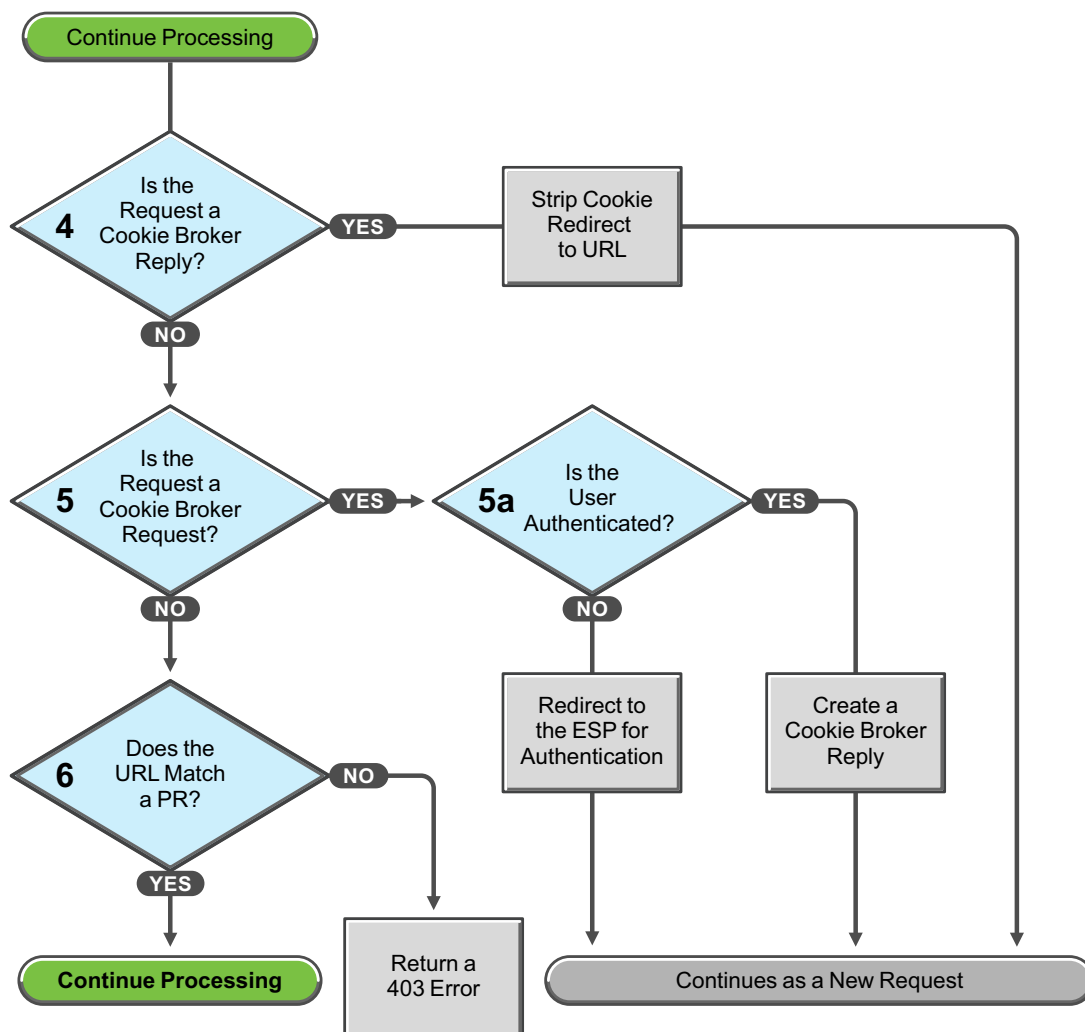
The Access Gateway examines the request to determine what type of request it is.

If the request is a cookie broker reply, the Access Gateway strips the cookie from the URL and redirects the request to the URL. The redirect is handled as a new request, and this new request flows to the task in decision point 6, where the URL is examined.

If the request isn't a cookie broker reply, the Access Gateway examines the request to see if it is a cookie broker request. If it is a cookie broker request, the Access Gateway determines whether the user is authenticated with the contract required by the protected resource.

- If the user is authenticated, the Access Gateway creates a cookie broker reply. This reply is handled as a new request, and flows to the task in decision point 4.

- If the user is not authenticated, the request is redirected to the Embedded Service Provider (ESP). The ESP interacts with the Identity Server to authenticate the user. The Identity Server, the ESP, and the reverse proxy all maintain authentication information. The ESP returns a new request, which flows to the task in decision point 6, where the URL is examined.

If the URL does not match a URL of a protected resource (PR), the Access Gateway returns an HTTP 403 error to the user.

If the URL in the request matches a URL of a protected resource, the Access Gateway needs to examine the protection type assigned to the resource. The Access Gateway continues with the tasks outlined in Figure 8-6 on page 200.

Figure 8-6 *Determining the Protection Type Assigned to the Resource*
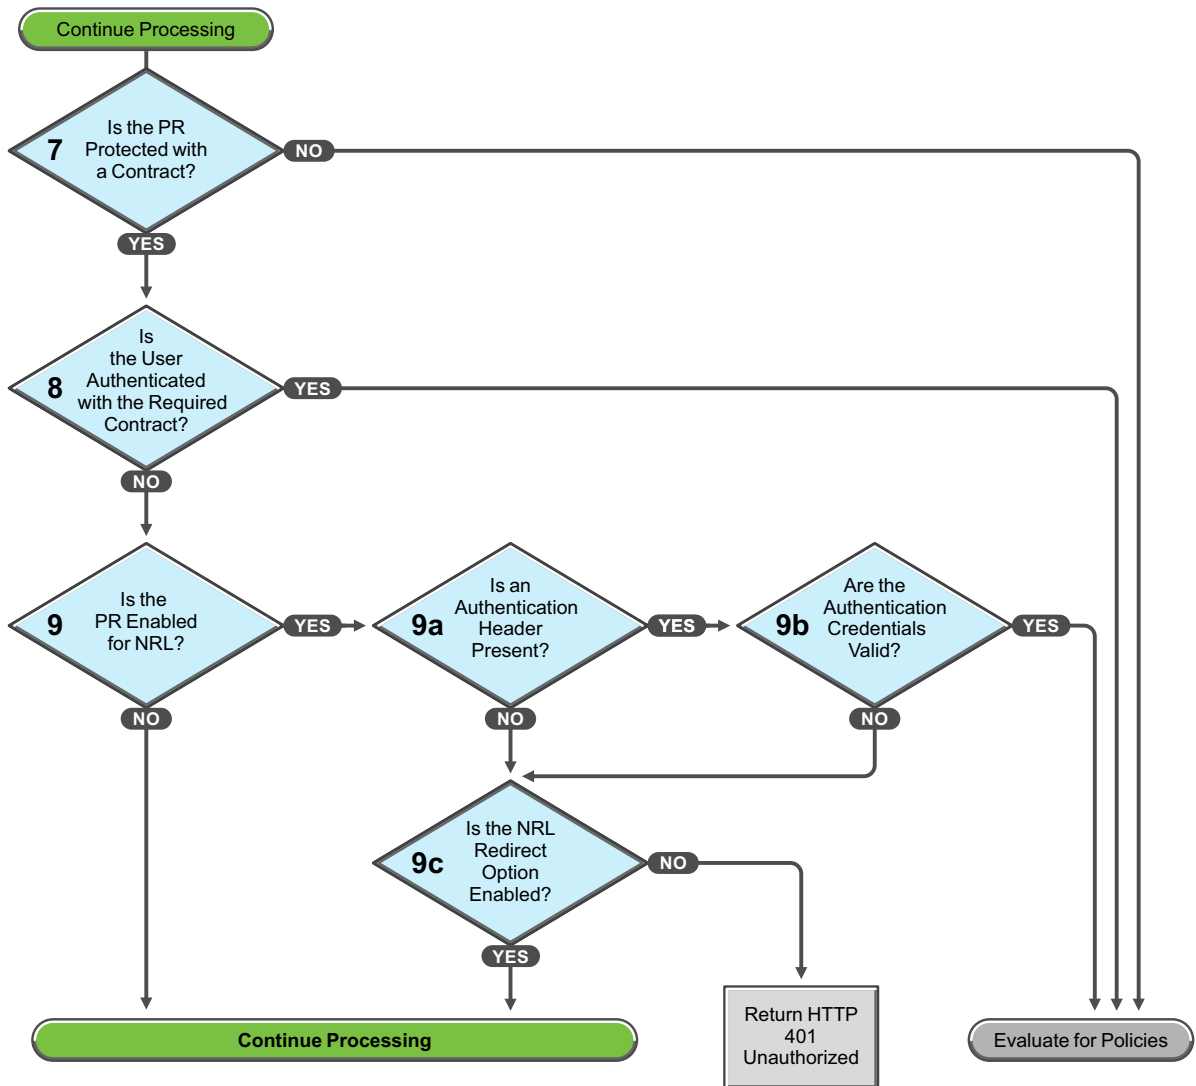
You configure a protected resource as a public resource when an authentication procedure/contract is not assigned to the protected resource. In decision point 7, the Access Gateway checks to see if a contract has been assigned to the protected resource.

  ◆ If the protected resource has not been assigned a contract, the Access Gateway is finished with its authentication checks and continues with policy evaluation.

  ◆ If the protected resource has been assigned a contract, the Access Gateway continues with the task in decision point 8.

For a user to gain access to a resource protected by a contract, the user must have authenticated with that contract, or if the contract is configured for it, the user can authenticate with another contract as long as the contract is of a equal or higher level.

  ◆ If the user is authenticated with the required contract, the Access Gateway is finished with its authentication checks and continues with policy evaluation.

  ◆ If the user is not authenticated with the required contract, the Access Gateway continues with the task in decision point 9.

Before the user is prompted for credentials, the Access Gateway needs to know whether the protected resource has been enabled for non-redirected login (NRL).

- ◆ If the resource has not been configured for non-redirected login, the Access Gateway continues with the tasks outlined in .
- ◆ If the resource has been configured for non-redirected login, the Access Gateway needs to examine the request for an authentication header and determine whether the header is valid. Processing continues with the tasks outlined in decision points 9a, 9b, and 9c.

If the request does not contain an authentication header, the Access Gateway needs to determine how non-redirected login has been configured. On the Authentication Procedure configuration page, you can select to enable the **Redirect to Identity Server When No Authentication Header Is Provided** option.

- ◆ If this option is enabled, the Access Gateway continues with the tasks outlined in .
- ◆ If this option is disabled, the Access Gateway returns an HTTP 401 unauthorized message.

If the request does contain an authentication header, the Access Gateway must verify that the credentials are valid.

- ◆ If the authentication credentials are valid, the Access Gateway is finished with its authentication checks and continues with evaluating the protected resource for policies.
- ◆ If the authentication credentials are not valid, the process is the same as if the request did not contain an authentication header and continues with the task in decision point 9c.

Figure 8-7 *Evaluating the Cookie Domain*
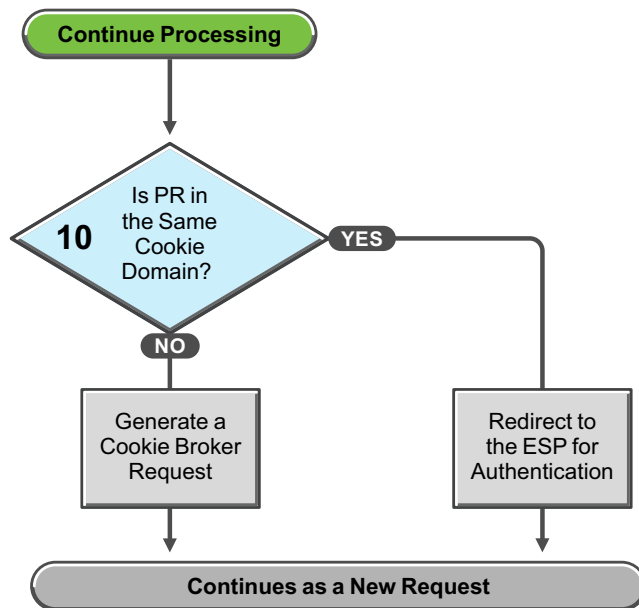
If you have configured your Access Gateway to use multiple domain-based proxy services, you can configure them to share the same cookie domain (domains of `development.novell.com` and `support.novell.com` can share the cookie domain of `novell.com`) or configure them so that they cannot share a cookie domain (domains of `a.slc.com` and `b.provo.com` cannot share a cookie domain).

When the Access Gateway reaches the task in decision point 10, it has determined that the protected resource requires a contract and that user is not authenticated with that contract.

  ◆ If the protected resource is in the same cookie domain, the Access Gateway redirects the request to the Embedded Service Provider (ESP). The ESP interacts with the Identity Server to authenticate the user. The ESP returns a new request, which flows to the task in decision point 6, where the URL is examined.

  ◆ If the protected resource is in a different cookie domain, the Access Gateway generates a cookie broker request. This new request flows to the task in decision point 5.

8.9 Enabling Caching of Audit Events for Apache Gateway Service

1 Open the `log4j.xml.base` file.

```
/etc/opt/novell/amlogging/config/log4j.xml.base
```

2 Look for Audit Server entry. By default the EnableCaching value is set to false. The xml entry for Audit server looks as below:

```
<appender name="AMAuditNSureAuditAppender"
class="com.novell.nacm.logging.audit.AMNSureAuditAppender">
        <param name="AppendMode" value="DIRECT"/>
        <param name="ErrorHandling" value="DISCARD"/>
        <param name="CertificatePath"
value="/etc/opt/novell/amlogging/certs/amnacert.pem"/>
        <param name="PrivateKeyPath"
value="/etc/opt/novell/amlogging/certs/amnapkey.pem"/>
        <param name="EnableCaching" value="false"/>
        <param name="ServerCheckInterval" value="4"/>

        <filter class="com.novell.nacm.logging.audit.AMNSureAuditFilter">
        </filter>
    </appender>
```

3 Modify the EnableCaching value from **false** to **true**.

4 In the Administration Console, click **Devices** > **Access Gateways** > **Edit** > **Auditing**. If any of the events are enabled, then disable all the events by deselecting them. Click **OK** twice. On the Access Gateways page, click **Update**.

8.10 Issue While Accelerating the Ajax Applications

If you are accelerating an Ajax application that cannot handle redirect and uses an authentication contract of 5 or 10 min, then increase the contract time out. Ensure that your Ajax application refreshes at an interval of 2 or 5 min. As a best practice, ensure that the Ajax application refresh interval is less than 2/3 of the contract time out.

8.11 Accessing Lotus-iNotes through the Access Gateway Asks for Authentication

This issue in not related to Access Manager. You need to configure authentication in Lotus-iNotes.

For more information about configuring Lotus-iNotes, see section 2.1 Authentication in the *iNotes Web Access Deployment and Administration guide* *(http://www.redbooks.ibm.com/redbooks/pdfs/ sg246518.pdf).*

8.12 Configuration Issues

If you get pending configuration issues when you apply changes on the device, one of the reasons could be that the soft link for the `certs` folder does not exist.

Enter the following command to check if the soft link exists for the `certs` folder:

```
ls -ltrh /opt/novell/apache2/
```

The following output is displayed:

```
lrwxrwxrwx 1 root root   34 2012-03-09 19:43 certs -> /etc/opt/novell/apache2/conf/
certs
```

If the soft link does not exist, perform the following steps:

1 Enter the following command:

```
ln -sf  /etc/opt/novell/apache2/conf/certs opt/novell/apache2/conf/certs
```

2 Click **Auditing** > **Troubleshooting** > **Certificates**.

3 Select the store that is reporting errors, then click **Re?push certificates**. You can select multiple stores at the same time.

4 (Optional) To verify that the re?push of the certificates was successful, click **Security** > **Command**

8.13 Keepalive timeout in the Reverse Proxy Configuration Does Not Get Applied to the Access Gateway

This will create issue for a customer with high traffic load. The slots will be fully occupied in apache and no more new requests will be entertained. To find the server load status, check http://localhost:8181/server-status. If there are zero idle workers and lot of slots shows 'K' instead of '.', then decrease the idle timeout or disable it if required.

8.14 Cannot Inject a Photo into HTTP Headers

You can use the jpegPhoto LDAP attribute to store your photo in JPEG format. This LDAP attribute is not injecting the image into a custom HTTP header and returns a 400 Bad Request error.

Edit the `index.php` file and add the following line:

```
<img src="data:image/jpeg;base64,/9j/4....Base64 value got from custom

header......."/>
```

8.15 Access Gateway Caching Issues

If you have caching issues with inodes, disk space, and cache corruption in the Access Gateway, use Apache htcacheclean tool which is used to keep the size of mod_disk_cache's storage within a certain limit. This tool can run either manually or in daemon mode. When running in daemon mode, it sleeps in the background and checks the cache directories at regular intervals for cached content to be removed. For more information on Apache htcacheclean tool, see (https://httpd.apache.org/docs/2.2/programs/htcacheclean.html).

8.16 Issues while Changing Management IP Address on an Access Gateway Appliance

If the Access Gateway Appliance has two NICs, a public and a private NIC, it is unable to change the Management IP address of the Access Gateway Appliance to a new value. The Administration Console connects to the changed IP address, but also tries to connect to the old IP address. The following procedure allows you to change the IP address manually.

1 Stop the AG service /etc/init.d/novell-mag stop

2 Stop the JCC service by using the `/etc/init.d/novell-jcc stop` command.

3 Change the IP address in `/opt/novell/devman/jcc/conf/settings.properties` file.

4 Change the IP address in `/opt/novell/nam/mag/webapps/agm/WEB-INF/config/current/config.xml` file.

5 Change the IP address in `/etc/opt/novell/apache2/conf/listen.conf` file.

6 Using YaST, change the network IP address.

7 In the Administration Console Edir, edit and change the IP address in the following attributes:

- For a specific Access Gateway device entry:

  1. romaAGDeviceXMLDoc of ou=ag-xxxxxx, ou=AppliancesContainer, ou=Partition, ou=PartitionsContainer, ou=VCDN_Root, ou=accessManagerContainer, o=novell

  2. romaAGDeviceSAXMLDoc of ou=ag-xxxxxx, ou=AppliancesContainer, ou=Partition, ou=PartitionsContainer, ou=VCDN_Root, ou=accessManagerContainer, o=novell

  3. romaAGConfigurationXMLDoc of ou=WorkingConfig, ou=ag-xxxx, ou=AppliancesContainer, ou=Partition, ou=PartitionsContainer, ou=VCDN_Root, ou=accessManagerContainer, o=novell

  4. romaAGConfigurationXMLDoc of ou=CurrentConfig ou=ag-xxxx, ou=AppliancesContainer, ou=Partition, ou=PartitionsContainer, ou=VCDN_Root, ou=accessManagerContainer, o=novell

- For the specific Access Gateway (esp)-identity server entry:

  1. romaIDPDeviceSAXMLDoc of ou=idp-esp-xxxxxx, ou=AppliancesContainer, ou=Partition, ou=PartitionsContainer, ou=VCDN_Root, ou=accessManagerContainer, o=novell

  2. romaIDPDeviceXMLDoc of ou=idp-esp-xxxxxx, ou=AppliancesContainer, ou=Partition, ou=PartitionsContainer, ou=VCDN_Root, ou=accessManagerContainer, o=novell

- For other Access Gateway device entry (if they are in a cluster):

  1. romaAGConfigurationXMLDoc of ou=WorkingConfig, ou=ag-yyyy, ou=AppliancesContainer, ou=Partition, ou=PartitionsContainer, ou=VCDN_Root, ou=accessManagerContainer, o=novell

  2. romaAGConfigurationXMLDoc of ou=CurrentConfig, ou=ag-yyyy, ou=AppliancesContainer, ou=Partition, ou=PartitionsContainer, ou=VCDN_Root, ou=accessManagerContainer, o=novell

- For tmp folder entry:

  1. romaAGConfigurationXMLDoc of ou=CurrentConfig, ou=tmp_zzz, ou=AppliancesContainer, ou=Partition, ou=PartitionsContainer, ou=VCDN_Root, ou=accessManagerContainer, o=novell

  2. romaAGConfigurationXMLDoc of ou=WorkingConfig, ou=tmp_zzz, ou=AppliancesContainer, ou=Partition, ou=PartitionsContainer, ou=VCDN_Root, ou=accessManagerContainer, o=novell

8 Start the Access Gateway service by using the `/etc/init.d/novell-mag start` command.

9 Start the JCC service by using the `/etc/init.d/novell-jcc start` command.

10 Restart the Administration Console, Identity Server and other Access Gateways in cluster.

8.17 Clustered Nodes Looping Due to JGroup Issues

In a cluster when multiple nodes are down, failover does not occur and user experiences looping due to jgroups issues.

Workaround: Modify the `/opt/novell/nids/lib/webapp/WEB-INF/web.xml` file on the Identity Server as follows to increase the jgroup timeouts.

```
<context-param>

 <param-name>JGroupsConfiguration</param-name>

 <param
value>TCP(start_port=[nidp:ClusterPort];end_port=[nidp:ClusterPort][nidp:IfExterna
lAddress];external_addr=[nidp:ExternalAddress][nidp:EndIf]):TCPPING(initial_hosts=
[nidp:ClusterMembers];port_range=1;timeout=3500;num_initial_members=2;up_thread=tr
ue;down_thread=true):MERGE2(min_interval=10000;max_interval=30000):FD_SOCK([nidp:I
fExternalAddress]bind_addr=[nidp:ExternalAddress][nidp:EndIf]):FD(shun=true;timeou
t=5000;max_tries=5;up_thread=true;down_thread=true):VERIFY_SUSPECT(timeout=2000;do
wn_thread=false;up_thread=false):pbcast.NAKACK(down_thread=true;up_thread=true;gc_
lag=100;retransmit_timeout=3000):pbcast.STABLE(desired_avg_gossip=20000;down_threa
d=false;up_thread=false):pbcast.STATE_TRANSFER():pbcast.GMS(merge_timeout=10000;jo
in_timeout=5000;join_retry_timeout=2000;shun=true;print_local_addr=[nidp:DebugOn];
down_thread=true;up_thread=true)</param-value>

 </context-param>
```

For more information about the timeout options, see the following links:

- TCPPING
- MERGE2
- FD_SOCK
- FD
- VERIFY_SUSPECT (http://www.jgroups.org/manual/html/protlist.html#d0e5762)
- NAKACK
- pbcast.STATE_TRANSFER
- pbcast.GMS