AUTHORIZED DOCUMENTATION

Access Gateway Guide

Novell® Access Manager

3.1 SP1

April 5, 2010

www.novell.com

Legal Notices

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

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

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

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

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

Novell Trademarks

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

Third-Party Materials

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

Contents

About This Guide

This guide describes the following features of Novell® Access Gateways:

This administration guide is intended to help you understand and configure all of the features provided by Access Manager, and includes advanced topics.

It is recommended that you first become familiar with the information in the *Novell Access Manager 3.1 SP1 Setup Guide*, which helps you understand how to perform a basic Identity Server configuration, set up a resource protected by an Access Gateway, and configure SSL.

The basic setup and the administration guides are designed to work together, and important information and setup steps are not always repeated in both places.

Audience

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

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

Feedback

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

Documentation Updates

For the most recent version of the *Access Manager Administration Guide*, visit the Novell Access Manager Documentation Web site (http://www.novell.com/documentation/novellaccessmanager).

Additional Documentation

Before proceeding, you should be familiar with the *Novell Access Manager 3.1 SP1 Installation Guide* and the *Novell Access Manager 3.1 SP1 Setup Guide*, which provides information about setting up the Access Manager system.

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

- *Novell Access Manager 3.1 SP1 Administration Console Guide*
- *Novell Access Manager 3.1 SP1 Identity Server Guide*
- *Novell Access Manager 3.1 SP1 Policy Management Guide*
- *Novell Access Manager 3.1 SP1 Agent Guide*
- *Novell Access Manager 3.1 SP1 SSL VPN Server Guide*
- *Novell Access Manager 3.1 SP1 Event Codes*

Documentation Conventions

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

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

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

Configuring the Access Gateway to Protect Web Resources

1

The Novell® Access Gateway is a reverse proxy server (protected site server) that restricts access to Web-based content, portals, and Web applications that employ authentication and access control policies. It also provides single sign-on to multiple Web servers and Web applications by securely providing the credential information of authenticated users to the protected servers and applications. The Access Gateway lets you simplify, secure, and accelerate your Internet business initiatives.

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

Figure 1-1 *Accessing a Web Resource*

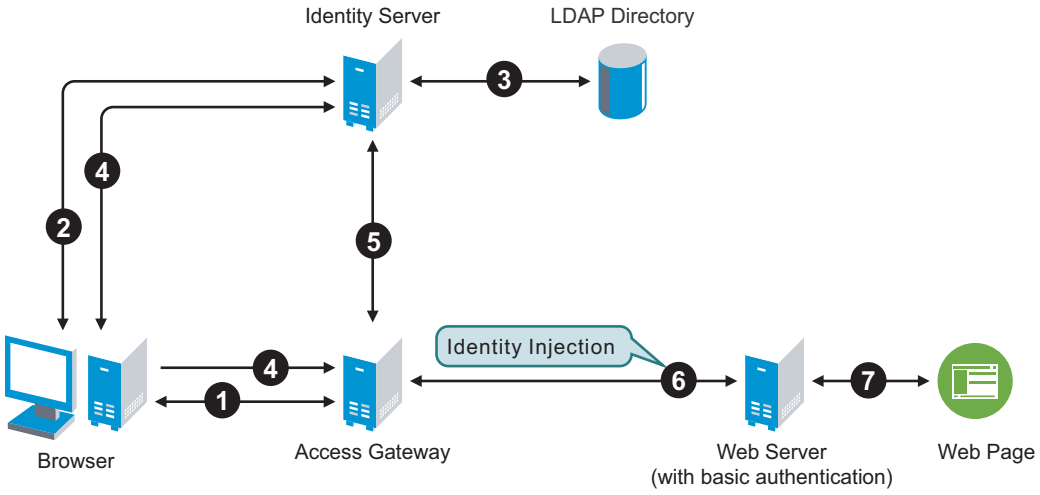

1. The user requests access to a resource protected by the Access Gateway.
2. The Access Gateway redirects the user to the Identity Server, which prompts the user for a username and password.
3. The Identity Server verifies the username and password against an LDAP directory (eDirectory™, Active Directory, 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 1-2 *Access Gateway Modules and Their Configuration Options*

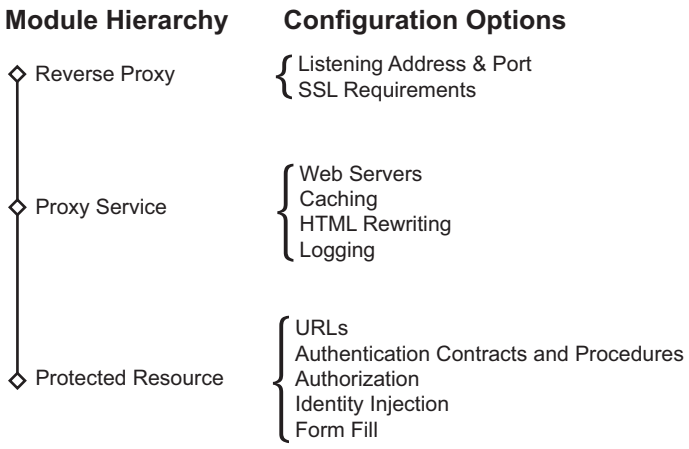

This hierarchy allows you to have precise control over what is required to access a particular resource, and also allows you to provide a single sign-on solution for all the resources protected by the Access Gateway. The authentication contract, authentication procedure, Authorization policy, Identity 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.

This section describes the following tasks:

- Section 1.1, “Creating a Reverse Proxy and Proxy Service,” on page 12
- Section 1.2, “Configuring a Proxy Service,” on page 16
- Section 1.3, “Configuring the Web Servers of a Proxy Service,” on page 18
- Section 1.4, “Configuring Protected Resources,” on page 19
- Section 1.5, “Configuring Protected Resources for Specific Applications,” on page 30
- Section 1.6, “Configuring HTML Rewriting,” on page 41
- Section 1.7, “Configuring Connection and Session Limits,” on page 59

1.1 Creating a Reverse Proxy and Proxy Service

A reverse proxy acts as the front end to your Web servers on your Internet or intranet and off-loads frequent requests, thereby freeing up bandwidth. The proxy also increases security because the IP addresses of your Web servers are hidden from the Internet.

To create a reverse proxy, you must create at least one proxy service with a protected resource. You must supply a name for each of these components. Reverse proxy names and proxy service names must be unique to the Access Gateway because they are configured for global services such as IP

addresses and TCP ports. For example, if you have a reverse proxy named `products` and another reverse proxy named `library`, only one of these reverse proxies can have a proxy service named `corporate`.

Protected resource names need to be unique to the proxy service, but they don't need to be unique to the Access Gateway because they are always accessed through their proxy service. For example, if you have a proxy service named `account` and a proxy service named `sales`, they both can have a protected resource named `public`.

The first reverse proxy and proxy service you create is 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*.

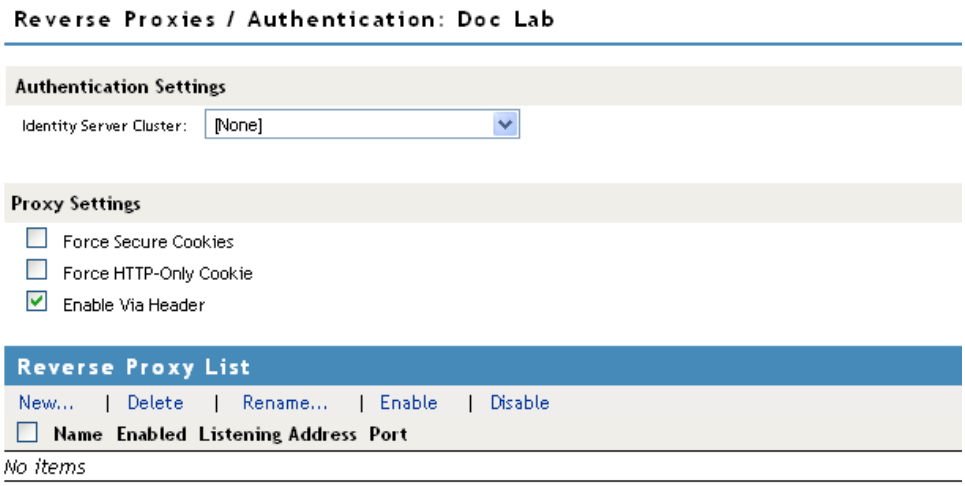

3 Configure the authentication settings.

Identity Server Cluster: Specifies the Identity Server you want the Access Gateway to trust for authentication. Select the configuration you have assigned to the Identity Server.

Whenever an Identity Server is assigned to a new trust relationship, the Identity Server needs to be updated. This process is explained following the step that saves this configuration setting (see Step 5 on page 17 and Step 6 on page 17).

4 (Optional) Configure the proxy settings.

Use the options in this section to control how all reverse proxies handle the security issues involving secure cookies and the Via header

Force Secure Cookies: Forces the Access Gateway to set the secure keyword for the proxy authentication cookie, regardless of whether the services hosted are all based on HTTPS. You should enable this option for either of the following conditions:

- All services that require the proxy authentication cookie (such as identity based policies) are hosted as HTTPS services
- An SSL accelerator, such as the Cisco* SSL accelerator, is placed between the Access Gateway and the browsers and the browser receives only HTTPS links.

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.

For more information and other options for securing Access Manager cookies, see Section 2.5, "Enabling Secure Cookies," on page 71.

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: 1.0 www.mylag.com (Access Gateway 3.1.1-72-D06FBFA8CF21AF45)
```

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

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

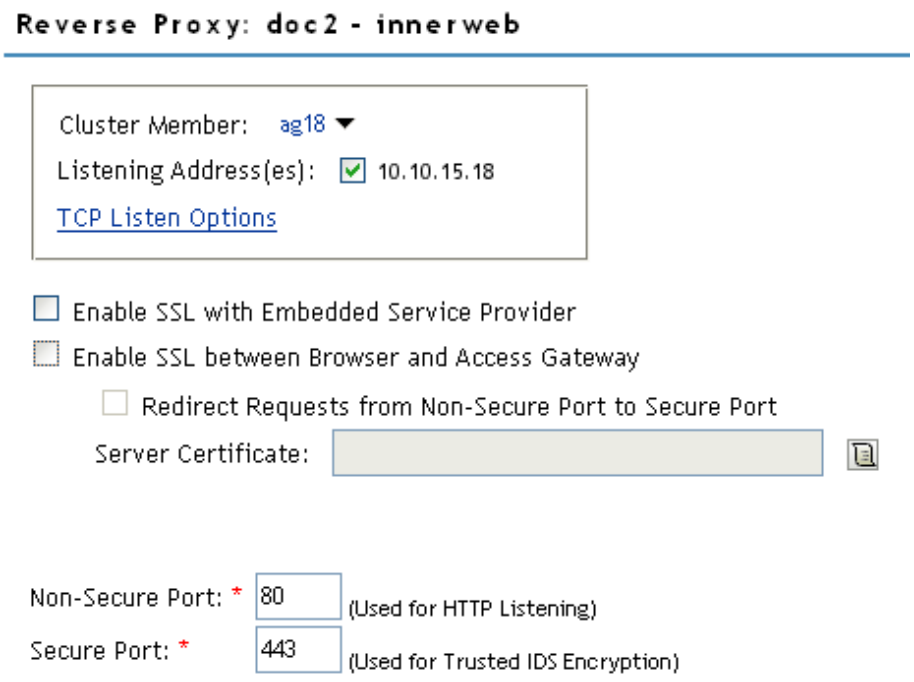

6 Enable a listening address. Fill in the following fields:

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 1.7.1, "Configuring TCP Listen Options for Clients," on page 59.

7 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 2, "Configuring the Access Gateway for SSL," on page 63.

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

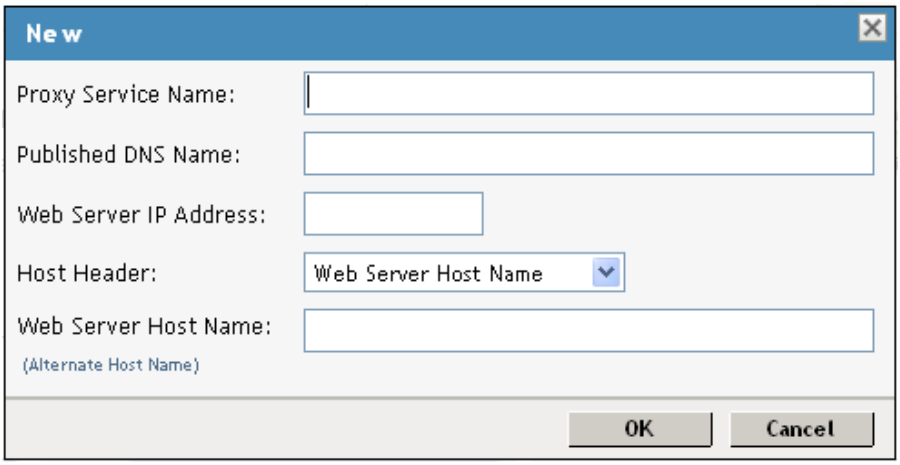

The first proxy service of a reverse proxy is considered the master (or parent) proxy. Subsequent proxy services can use domain-based, path-based, or virtual multi-homing, relative to the published DNS name of the master proxy service. If you are creating a second proxy service for a reverse proxy, see Section 6.2, "Using Multi-Homing to Access Multiple Resources," on page 155.

9 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.

---

10 Click *OK*.

11 Continue with Section 1.2, "Configuring a Proxy Service," on page 16 or select one of the following tasks:

- For instructions on creating multiple reverse proxies, see Section 6.3, "Managing Multiple Reverse Proxies," on page 164.
- For instructions on creating multiple proxy services for a reverse proxy, see Section 6.2, "Using Multi-Homing to Access Multiple Resources," on page 155.

1.2 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]*.

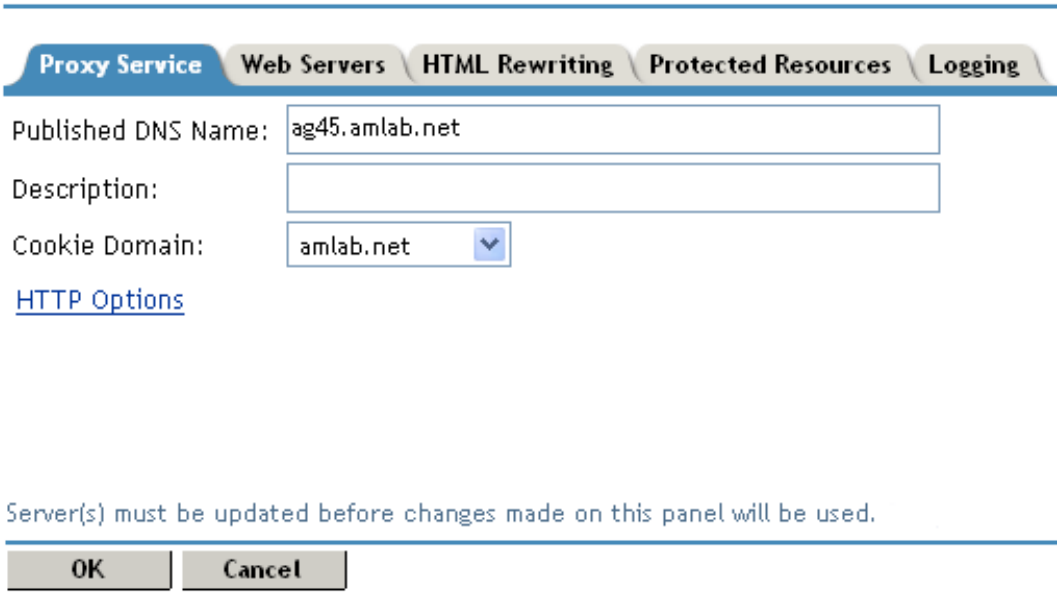

2 Fill in the following fields:

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.

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

- Section 5.2, "Controlling Browser Caching," on page 143
- Section 5.3, "Configuring Custom Cache Control Headers," on page 144

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 1.3, "Configuring the Web Servers of a Proxy Service," on page 18.
- 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 1.4, "Configuring Protected Resources," on page 19.

1.3 Configuring the Web Servers of a Proxy Service

The Web server configuration determines how the Access Gateway handles connections and packets between itself and the Web servers.

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

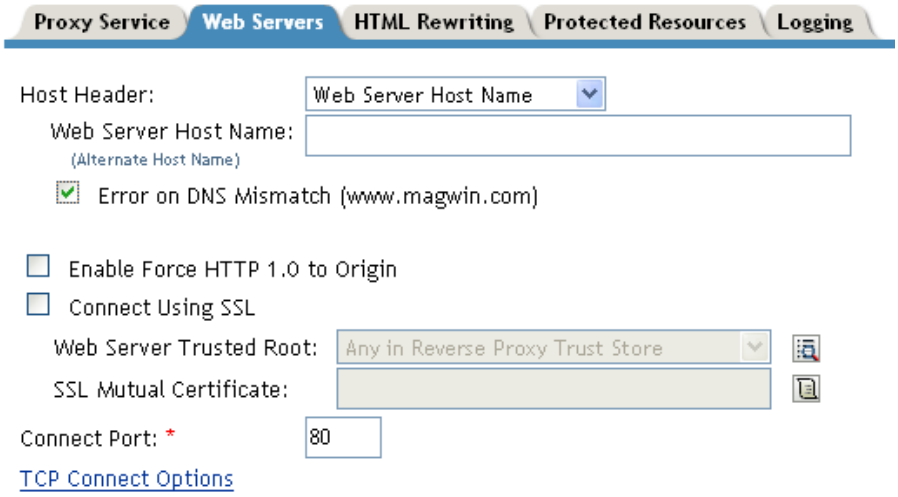

2 Specify the hostname that is placed in the HTTP header of the packets being sent to the Web servers. In the *Host Header* field, select one of the following:

- **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.

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.

4 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.

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 2.4, “Configuring SSL between the Proxy Service and the Web Servers,” on page 68.

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 1.7.2, “Configuring TCP Connect Options for Web Servers,” on page 60.

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 6.1, “Setting Up a Group of Web Servers,” on page 154.

9 To delete a Web server, select the Web server, then click *Delete*.

This deletes the Web server from the list so that the Access Gateway no longer sends requests to the deleted Web server. At least one Web server must remain in the list. You must delete the proxy service to remove the last server in the list.

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

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

1.4 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 protections a resource requires depends upon the resource, the Web server, and the conditions you define for the resource.

You can select from the following types of protection:

Authentication 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 contract, 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.

This section describes the following tasks:

1.4.1 Setting Up a Protected Resource

To configure a protected resource:

1. Click *Access Gateways* > *Edit* > *[Name of Reverse Proxy]* > *[Name of Proxy Service]* > *Protected Resources*.
2. Either click the name of an existing resource or click *New*, then specify a display name for the resource.

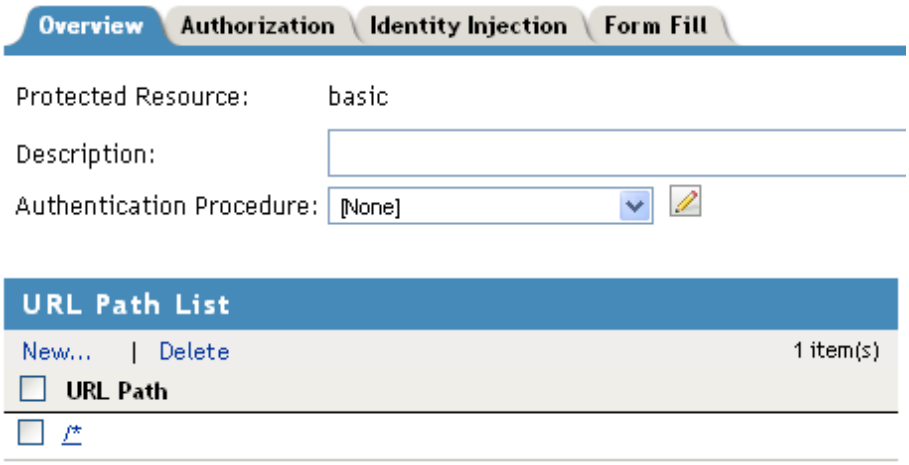

3 (Optional) Specify a description for the protected resource. You can use it to briefly describe the purpose for protecting this resource.

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, allows any contract defined for the Identity Server to be valid, or if the user has not authenticated, prompts the user to authenticate, using the default contract assigned to the Identity Server configuration.
- **Name/Password - Basic:** Specifies basic authentication over HTTP, using a standard login pop-up 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.

You can configure other types of contracts. For more information, see "Configuring Authentication Contracts" in the *Novell Access Manager 3.1 SP1 Identity Server Guide*.

If these default contracts are not available, you have not configured a relationship between the Access Gateway and the Identity Server. See Section 1.1, "Creating a Reverse Proxy and Proxy Service," on page 12.

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 1.4.4, "Modifying Authentication Procedures," on page 24.

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 on how a user's request is matched to a protected resource, see Section 1.4.2, "Understanding URL Path Matching," on page 23.

  For information on using a query string, see Section 1.4.3, "Using a Query String in the URL Path," on page 23.

- **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 25
- "Assigning an Identity Injection Policy to a Protected Resource" on page 26
- "Assigning a Form Fill Policy to a Protected Resource" on page 27
- "Assigning a Policy to Multiple Protected Resources" on page 29

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

1.4.2 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.ph`) 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).

URL paths are case insensitive. If your Web server has two paths (`/public/current` and `/public/Current`), a URL path of `/public/current` matches both.

1.4.3 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 Appliance allows you to configure the URL matching process so that it ignores the query string in the URL path.

On the Access Gateway Appliance, you can remove the query string from the URL path or you can create the following touch file:

```
/var/novell/.prWithOutQuestionMark
```

You need to then restart the Access Gateway Appliance to activate the touch file. When this touch file is used, the Access Gateway Appliance ignores the query string and uses just the path to find a match.

1.4.4 Modifying Authentication Procedures

When a user requests access to a protected resource that is protected by a contract, the default authentication procedure is to redirect the request to the Identity Server for the following conditions:

- When a user attempts to connect to a protected resource for the first time.
- When the user's session reaches a soft timeout.
- When the user's session reaches a hard timeout.

The session hard timeout (*Devices > Identity Servers > Edit > Session timeout*) is a global setting that applies to all users. The default value is 60 minutes. The Identity Server passes this value to the Embedded Service Providers (Access Gateway, SSL VPN, or J2EE agent) and service providers.

When the Access Gateway receives the session hard timeout, it uses the value to calculate a separate soft timeout that is 66% of the hard timeout. The Access Gateway uses the soft timeout as a trigger to inform the Identity Server that the session is still active.

- When the Access Gateway gets a request from a browser after the soft timeout expires, but before the hard session timeout, the Access Gateway attempts to renew the session with the Identity Server. This is done by redirecting the browser to the Identity Server. After the session renewal request, the Identity Server redirects back to the Access Gateway and the session has new soft and hard timeout values.
- When the Access Gateway receives a request after the hard timeout has expired, the Access Gateway allows the user to create a new session by redirecting the browser to the Identity Server, where the user is prompted to re-authenticate. After this re-authentication, the browser is redirected back to the Access Gateway and the session has new timeout values.

Some applications, such as AJAX and WebDAV applications, do not allow redirection for authentication. You can use non-redirected login to change the default authentication behavior of Access Manager 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 then used to complete the authentication on the user's behalf. The SOAP back channel, rather than a redirect, 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*.

To modify the authentication procedures:

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.

3 Click *New*, specify a name, then click *OK*.

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

Contract: Select the contract that you have configured for this protected resource. This needs to be a contract that supports name and password credentials.

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. This is what displays 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 (Conditional) If you have more than one realm to query for credentials, repeat Step 4 and Step 5.

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 30
- "Configuring a Protected Resource for a SharePoint Server with an ADFS Server" on page 31
- "Configuring a Protected Resource for Outlook Web Access" on page 34
- "Configuring a Protected Resource for a Novell Teaming 2.0 Server" on page 36

1.4.5 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*.

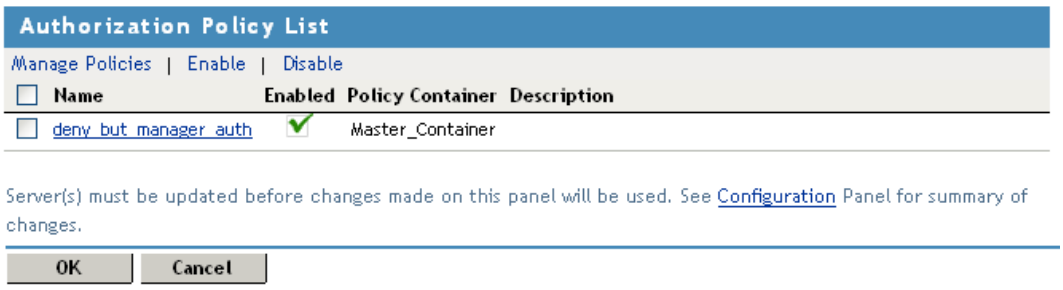

The *Authorization Policy List* contains all the Access Gateway Authorization policies that have been created on this Administration Console for the selected policy container.

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 *Novell Access Manager 3.1 SP1 Policy Management 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 *Novell Access Manager 3.1 SP1 Policy Management 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 browser cache, click *OK*.

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

1.4.6 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*.

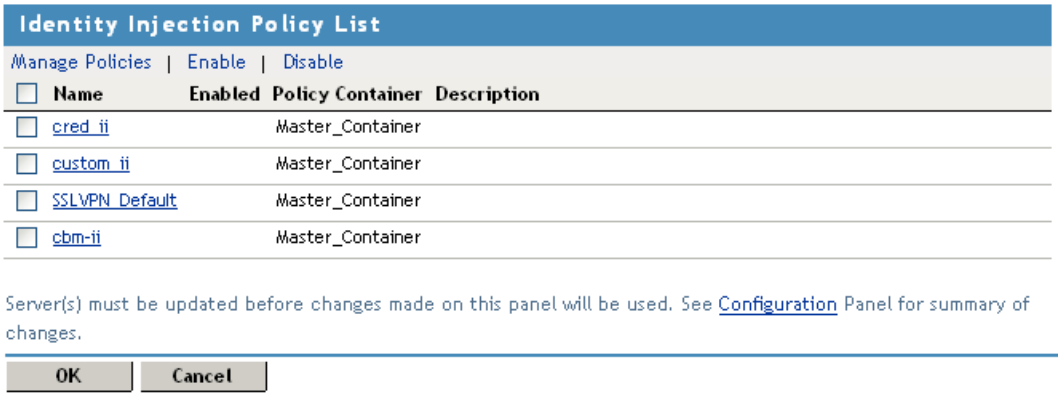

The *Identity Injection Policy List* contains all the Identity Injection policies that have been created on this Administration Console for the selected policy container.

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 *Novell Access Manager 3.1 SP1 Policy Management 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 *Novell Access Manager 3.1 SP1 Policy Management 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 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.

1.4.7 Assigning a Form Fill Policy to a Protected Resource

Some client requests cause the Web server to return a form. Sometimes this form contains a request to log in. If you create a Form Fill policy, you can have the Access Gateway fill in the form. When a user first logs in, the Access Gateway prepopulates some fields and prompt the users for the others. The Access Gateway securely saves the information, so that on subsequent logins, the Access Gateway can fill in the form. The user is only prompted to fill in the form when something changes, such as a password expiring.

Form Fill uses two components: the HTML form and the Form Fill policy. The HTML form is created with HTML tags and consists of form elements such as fields, menus, check boxes, and buttons. The Form Fill policy is created by specifying the following:

- Which information is entered automatically and not displayed to the user.
- Which information is displayed so that the user, at least the first time, can enter the information.
- What is done with the information (for example, is it saved so that the user doesn't need to enter it when accessing the form again).

You must create the policy before you can assign it to a resource (see "Creating Form Fill Policies" in the *Novell Access Manager 3.1 SP1 Policy Management 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 at all 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 on the performance problems this can cause, see "Creating a Form Matching Rule" in the *Novell Access Manager 3.1 SP1 Policy Management Guide*.

   ---

3. (Conditional) If the URL is not specific, click the name of the path and modify it.
4. Click *Form Fill*.

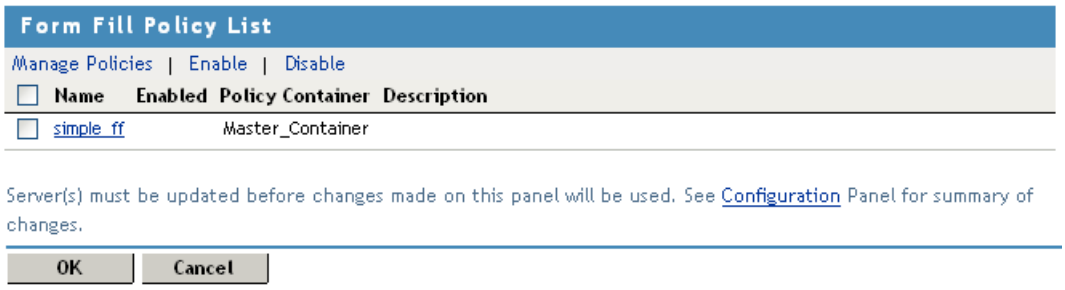

   The *Form Fill Policy List* contains all the Form Fill policies that have been created on this Administration Console for the selected policy container.

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 *Novell Access Manager 3.1 SP1 Policy Management 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 *Novell Access Manager 3.1 SP1 Policy Management 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 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 and how to download and configure it, see Access Management Authentication Class Extension to Retrieve Password for Single Sign-on (http://www.novell.com/communities/node/4556).

1.4.8 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. The one limitation is that 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*.

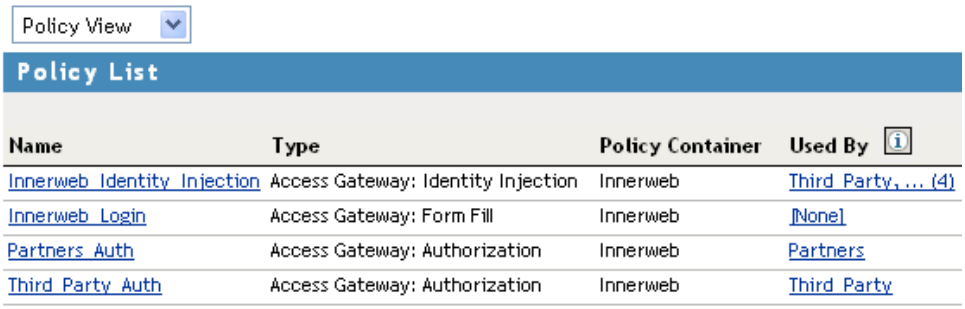

3 Select the *Used By* link of the policy you want to assign to multiple resources.

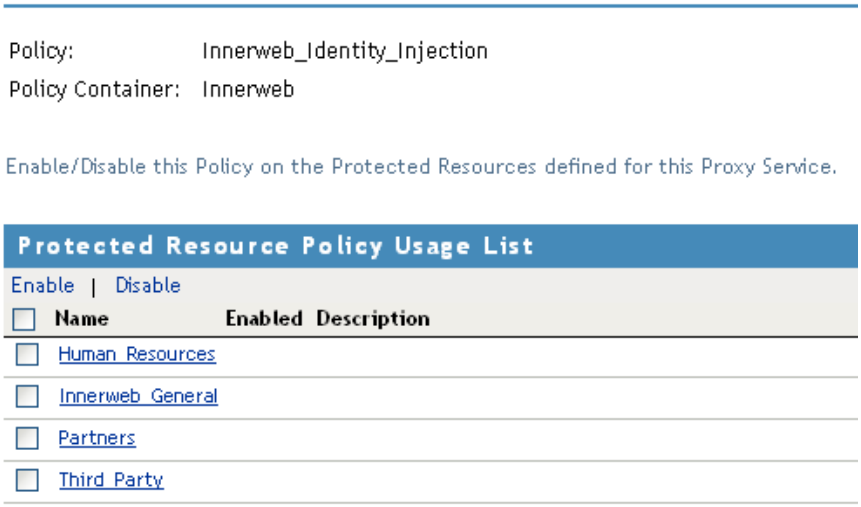

The *Policy* and *Policy Container* fields identify the policy. The *Protected Resource Policy Usage List* displays the protected resources defined for this proxy service and indicates which resources the policy has been enabled on.

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*.

1.5 Configuring Protected Resources for Specific Applications

- Section 1.5.1, “Configuring Protected Resource for a SharePoint Server,” on page 30
- Section 1.5.2, “Configuring a Protected Resource for a SharePoint Server with an ADFS Server,” on page 31
- Section 1.5.3, “Configuring a Protected Resource for Outlook Web Access,” on page 34
- Section 1.5.4, “Configuring a Protected Resource for a Novell Teaming 2.0 Server,” on page 36

1.5.1 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 Linux Access Gateway Appliance. When you protect a SharePoint server on Linux 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 the non-browser clients such as Microsoft Network Place, Nautilus in SLES 10 SP2 or MAC finder, you might see issues as these WebDAV clients do not support 302 redirection for authentication. You must modify the authentication procedure to prevent redirection when the user session expires.

For more information on how to configure a protected resource for a SharePoint server, see Novell Cool Solutions (http://www.novell.com/communities/node/8346/sharepoint-integration-linux-access-gateway).

1.5.2 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 following sections assume that you have a functioning SharePoint server and a functioning Access Manager 3.1 SP1 system:

- "Configuring a Custom Contract" on page 31
- "Creating a Reverse Proxy Service" on page 32
- "Configuring Multiple Protected Resources" on page 32

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 created from the same method as the default 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 on the other options, see "Configuring Authentication Contracts" in the *Novell Access Manager 3.1 SP1 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. Select the 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" on page 32.

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, which 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 on creating a reverse proxy, see Section 1.1, "Creating a Reverse Proxy and Proxy Service," on page 12

4 Click *OK*.

5 Continue with "Configuring Multiple Protected Resources" on page 32.

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 1.4.4, “Modifying Authentication Procedures,” on page 24.
5. Click *OK* three times.

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

   For information on configuring protected resources, see Section 1.4.1, “Setting Up a Protected Resource,” on page 20.
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*, update the Identity Server.

1.5.3 Configuring a Protected Resource for Outlook Web Access

To protect your Outlook Web Access server with the Access Gateway Appliance, configure some of the Access Manager features. The following sections assume that you have a functioning Outlook Web Access server and a functioning Access Manager 3.1 SP1 system:

Configuring a Protected Resource for Outlook Web Access

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. If you do not already have a reverse proxy service created, create one. For more information on creating a reverse proxy, see Section 1.1, "Creating a Reverse Proxy and Proxy Service," on page 12.
3. In the *Reverse Proxy List*, click *New*, specify a display name for the reverse proxy, then click *OK*.
4. In the *Proxy Service List* section, click *New*.
5. Specify a name for the proxy service, then click *OK*.
6. 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.
7. Click *OK*.
8. Continue with "Configuring an Authentication Procedure" on page 34.

Configuring an Authentication Procedure

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 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 To create protected resource as follows:

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 To create protected resource as follows:

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 as 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, then 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 "Configure Identity Injection" on page 36

11 Continue with "Configuring a Rewriter Profile" on page 35.

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 as follows:

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

3b Click the newly added word profile.

3c Click *New* in the *Variable or Attribute Name to Search for Is* section, then specify a variable to search for. For example, `value`.

3d Click *OK*.

3e Select *Rewrite Inbound Query String Data*.

3f Select *Rewrite Inbound Post Data.*

3g Select *Rewrite Inbound Headers*.

3h Make sure that *Enable Rewrite Actions* remains selected.

3i Click *OK*.

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

Configure Identity Injection

You must configure an Identity Injection policy in order to enable single sign on with the Outlook Web Access server which has basic authentication configured. This Identity Injection policy should be configured to inject an authentication header. For information on creating such a policy, see "Configuring an Authentication Header Policy" in the *Novell Access Manager 3.1 SP1 Policy Management Guide*.

For the Linux Access Gateway Appliance, you must do the following to ensure that when a browser sends an authentication header, the Linux Access Gateway Appliance overwrites this with the authentication header configured in the Identity Injection policy:

1.5.4 Configuring a Protected Resource for a Novell Teaming 2.0 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 Teaming 2.0 server on Linux and a functioning Access Manager 3.1 SP1 IR1 system with a reverse proxy configured for SSL communication between the browsers and the Access Gateway.

The Teaming 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 Teaming Server to Trust the Access Gateway" on page 37
- "Configuring a Domain-Based Multi-Homing Service for Novell Teaming" on page 38

- "Creating a Pin List" on page 40
- "Configuring Single Sign-On" on page 40

For information on other possible Access Gateway configurations, see "How-to: Integrating Access Manager 3.1.1 IR1 Linux Access Gateway with Teaming 2.0" (http://www.novell.com/support/php/search.do?cmd=displayKC&docType=kc&externalId=7004284&sliceId=1&docTypeID=DT_TID_1_1&dialogID=49882022&stateId=1%200%2049878366).

Configuring the Teaming Server to Trust the Access Gateway

To use Teaming as a protected resource of a Novell 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 Teaming 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, Teaming 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 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 teaming.doc.provo.novell.com, specify the following URL:

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

9. When prompted to use the Access Gateway for WebDAV connections, type `No`.
10. Follow the prompts to complete the reconfiguration process.
11. Start the Teaming server with the following command:

    ```
    /etc/init.d/teaming start
    ```

12. Continue with "Configuring a Domain-Based Multi-Homing Service for Novell Teaming" on page 38.

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

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

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

- "Configuring the Domain-Based Proxy Service" on page 38
- "Configuring Protected Resources" on page 38
- "Configuring a Rewriter Profile" on page 39

Configuring the Domain-Based Proxy Service

1. In the Administration Console, click *Devices* > *Access Gateways* > *Edit* > *[Name of Reverse Proxy]*.
2. In the *Reverse Proxy List*, click *New*, then fill in the following fields:

   Proxy Service Name: Specify a display name for the proxy service, which 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, teaming.doc.provo.novell.com.

   Web Server IP Address: Specify the IP address of the Novell Teaming server.

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

   Web Server Host Name: Specify the DNS name of the Teaming 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 Teaming 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 1200 seconds.

   This longer timeout is needed for file uploads.
8. Click *OK*.
9. Continue with "Configuring Protected Resources" on page 38.

Configuring Protected Resources

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

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

2 Create a protected resource for HTML content:

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

2b (Optional) Specify a description for the protected resource. You can use it to briefly describe the purpose for protecting this resource.

2c Specify a value for *Authentication Procedure.* For example, select the *Secure Name/Password - Form* contract.

2d In the *URL Path List*, remove the `/*` path and add the following two paths:

```
/teaming/*
/ssf/*
```

2e Click *OK*.

3 Create a protected resource for WebDAV and AJAX content:

3a In the *Protected Resource List*, click *New*, specify a unique name, then click *OK*.

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

3c Click the *Edit Authentication Procedure* icon.

3d In the *Authentication Procedure List*, click *New*, specify a name, then click *OK*.

3e Fill 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 Teaming configuration option. It appears when the user is prompted for credentials.

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

3f Click *OK* twice.

3g For the Authentication Procedure, select the procedure you just created.

3h In the *URL Path List*, remove the `/*` path and add the following two paths:

```
/ssfs/*
/ssf/a/do?*
/ssf/rss/*
```

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

3i Click *OK*.

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

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

6 Continue with "Configuring a Rewriter Profile" on page 35.

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 the *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. Make sure that *Enable Rewrite Actions* remains selected.
8. Click *OK*.
9. In the *HTML Rewriter Profile List*, move the Word profile you created to be the first profile in the list, and the default profile to be the second profile in the list.
10. Click *OK*.
11. To apply your changes, click *Devices* > *Access Gateways*, then click *Update*.
12. Continue with "Creating a Pin List" on page 40.

Creating a Pin List

The Access Gateway needs to be configured to bypass the published URL of 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, teaming.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*.
7. Continue with "Configuring Single Sign-On" on page 40.

Configuring Single Sign-On

You must configure an Identity Injection policy to enable single sign on with the Novell Teaming 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*.

8 To save the policy, click *OK*, then click *Apply Changes*.

For more information on creating such a policy, see "Configuring an Authentication Header Policy" in the *Novell Access Manager 3.1 SP1 Policy Management Guide*.

9 Assign this policy to the protected resources:

9a Click *Devices* > *Access Gateways* > *Edit* > *[Name of Reverse Proxy]*.

9b Click the name of the teaming proxy service, then click the *Protected Resources* tab.

9c For each teaming protected resource, click the *Identity Injection* link, select the Identity Injection policy, click *Enable*, then click *OK*.

9d Click *OK*.

9e To save the configuration changes, click *Devices* > *Access Gateways*, then click *Update*.

1.6 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 1-3 illustrates these processes.

Figure 1-3 *HTML Rewriting*

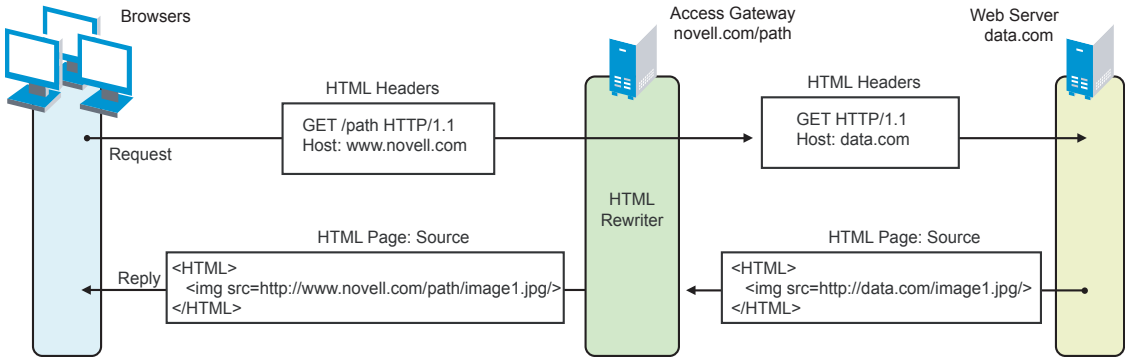

The following sections describe the HTML rewriting process:

- Section 1.6.1, "Understanding the Rewriting Process," on page 42
- Section 1.6.2, "Specifying the DNS Names to Rewrite," on page 43
- Section 1.6.3, "Defining the Requirements for the Rewriter Profile," on page 46
- Section 1.6.4, "Configuring the HTML Rewriter and Profile," on page 52
- Section 1.6.5, "Disabling the Rewriter," on page 57

1.6.1 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 this link, the browser bounces between HTTP and HTTPS to 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/imagel.jpg. Because Internet users are unable to resolve data.com/imagel.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 1-3 on page 41, 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 "Configuring the Web Servers of a Proxy Service" on page 18.

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`

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.

1.6.2 Specifying the 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" on page 44
- "Determining Whether You Need to Exclude DNS Names from Being Rewritten" on page 45

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 1-4 illustrates a scenario that requires an entry in the *Additional DNS Name List*.

Figure 1-4 *Rewriting a URLs for Web Servers*

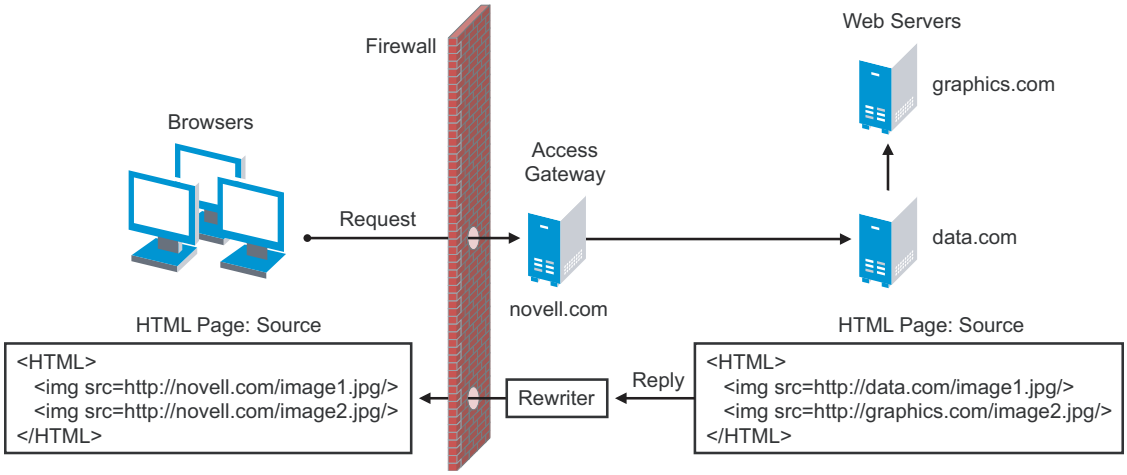

The page on the data.com Web server contains two links, one to an image on the data.com server and one to an image on the graphics.com server. The link to the data.com server is automatically rewritten to novell.com, when rewriting is enabled. The link to the image on graphics.com is not rewritten, until 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 to this additional Web server needs to rewritten.

- If the Web server listens on one port (for example, 80), and redirects the request to a secure port (for example, 443). The response to the user comes back on https://<*DNS_name*>:443. This does not match the request which was sent on http://<*DNS_name*>:80. If you add the DNS name to the list, the response can be sent in the format that the user expects.
- If an application is written to use a private hostname. For example, assume that an application URL reference contains the hostname of home (`http://home/index.html`). This hostname would need 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 acme.com.uk, the response is returned with references to acme.com.uk and not acme.com.usa. If you have a third reverse proxy protecting a Web server, the rewriting rules can become ambiguous. For example, consider the configuration illustrated in Figure 1-5.

Figure 1-5 *Excluding URLs*

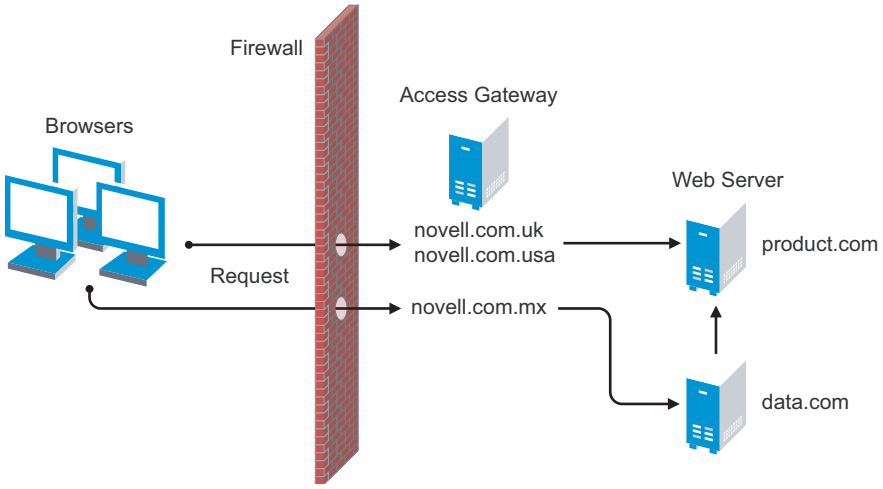

A user accesses data.com through the published DNS name of novell.com.mx. The data.com server has references to product.com. The novell.com.mx proxy has two ways to get to the product.com server because this Web server has two published DNS names (novell.com.uk and novell.com.usa). The rewriter could 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*.

1.6.3 Defining the Requirements for the Rewriter Profile

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

- "Types of Rewriter Profiles" on page 46
- "Page Matching Criteria for Rewriter Profiles" on page 47
- "Possible Actions for Rewriter Profiles" on page 48
- "String Replacement Rules for Word Profiles" on page 50
- "String Replacement Rules for Character Profiles" on page 50
- "Using $path to Rewrite Paths in JavaScript Methods or Variables" on page 51

Types of Rewriter Profiles

The Access Gateway allows you to define two types of profiles:

- "Word Profile" on page 47
- "Character Profile" on page 47

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.”

The Access Gateway has a default Word profile. It is not specific to a reverse proxy or its proxy services. When you modify its behavior, remember its scope.

If you enable HTML rewriting, but do not define a 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

If this default behavior does not match your requirements for a particular page, create your own Word profile and 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.

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

- “String Replacement Rules for Word Profiles” on page 50
- “Using $path to Rewrite Paths in JavaScript Methods or Variables” on page 51

Character Profile

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 need functionality not provided by the default profile, create a Character profile. If you create multiple Character profiles, order is important. The first Character profile that matches the page is applied. Profiles lower in the list are ignored.

For information on how strings are replaced in a Character profile, see “String Replacement Rules for Character Profiles” on page 50.

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.

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 request URL does not contain a trailing slash.
http://www.a.com/content/	Matches pages only if the request 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 request 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 that matches 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-type, the page matches if the file extension is `html`, `htm`, `shtml`, `jhtml`, `asp`, or `jsp`.

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 following options. 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 unselected 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 43. 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 58.
  - **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 58.

- **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 50
- "String Replacement Rules for Character Profiles" on page 50

For information on 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 51.

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 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 proxy 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 1-6 *Rewriting with a Multi-homing Path*

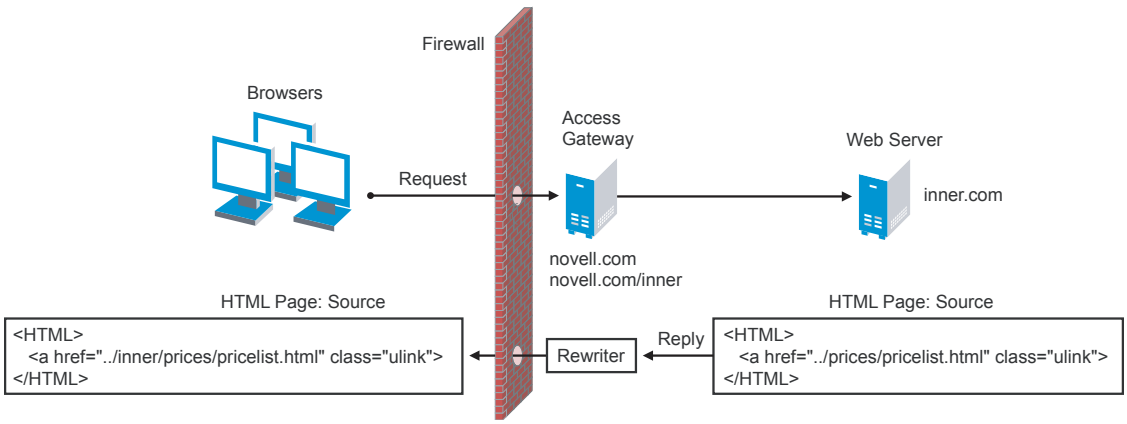

To make sure all the paths generated by JavaScript are rewritten, you must search the Web pages of the application. You can then either list all the JavaScript methods and variables in the *Additional Names to Search for URL Strings to Rewrite with Host Name* section of the rewriter profile, or you can 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.

Table 1-1 *Rewriting Strings Sent from the Browser to the Web Server*

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`

1.6.4 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*.

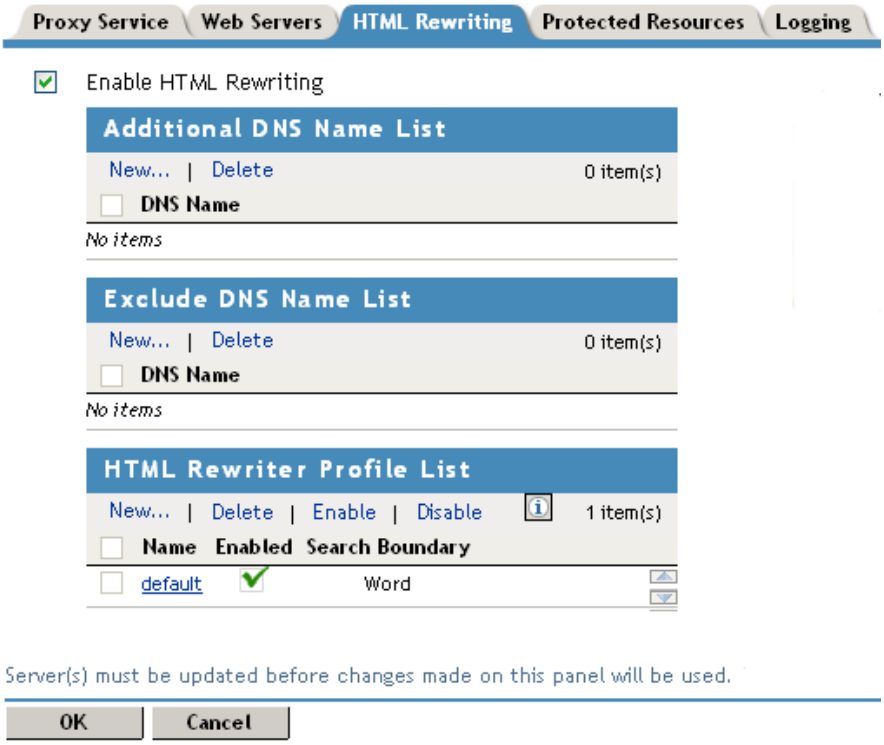

The HTML Rewriting page specifies which DNS names are to be rewritten. The HTML Rewriter Profile specifies which pages to search for DNS names that need to be rewritten.

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 "Determining Whether You Need to Specify Additional DNS Names" on page 44.

4 In the *Exclude DNS Name List* section, click *New*, specify a DNS name that appears on the Web pages of your server that you do not want rewritten, then click *OK*.

For more information, see "Determining Whether You Need to Exclude DNS Names from Being Rewritten" on page 45.

5 Use the *HTML Rewriter Profile List* to configure a profile. Select one of the following actions:

- **New:** To create a profile, click *New*. Specify a display name for the profile and select either a *Word* or *Character* for the *Search Boundary*. Continue with Step 6.
  - **Word:** A Word profile searches for matches on words. For example, “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 executed. 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 executed. Profiles lower in the list are ignored.
- **Delete:** To delete a profile, select the profile, then click *Delete*. Continue with Step 13.
- **Enable:** To enable a profile, select the profile, then click *Enable*. Continue with Step 13.
- **Disable:** To disable a profile, select the profile, then click *Disable*. Continue with Step 13.
- **Modify:** To view or modify the current configuration for a profile, click the name of the profile. Continue with Step 6.

  The default profile is designed to be applied to all pages protected by the Access Gateway. It is not specific to a reverse proxy or its proxy services. If you modify its behavior, remember its scope. Rather than modify the default profile, you should create your own customized Word profile and enable it

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

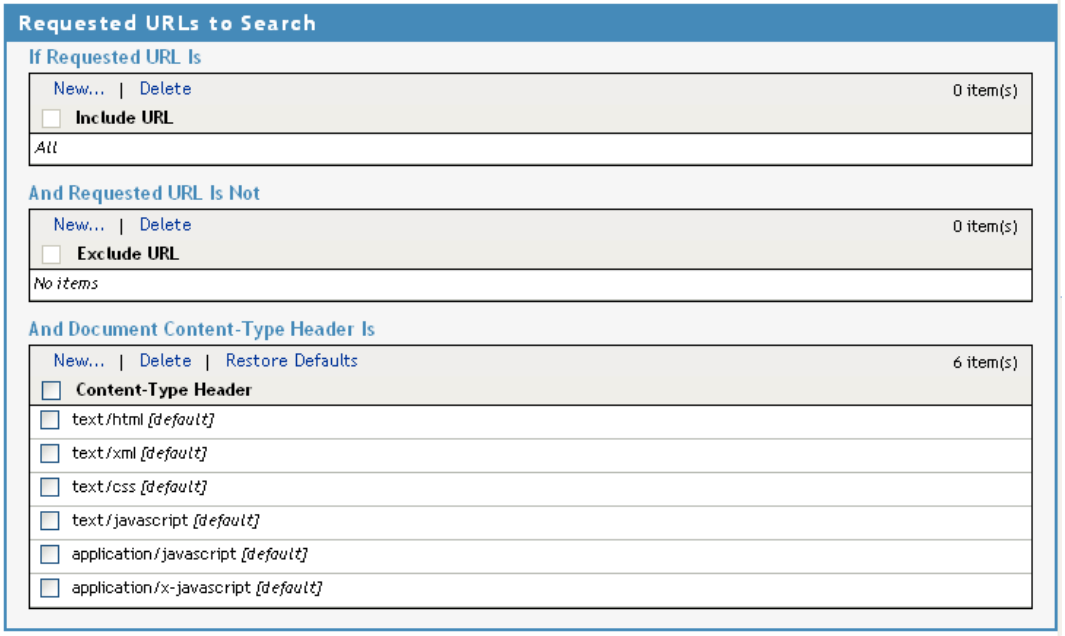

Fill in the following fields:

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, 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 on how to use these options, see "Page Matching Criteria for Rewriter Profiles" on page 47.

7 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.

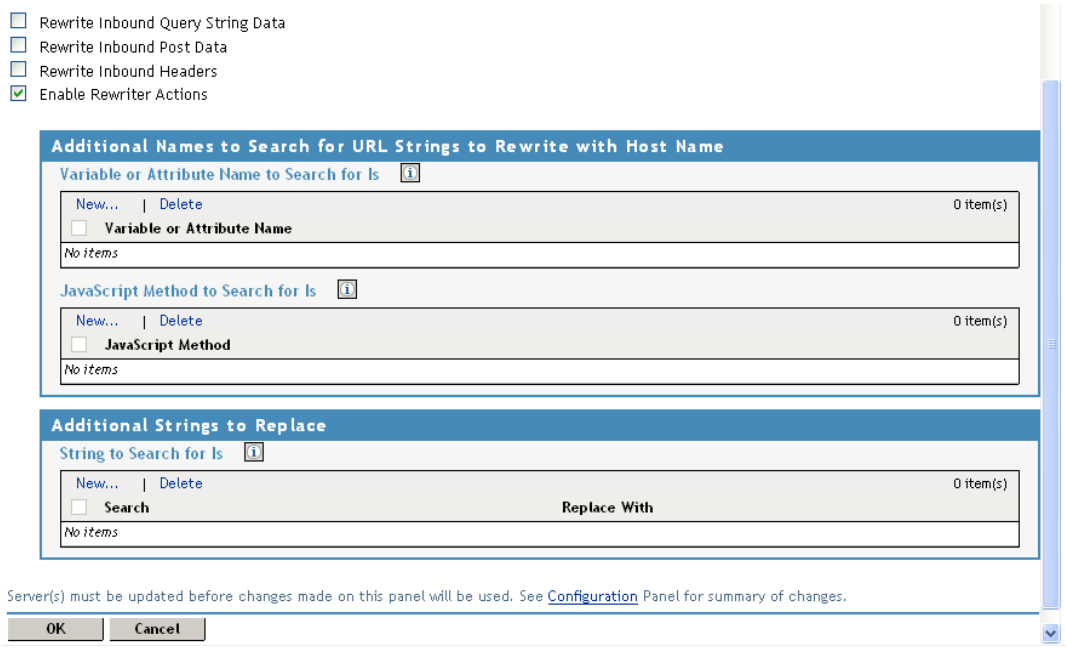

Configure the following actions:

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 48.

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 48.

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.

8 (Not available for Character profiles) If your pages contain JavaScript, use the *Additional Names to Search for URL Strings to Rewrite with Host Name* section to specify JavaScript variables or methods. You can also add HTML attribute names. (For the list of attribute names that are automatically searched, see "HTML Tags" on page 43.)

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.

9 Use the *Additional Strings to Replace* section to specify a string to search for and specify the text it should be replaced with. The search boundary (word or character) that you specified when creating the profile is used when searching for the string.

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 50
- "String Replacement Rules for Word Profiles" on page 50
- "Using $path to Rewrite Paths in JavaScript Methods or Variables" on page 51

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

10 Click *OK*.

11 If you have more than one profile in the *HTML Rewriter Profile List*, use the 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, only a maximum of one Word profile and one Character profile is executed per page. The first one in the list that matches a page is executed, and the others are ignored.

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

The default profile cannot be disabled. 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.

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

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

15 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 update them from a browser. Access the page, then press Ctrl+Shift, then click *Refresh* to force a refresh of the page.

1.6.5 Disabling the Rewriter

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

- "Disabling per Proxy Service" on page 57
- "Disabling per URL" on page 57
- "Disabling with Page Modifications" on page 58

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.

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 initial 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
```

1.7 Configuring Connection and Session Limits

The Access Gateway establishes connections with clients and with Web servers. The Identity Server establishes the session and sets the session timeout. For most networks, the default values for the connection and session limits provide adequate performance, but you can fine-tune the options to match for your network, its performance requirements, and your users:

- Section 1.7.1, "Configuring TCP Listen Options for Clients," on page 59
- Section 1.7.2, "Configuring TCP Connect Options for Web Servers," on page 60
- Section 1.7.3, "Configuring Connection and Session Persistence," on page 62
- Section 1.7.4, "Configuring the Session Timeout," on page 62

1.7.1 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*.

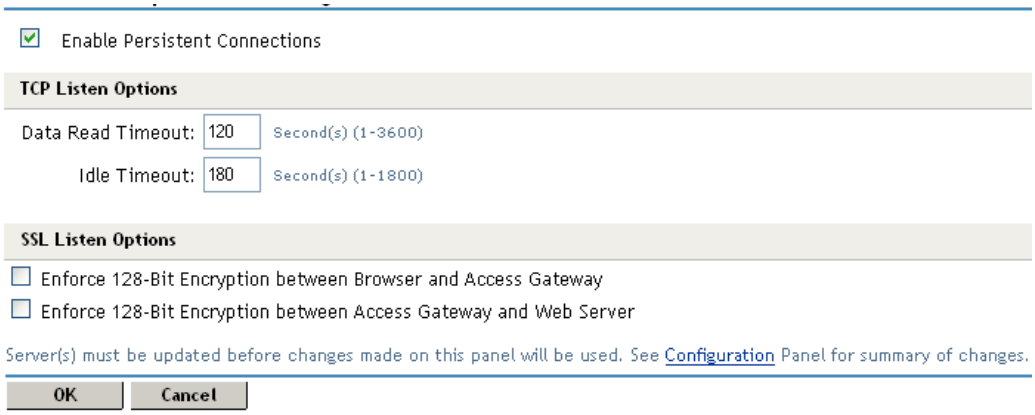

2 Select *Enable Persistent Connections* to allow the Access Gateway to establish a persistent HTTP connection between the Access Gateway and the browser. Usually, HTTP connections service only one request and response sequence. A persistent connection allows multiple requests to be serviced before the connection is closed.

   This option is enabled by default.

3 Specify values for the *TCP Listen Options*:

   Data Read Timeout: Determines when an unresponsive connection is closed. When exchanging data, if an expected response from the connected device is not received within this amount of time, the connection is closed. This value might need to be increased for slow or congested network links. The value can be set from 1 to 3600 seconds (1 hour). The default is 120 seconds (2 minutes).

   Idle Timeout: Determines when an idle connection is closed. If no application data is exchanged over a connection for this amount of time, the connection is closed. This value limits how long an idle persistent connection is kept open. This setting is a compromise between freeing resources to allow additional inbound connections, and keeping connections established so that new connections from the same device do not need to be re-established. The value can be set from 1 to 1800 seconds (30 minutes). The default is 180 seconds (3 minutes).

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.

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

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

1.7.2 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.

Make Outbound Connection Using: Specifies which IP address the proxy service should use when establishing connections with the back-end Web servers.

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: 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*.

1.7.3 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 to the browser connections and the Access Gateway to the Web server connections involve setting up a TCP connection for an HTTP request. HTTP connections usually service only one request and response sequence, and the TCP connection is opened and closed during the sequence. A persistent connection allows multiple requests to be serviced before the connection is closed and saves a significant amount of processing time. To configure this type of persistence, see the following:

- **Access Gateway to Browser:** Click *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.

1.7.4 Configuring the Session Timeout

When a user logs in and authenticates to the Identity Server, the Identity Server establishes a session for the user and sets an inactivity timeout for the session. If the user's session becomes inactive and reaches this time limit, the session becomes invalid. If the user tries to access a resource from an invalid session, the user is prompted to log in again.

The session timeout is a global value, affecting all users who authenticate to the Identity Server and all resources protected by Access Manager. The default value for the session timeout is 60 minutes.

To modify this value:

1. In the Administration Console, click *Devices > Identity Servers > Edit*.
2. For the *Session timeout* option, use the up-arrow button to increase the timeout and the down-arrow button to decrease the timeout.
3. Click *OK*, then update the Identity Server.

Configuring the Access Gateway for SSL

2

SSL provides the following security features:

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

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

To ensure the validity of X.509 certificates, Access Manager supports both Certificate Revocation Lists (CRLs) and Online Certificate Status Protocol (OCSP) methods of verification.

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

2.1 Using SSL on the Access Gateway Communication Channels

You can configure the Access Gateway to use SSL in its connections to the Identity Server, to the browsers, and to its Web servers. Figure 2-1 illustrates these communication channels.

Figure 2-1 *Setting Up SSL for the Access Gateway Communication Channels*

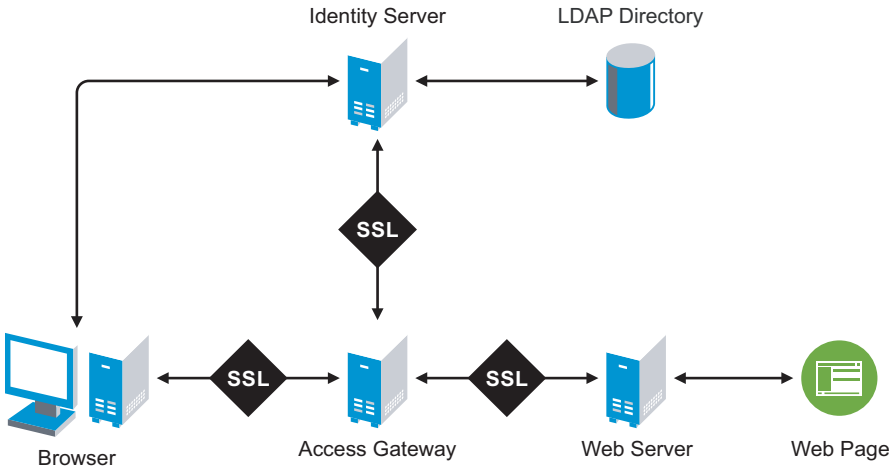

This section only describes how to set up SSL for the Access Gateway communication channels. The Identity Server needs to be configured for SSL before the Access Gateway can be configured for SSL. See "Configuring Secure Communication on the Identity Server" in the *Novell Access Manager 3.1 SP1 Setup Guide*.

When the user logs in to the Identity Server, the Identity Server verifies the user's credentials, usually with the credentials stored in an LDAP directory, but other methods are available. If the login is successful, the Identity Server sends an artifact to the browser, and the browser forwards it to the Access Gateway. The Access Gateway uses the artifact to retrieve the user's name and password from the Identity Server. The Access Gateway and Identity Server channel is probably the first communication channel you should enable for SSL. The Access Gateway uses an Embedded Service Provider to communicate with the Identity Server. When you enable SSL between the two, the Access Manager distributes the necessary certificates to set up SSL. However, if you have configured the Identity Server to use certificates from an external 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.

2.2 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 *Novell Access Manager 3.1 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 *Novell Access Manager 3.1 SP1 Administration Console Guide*.

2.2.1 Prerequisite for SSL Communication between the Identity Server and the Access Gateway

If you are going to set up SSL communication between the Identity Server and the Access Gateway for authentication and you have configured the Identity Server to use certificates created by an external CA, you need to import the public certificate of this CA into the trusted root keystore of the Access Gateway.

1. If you haven't already imported the public certificate of this CA into the trusted root store of the Identity Server, do so now. For instructions, see "Importing Public Key Certificates (Trusted Roots)" in the *Novell Access Manager 3.1 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 on the Access Gateways page, click *Update*.

2.2.2 Prerequisites for SSL Communication between the Access Gateway and the Web Servers

If you are going to set up SSL between the Access Gateway and the Web servers, you need to configure your Web servers for SSL. Your Web servers must supply a certificate that clients (in this case, the Access Gateway) can import. See your Web server documentation for information on how to configure the Web server for SSL.

For mutual SSL, the proxy service must supply a certificate that the Web server can trust. This certificate can be the same one you use for SSL between the browsers and the reverse proxy.

2.3 Configuring SSL Communication with the Browsers and the Identity Server

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

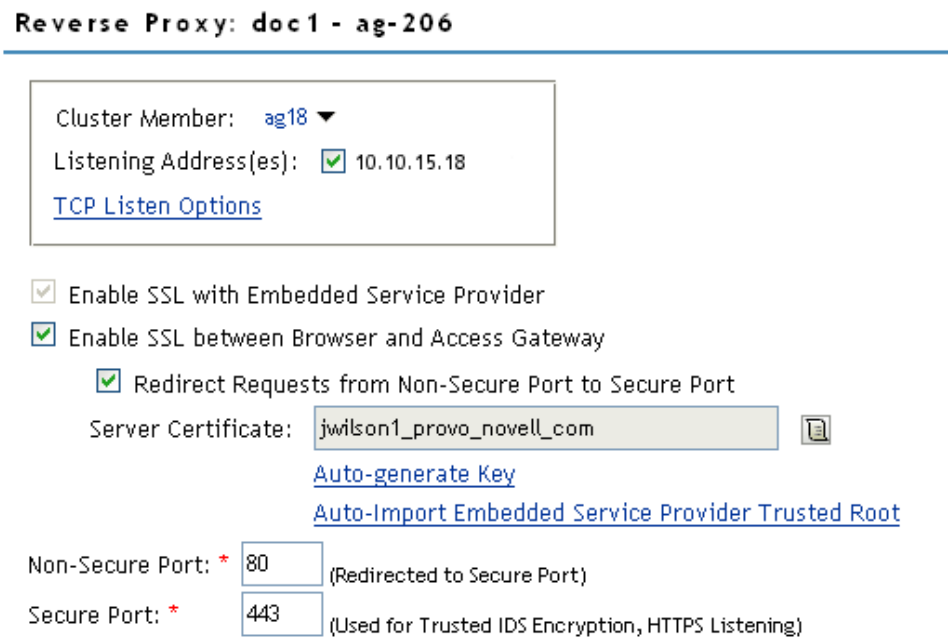

2 Configure the reverse proxy for SSL. Fill in the following fields:

Enable SSL with Embedded Service Provider: Select this option to encrypt the data exchanged for authentication (the communication channel between the Identity Server and the Access Gateway). This option is only available for the reverse proxy 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 to require SSL connections between your clients and the Access Gateway. SSL must be configured between the browsers and the Access Gateway before you can configure SSL between the Access Gateway and the Web servers.

Redirect Requests from Non-Secure Port to Secure Port: Determines whether browsers are redirected to the Secure Port and allowed to establish an SSL connection. If this option is not selected, browsers that connect to the non-secure port are denied service.

This option is only available if you have selected *Enable SSL with Embedded Service Provider*.

3 Select the certificate to use for SSL between the Access Gateway and the browsers. Select one of the following methods:

- To auto-generate a certificate key 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 (Conditional) If you selected a certificate in Step 3 that was created by an external CA, click *Auto-Import Embedded Service Provider Trusted Root*, click *OK*, specify an alias name, click *OK*, then click *Close*.

This option imports the public key from the Embedded Service Provider into the trust store of the Identity Servers in the selected Identity Server Configuration. This sets up a trusted SSL relationship between the Identity Server and the Embedded Service Provider.

If you are using certificates signed by the Novell Access Manager CA, the public key is automatically added to this trust store.

5 Configure the ports for SSL:

Non-Secure Port: Specifies the port on which to listen for HTTP requests. The default port for HTTP is 80.

- If you 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.

6 Click *OK*.

7 On the *Configuration* page, click *Reverse Proxy / Authentication*.

8 (Conditional) If you are using an externally signed certificate for the Identity Server cluster, you need to import the public key of the CA:

8a In the *Embedded Service Provider* section, click *Auto-Import Identity Server Trusted Root*, then click *OK*

8b Specify an alias, click *OK* twice, then click *Close*.

This option imports the public key of the Identity Server into the trust store of the Embedded Service Provider. This sets up a trusted SSL relationship between the Embedded Service Provider and the Identity Server.

The configCA public key certificate of the Access Manager CA is automatically added to the ESP Trust Store. If you are using Access Manager CA certificates for the Identity Server, you do not need to import the configCA certificate unless someone has deleted it from this trust store.

9 Click *OK*.

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

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

The Embedded Service Provider is restarted during the update.

12 Update the Identity Server so that it uses the new SSL configuration. Click *Identity Servers* > *Update*.

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

13a Enter the URL to a protected resource on the Access Gateway.

13b 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 on solving this problem, see "Troubleshooting 100101043 and 100101044 Liberty Metadata Load Errors" in the *Novell Access Manager 3.1 SP1 Identity Server Guide*.

2.4 Configuring SSL between the Proxy Service and the Web Servers

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

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

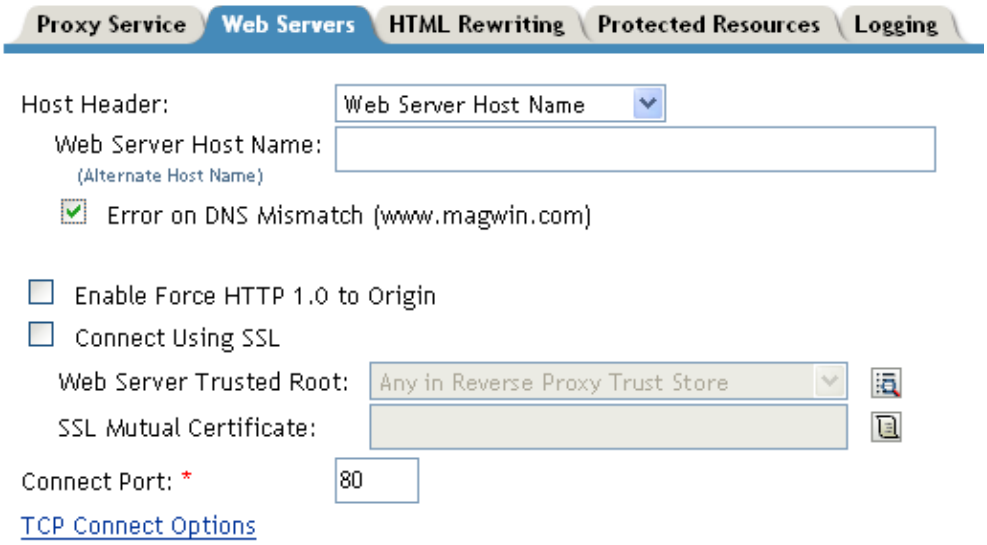

2 To configure SSL, select *Connect Using SSL*.

This option is not available if you have not set up SSL between the browsers and the Access Gateway. See Section 2.3, "Configuring SSL Communication with the Browsers and the Identity Server," on page 66 and select the *Enable SSL between Browser and Access Gateway* field.

3 Configure how you want the proxy service to verify the Web server certificate:

3a Select one of the following options:

- To not verify this certificate, select *Do not verify* for the *Web Server Trusted Root* option.

  Use this option when you want the information between the Access Gateway and the Web server encrypted, but you don't need the added security of verifying the Web server certificate.

  Continue with Step 4.

- To verify the certificate authority of the Web server certificate, select *Any in Reverse Proxy Trust Store*. When this option is selected, the public certificate of the certificate authority must be added to the proxy trust store.

  Click the *Manage Reverse Proxy Trust Store* icon. The auto import screen appears.

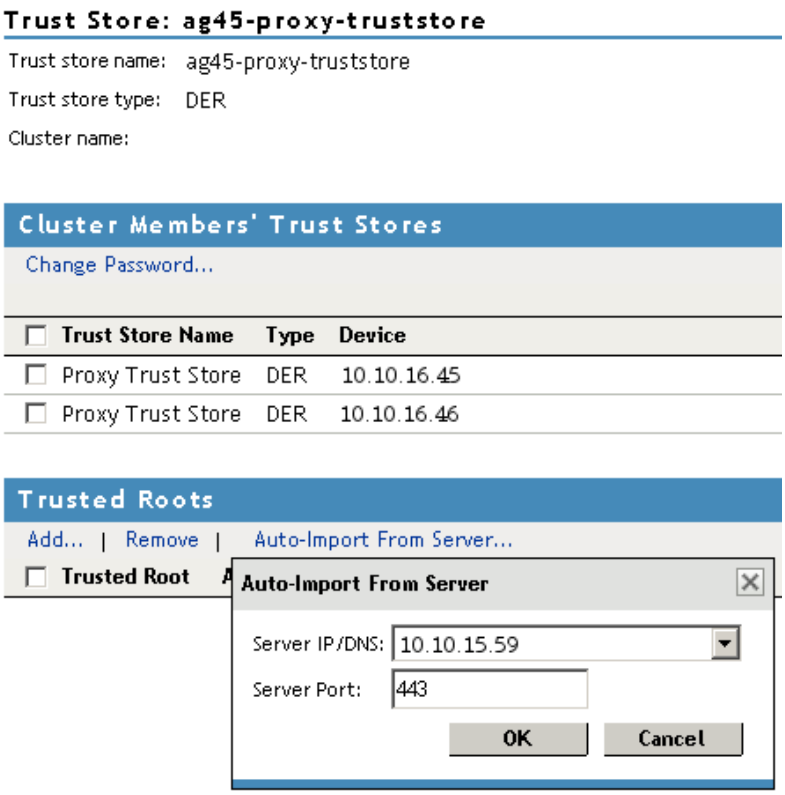

If the Access Gateway is a member of a cluster, the cluster members are listed. The Web server certificate is imported into the trust stores of each cluster member.

3b Ensure that the IP address of the Web server and the port match your Web server configuration.

If these values are wrong, you have entered them incorrectly on the Web server page. Click *Cancel* and reconfigure them before continuing.

3c Click *OK*.

The server certificate, the Root CA certificate, and any certificate authority (CA) certificates from a chain are listed.

If the whole chain is not displayed, import what is displayed. You then need to manually import the missing parents in the chain. A parent is missing if the chain does not include a certificate where the Subject and the Issuer have the same CN.

3d Specify an alias, then click *OK*.

All the certificates displayed are added to the trust store.

3e Click *Close*.

4 (Optional) To set up mutual authentication so that the Web server can verify the proxy service certificate:

4a Click the *Select Certificate* icon,

4b 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.

5 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

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

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

2.5 Enabling Secure Cookies

The Embedded Service Provider of the Access Gateway and 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 2.5.1, "Securing the Embedded Service Provider Session Cookie," on page 71
- Section 2.5.2, "Securing the Proxy Session Cookie," on page 72

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)

2.5.1 Securing the Embedded Service Provider Session Cookie

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 from happening, you must first configure Access Gateway to use SSL. See Section 2.3, "Configuring SSL Communication with the Browsers and the Identity Server," on page 66. After you have SSL configured, create a touch file as follows:

1 On the Linux Access Gateway Appliance, log in as `root`.

2 Specify the following command to create the `.setsecureESP` file:

touch `/var/novell/.setsecureESP`

3 Specify the following command to restart Linux Access Gateway:

```
/etc/init.d/novell-vmc stop
/etc/init.d/novell-vmc start
```

2.5.2 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 from happening, you can use the following configuration options:

- **Set 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. Enable the *Force 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

- **Prevent 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.

- **Prevent the browser from sending cookies on a non-HTTPS channel:** You can configure the Linux Access Gateway Appliance to set its authentication cookie with the secure keyword in order to prevent the browser from sending this cookie on a non-HTTPS channel. To enable this, use the following touch file:

  ```
  /var/novell/.EnableSecureCookie
  ```

  This file works when the *Force Secure Cookie* option is disabled in the Administration Console.

NOTE: This works only for HTTPS services. When this setting is enabled, you cannot configure the Access Gateway to have an HTTP service that requires authentication, or create a policy that depends on the authentication cookie.

2.6 Managing Access Gateway Certificates

2.6.1 Managing Embedded Service Provider Certificates

The Access Gateway uses an Embedded Service Provider to communicate with the Identity Server. The Service Provider Certificates page allows you to view the private keys, certificate authority (CA) certificates, and certificate containers associated with this module. These keystores do not contain the certificates that the Access Gateway uses for SSL connections to browsers or to back-end Web servers.

To view or modify these certificates:

1. In the Administration Console, click *Devices* > *Access Gateways* > *Edit* > *Service Provider Certificates*.
2. Configure the following:

   Signing: The signing certificate keystore. Click this link to access the keystore and replace the signing certificate as necessary. The signing certificate is used to sign the assertion or specific parts of the assertion.

   Trusted Roots: The trusted root certificate container for the CA certificates associated with the Access Gateway. Click this link to access the trust store, where you can change the password or add trusted roots to the container.

   The Embedded Service Provider must trust the certificate of the Identity Server that the Access Gateway has been configured to trust. The public certificate of the CA that generated the Identity Server certificate must be in this trust store. If you configured the Identity Server to use a certificate generated by a CA other than the Access Manager CA, you must add the public certificate of this CA to the Trusted Roots store. To import this certificate, click *Trusted Roots*, then in the *Trusted Roots* section, click *Auto-Import From Server*. Fill in the IP address or DNS name of your Identity Server and its port, then click *OK*.

   You can also auto import the Identity Server certificate by select the *Auto-Import Identity Server Configuration Trusted Root* option on the *Reverse Proxies / Authentication* page (click *Devices* > *Access Gateways* > *Edit* > *Reverse Proxies / Authentication*). With this option, you do not need to specify the IP address and port of the Identity Server.
3. To save your changes to browser cache, click *OK*.
4. To apply your changes, click the *Access Gateways* link, then click *Update* > *OK*.

2.6.2 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 configuring 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.

Server Configuration Settings

3

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 Section 4.2, “Configuring Proxy Service Logging,” on page 110 and Section 4.5, “Enabling Access Gateway Audit Events,” on page 133.

For cache management, see Chapter 5, “Configuring the Content Settings,” on page 141.

3.1 Viewing and Updating the Configuration Status

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

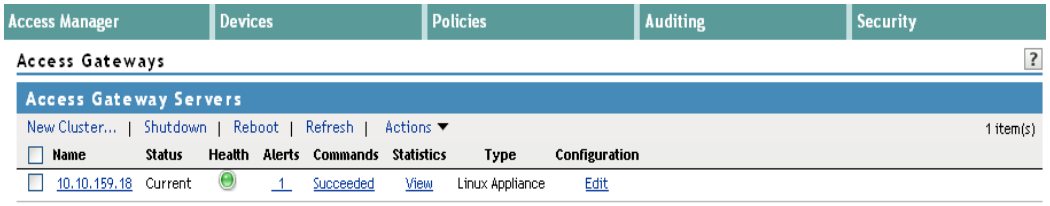

2 View the *Status* column.

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. Click this link to apply the changes. • **All Configuration:** You can select to have the server read its complete configuration file. Depending upon what has been modified, updating the complete configuration might cause logged-in users to lose data and connections. • **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 that the server has enabled for a protected resource 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.
Update !	If the configuration update contains a configuration error, the *Update* option 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	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 that might interrupt service, we recommend that you update the servers one at a time. 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 6.4.3, “Applying Changes to Cluster Members,” on page 168.
Update All !	If the configuration update contains a configuration error, the *Update All* and the member *Update* options 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.

Status	Description
Pending	Indicates that the server is processing a configuration change, but has not completed the process.

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 (*Devices* > *Access Gateways* > *Edit*) allows you to control how your changes are saved so they can be applied with the update options (see Section 3.1, "Viewing and Updating the Configuration Status," on page 75).

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* on the Configuration page, 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* on the Configuration page, 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* on the Configuration page, 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* on the Configuration page, 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 cancel 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 cancel 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 cancel 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 Starting and Stopping the Access Gateway

The Access Gateway has two processes that can be stopped and started: the Access Gateway and the Embedded Service Provider within the Access Gateway. Normally, you do not need to stop and start these services. However, if you need to change certain configuration options, you can be prompted to update the Access Gateway or to restart the Embedded Service Provider.

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

- Section 3.3.1, “Updating the Access Gateway,” on page 79
- Section 3.3.2, “Restarting the Access Gateway Service Provider,” on page 79
- Section 3.3.3, “Starting the Access Gateway Service Provider,” on page 80
- Section 3.3.4, “Stopping the Access Gateway Service Provider,” on page 80
- Section 3.3.5, “Restarting the Access Gateway Appliance,” on page 80
- Section 3.3.6, “Stopping the Access Gateway Appliance,” on page 81

3.3.1 Updating the Access Gateway

When a configuration change has been made, but not applied, the Access Gateway is in an *Update* status on the Access Gateways page. If the Access Gateway is a member of a cluster, the cluster is in an *Update All* status. You can click *Update* to apply the configuration change to a single Access Gateway or *Update All* to apply the configuration change to all members of a cluster.

If the changes have been saved to browser cache, but not to the configuration store, the changes are lost if your session times out before you apply the changes. The Access Gateway remains in an *Update* status, but when you click *Update*, there are no changes to apply. If you prefer to update members of a cluster one at a time, it is best to save the changes to the configuration datastore before applying them. Click *Edit*, then click *Save*.

When you click *Update*, three options are displayed:

- When you have modified services of the Access Gateway, the update option for *All Configuration* is available. Depending upon what has been modified, updating might cause logged in users to lose data and their connections.
- 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.
- If a policy is modified that the server has enabled for a protected resource or a protected resource has a policy enabled or disabled and the policy changes are the only modifications that have occurred, the update option for *Policy Settings* is available. The Policy Settings option causes no interruption in services.

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. To access this option, click *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. To access this option, click *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. To access this option, click *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. To access this option, click *Access Gateways > Edit > Reverse Proxy/Authentication > [Name of Reverse Proxy] > [Name of First Proxy Service]*, then modify the *Published DNS Name* option.

3.3.2 Restarting the Access Gateway Service Provider

To stop and start the Access Gateway service provider:

1. In the Administration Console, click *Access Manager > Access Gateways*, select the Access Gateway, then click *Actions*.
2. Click *Service Provider > Restart Service Provider*, then click *OK*.

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

3.3.3 Starting the Access Gateway Service Provider

When an Access Gateway is removed from a cluster configuration, the Embedded Service Provider is stopped. It should remain stopped until you have reconfigured the Access Gateway. When you have finished the reconfiguration, you should start the Embedded Service Provider.

1. In the Administration Console, click *Devices* > *Access Gateways*, select the Access Gateway, then click *Actions*.
2. Click *Service Provider* > *Start Service Provider*, then click *OK*.

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

3.3.4 Stopping the Access Gateway Service Provider

Stopping the Embedded Service Provider is a quick way to make the Access Gateway inaccessible to users.

1. In the Administration Console, click *Devices* > *Access Gateways*, select the Access Gateway, then click *Actions*.
2. Click *Service Provider* > *Stop Service Provider*, then click *OK*.

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

3.3.5 Restarting the Access Gateway Appliance

For a Gateway Appliance, the Restart option is really a reboot option. The Access Gateway is stopped, the operating system is rebooted, then the Access Gateway is started.

- "Immediately Rebooting the Gateway Appliance" on page 80
- "Scheduling a Reboot of the Gateway Appliance" on page 80

Immediately Rebooting the Gateway Appliance

1. In the Administration Console, click *Devices* > *Access Gateways*, select the Access Gateway.
2. Click *Restart*.

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

Scheduling a Reboot of the Gateway Appliance

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

1. In the Administration Console, click *Devices* > *Access Gateways*, select the Access Gateway, then click *Actions*.
2. Click *Schedule Reboot*.

   The following field displays information about the command you are scheduling.

   Type: Displays the type of command that is being scheduled, such as *Access Gateway Shutdown, Access Gateway Reboot, Access Gateway Upgrade, Device Configuration*.
3. Fill in the following fields:

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

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

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

4 Click *OK*.

3.3.6 Stopping the Access Gateway Appliance

You should stop the Access Gateway Appliance only when you plan to turn off the power. After you have stopped the Access Gateway Appliance, you must have physical access to the machine to start it.

- "Immediately Stopping the Gateway Appliance" on page 81
- "Scheduling the Shutdown of the Gateway Appliance" on page 81

Immediately Stopping the Gateway Appliance

1 In the Administration Console, click *Devices* > *Access Gateways*, select the Access Gateway, then click *Stop*.

2 To confirm the shutdown, click *OK*.

The machine is physically turned off.

Scheduling the Shutdown of the Gateway Appliance

Scheduling a shutdown allows you to pick the best time for the Access Gateway to be unavailable.

1 In the Administration Console, click *Devices* > *Access Gateways*, select the Access Gateway, then click *Actions*.

2 Click *Schedule Shutdown*.

The type field displays information about the command you are scheduling, such as *Access Gateway Shutdown, Access Gateway Restart, Access Gateway Upgrade, Device Configuration*

3 Fill in the following fields:

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

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

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

4 Click *OK*.

The machine is turned off when the scheduled command executes.

3.4 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 identity one Access Gateway from another.

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

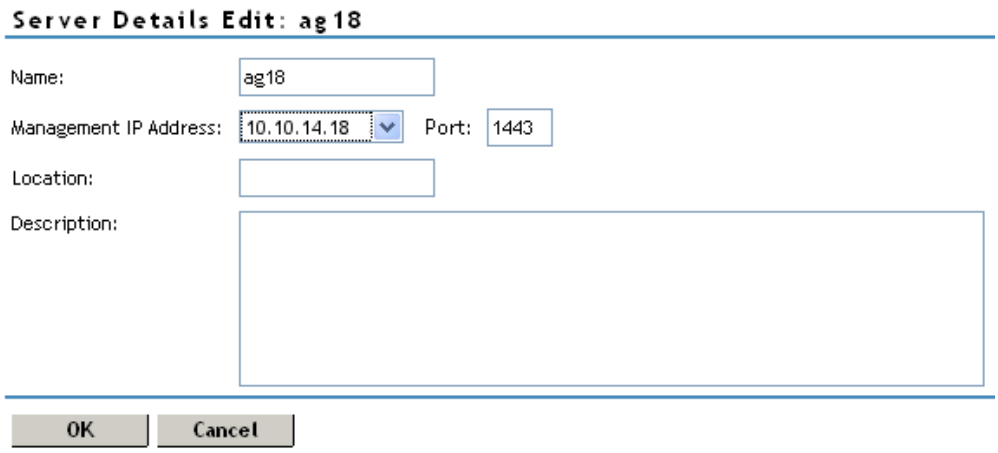

2 Modify the values in the following fields:

Name: Specifies 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: Specifies the IP address used to manage the Access Gateway. Select an IP address from the list. For information on changing the *Management IP Address*, see "Changing the IP Address of the Access Gateway Appliance" in the *Novell Access Manager 3.1 SP1 Administration Console Guide*.

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

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

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

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

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

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*.

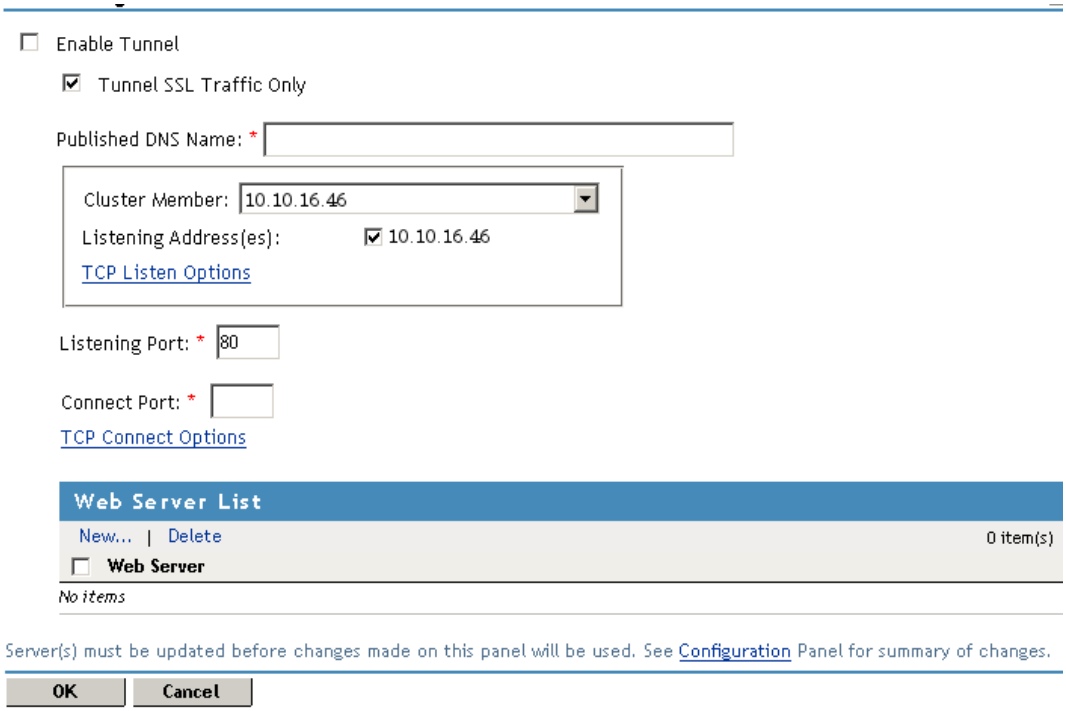

3. Fill in the following fields:

   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 1.7.1, "Configuring TCP Listen Options for Clients," on page 59. 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 1.7.2, "Configuring TCP Connect Options for Web Servers," on page 60.

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*.

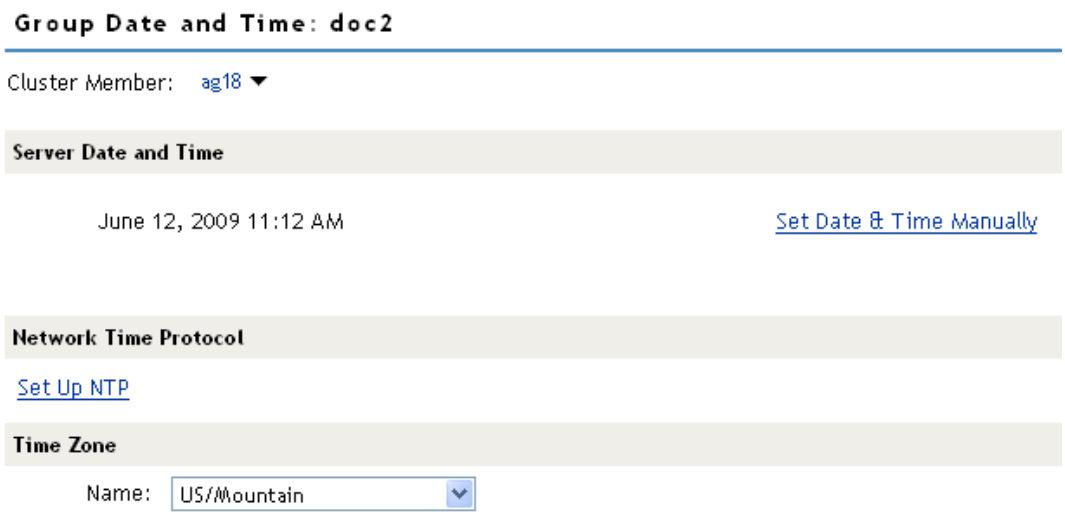

2 (Conditional) If the Access Gateway belongs to a cluster of Access Gateways, select the Access Gateway from the list displayed in the *Cluster Member* field. The modifications you make on this page apply only to the selected Access Gateway.

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

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

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

3.7 Customizing Error Pages on the Gateway Appliance

The Access Gateway Appliance uses the custom error page template to rebrand and localize the language of error pages that are published to the browser.

By default, the Gateway Appliance contains the following files to help customize and localize the error messages:

- The error page configuration file, `ErrorPagesConfig.xml`
- The error page template file, `ErrorPageTemplate.htm.en`
- The error messages file, `ErrorMessages.xml.en`

NOTE: If you are modifying any of the above files, ensure that you retain the original filenames.

The Access Gateway Appliance maintains three directories to save files that are used for error page configuration:

```
/var/novell/errorpagesconfig/.factory
/var/novell/errorpagesconfig/.backup
/var/novell/errorpagesconfig/current
```

During the initial installation, the default template files are copied to the `.factory` and the `current` directories. If you have not customized the files in the `current` directory, subsequent installations do not overwrite these files.

When the next version of the files is installed, the files in the `current` directory are copied to the `.backup` directory with the format `<filename>.oldBuildNo`. This ensures that the old build files and customized files are always available in the `.backup` directory.

You can customize and localize the error template and the error messages:

- Section 3.7.1, “Customizing the Error Pages by Using the Default Template,” on page 86
- Section 3.7.2, “Customizing and Localizing Error Messages,” on page 88

3.7.1 Customizing the Error Pages by Using the Default Template

To customize the default error page template, you must edit the `ErrorPageTemplate.htm.en` file as follows:

1. Log in to the Access Gateway Appliance machine.
2. Back up the `ErrorPageTemplate.htm.en` file in the `/var/novell/errorpagesconfig/current` directory.
3. Open the `ErrorPageTemplate.htm.en` file located in the `/var/novell/errorpagesconfig/current` directory.

   A sample error page template looks similar to the following:

```
<html>
   <head><title>Information Alert</title></head>
   <body bgcolor="white">
   <div align="center">
   <center>
  <table border="0" cellpadding="2" frame height="199" style="margin-top:
1px; margin-bottom: 1px; padding-top: 1px; padding-bottom: -1px">
         <tr>
           <td height="34" align="center"><font color="black" face="Arial
Bold" size="4"><b><p align="center"></b></font>
```

```
          <font face="Intrepid" size="6" color="#000080">
<strong>Information Alert </strong></font>
            </td>
          </tr>
          <tr>
                <td height="20" align="center"><img height="8" width="445"
src="<PROXY_ADDRESS>/images/Odyssey_LoginHead.gif"></td>
              </tr>
              <tr>
    <td height="24" width="444" bgcolor="white" align="center">
       <p align="left">
         <b><br><font color="black" face="Comic Sans MS">Status</font></b>
            <font color="#ff0033" face="Comic Sans MS"><b>: </b></font>
           <font color="black" face="Comic Sans MS"><ERROR_STATUS> </font>
       </p>
       <p align="left">
          <font color="black" face="Comic Sans MS"><b>Description</b></
font>
          <font color="#ff0033" face="Comic Sans MS"><b>: </b></font>
          <font color="black" face="Comic Sans MS"><ERROR_DESCRIPTION></
font>
       </p>    <br>    <br>
     </font></td>
      </tr>
     <tr><td width="444" height="10" align="center"><img height="8"
width="445" src="<PROXY_ADDRESS>/images/
LAP_interoperable_logo_100.gif"></td></tr>
</table>
</center>
</div>
</body>
</html>
```

4 Modify the error page template. You can edit the default template to customize the user interface, to modify embedded images, and to provide localization. However, `<ERROR_STATUS>` and `<ERROR_DESCRIPTION>` tags should not be removed because the following actions take place when the error page is served to the browser:

- `<ERROR_STATUS>`: When the error page is served to the browser, `<ERROR_STATUS>` is replaced with the HTTP status code description.
- `<ERROR_DESCRIPTION>`: When the error page is served to the browser, `<ERROR_DESCRIPTION>` is replaced with the detailed error description.

5 If you have changed the file to use a new image:

- All the images must be linked to the *<PROXY_ADDRESS>*`/images/` directory.
- All the images must be copied to Tomcat in the path `/var/opt/novell/tomcat5/webapps/LAGERROR/images`.

If you have changed an image but retained the filename, press Ctrl+F5 in the browser to refresh the Access Gateway cache.

6 Save the file.

7 Enter the following commands to restart the machine:

```
/etc/init.d/novell-vmc stop
/etc/init.d/novell-vmc start
```

8 If the Access Gateway belongs to a cluster, copy the modified file and images to each member in the cluster, then restart that member.

3.7.2 Customizing and Localizing Error Messages

When the Access Gateway Appliance 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 88.
- Modify the `ErrorPagesConfig.xml` file with the header value and the template mapping information. For more information, see "Modifying the ErrorPagesConfig.xml File" on page 89.

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 Enter the following commands to restart the Linux Access Gateway:

```
/etc/init.d/novell-vmc stop
/etc/init.d/novell-vmc start
```

6 If the Access Gateway belongs to a cluster, copy the modified file to each member of the cluster, then restart that member.

Modifying the ErrorPagesConfig.xml File

The `ErrorPagesConfig.xml` file stores the header value and the template mapping information. You must edit the `ErrorPagesConfig.xml` file to provide localization for error messages in various languages. In the `ErrorPagesConfig.xml` file, each `<Profile>` element corresponds to a template file `ErrorPageTemplate.htm.<lang>` and a messages file `ErrorMessages.xml.<lang>`, where <*lang*> is the `fileXn` attribute value. For example, if the `fileXn` attribute value is `de`, the `ErrorPagesTemplate.htm.de` file is served to the browser.

To map a list of `Accept-Language` header values to the template, you must add the header value as the `<header>` element under the corresponding `<Profile>` element.

To modify the `errorpagesconfig.xml` file:

1 Log in to the Access Gateway Appliance machine.

2 Open the `ErrorPagesConfig.xml` file located in the `/var/novell`/errorpagesconfig/current directory.

3 Add the language information within the `<profile>` tag as follows:

```
<ErrorPageConfiguration>
      <Profile name = "English"  enable = "1"  fileXn = "en">
               <header value = "en-us" />
                <header value = "en-uk" />
               <header value = "en-any" />
               <header value = "any" />
       </Profile>
       <Profile name = "German" enable = "1" fileXn = "de">
               <header value = "de-CH" />
               <header value = "de-any" />
       </Profile>
</ErrorPageConfiguration>
```

This file serves the error messages from:

- The English profile, if the header value is en-us or en-uk or en-*
- The German profile, if the header value is de-CH or de-*
- The default profile, if the header value is not any of the above, or if it is defined as `any`.

When the header value is defined as `any`, the default profile is served. This profile matches any header value that did not have a matching profile. For example, if the header value entry is `en-any`, and the `Accept-Language` header value of the browser is `en-xyz` (for which there is no proper match), then the profile with the entry `en-any` would be a match.

If `any` is used to search for any language-specific files, then the word `any` must be preceded by the hyphen (-). For example, you must not specify `en-cany` as the header value entry to match en-c* header values.

4 Save the file.

5 Enter the following commands to restart the machine:

```
/etc/init.d/novell-vmc stop
/etc/init.d/novell-vmc start
```

6 If the Access Gateway belongs to a cluster, copy the modified file to each member of the cluster, then restart that member.

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.

This section describes the following tasks:

- Section 3.8.1, "Viewing and Modifying Adapter Settings," on page 90
- Section 3.8.2, "Viewing and Modifying Gateway Settings," on page 92
- Section 3.8.3, "Viewing and Modifying DNS Settings," on page 94
- Section 3.8.4, "Configuring Hosts," on page 96
- Section 3.8.5, "Adding New Network Interfaces to the Gateway Appliance," on page 97

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 machine and manage the IP addresses that are assigned to them.

- If you want to configure an adapter to use more than one IP address, you can use these settings to add them.
- If you have multiple adapters installed on a Gateway Appliance machine, you can only configure eth0 during installation. Use the procedure described in this section to configure the others.
- If you have added an adapter to the machine after installing the Access Gateway, you need to use the *New NIC* option before it can appear in the adapter list. See Section 3.8.5, "Adding New Network Interfaces to the Gateway Appliance," on page 97.

To view or modify your current adapter settings:

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

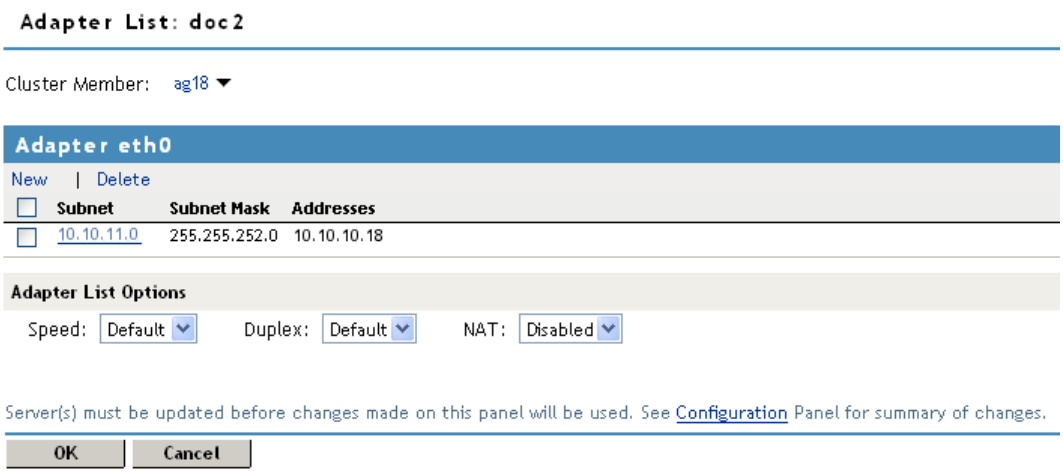

2 (Conditional) If the Access Gateway is a member of a cluster, select the server you want to configure from the list of servers in the *Cluster Member* field. All changes made to this page apply to the selected server.

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:

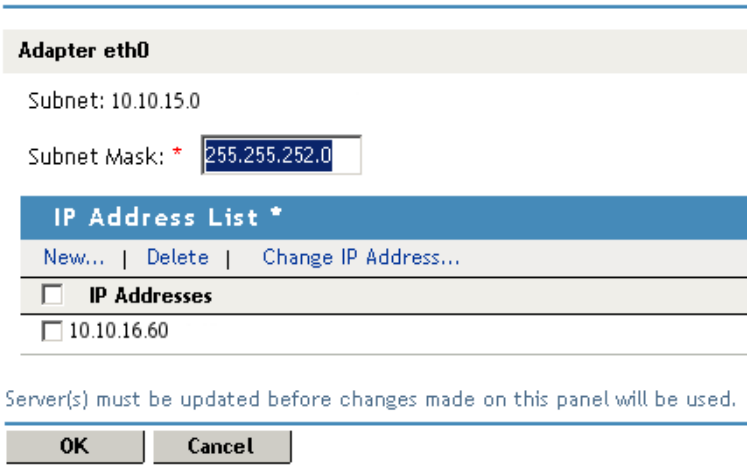

Subnet: Displays the address of the subnet that you are modifying. This is empty if you are creating a new one.

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, see "Changing the IP Address of the Access Gateway Appliance" in the *Novell Access Manager 3.1 SP1 Administration Console Guide*.

5 Click *OK*.

6 Configure the *Adapter List Options*.

These options let you change settings for the network adapters on the Access Gateway to ensure compatibility with an existing LAN. Modify the default settings only if your LAN requires specialized adapter card changes.

- **Speed:** Select *Default*, *10 MB*, *100 MB*, or *1000 MB*.
- **Duplex:** Select *Default*, *Half*, or *Full*.

IMPORTANT: Some network adapter drivers do not correctly detect duplex settings. This is a general industry problem with Fast Ethernet technology.

If your Access Gateway isn't performing as expected, check to ensure that the duplex settings for its network adapters match your network configuration. It might be necessary to manually configure the duplex settings on both your Access Gateway and your Ethernet switch or hub.

- **NAT:** Select *Dynamic* or *Disabled*.

  If the Access Gateway is serving as a router, and your network employs non-unique private IP addresses, you can configure the Access Gateway to provide Network Address Translation (NAT) services.

  For example, if you have a 10.0.0.0 private network on eth0 and a registered public network such as 130.0.0.0 on eth1, the clients on the private network can access the Internet through the Access Gateway, provided that the *Dynamic* option is selected in the NAT drop-down list for the eth1 adapter.

  The Access Gateway then functions as a network address translator and dynamically maps the private, non-routable 10-net addresses to the registered public address assigned to eth1.

  IMPORTANT: You cannot configure a reverse proxy on an IP address assigned to an adapter that has the *Dynamic* option set for NAT. NAT and a reverse proxy cannot coexist on the same adapter.

  ---

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

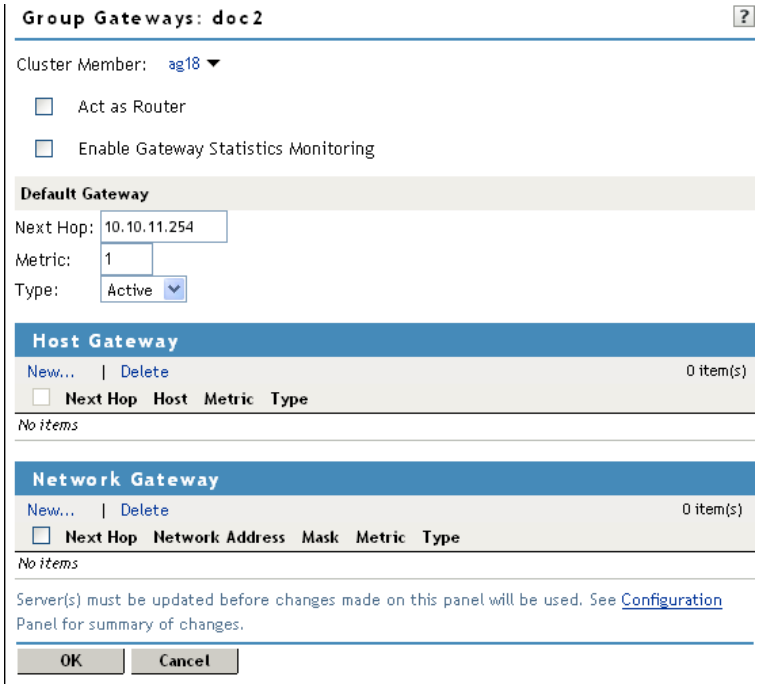

2 (Conditional) If the Access Gateway is a member of a cluster, select the server you want to configure from the list of servers in the *Cluster Member* field. All changes made to this page apply to the selected server.

3 Fill in the following fields:

Act as Router: Select this option if the Access Gateway functions as the default gateway for clients on the network. If you select this option, you can specify additional gateways.

Enable Gateway Statistics Monitoring: Select this option if you want to gather statistics and monitor the traffic on the gateways.

4 Configure your default gateway, which specifies the gateway to use when no other routes apply. 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.

5 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.

6 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 can also enter a specific IP address on a given subnet, and the Access Gateway calculates the subnet address using the mask.

Mask: The subnet mask for the subnet or IP address above. A valid entry must be at least as large as a class mask where 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.

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

8 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 and allows you to modify it.

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

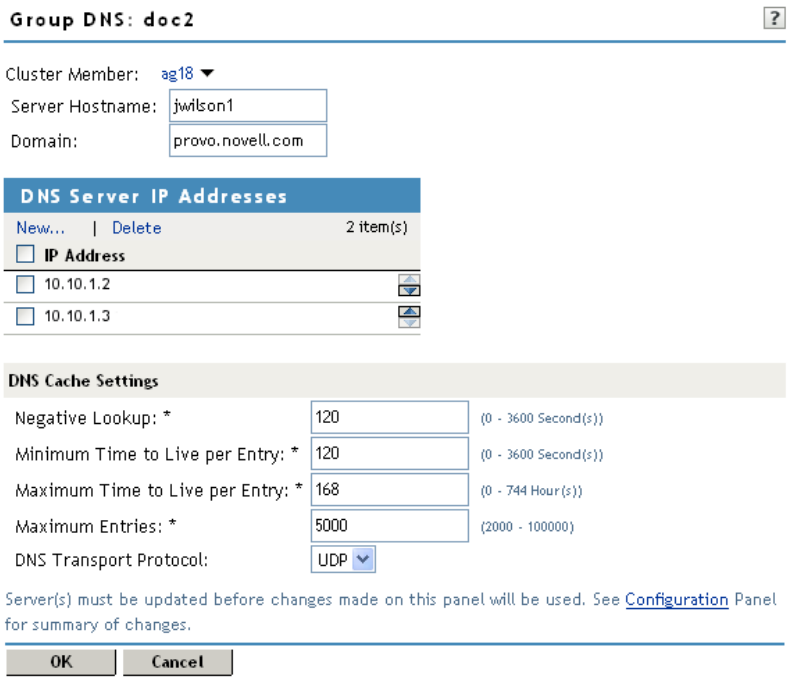

2 (Conditional) If the Access Gateway is a member of a cluster, select the server you want to configure from the list of servers in the *Cluster Member* field. All changes made to this page apply to the selected server.

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 Configure the DNS Cache Settings. These options allow you to control the refresh of DNS information. These are all standard DNS options.

Negative Lookup: Specifies how long a failed DNS lookup domain name remains in cache. If the Access Gateway cannot resolve a domain name, it stores that information in its cache for the specified amount of time. If the Access Gateway receives requests for that domain name within this period, it sends a "Bad Gateway" error message to the browser and does not resolve the domain name again. Valid field values include 0–3600 seconds. The default is120 seconds.

Minimum Time To Live per Entry: Specifies the minimum amount of time that DNS entries remain in cache before they expire. This is the minimum value the Access Gateway uses regardless of the value the DNS server returns. Valid field values include 0–3600 seconds. The default is 120 seconds.

Maximum Time To Live per Entry: Specifies the maximum amount of time that DNS entries remain in cache before they expire. This is the maximum value the Access Gateway uses regardless of the value the DNS server returns. Valid field values include 0–744 hours. The default is 168 hours.

Maximum Entries: Specifies the maximum number of DNS cache entries. When this number is reached, the Access Gateway deletes old entries to make room for newer ones. Valid field values include 2000–100000. The default is 5000.

DNS Transport Protocol: Specifies the transport protocol that DNS uses on the network where the Access Gateway is installed. Valid values are UDP and TCP. The default is UDP.

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

6 On the Server Configuration page, click *OK*, then click *Update* > *OK*.

3.8.4 Configuring Hosts

You can configure the Access Gateway 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 In the Administration Console, click *Devices* > *Access Gateways* > *Edit* > *Hosts*.

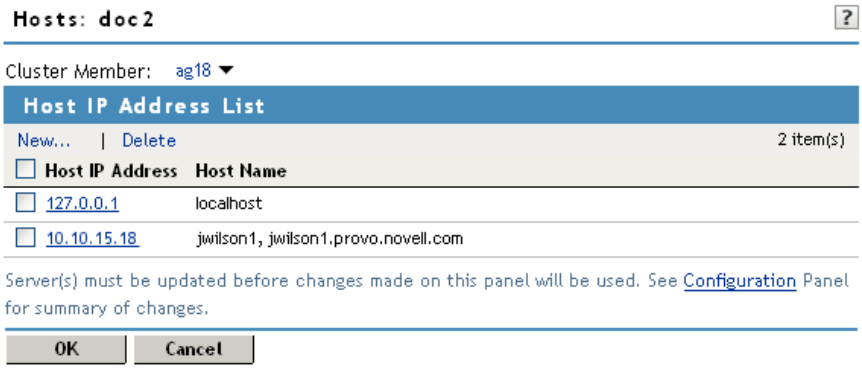

This page displays a list of host IP addresses.

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*.

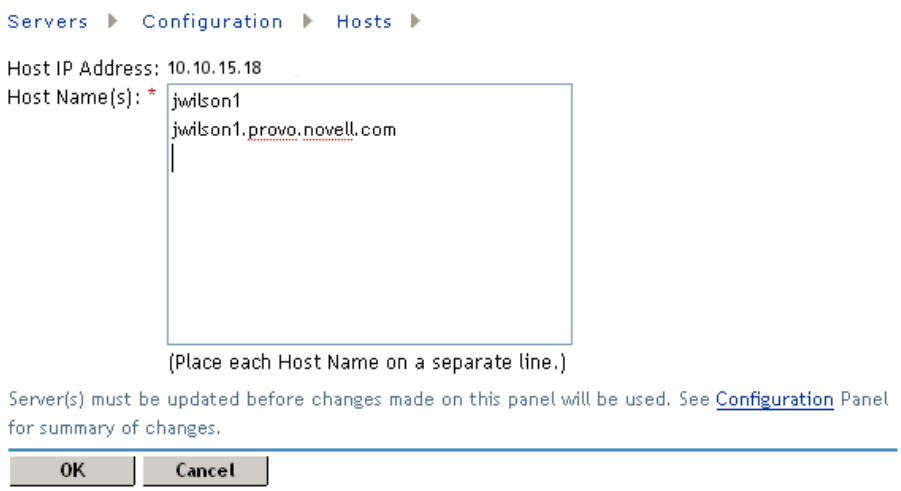

4 In the *Host Name(s)* text box, specify a name for the host. Place each hostname on a separate line, then click *OK*.

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

8 On the Server Configuration page, click *OK*, then click *Update* > *OK*.

3.8.5 Adding New Network Interfaces to the Gateway Appliance

If you add new network interface cards to the Gateway Appliance machine after installation, you need to scan for these cards. Then you can configure them.

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

2 Click the name of the Access Gateway (this is usually the IP address) that you want to add a NIC to.

3 On the Server Details page, click *New NIC* to scan for new network interface, then click *OK* to confirm.

You can click the *Command Status* tab to check if the scan has completed.

4 Click *Access Gateways*, then click *Edit* for the cluster or server that has the new card.

5 Click *Adapter List*. If the server is a member of a cluster, select the cluster member you want to configure.

The newly added network interface is displayed here.

6 In the newly added adapter section, click *New*, then configure the subnet mask and IP address.

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

8 On the Server Configuration page, click *OK*, then click *Update* > *OK*.

3.9 Customizing Logout Requests

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 re-authenticates 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. For more information on changing the authentication proxy, see Section 6.3.2, “Changing the Authentication Proxy Service,” on page 165.
3. Use this path to redirect application logout requests to this 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 `/opt/novell/nesp/lib/webapp/jsp` directory.

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

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 on how to modify `nidp.jsp` for logos, titles, and colors, see "Rebranding the Header" in the *Novell Access Manager 3.1 SP1 Identity Server Guide*.

IMPORTANT: Save a copy of your modified `nidp.jsp` file. Every time you upgrade your Access Gateway, you will need to restore this file.

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: Save a copy of your modified `logoutSuccess.jsp` file. Every time you upgrade your Access Gateway, you will need to restore this file.

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 Section 3.9.1, "Customizing Applications to Use the Access Gateway Logout Page," on page 98). The AGLogout page redirects to the plogout page, which calls the `logoutSuccess.jsp`. You cannot modify the AGLogout URL for parameters, because they are discarded. However, any parameter added to the plogout URL is saved and passed to the `logoutSuccess.jsp` file.

The parameter passed to the `logoutSuccess.jsp` file can be use 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
```

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. This option is disabled by default.

To enable it:

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

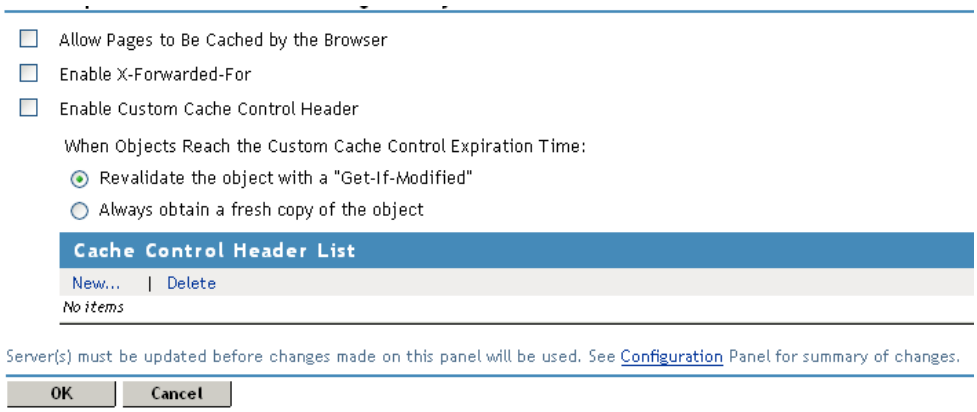

2. Select the *Enable X-Forwarded-For* option.

   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. To save your changes to browser cache, click *OK*.
4. To apply your changes, click the *Access Gateways* link, then click *Update* > *OK*.

3.11 Upgrading the Access Gateway Software

You can upgrade the software currently running on Gateway Appliance to a newer version without losing configuration information and with down time limited to the time it takes the Access Gateway to restart. See "Upgrading the Linux Access Gateway Appliance" in the *Novell Access Manager 3.1 SP1 Installation Guide*.

3.12 Exporting and Importing an Access Gateway Configuration

You can export an existing Access Gateway configuration as well as its dependent policies, and then import this configuration to a new machine. 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 Components" in the *Novell Access Manager 3.1 SP1 Administration Console Guide*

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 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.

The following sections explain this process:

- Section 3.12.1, "Exporting the Configuration," on page 101
- Section 3.12.2, "Importing the Configuration," on page 102
- Section 3.12.3, "Cleaning Up and Verifying the Configuration," on page 103

3.12.1 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 importing 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. (Conditional) If you want to change the names of the proxy services and their DNS names from a staging name to a production name, complete the following:

   5a Open the file in a text editor.

   5b Search and remove the staging suffix.

   If you have specified DNS names with a staging suffix (for example, `innerwebstaging.provo.novell.com`), you can search for `staging.provo.novell.com` and remove `staging` from the name.

   In particular, you need to change the following:

   - Any fully qualified DNS names from the staging name to the production name (DNSName elements in the file).
   - The cookie domains associated with each proxy service (AuthenticationCookieDomain elements in the file)
   - The URL masks in Pin Lists that contain fully qualified names (URLMask elements in the file.

Depending upon your naming standards, you might want to change the names of the following:

- UserInterfaceID elements (proxy service, pin list, and protected resource user interface ID's)
- Description elements (proxy service, pin list, and protected resource descriptions)
- Name (proxy service, pin list, and protected resource names)
- SubServiceID elements
- MultiHomeMasterSubserviceIDRef elements
- LogDirectoryName elements
- ProfileIDRef elements
- ProtectedResourceID elements
- ProfileID elements (TCP Listen options name)

5c (Conditional) If your Web servers in the staging area have different IP addresses and hostnames than the Web Servers in the production area, you can search and replace them in the configuration file or wait until after the import and modify them in the Administration Console.

6 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.

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

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

9 (Conditional) If you have created multiple policy containers, select the next policy container in the list, and repeat Step 6 through Step 8.

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

10 (Conditional) If your policies redirect users to staging URLs when they are denied access, search and replace these URLs with the production URLs. Open the policy file with a text editor and search for your staging name.

11 Copy the Access Gateway and policy configuration files to a place accessible by the new Access Gateway.

12 Continue with Section 3.12.2, “Importing the Configuration,” on page 102.

3.12.2 Importing the Configuration

1 Verify that the Access Gateway meets the conditions for an import:

- The Access Gateway should not be a member of a cluster. If it is a member of a cluster, remove it from the cluster before continuing.

  In the Administration Console, click *Devices* > *Access Gateways*, select the Access Gateway, then click *Actions* > *Remove from Cluster*.

  You can create a cluster and add this machine to the cluster as the primary server after you have completed the import.

- The Access Gateway should be an unconfigured machine. If it contains reverse proxies, delete them before continuing.

  In the Administration Console, click *Devices* > *Access Gateways* > *Edit* > *Reverse Proxies / Authentication*. In the *Reverse Proxy List*, select *Name*, then click *Delete*. Update the Access Gateway and the Identity Server.

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

  The policies that the Access Gateway is dependent upon must be imported first.

3 (Conditional) If you have exported policies from more than one container, create the policy containers. Click the *Containers* tab; in the *Container List*, click *New*, specify the name for the container, then click *OK*.

4 (Conditional) If your system already contains policies, delete them if they aren't being used.

  If they are in use and you have policies with the same names as the policies you are going to import, you need to manually reconcile the duplicate policies. See Step 5 in Section 3.12.3, "Cleaning Up and Verifying the Configuration," on page 103.

5 In the Policy List, click *Import*.

6 Browse to the location of the policy configuration file, select the file, then click *OK*.

7 (Conditional) If you exported multiple policy configuration files, repeat Step 5 and Step 6.

8 Enable all new Role policies. Click *Identity Servers* > *Edit* > *Roles*.

9 Either select *Name* to enable all policies or individually select the policies, then click *Enable*.

10 Click *OK*, then click *Update*.

11 To import the Access Gateway configuration, click *Access Gateways* > *[Name of Access Gateway]* > *Configuration* > *Import*.

12 Browse to the location of the file, select the file, enter a password if you specified one on export, then click *OK*.

13 Continue with Section 3.12.3, "Cleaning Up and Verifying the Configuration," on page 103.

3.12.3 Cleaning Up and Verifying the Configuration

1 When the configuration import has finished, verify the configuration for your reverse proxies.

  1a Click *Access Gateways* > *Edit* > *[Name of Reverse Proxy]*.

  1b Verify the listening address.

  This is especially important if your Access Gateway has multiple network adapters. By default, the IP address of eth0 is always selected as the listening address.

  1c Verify the certificates assigned to the reverse proxy.

  The Subject Name of the certificate should match the published DNS name of the primary proxy service in the *Proxy Service List*.

  1d Verify the Web Server configuration. In the *Proxy Service List*, click the *Web Server Addresses* link. Check the following values:

  - **Web Server Host Name.** If this name has a staging prefix or suffix, remove it.
  - **IP addresses in the Web Server List.** If the IP addresses in the production area are different from the IP addresses in the staging area, modify the IP addresses to match the production area.

- **Certificates.** If you have configured SSL or mutual SSL between the proxy service and the Web servers, configure the *Web Server Trusted Root* and *SSL Mutual Certificate* options. The export and import configuration option does not export and import certificates.

1e Click *OK* twice.

2 (Conditional) If you have multiple reverse proxies, repeat Step 1 for each proxy service.

3 On the Configuration page, click *Reverse Proxy / Authentication*, then select the *Identity Server Cluster* configuration.

4 If you have multiple reverse proxies, verify that the Reverse Proxy value in the *Embedded Service Provider* section is the reverse proxy you want to use for authentication, then click *OK* twice.

5 (Conditional) If the Administration Console already contained some policies, verify that you do not have policies with duplicate names. Click *Policies* > *Policies*.

Policies with duplicate names have Copy-*n* appended to the end of the name, with *n* representing a number. If you have duplicates, reconcile them:

- If they contain the same rules, you need to reconfigure the resources using one policy to use the other policy before you can delete the duplicate policy.
- If they contain different rules, rename the duplicate policies.

6 (Conditional) Apply any policy configuration changes.

7 Click *Access Gateways* > *Update*.

8 Click *Identity Servers* > *Update*.

If your Identity Server does not prompt you for an update, complete the following steps to trigger the update.

8a In the Administration Console, click *Devices* > *Access Gateways* > *Edit* > *Reverse Proxy / Authentication*.

8b Set the Identity Server Cluster field to *None*, then click *OK*.

8c Click *Reverse Proxy / Authentication*.

8d Set the Identity Server Cluster field to the correct value, then click *OK*.

8e Update the Access Gateway.

8f Update the Identity Server.

9 Configure the keystores for the Access Gateway.

If you have configured the Access Gateway for SSL between the Identity Server and the Access Gateway and between the Access Gateway and the browsers, verify that the trust stores and the keystores contain the correct certificates.

9a In the Administration Console, click *Security* > *Certificates*.

9b Find the certificate for the Access Gateway.

The subject name of this certificate should match the DNS name of the Access Gateway. If this certificate is not in the list, you need to create it or import it.

This certificate should be in use by the ESP Mutual SSL and Proxy Key Store of the Access Gateway.

9c If the certificate is not in use by the required keystores, select the certificate, then click *Actions* > *Add Certificate to Keystores*.

9d Click the *Select Keystore* icon, select ESP Mutual SSL and Proxy Key Store of the Access Gateway, then click *OK* twice.

10 Configure the trust stores for the Access Gateway.

10a In the Administration Console, click *Security* > *Certificates* > *Trusted Roots*.

The trusted root certificate of the CA that signed the Access Gateway certificate needs to be in the NIDP-truststore.

The trusted root certificate of the CA that signed the Identity Server certificate, needs to be in the ESP Trust Store of the Access Gateway.

10b If you need to add a trusted root to a trust store, select the trusted root, click *Add Trusted Roots to Trust Stores*.

10c Click the *Trust Store* icon, select the required trust store, then click *OK* twice.

11 If you made any keystore or trust store modifications, update the Access Gateway and the Identity Server.

12 (Optional) Create a cluster configuration and add this server as the primary server.

Access Gateway Maintenance

4

4.1 Gateway Appliance Logs

This section contains the following information about the Access Gateway Appliance logs:

4.1.1 Configuring Log Levels

You can use the following procedure to set the level of information logged to the `ics_dyn.log` file in the `/var/log` directory.

1. At the command prompt, enter the following command:

   ```
   nash
   ```

2. At the `nash` shell prompt, enter the following command:

   ```
   configure .current
   ```

3. To change the log level, enter the following command:

   ```
   log-conf log-level <log level>
   ```

   Replace <*log level*> with the new log level that you want to set.

Level	Description
LOG_EMERG	Sends only messages that render the system unusable, if they are not resolved.
LOG_ALERT	Sends only messages that require immediate action.
LOG_CRIT	Sends only messages about critical situations.
LOG_ERR	Sends warning messages about recoverable errors.
LOG_WARNING	Sends warning messages.
LOG_NOTICE	Sends information about the status of a service to the service configuration logs.

Level	Description
LOG_INFO	Sends informational messages such as requests sent to Web servers and the results of authentication requests.
LOG_DEBUG	Sends debug messages.

When you run the `/etc/init.d/novell-vmc start` command, the default log level is set to LOG_NOTICE. You can change the log level to any level from LOG_EMERG to LOG_INFO.

4 To apply changes, enter the following command:

```
apply
```

5 To exit from the configuration mode, enter the following command:

```
exit
```

6 To exit from the nash shell, enter the following command:

```
exit
```

4.1.2 Interpreting Log Messages

The entries in the `ics_dyn.log` file have the following format:

```
<time-date-stamp> <hostname> : <AM#event-code> : <AMDEVICE#device-id> :
<AMAUTHID#auth-id> : <AMEVENTID#event-id> :<supplementary log entry data and
text>
```

A sample log message is given below:

```
Aug  3 14:35:41 c1h : AM#504503000: AMDEVICEID#ag-0BDF41AAC4CDCBE5 :
AMAUTHID#0: AMEVENTID#74: Process request 1 'www.lag-202.com' '/AGLogout'
[192.10.100.111:38091 -> 192.10.106.2:80]
```

The fifth and sixth digits in the `<AMEVENTID#event-id>` refer to the Access Gateway components. The following table list the numbers and the components which they denote.

Table 4-1 *Linux Access Gateway Components*

Number	Component
01	If the fifth and sixth digits are 01, the Multi-Homing component
02	Service Manager
03	Request Processing
04	Authentication
05	Authorization
06	Identity Injection
07	Form Fill
08	Caching
09	Response Processing

Number	Component
11	Rewriting
12	Soap Channel
14	IVM
15	Connection Manager.
16	VXE
17	DataStream

4.1.3 Configuring Logging of SOAP Messages and HTTP Headers

1. At the command prompt, enter the following command:

   ```
   nash
   ```

2. To enter the configuration mode, enter the following command:

   ```
   configure .current
   ```

3. Enter one of the following commands to configure logging:

Command	Purpose
`log-conf debug-soap-messages enable`	Logs all the SOAP messages between the Access Gateway and the Embedded Service Provider to the `/var/log/lagsoapmessages` file.
`log-conf no debug-soap-messages enable`	Disables the logging of SOAP messages between the Access Gateway and the Embedded Service Provider.
`log-conf debug-http-headers enable`	Logs all the HTTP headers between the browsers and the Access Gateway and between the Access Gateway and the Web servers to the `/var/log/laghttpheaders` file.
`log-conf no debug-http-headers enable`	Disables the logging of HTTP headers to the `/var/log/laghttpheaders` file.

4. To apply changes, enter the following command:

   ```
   apply
   ```

5. To exit from the configuration mode, enter the following command:

   ```
   exit
   ```

6. To exit from the nash shell, enter the following command:

   ```
   exit
   ```

4.2 Configuring Proxy Service Logging

Logging HTTP transactions has associated costs. The Access Gateway is capable of handling thousands of transactions per second. If transaction volume is high and each log entry consumes a few hundred bytes, the Access Gateway can fill up the available disk space in a matter of minutes. HTTP logging also increases system overhead, which causes some degradation in performance. By default, the logging of HTTP transactions is turned off. Before enabling logging, you need to determine what needs to be logged and then plan a logging strategy.

4.2.1 Determining Logging Requirements

Because logging requirements and transaction volume vary widely, Novell cannot make recommendations regarding a specific logging strategy. The following tasks guide you through the process of creating a strategy that fits your business needs.

1. Identify the reasons for tracking transactions such as customer billing, statistical analysis, or growth planning.
2. Determine which resources need logging.

   You enable logging at the proxy service level. If you have a proxy service protecting resources whose 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. You can also select to copy the files to another server and then delete them. 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, “Monitoring 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 three or more.

The following formulas can help you estimate when the system would run out of resources:

Calculating diskfull_time

Use the following formula to calculate how long it will take 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 will fill up every 17.47 minutes.

To calculate the diskfull_time for your Access Gateway:

1. Determine the values of the four variables listed above.
2. 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 111 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*.

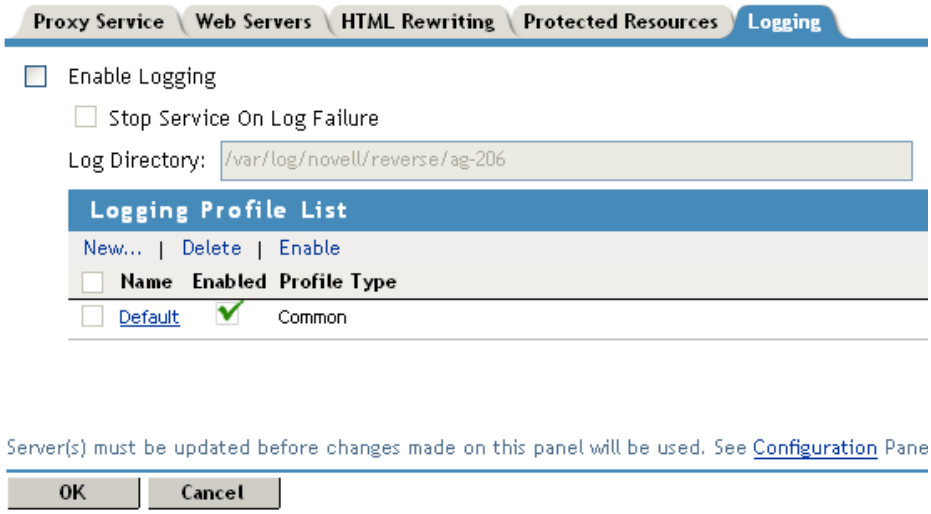

2. Fill in the following fields:

   Enable Logging: Select this field to enable logging.

   Stop Service On Log Failure: Select this field if you want the Access Gateway to deny requests to this proxy service because the Access Gateway cannot log entries for it.

   Log Directory: Displays the default location for the log files for this proxy service.

3. In the *Logging Profile List*, click one of the following options:

   - **New:** Click this option to create a new logging profile. Then specify a name and select either *Common* or *Extended*.

- **Default:** Click *Default* to modify or view the settings for the *Default* profile. The *Default* profile uses the common log options.

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]*.

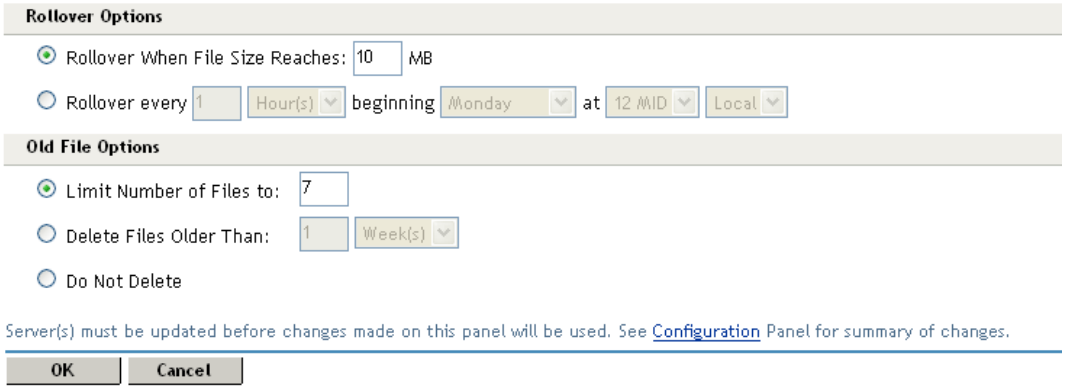

2 Select one of the following roll over options:

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]*.

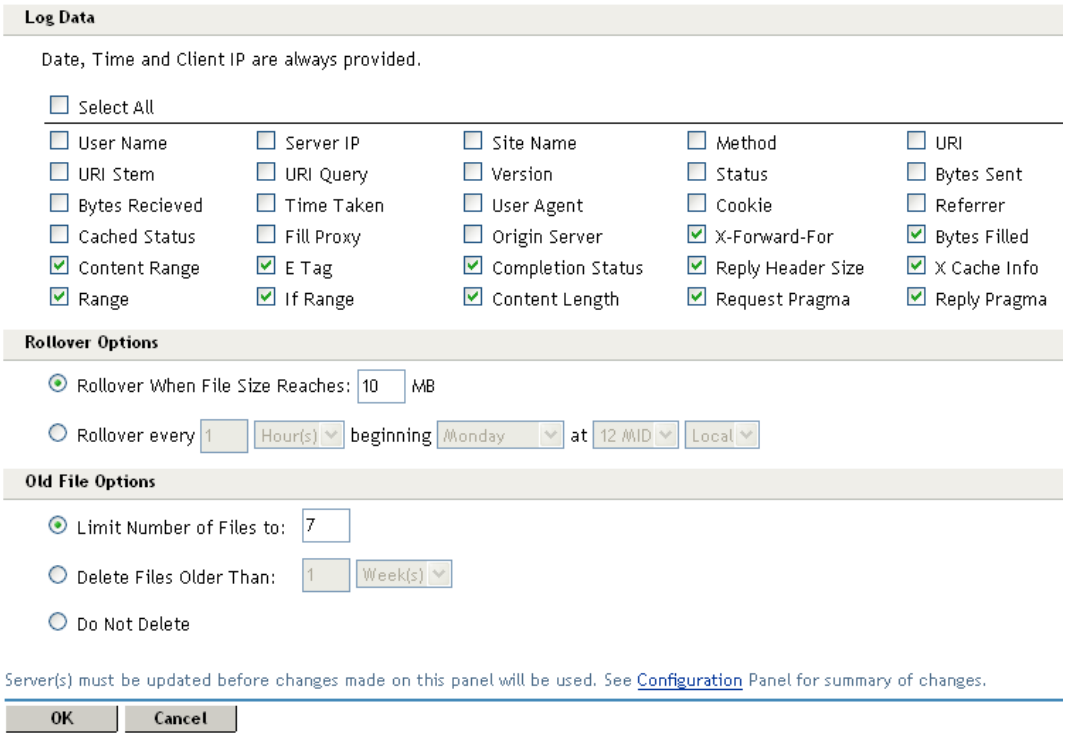

2 Select one or more of the log data options:

Name	Description
User Name	The name of the user sending the request.
Server IP	The IP address of the Access Gateway.
Site Name	The name of the reverse proxy.
Method	The HTTP method the browser sent to the Access Gateway.
URI	The HTTP URL the browser sent to the Access Gateway.
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.
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.
Version	The HTTP version specified in the URL the browser sent to the Access Gateway.
Status	The HTTP status code the Access Gateway sent to the browser.
Bytes Sent	The number of bytes of HTTP response data the Access Gateway sent to the browser.
Bytes Received	The number of bytes of HTTP request data the proxy service received from the browser.

Name	Description
Time Taken	The time in seconds it took the Access Gateway resources to deal with the request.
U*ser Agent*	The User-Agent HTTP request header value the browser sent to the Access Gateway.
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.
Referer	The Referer HTTP request header value the browser sent to the Access Gateway.
Cached Status	The value indicates whether the request was filled from cache. 1 = filled from cache 0 = not filled from cache
Fill Proxy	The IP address of the upstream proxy.
Origin Server	The IP address of the Web server. This assumes the Access Gateway retrieved the requested information directly from the Web server.
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.
Bytes Filled	The total bytes filled in response to the request.
Content Range	The byte ranges sent from the Access Gateway to a requesting browser.
E Tag	The tag sent from the Access Gateway to a requesting browser.
Completion 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).
Reply Header Size	The size in bytes of the HTTP header associated with a response to a client.
X Cache Info	Brief status statement for cached objects; brief reasons why an object was not cached.
Range	The Range header value.
If Range	The If Range header value, which indicates whether the browser request was a conditional range request.
Content Length	The size in bytes of the entire object delivered to a requesting browser.
Request Pragma	The pragma value associated with a browser request.
Reply Pragma	The pragma value associated with a server response to a requesting browser.

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 Appliance logs are stored in `/root` partition by default. You can create a `/var` partition to store the logs. The size of this partition depends on your requirements. For more information on creating the `/var` partition, see "Customizing the Partitions" in the *Novell Access Manager 3.1 SP1 Installation Guide*.

4.3 Monitoring Access Gateway Statistics

Access Gateway statistics are available for each Access Gateway and for clusters:

- Section 4.3.1, "Viewing Access Gateway Statistics," on page 118
- Section 4.3.2, "Viewing Cluster Statistics," on page 127

4.3.1 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*.

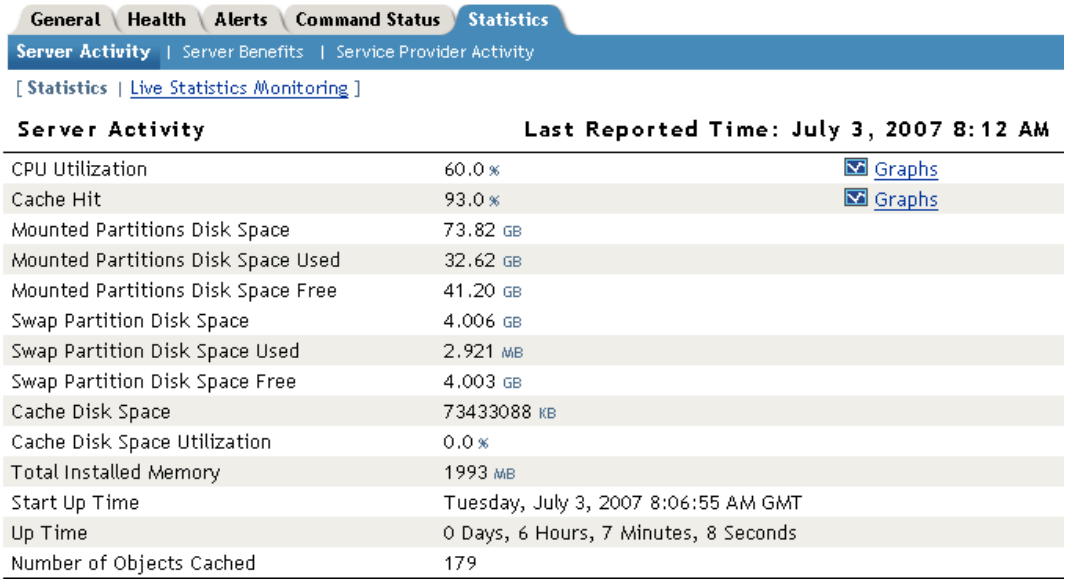

2 Select from the following types:

- "Server Activity" on page 119
- "Server Benefits" on page 123
- "Service Provider Activity" on page 123

3 Click *Close*.

Server Activity

Access Gateways > [Name of Server] > 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" on page 119
- "Connections" on page 120
- "Bytes" on page 121
- "Requests" on page 122
- "Cache Freshness" on page 122

Server Activity

The Server Activity section displays general server utilization statistics.

Table 4-2 *Server Activity*

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.
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 for the swap partition.
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	Reserved. Not currently used.

Connections

The connection statistics show the current and peak levels of usage in terms of TCP connections.

Table 4-3 *Connections*

Statistic	Description
Current Connections to Origin Server	Displays the current number of connections that the Access Gateway has established with Web servers.

Statistic	Description
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.
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

Table 4-4 *Bytes*

Statistic	Description
Bytes per Second from Origin Server	Displays the number of bytes of data being sent each second from the Web servers to the Access Gateway.
Bytes per Second to Browsers	Displays the number of bytes of data being sent each second from the Access Gateway to the browsers.
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.
Bytes Received from Origin Server	Displays the total number of bytes of data sent to the Access Gateway from the Web servers since the Access Gateway last started.
Bytes Sent to Browser	Displays the total number of bytes of data sent to the browsers from the Access Gateway since the Access Gateway last started.

Statistic	Description
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.

Table 4-5 *Requests*

Statistic	Description
Current Requests to Origin Server	Displays the current number of requests that the Access Gateway has made to the Web servers.
Current Requests from Browsers	Displays the current number of requests that the browsers have made to the Access Gateway.
Total Current Requests	Displays the total number of current requests that the Access Gateway has received from the browsers and that the Access Gateway has sent to the Web servers.
Successful Requests to Origin Server	Displays the total number of successful requests that the Access Gateway has sent to the Web servers since the Access Gateway last started.
Failed Requests to Origin Server	Displays the total number of failed requests that the Access Gateway has sent to the Web servers since the Access Gateway last started.
Cumulative Requests to Origin Server	Displays the total number of requests that the Access Gateway has sent to the Web servers since the Access Gateway last started.
Cumulative Requests to Browsers	Displays the total number of requests that the browsers have sent to the Access Gateway since the Access Gateway last started.
Total Cumulative Requests	Displays the total number of cumulative requests that the Access Gateway has processed since the Access Gateway last started.
Requests per Second to Origin Server	Displays the number of requests that are being sent each second from the Access Gateway to the Web servers.
Requests per Second from Browsers	Displays the number of requests that are being sent each second from the browsers to the Access Gateway.
Total Requests per 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.

Table 4-6 *Cache Freshness*

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.
Oldest Object in Memory	Displays how long the oldest cache object has been cached.

Server Benefits

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:

Table 4-7 *Server Benefits*

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.
Bandwidth Saved	Displays the amount of bandwidth saved by using data cached by the Access Gateway rather than requesting the data from the Web servers.
Total DNS Lookups Saved	Displays the number of DNS requests that the Access Gateway could solve locally without performing a DNS lookup.
DNS “Modified Since” Queries Returning False	Displays the number of DNS Modified Since queries that the Access Gateway was able to service with a false value.
Total Number of Connections Saved	Displays the number of connections that the Access Gateway has with clients minus the number of connections that the Access Gateway has with Web servers. This statistic indicates the number of connections that the Access Gateway is off loading from the Web servers.

Service Provider Activity

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 module (ESP) and the Identity Server. These statistics are grouped into the following categories:

- Application
- Authentications
- Incoming HTTP Requests
- Outgoing HTTP Requests
- Liberty
- Clustering

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 memory 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 memory that is free for the selected time period.

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.
Logouts	The number of explicit logouts performed by users. This does not include logouts where an inactive session was destroyed.
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 session 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.

Statistic	Description
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 request that were processed during the last 60 second interval.
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 request that were processed during the last 60 second interval.
Historical Maximum Request Duration (Milliseconds)	The age, since the Identity Server was started, of the longest outgoing HTTP request that was processed.
Historical Mean Request Duration (Milliseconds)	The mean age, since the Identity Server was started, of all outgoing HTTP requests that were processed.

Liberty

Statistic	Description
Liberty Federation	The number, since the Identity Server was started, of Liberty protocol federations performed.
Liberty De-Federations	The number, since the Identity Server was started, of Liberty protocol de-federations performed.
Liberty Register-Names	The number, since the Identity Server was started, of Liberty protocol register names performed.

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, since the Identity Server was started, of proxied requests that have been processed. This is the case where the request hits a non-authoritative (wrong) box.
Total Non-Proxied Requests	The total number, since the Identity Server was started, of non-proxied requests that have been processed. This is the case where the request hits the authoritative (correct) box.
Authoritative Server Obtained from URL Parameter	The total number, since the Identity Server was started, of authoritative servers identified using the parameter from the URL query string.
Authoritative Server Obtained from Cookie	The total number, since the Identity Server was started, of authoritative servers identified using the HTTP cookie.
Payload Examinations	The total number, since the Identity Server was started, of attempted payload examinations to identify the authoritative server.
Successful Payload Examinations	The total number, since the Identity Server was started, of successful payload examinations to identify the authoritative server.
Identity ID Broadcasts	The total number, since the Identity Server was started, of attempted Identity ID Broadcasts to identify the authoritative server.
Successful Identity ID Broadcasts	The total number, since the Identity Server was started, of successful Identity ID Broadcasts to identify the authoritative server.
Session ID Broadcasts	The total number of attempted Session ID Broadcasts to identify the authoritative server.
Successful Session ID Broadcasts	The total number, since the Identity Server was started, of successful Session ID Broadcasts to identify the authoritative server.

4.3.2 Viewing Cluster Statistics

To view general performance statistics of the servers assigned to the selected cluster:

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

2 To determine performance, analyze the following statistics:

Column	Description
Server Name	Lists the name of the Access Gateways that belong to the group. To view additional statistical information about a specific Access Gateway, click the name of an Access Gateway.
CPU %	Displays the current CPU utilization rate. Use this statistic for capacity planning.
Cache Hit Rate %	Displays the current cache hit rate. A high cache hit rate indicates that the caching system is off-loading significant request processing from the Web server whose objects have been cached. If the percentage is low, you might want to configure a pin list. For this and other caching options, see Chapter 5, "Configuring the Content Settings," on page 141.
Bytes per second to/from Server	Displays the rate at which the Access Gateway is requesting Web objects from the Web servers it is protecting.
Bytes per second to/from Browser	Displays the rate at which browser clients are requesting Web objects.
Current Connections	Displays the total number of TCP connections that are active, idle, or closing.
Statistics	Allows you to view all the statistics for a selected server. Click *View* to see these additional statistics. For more information, see Section 4.3, "Monitoring Access Gateway Statistics," on page 118.

3 Click *Close*.

4.4 Monitoring Access Gateway Alerts

The Access Gateway has been programmed to issue events to various types of systems (such as a Novell® Audit server or a Syslog server) so that the administrator can be informed when significant changes occur that modify how the Access Gateway is performing. For information about auditing and audit events, see Section 4.5, "Enabling Access Gateway Audit Events," on page 133. This section describes how to use the following types of alerts:

- Section 4.4.1, "Reviewing Java Alerts," on page 128
- Section 4.4.2, "Configuring Access Gateway Alerts," on page 129

4.4.1 Reviewing Java Alerts

The Alerts page allows you to view information about current Java alerts and to clear them. An alert is generated whenever the Access Gateway detects a condition that prevents it from performing normal system services.

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

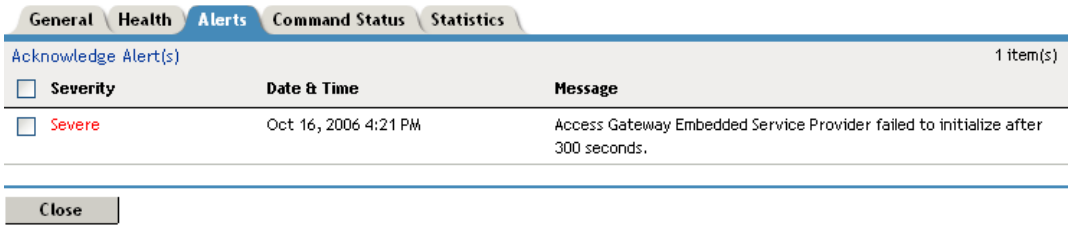

2 To delete an alert from the list, select the check box for the alert, then click *Acknowledge Alert(s)*. To remove all alerts from the list, click the *Severity* check box, then click *Acknowledge Alert(s)*.

3 Click *Close*.

4 (Optional) To verify that the problem has been solved, click *Access Gateways > [Server Name] > Health > Update from Server*.

4.4.2 Configuring Access Gateway Alerts

You can configure alerts for individual Access Gateways and for Access Gateway clusters. To set up notification for these types of alerts, see the following sections:

- "Access Gateway Alerts" on page 129
- "Access Gateway Cluster Alerts" on page 132

Access Gateway Alerts

For a 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 To add a new profile, click *New*.

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

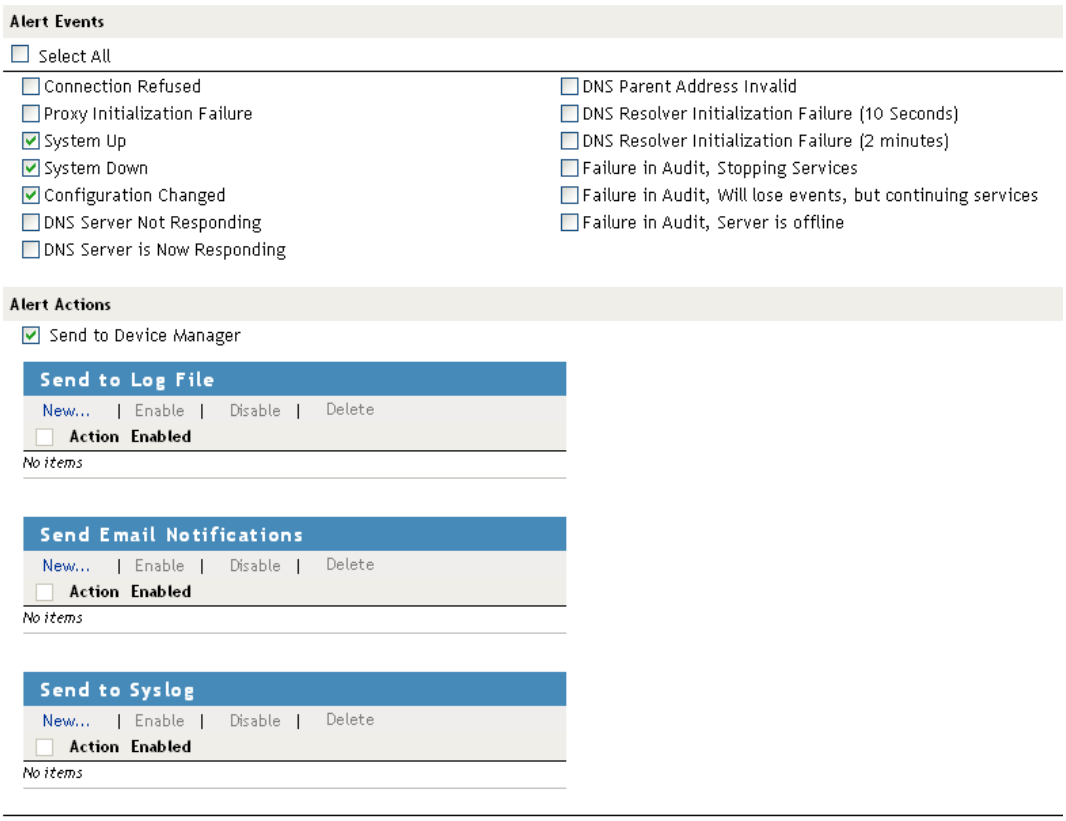

4 To select the alerts for notification, select one or more of the following:

Alert	Description
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.
DNS Server Not Responding	Generated each time the DNS stops responding.
DNS Server Is Now Responding	Generated each time the DNS server comes up.
DNS Parent Address Invalid	Generated when the IP address of DNS parent is invalid.
DNS Resolver Initialization Failure (10 seconds)	Generated when the DNS resolver initialization fails.
DNS Resolver Initialization Failure (2 minutes)	Generated when the DNS resolver initialization fails.

Alert	Description
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 Proxy Service Logging," 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).

5 To send alerts to the Administration Console, select the *Send to Device Manager* check box.

6 To send alerts to a log file, click *New* in the *Send to Log File* section, specify a name for the log profile, then click *OK*.

6a Configure the following Log File details:

- **Log File Name:** Specify a name for the log file and a path where the file should be stored.
- **Max File Size:** Specify a maximum size in KB for the log file. The size can be from 50 to 100000 KB. Specify 0 to indicate that there is no maximum file size.

6b Click *OK*.

7 To enable e-mail notification click *New* in the Send Email Notifications section, specify a name for the e-mail profile, then click *OK*.

7a Configure the following e-mail details:

- **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*.

7b Click *OK*.

8 To enable syslog alerts click *New* in the Send to Syslog section, specify a name for the Syslog profile, then click *OK*.

8a Configure the following syslog details:

- **Facility Name:** Specify a facility name for the Syslog server. It can be any name from local0 to local7. If you specify local0 as your facility name, the alerts are stored at `\var\logs\ics_dyn.log`. The Access Gateway Appliance uses local0 for normal logging information. Therefore, it is not recommended to specify local0 as your facility name.

8b Click *OK*.

9 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 email addresses, or to a syslog file is not performed until the action profile is enabled.

10 On the Alert Profiles page, verify that the Alert Profile you have created is enabled.

11 To save your modifications, click *OK* twice.

12 On the *Access Gateways* page, click *Update*.

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*.

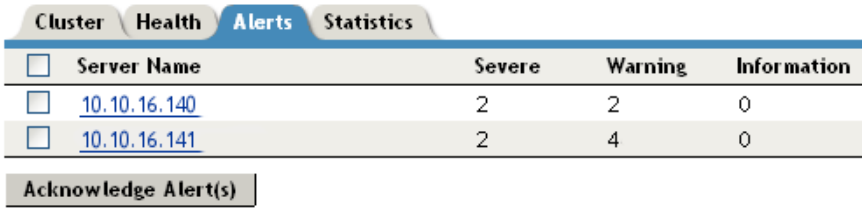

2 Analyze the data displayed in the table.

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.
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.

4.5 Enabling Access Gateway Audit Events

The *Novell Audit* option in the Access Gateway allows you to configure the events you want audited. The following steps assume that you have already set up Novell Audit on your network. For more information, see "Configuring Access Manager for Novell Auditing" in the *Novell Access Manager 3.1 SP1 Administration Console Guide*.

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

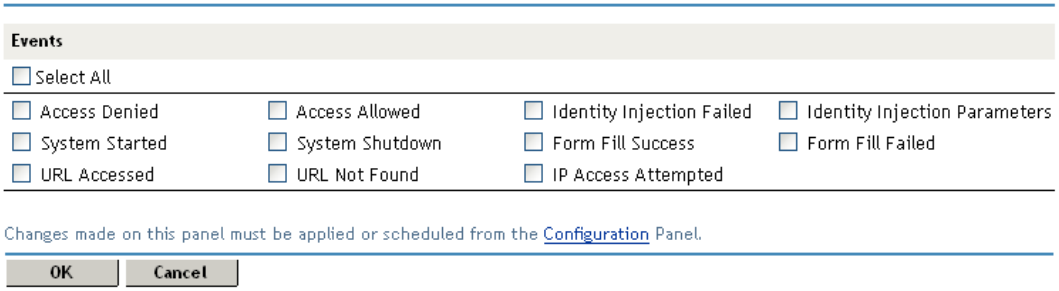

2 Select the events for notification.

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.

Event	Description
Form Fill Failed	Generated when a Form Fill policy fails to successfully fill in a form.

3 To save your modifications, click *OK* twice.

4 On the Access Gateways page, click *Update*.

4.6 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:

- Section 4.6.1, “Health States,” on page 134
- Section 4.6.2, “Monitoring the Health of an Access Gateway,” on page 135
- Section 4.6.3, “Viewing the Health of an Access Gateway Cluster,” on page 138

4.6.1 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.

Icon	Description
	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.

4.6.2 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.
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.

Service Category	If not healthy
Time: Indicates the type of time configuration. Time must be configured so that it remains synchronized with the other servers in the configuration (the Identity Server, SSL VPN server, J2EE agents, Web servers, etc.).	See Section 3.6, “Setting the Date and Time,” on page 84.
Gateway: Specifies the type of routing that is configured for the gateway.	See Section 3.8.2, “Viewing and Modifying Gateway Settings,” on page 92.
DNS: Specifies whether a domain name server has been configured and is active.	Displays the IP address of the each configured DNS server and when the server last responded. See Section 3.8.3, “Viewing and Modifying DNS Settings,” on page 94.
Services: Indicates the general health of all configured services.	Displays messages about the health of the reverse proxy, the back-end Web servers, and internal services (the SOAP back channel and the communication module).
Address: Indicates whether an IP address has been configured for the reverse proxy to listen on. This is required for the Access Gateway to function.	See Section 1.1, “Creating a Reverse Proxy and Proxy Service,” on page 12.
Embedded Service Provider Communication: Indicates whether the Embedded Service Provider can communicate with the Identity Server. 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.	Restart the Embedded Service Provider. If restarting the Embedded Service Provider fails, try restarting Tomcat.

Service Category	If not healthy

L4 and Cache: The L4 status indicates whether the Access Gateway is responding to health checks from the L4 switch. The number increments with each health check for which the Access Gateway does not send a response.
◆ When it reaches 13, the health is changed to yellow.
◆ When it reaches 31, the health is changed to red.
If the Access Gateway recovers and starts responding, the health turns green after 20 seconds and the unresponsive count is reset to 0.
To fix the problem if it does not resolve itself, restart the Access Gateway Appliance.
The cache status indicates the current number of delayed cache requests and whether enough memory is available to process new requests.
◆ When this number reaches 101, the health is changed to yellow.
◆ When this number reaches 151, the health changes to red. To solve the problem, you need to restart the Access Gateway Appliance.
	Restart the Access Gateway Appliance by entering the following commands:
	`/etc/init.d/novell-vmc stop`
	`/etc/init.d/novell-vmc start`

Embedded Service Provider Configuration: Specifies 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 "Configuring an Identity Server" in the *Novell Access Manager 3.1 SP1 Identity Server Guide* for information on configuring an Identity Server. See Section 1.1, "Creating a Reverse Proxy and Proxy Service," on page 12 for information on assigning an Identity Server configuration to the Access Gateway.
Configuration Datastore: Indicates whether the configuration data store is functioning correctly.	See "Repairing the Configuration Datastore" in the *Novell Access Manager 3.1 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.
Signing and Encryption Keys: Indicates whether the Signing keystore contains a key.	Click *Access Gateways* > *Edit* > *Service Provider Certificates* > *Signing* and replace signing key in this keystore.

Service Category	If not healthy
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 Access Gateway and the Embedded Service Provider are communicating.	Restart the Access Gateway. See Section 3.3.5, "Restarting the Access Gateway Appliance," on page 80.
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. See "Assigning an Identity Server to a Cluster Configuration" in the *Novell Access Manager 3.1 SP1 Identity Server Guide*. Reconfigure the Access Gateway to trust a different Identity Server configuration. See Section 1.1, "Creating a Reverse Proxy and Proxy Service," on page 12.

4 Click *Close*.

4.6.3 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. 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.

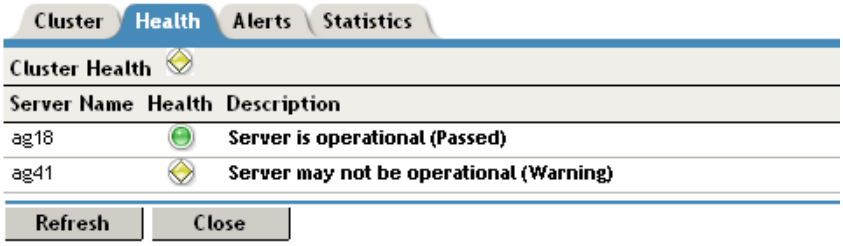

3 To ensure that the information is current, click *Refresh*.

4 To view specific information about the status of an Access Gateway, click the Health icon in the Access Gateway row.

4.7 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 4.7.1, “Viewing the Status of Current Commands,” on page 139
- Section 4.7.2, “Viewing Detailed Command Information,” on page 140

4.7.1 Viewing the Status of Current Commands

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

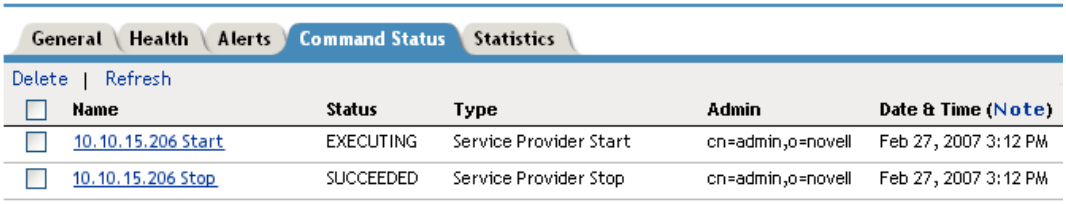

This page lists the current commands and the following information about the commands:

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 4.7.2, “Viewing Detailed Command Information,” on page 140.
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.

Column Name	Description
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*.

4.7.2 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. The following information is displayed:

Name: The Identity Server name.

Type: The type of command issued to the server.

Admin: The distinguished name of the admin user of the LDAP directory who performed the command.

Status: The status of the server command.

Last Executed On: The date and time that the command was executed.

3 To determine if any problems occurred, view the *Command Execution Details* section.

For a command that fails because the Administration Console cannot communicate with the Identity Server, 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.

4 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*.

5 Click *Close* to return to the Command Status page.

Configuring the Content Settings

5

One of the major benefits of using an Access Gateway to protect Web resources is that it can cache the requested information and send it directly to the client browser rather than contacting the origin Web resource and waiting for the requested information to be sent. This can significantly accelerate access to the information.

The object cache on an Access Gateway is quite different from a browser's cache, which all users access when they click the *Back* button and which can serve stale content that doesn't accurately reflect the fresh content on the origin Web server.

The Access Gateway caching system uses a number of methods to ensure cache freshness. Most time-sensitive Web content is flagged by Webmasters in such a way that it cannot become stale unless a caching system ignores the Webmaster's settings. The Access Gateway honors all flags that affect cache freshness, including Time to Expire, Don't Cache, and Must Revalidate directives.

In addition, the Access Gateway can be fine-tuned for cache freshness in the following ways:

- Accelerated checking of objects that have longer than desirable Time to Expire headers
- Delayed checking of objects that have shorter than desirable Time to Expire headers
- Checking for freshness of objects that do not include Time to Expire headers

The following sections describe the features available to fine-tune this process for your network:

5.1 Configuring Caching Options

The Cache Options allow you to control how the Access Gateway caches objects.

1 Click *Access Gateways* > *Edit* > *Cache Options*.

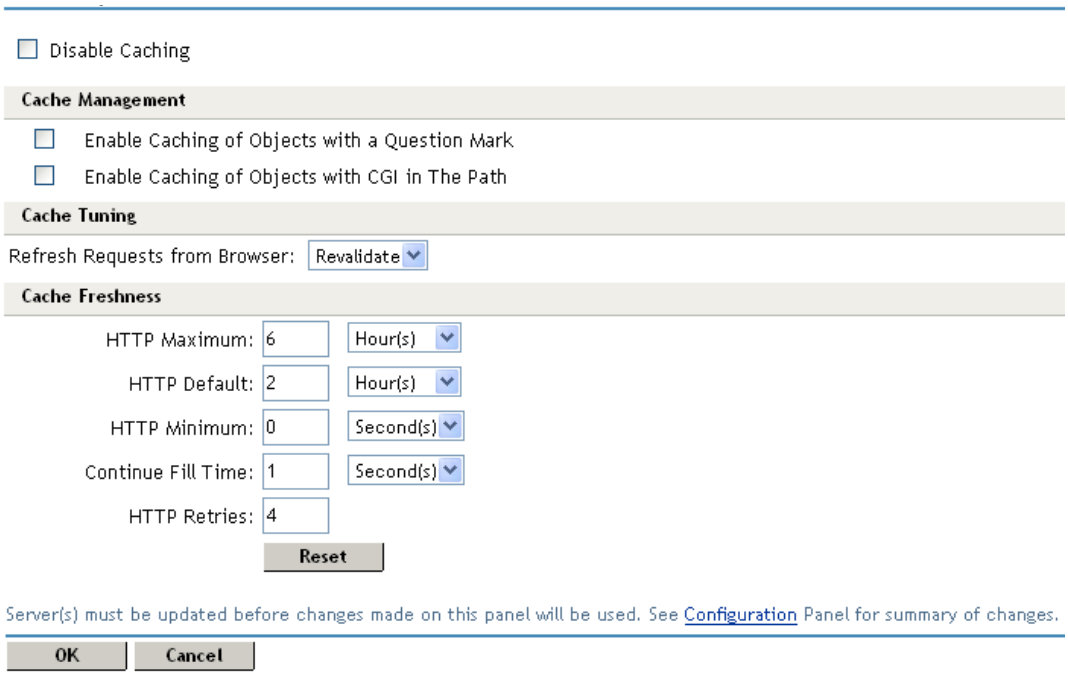

2 To disable caching of all Web server content, select *Disable Caching*.

When this option is selected, the other caching options for management, tuning, and freshness are disabled.

3 Configure the *Cache Management* options:

Enable Caching of Objects with a Question Mark: If this option is selected, a cacheable object is cached if it has a question mark in the URL.

Enable Caching of Objects with CGI in the Path: If this option is selected, a cacheable object is cached if it has `/cgi` in its URL.

Objects that meet these criteria are only cached if they are also cacheable objects. Web server administrators can mark objects as non-cacheable. When so marked, these objects are not cached, even when the above options are selected.

If you disable both of these options, it does not mean that objects with question marks or cgi in their paths cannot be cached. These objects can match some other criteria and be cached.

4 Configure the *Cache Tuning* options.

These options restrict or enable functionality that affects all the resources protected by the Access Gateway.

Refresh Requests from Browser: When a user clicks *Refresh* or *Reload* in the browser, this action sends a new request to the Web server. Select one of the following options to control how the Access Gateway handles the request:

- **Refill:** Causes the proxy service to send the request to the Web server
- **Revalidate:** Causes the proxy service to check whether the current information is valid. If it is, the currently cached information is returned. If it isn't valid, the request is forwarded to the Web server.
- **Ignore:** Causes the proxy service to ignore the request and send the data from cache without checking to see if the cached data is valid.

5 Modify the Cache Freshness settings for the Gateway Appliance. Use the *Reset* button to return these settings to their default values.

These options govern when the proxy service revalidates requested cached objects against those on their respective origin Web servers. If the objects have changed, the proxy service re-caches them.

HTTP Maximum: Specifies the maximum time the proxy service serves HTTP data from cache before revalidating it against content on the origin Web server. No object is served from cache after this value expires without being revalidated.

This overrides a freshness or Time to Expire directive specified by the 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.

HTTP Minimum: Specifies the minimum time the proxy service serves HTTP data from cache before revalidating it against content on the origin Web server. No requested object is revalidated sooner than specified by this value.

This overrides the freshness or Time to Expire directive specified by the Webmasters if they specified a shorter time.

You can use this value to increase the minimum time the proxy service waits before checking whether requested objects need to be refreshed. This parameter does not override No Cache or Must Revalidate directives from the origin Web server.

The default value is 0, which allows the proxy service to honor the Time To Expire directive of each object (unless it is longer than the *HTTM Maximum* option). If the *HTTP Minimum* option is set to a value other than 0, the value overrides any object's Time to Expire directive that is shorter than the value set. The default is 0.

Continue Fill Time: Specifies the how long the proxy service ignores browser request cancellations and continues downloading objects from the target Web server until the download is complete. The default is 1 second.

HTTP Retries: Specifies the number of retry requests to issue to a Web server. The default is 4 retries.

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

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

5.2 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*.

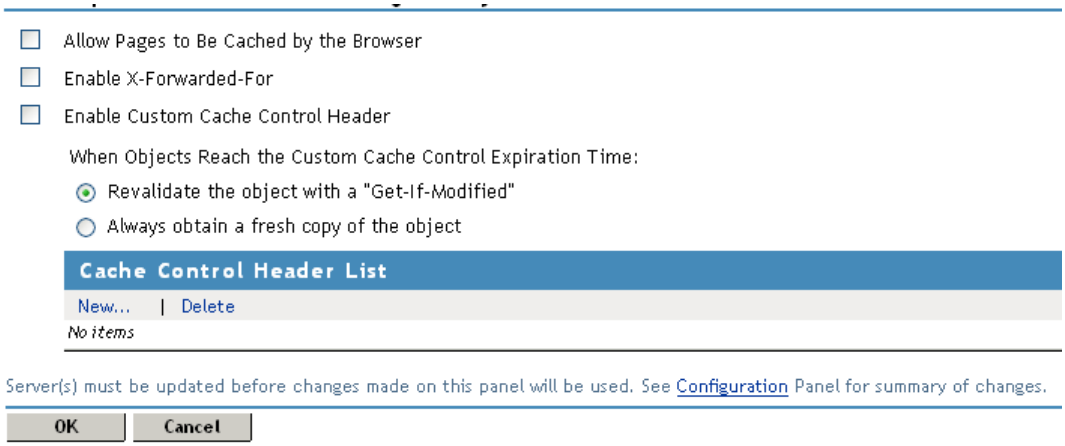

2 To mark all pages coming through this host as cacheable on the browser, select *Allow Pages to be Cached by the Browser*.

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 information about the *Enable X-Forwarded-For* option, see Section 3.10, “Configuring X-Forwarded-For Headers,” on page 100.

4 To configure custom caching instructions, see Section 5.3, “Configuring Custom Cache Control Headers,” on page 144.

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

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

5.3 Configuring Custom Cache Control Headers

In addition to fine-tuning cache freshness by using the HTTP timers, as explained in Section 5.1, “Configuring Caching Options,” on page 141, 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 5.3.1, “Understanding How Custom Cache Control Headers Work,” on page 145
- Section 5.3.2, “Enabling Custom Cache Control Headers,” on page 146

5.3.1 Understanding How Custom Cache Control Headers Work

Only the proxy service containing the custom header definition follows the cache policies specified in the custom headers.

All other proxy services, requesting browsers, and external proxy caches (transparent caches, client accelerators, etc.), do not recognize the custom headers. They follow only the cache policies specified by the standard cache control headers.

This means that you have the following options for configuring your Web server:

- You can specify that browsers and/or external caches cannot cache the objects, but the proxy service can.

  This lets you offload request processing from the origin Web server while still requiring that users return to the site each time they request an object.
- You can also specify separate cache times for browsers, external caches, and the proxy service.

To implement custom cache control headers, you must do the following:

- Configure a proxy service to use custom cache control headers by enabling the feature and specifying a header string such as MYCACHE (see Section 5.3.2, “Enabling Custom Cache Control Headers,” on page 146).
- 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
```

5.3.2 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* > *Header Options*.

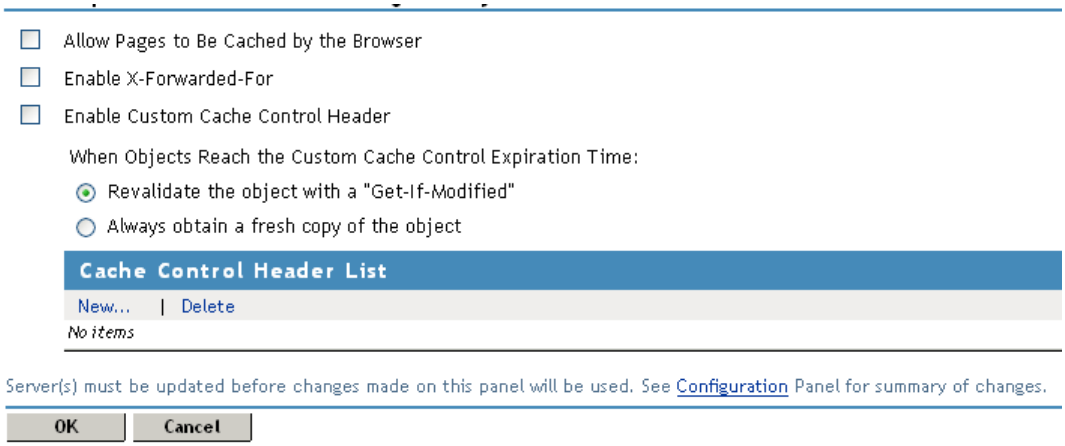

2 To enable the use of custom headers, select *Enable Custom Cache Control Header*.

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 5.3.1, "Understanding How Custom Cache Control Headers Work," on page 145).

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 due to their relatively brief time-to-live value.

5.4 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*.

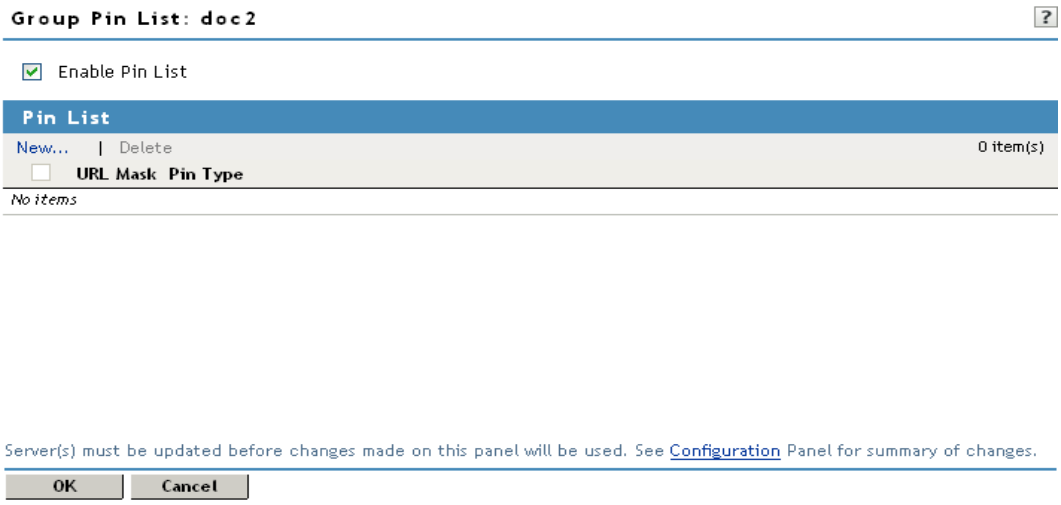

2 Select the *Enable Pin List* option to enable the use of pinned objects. If this option is not selected, the pinned objects in the pin list are not used.

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 Section 5.4.1, "URL Mask," on page 148.

Pin Type: Specifies how the URL is to be used to cache objects. Select from *Normal* and *Bypass*. For more information, see Section 5.4.2, "Pin Type," on page 150.

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*.

5.4.1 URL Mask

The URL mask can contain complete or partial URL patterns. A single URL mask might apply to a large set of URLs, or it might be so specific that only a single file on the Web matches it.

The Access Gateway processes the masks in the pin list in order of specificity. A mask containing a 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.

Level	Examples

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`
	
	Filenames are processed after paths. A leading forward slash must always be used when specifying a filename. If a path is included with a filename, the path must start with the root directory of the Web server, and the entry is processed as a path entry, not as a filename entry.

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.

Specific rules have precedence over less specific rules. Thus, objects matched by a more specific rule are always processed according to its conditions. If a less specific rule also matches the object, the less specific rule is ignored for the object. For example, assume the following two entries in the pin list:

URL Mask	Pin Type	Pin Links
`http://www.foo.gov/documents/*`	cache	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 (.com, .net, .org, etc.) whose DNS names begin with www.foo from being cached.

5.4.2 Pin Type

The pin type specifies how the Access Gateway caches objects that match the URL mask.

- **Normal:** The Access Gateway handles objects matching the mask in the same way it handles any other requested objects. In other words, the objects are cached but not pinned.

  Administrators often use this pin type in combination with a broad URL mask that has a bypass pin type. This allows them to insulate specific objects from the effects of the bypass rule.

  For example, you could specify a URL mask of `/*.jpg` with a pin type of bypass and a second URL mask of `www.foo.gov/graphics/*` 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.

5.5 Configuring a Purge List

The purge list is global to the Access Gateway and affects all protected resources. This option allows you to specify URL patterns or masks for the pages and sites whose objects you want to purge from cache.

When defining the masks, keep in mind that the Access Gateway interprets everything in the URL mask between the asterisk wildcard (*) and the following delimiter as a wildcard. Delimiters include the forward slash (/), the period (.), and the colon (:) characters. For example:

URL Mask	Effects
/*.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. You can configure the Access Gateway to use the pin list to add objects to the cache and to use the purge list to remove the same objects.

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

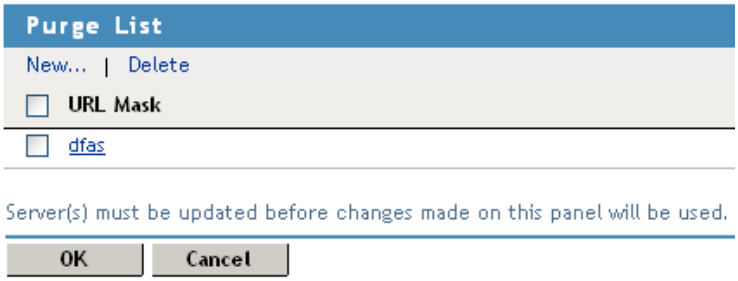

2 Click *New,* enter a URL pattern, then click *OK*.

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*.

5.6 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.

Protecting Multiple Resources

6

This section describes how to create multiple resources for the various Access Gateway components, including a cluster of Access Gateways. Figure 6-1 illustrates the relationships that Access Gateways, reverse proxies, proxy services, Web servers, and protected resources have with each other when two Access Gateways are members of a cluster.

Figure 6-1 *Hierarchical View of the Access Gateway Configured Objects*

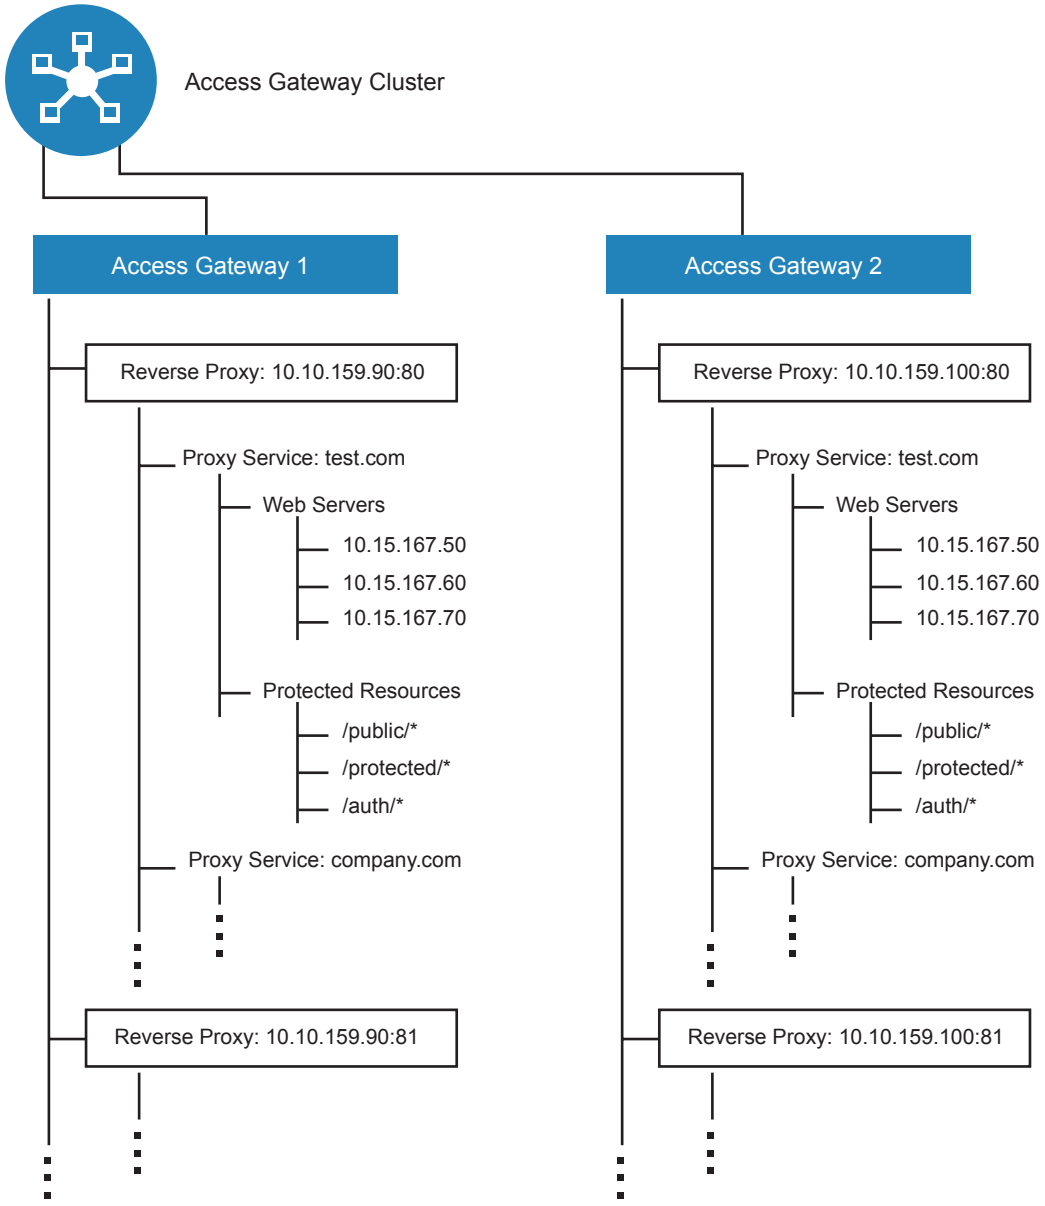

In Figure 6-1, Access Gateway 1 and Access Gateway 2 have the same configuration except for the reverse proxy listening address. They share the other configuration settings because they are members of an Access Gateway cluster. This section explains how to create a group of Web servers, how to add multiple proxy services and reverse proxies to an Access Gateway, and how to manage a cluster of Access Gateways.

- Section 6.1, “Setting Up a Group of Web Servers,” on page 154
- Section 6.2, “Using Multi-Homing to Access Multiple Resources,” on page 155
- Section 6.3, “Managing Multiple Reverse Proxies,” on page 164
- Section 6.4, “Managing a Cluster of Access Gateways,” on page 166

6.1 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 6-2 *Adding Redundant Web Servers*

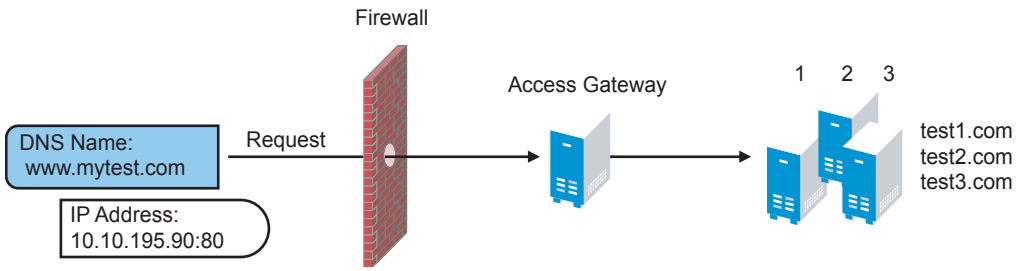

For this configuration, you use a single reverse proxy and proxy service. To add multiple Web servers to a host:

1. In the Administration Console, click *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 Steps 2 and 3 to add additional Web servers to the group.
5. To save your changes to browser cache, click *OK*.
6. To apply the changes, click the *Access Gateways* link, then click *Update* > *OK*.

The Access Gateway can perform a round robin, or it can be configured to perform a simple failover, sending 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 1.7.2, “Configuring TCP Connect Options for Web Servers,” on page 60.

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 1.7.2, “Configuring TCP Connect Options for Web Servers,” on page 60.

Session persistence is enabled whenever a second Web server is added to the list. This allows a browser to persistently use the same Web server after an initial connection has been established. This type of persistence is not configurable. For more information on persistent connections, see Section 1.7.3, “Configuring Connection and Session Persistence,” on page 62.

6.2 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 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 6.2.1, “Domain-Based Multi-Homing,” on page 155
- Section 6.2.2, “Path-Based Multi-Homing,” on page 157
- Section 6.2.3, “Virtual Multi-Homing,” on page 159

This section describes these multi-homing methods, then explains the following:

- Section 6.2.4, “Creating a Second Proxy Service,” on page 160
- Section 6.2.5, “Configuring a Path-Based Multi-Homing Proxy Service,” on page 162

6.2.1 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 6-3 illustrates this type of configuration for three proxy services.

Figure 6-3 *Using a Base Domain Name with Host Names*

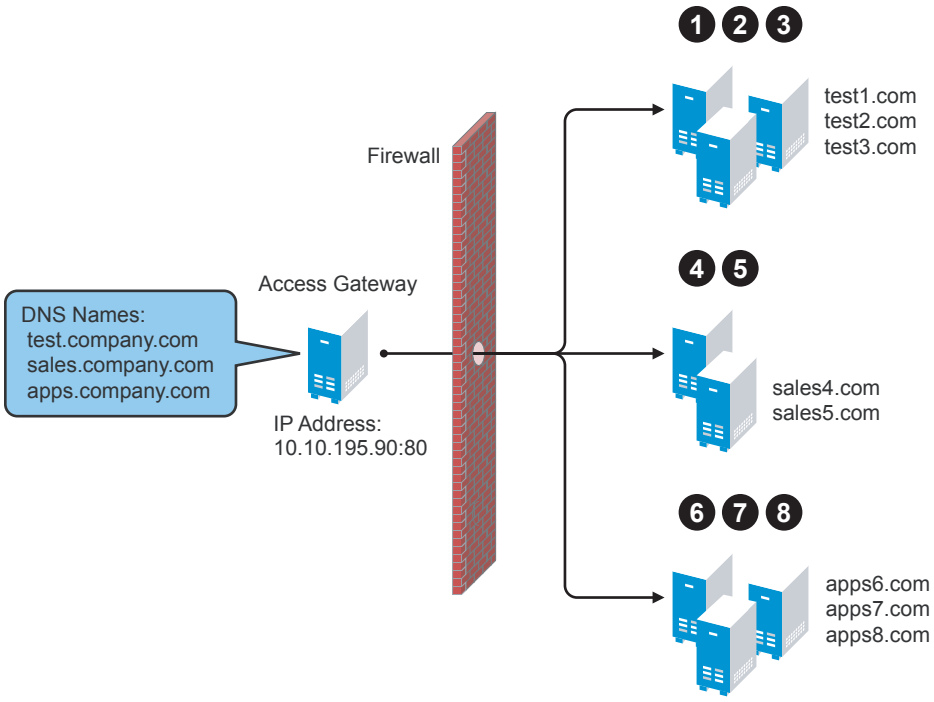

Domain-based multi-homing has the following characteristics:

- If you are using SSL, the back-end servers can all listen on the same SSL port (default for HTTPS is 443).
- If you are using SSL, the back-end servers can share the same SSL certificate. Instead of using a specific 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 and then 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.

To create a domain-based multi-homing proxy service, see Section 6.2.4, “Creating a Second Proxy Service,” on page 160, and select domain-based for the multi-homing type.

6.2.2 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 6-4 illustrates this type of configuration.

Figure 6-4 *Using a Domain Name with Path Elements*

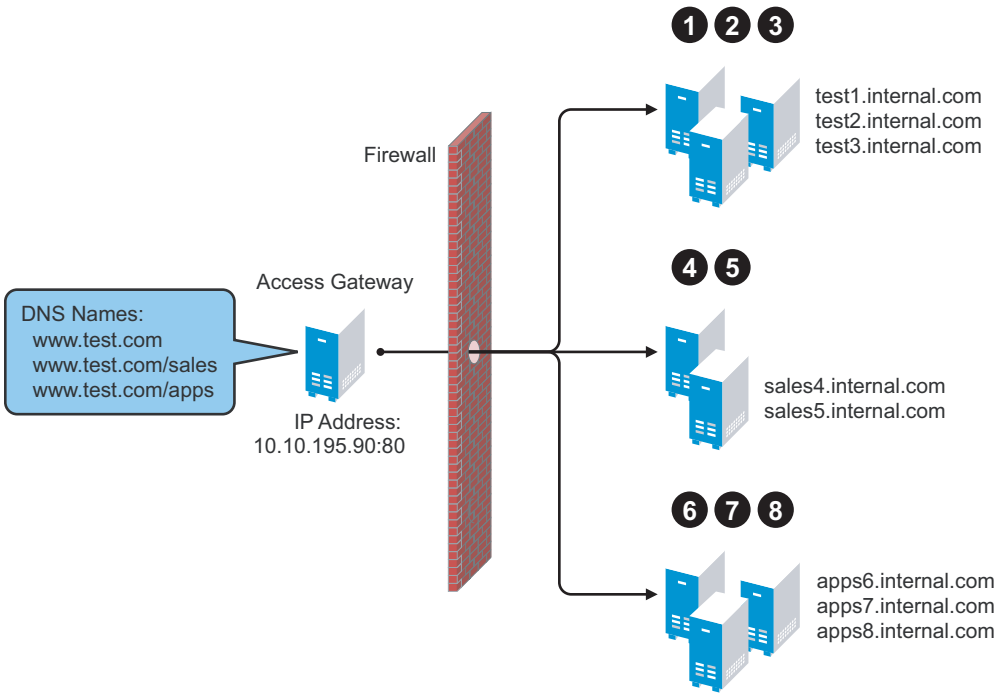

Path-based multi-homing has the following characteristics:

- It is considered to be more secure than domain-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 158
- “Configuring the Host Header Option” on page 158
- “Configuring for Path-Based Multi-Homing” on page 159

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 6-4 on page 157 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 6-4 on page 157 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 44.

Configuring 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.

To create a path-based multi-homing proxy service, see Section 6.2.4, “Creating a Second Proxy Service,” on page 160, and select path-based for the multi-homing type.

6.2.3 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 6-5 illustrates this type of configuration.

Figure 6-5 *Using Multiple DNS Names*

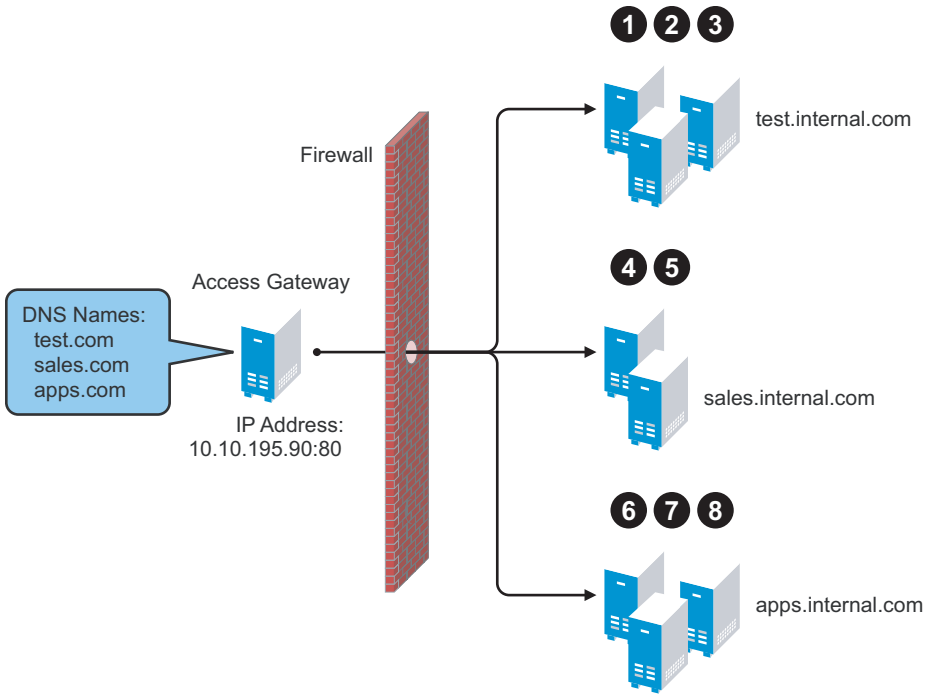

Virtual multi-homing cannot be used with SSL. You should use this configuration with resources that need to be protected, but the information exchanged should be public information that does not need to be secure. For example, you could use this configuration to protect your Web servers that contain the catalog of your shipping products. It isn't until the user selects to order a product that you need to switch the user to a secure site.

Whether a client can use one DNS name or multiple DNS names to access the Access Gateway depends upon the 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 6.2.4, "Creating a Second Proxy Service," on page 160, and select *Virtual* for the multi-homing type.

6.2.4 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*.

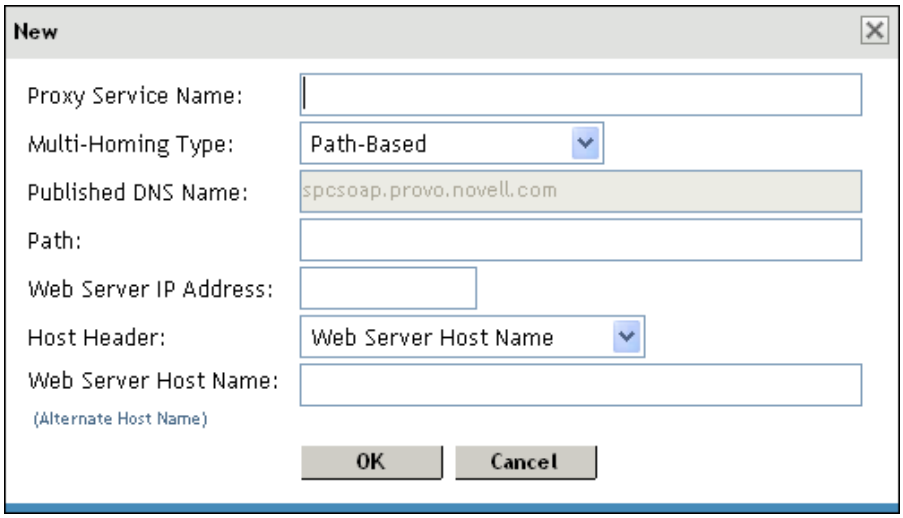

3 Fill in the fields.

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 6.2.1, "Domain-Based Multi-Homing," on page 155.
- **Path-Based:** Uses the published DNS name (www.test.com) with a path (www.test.com/path). For more information, see Section 6.2.2, "Path-Based Multi-Homing," on page 157.
- **Virtual:** Uses a unique DNS name (www.newsite.newcompany.com). Virtual multi-homing cannot be used with SSL. For more information, see Section 6.2.3, "Virtual Multi-Homing," on page 159. If you need a unique DNS name and SSL, you need to create a reverse proxy rather than a proxy service. For information on creating a second reverse proxy, see Section 6.3, "Managing Multiple Reverse Proxies," on page 164.

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 158.

Web Server Host Name: Specify the DNS name of the Web server that the Access Gateway should forward to the Web server. If you have set up a DNS name for the Web server and the Web server requires its DNS name in the HTTP header, specify that name in this field. If you selected *Forward Received Host Name*, this option is not available.

NOTE: For iChain® administrators, the *Web Server Host Name* is the alternate 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 1.2, "Configuring a Proxy Service," on page 16.
- To configure a path-based proxy service, see Section 6.2.5, "Configuring a Path-Based Multi-Homing Proxy Service," on page 162.

6.2.5 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]*.

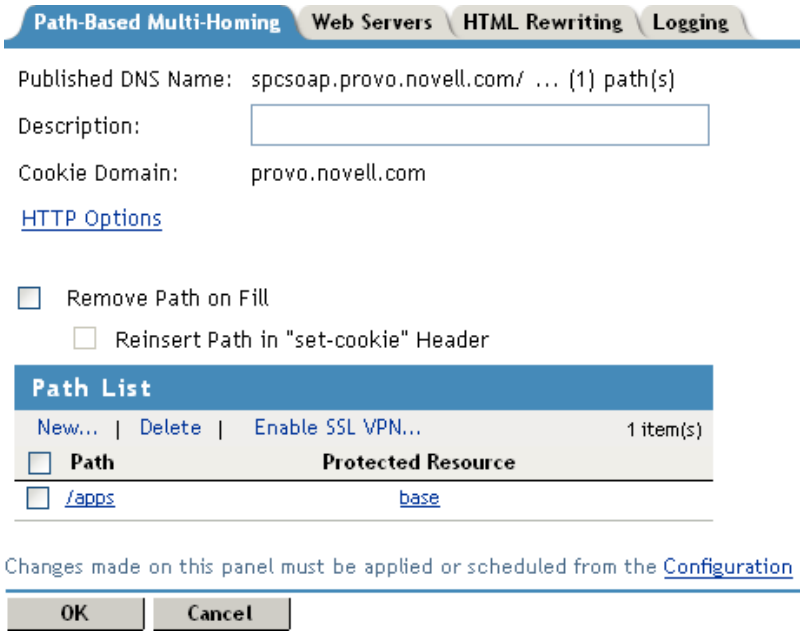

The following fields display information that must be configured on the parent proxy service (the first proxy service created for this reverse proxy).

- **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 5.3, "Configuring Custom Cache Control Headers," on page 144 and Section 5.2, "Controlling Browser Caching," on page 143.

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.

4c To return to the Path-Based Multi-Homing page, click the *Overview* tab, then click *OK*.

- If the protected resource meets your needs, continue with Step 5
- If it does not meet your needs, you must create a protected resource for the path-based proxy service. Continue with Step 4d.

4d Click *OK*, the name of the parent proxy service, then *Protected Resources*.

4e In the *Protected Resource List*, click *New*, specify a name, then click *OK*.

4f Assign a contract.

4g In the *URL Path List*, specify the path you used when creating the path-based proxy service. 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 To save your changes to browser cache, click *OK*.

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

6.3 Managing Multiple Reverse Proxies

Each reverse proxy must have a unique IP address and port combination. If your Access Gateway has only one IP address, you must select unique port numbers for each additional reverse proxy that you create. You can configure the Access Gateway to use multiple IP addresses. These addresses can be configured to use the same network interface card, or if you have installed multiple network cards, you can assign the IP addresses to different cards. To configure IP addresses and network interface cards, see Section 3.8.1, “Viewing and Modifying Adapter Settings,” on page 90.

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.

6.3.1 Managing Entries in the Reverse Proxy List

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

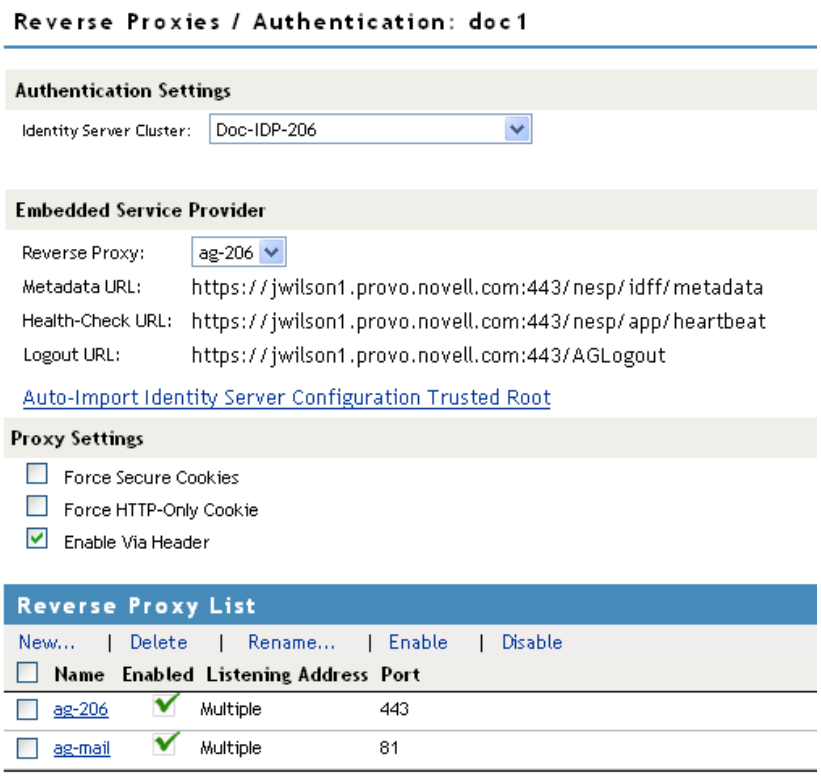

2 In the *Reverse Proxy List*, select one of the following actions:

- **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 1.1, "Creating a Reverse Proxy and Proxy Service," on page 12.

  Reverse proxy names and proxy service names must be unique to the Access Gateway. Protected resource names need to be unique to the proxy service, but they don't need to be unique to the Access Gateway.

- **Delete:** To delete a reverse proxy, select the check box by a specific reverse proxy, then click *Delete*. To delete all reverse proxies, select the check box by the *Name* column, then click *Delete*.

- **Enable:** To enable a reverse proxy, select the check box by a specific reverse proxy, then click *Enable*. To enable all reverse proxies, select the check box by the *Name* column, then click *Enable*.

- **Disable:** To disable a reverse proxy, select the check box by a specific reverse proxy, then click *Disable*. To enable all reverse proxies, select the check box by the *Name* column, then click *Disable*.

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

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

6.3.2 Changing the Authentication Proxy Service

If you have multiple reverse proxies, you can select the reverse proxy that users are redirected to for login and logout.

IMPORTANT: Changing the reverse proxy that is used for authentication is not a trivial task. For example, if you have customized the logout options on your Web servers to redirect the logout request to the Logout URL of the current authentication reverse proxy, you need to modify these options to point to a new Logout URL.

If you have set up SSL connections, you need to change your certificate configurations.

To select the reverse proxy to use for authentication:

1. In the Administration Console, click *Devices* > *Access Gateways* > *Reverse Proxy / Authentication*.
2. In the *Embedded Service Provider* section, select a value for the *Reverse Proxy* option. This is the reverse proxy that is used for authentication.

   The screen is refreshed and the *Metadata URL*, *Health-Check URL,* and *Logout URL* are rewritten to use the selected reverse proxy.
3. (Conditional) If your Access Gateway certificates were generated by a different certificate authority than your Identity Server certificates, you need to import the trusted root of the Identity Server into the trusted root keystore of the Embedded Service Provider. Click *Auto-Import Identity Server Configuration Trusted Root*, click *OK*, specify an alias, click *OK*, then click *Close*.

   If you don't know whether you need to import the trusted root, click the option. If the trusted root is already in the keystore, the duplicate key is not imported and you are informed of this condition.
4. In the *Reverse Proxy List*, click the name of the reverse proxy that you have selected for authentication.
5. If you have enabled SSL between the Embedded Service Provider and the Identity Server, you need to import the trusted root of the Embedded Service Provider into the trusted root keystore of the Identity Server. Click *Auto-Import Embedded Service Provider Trusted Root*, click *OK*, specify an alias, click *OK*, then click *Close*.

   If you don't know whether you need to import the trusted root, click the option. If the trusted root is already in the keystore, the duplicate key is not imported and you are informed of this condition.
6. To save your changes to browser cache, click *OK*.
7. To apply the changes, click the *Access Gateways* link, then click *Update* > *OK*.
8. (Conditional) If you have customized Web logout pages, update them to use the new Logout URL.

6.4 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. (For information on how to create a cluster of Access Gateways, see "Clustering Access Gateways" in the *Novell Access Manager 3.1 SP1 Setup Guide*.) This section describes the tasks that are specific to managing the servers of an existing cluster:

For information about monitoring the health or statistics of a cluster, see Section 4.6, “Managing Server Health,” on page 134 and Section 4.3, “Monitoring Access Gateway Statistics,” on page 118.

6.4.1 Managing the Servers in the Cluster

To view the servers that are currently members of clusters:

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

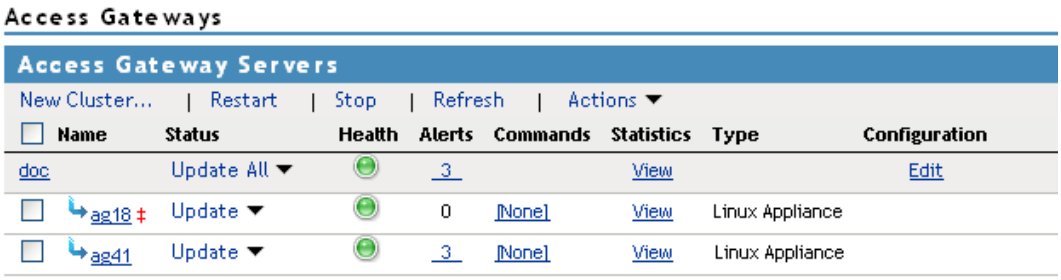

   The members of a cluster are listed under the cluster name. The red double dagger symbol identifies the server that is the primary cluster server.

2. To add a server to a cluster, select the server, then click *Actions* > *Assign to Cluster* > *[Name of Cluster]*.

3. To remove a server from a cluster, select the server, then click *Actions* > *Remove from Cluster*.

   Usually when you delete a server from a cluster, you have discovered that traffic is lighter than anticipated and that it can be handled with fewer machines while another cluster is experiencing higher traffic and can benefit from having another cluster member. When the server is removed, its configuration object maintains all the configuration settings from the cluster. When it is added to a new cluster, its configuration object is updated with the configuration settings of the new cluster. If your clusters are behind an L4 switch, you need to reconfigure the switch so that the server is assigned to the correct cluster.

   When a server is removed from a cluster, its Embedded Service Provider is stopped. If you are not going to assign it to another cluster, you need to reconfigure the server so that it is protecting resources other than the ones it protected in the cluster. When you apply the changes by clicking *Update*, the Embedded Service Provider is restarted.

   You cannot remove the primary cluster server unless it is the only server in the cluster. If you need to remove the primary cluster server from a multiple server cluster, you need to assign another the server to be the primary cluster server.

4. To modify which server is the primary cluster server, see Section 6.4.2, “Changing the Primary Cluster Server,” on page 168.

5. To view detailed information about a server in the group, click the name of the server.

6. To view detailed health information about a server, click the health icon of the server. For more information, see Section 4.6.2, “Monitoring the Health of an Access Gateway,” on page 135.

7. Click *Close*.

6.4.2 Changing the Primary Cluster Server

If the current primary cluster server is down and will be down for an extended period of time, you should select another server to be the primary cluster server

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

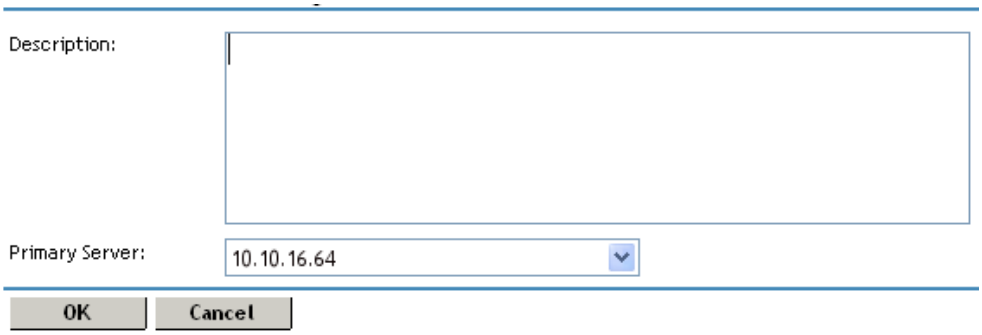

2 In the *Primary Server* drop-down list, select the name of a server, then click *OK*.

Please be patient. Wait until this configuration change has completed, before doing any other configuration updates.

3 To update the Identity Server, click *Identity Servers* > *Update*.

6.4.3 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.

When servers are in a cluster, you might want to update only one server in the cluster and verify that the changes are behaving as expected. If this is your plan, we highly recommend that you save the proposed changes to the configuration datastore so the changes are not lost. If 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.

After testing the configuration on one server, you can then apply the saved changes to the other servers in the cluster, either individually (with the *Update* link) or as group (with the *Update All* link).

If you discover that the configuration change is not behaving the way you want it to, you can revert back to the previous applied configuration by doing the following:

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.

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:

- 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).

Troubleshooting the Linux Access Gateway 7

For a solution to an Access Gateway problem, see the following sections:

For information about policy errors, see “Troubleshooting Access Manager Policies” in the *Novell Access Manager 3.1 SP1 Policy Management Guide*.

For XML validation errors, see “Troubleshooting XML Validation Errors” in the *Novell Access Manager 3.1 SP1 Administration Console Guide*.

For information about installation, reinstallation, and import issues, see “Troubleshooting a Linux Access Gateway Appliance Installation” and “Troubleshooting the Access Gateway Import” in the *Novell Access Manager 3.1 SP1 Installation Guide*.

For information on how to install security patches on your Linux Access Gateway, see “Installing or Updating the Latest Linux Patches” in the *Novell Access Manager 3.1 SP1 Installation Guide*.

7.1 Useful Tools for Troubleshooting the Linux Access Gateway

7.1.1 Useful Tools

Table 7-1 describes some of the tools available in the Linux operating system or installed by the Linux Access Gateway that can help you determine the cause of a problem.

Table 7-1 *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 select to push the current configuration in the Administration Console 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 4.6.2, "Monitoring the Health of an Access Gateway," on page 135.
`curl`	Use this command to view identity provider metadata from the Linux Access Gateway. See "Testing Whether the Provider Can Access the Metadata" in the *Novell Access Manager 3.1 SP1 Identity Server Guide*.
`tail -f`	Use this command to view real time activity in key log files. For information on useful files to tail, see "Useful Files for Troubleshooting the Access Gateway Appliance" on page 176.
proc	Use this command to check resources available on the system.
`netstat /ss`	Use this command to view statistics about the listeners on the Linux Access Gateway.
`netcat`	Use this command to access the Linux Access Gateway console, which displays statistics and information about various processes. For more information, see "The Linux Access Gateway Console" on page 173.
`tcpdump`	Use this command to capture data on standard and loopback interfaces and to view SSL data with imported keys.
`nash`	Use this command to manually configure log level verbosity and replace IP addresses. For log level information, see "Gateway Appliance Logs" on page 107.
`/etc/init.d/novell-vmc`	Use the `novell-vmc` command line options to restart the proxy and view status. For more information, see Table 7-2 on page 173.
The `/chroot/lag/opt/novell/bin` directory contains the following scripts:	
`getlaglogs.sh`	Generates a `/var/log/laglogs.tar.gz` file of the install and system log files. For more information, see "Linux Access Gateway Logs" on page 196.

Tool	Description
`lagupgrade.sh`	Use this script to apply patches. For more information, see "Upgrading the Linux Access Gateway Appliance" in the *Novell Access Manager 3.1 SP1 Installation Guide*.
`lagconfigure.sh`	Use this script to resolve auto-import issues. For more information, see "Triggering an Import Retry" in the *Novell Access Manager 3.1 SP1 Installation Guide*.

You can use the following commands to stop and start the Linux Access Gateway and to view its status.

Table 7-2 *novell-vcm Commands*

Command	Description
`/etc/init.d/novell-vmc start`	Starts the Linux Access Gateway.
`/etc/init.d/novell-vmc stop`	Stops the Linux Access Gateway.
`/etc/init.d/novell-vmc status`	Displays the Linux Access Gateway status.
`/etc/init.d/novell-vmc restart`	Stops and starts the Linux Access Gateway.

7.1.2 The Linux Access Gateway Console

1. To access the console, run the following command:

   ```
   netcat localhost 2300
   ```

2. Press Enter at the `Please enter terminal type` prompt.

   This displays the Linux Access Gateway console screens.

```
PLEASE NOTE:
Use of these screens is not offically supported.  Statistics contained herein
may not be accurate, and debugging options may affect system performance or
stability.  Use at your own risk.

1. Work Scheduler Screen
2. System Console
3. Callout Scheduler Console
4. Novell SSL Stack Screen
5. Novell SSL Server Handshake Screen
6. Novell SSL Client Handshake Screen
7. Novell SSL Performance Screen
8. CCAgent Console
9. Sockets Interface Screen
10. Sockets Interface Screen
11. USTL Console
12. Proxy Messages
13. Proxy Console
14. VXE Callout Scheduler

Pick a screen:
```

Most of the time, the Proxy Console screen is the one you should pick. The other screens are used mainly by the developers of the Linux Access Gateway. If you are having SSL connection problems, the SSL screens can help in diagnosing the problem.

3 To access the Proxy Console screen, enter 13.

```
Novell LAG Proxy Console

   1. Display current activity
   2. Display memory usage
   3. Display ICP statistics
   4. Display DNS options
   5. Display cache statistics
   6. Display not cached statistics
   7. Display HTTP server statistics
   8. Display HTTP client statistics
   9. Display connection statistics
  10. Display FTP client statistics
  11. Display GOPHER client statistics
  12. Display configured addresses and services
  13. Display SOCKS client statistics
  14. Application Proxies
  15. Transparent Proxy statistics
  16. Site download options
  17. Debug options
  18. Identity Agent Console

Enter option:
```

4 To access a specific screen, enter the number.

Screen	Description
1. Display current activity	Displays information about connections (server and client), cached objects, and HTTP requests.
2. Display memory usage	Displays information about memory pools and memory used and the types of objects stored in memory.
3. Display ICP statistics	Displays statistics for the Internet Cache Protocol.
4. Display DNS options	Displays statistics and information about the entries in the DNS table.
5. Display cache statistics	Displays information about cached objects and the COS partition. For more information, see “Checking if the COS Partition Is Mounted” on page 189.
6. Display not cached statistics	Displays statistics about requests for objects that cannot be cached.
7. Display HTTP server statistics	Displays statistics about the server handling of HTTP requests.
8. Display HTTP client statistics	Displays statistics about the client handling of HTTP requests.

Screen	Description
9. Display connection statistics	Displays general information about connections.
10. Display FTP client statistics	Displays statistics about FTP client requests.
11. Display GOPHER client statistics	Displays statistics about GOPHER requests.
12. Display configured addresses and services	Displays information about the IP addresses that the Access Gateway is using.
13. Display SOCKS client statistics	Displays statistics about SOCKS client requests.
14. Application Proxies	Displays proxy service statistics.
15. Transparent Proxy statistics	Displays transparent proxy statistics.
16. Site download options	Displays information about the last download and prompts for information to schedule a new download.
17. Debug options	Allows you to control cache purging.
18. Identity Agent Console	Displays user information. For more information about the user screen, see "User Details" on page 201.

5 To return to the opening page of the console from other console page, press Esc+Enter.

This keystroke works only on some pages.

6 To exit the console, press Ctrl+C.

7.1.3 Viewing Configuration Information

The configuration store maintains two versions of the Access Gateway configuration and the browser cache maintains one.

- **Current:** The current configuration is the version of the configuration that the Access Gateway is currently using.

  You can view this configuration in file format by clicking *Access Manager > Access Gateways > [Name of Server] > Configuration > Export*. Do not set a password to encrypt the file. The exported file contains the current configuration.
- **Working:** The working configuration is the version that you have saved by clicking the *OK* button on the Server Configuration page, but you have not applied the changes by clicking the *Update* or the *Update All* link on the Access Gateways page. This version is not viewable from the Administration Console.
- **Browser Cache:** All configuration changes are saved to browser cache when you click the *OK* button on a configuration page. To view the configuration currently in browser cache, click *Access Manager > Auditing > Troubleshooting*, scroll to the *Cached Access Gateway Configurations* section, then click *View*. You can view the cached configuration of an individual Access Gateway, or if the Access Gateway is a member of a cluster, you can view the cached configuration of the cluster and each member. The + and - buttons allow you to expand and collapse individual configurations.

7.2 Useful Files for Troubleshooting the Access Gateway Appliance

7.2.1 Viewing Log Files

Table 7-3 describes the Linux Access Gateway files that contain troubleshooting information.

Table 7-3 *Log Files with Troubleshooting Information*

Log File	Description
`catalina.out`	Located in the `/var/opt/novell/tomcat5/logs` directory and available from the General Logging page in the Administration Console. The Embedded Service Provider, which communicates with the Identity Server, writes to this log file. The log level is controlled by the Identity Server Configuration. For configuration information, see “Turning on Logging for Policy Evaluation” in the *Novell Access Manager 3.1 SP1 Policy Management Guide*. For information on how to use the entries for policy troubleshooting, see “Troubleshooting Access Manager Policies” in the *Novell Access Manager 3.1 SP1 Policy Management Guide*.
`ics_dyn.log`	Located in the `/var/log` directory and available from the General Logging page in the Administration Console. The proxy service writes to this log file. For information on enabling logging to this file, see “Gateway Appliance Logs” on page 107. For maximum verbosity, the proxy service must be started in debug mode. See Table 7-2, “novell-vcm Commands,” on page 173.
`lagsoapmessages`	Located in the `/var/log` directory and available from the General Logging page in the Administration Console. When enabled, this file contains a log of the SOAP messages between the Linux Access Gateway and the Embedded Service Provider for authentication (roles, contracts, and timeouts) and policy interaction (Authorization, Form Fill, and Identity Injection). For information on enabling logging to this file, see “Configuring Logging of SOAP Messages and HTTP Headers” on page 109.
`laghttpheaders`	Located in the `/var/log` directory and available from the General Logging page in the Administration Console. When enabled, this file contains a log of the HTTP headers to and from the Linux Access Gateway. For information on enabling logging to this file, see “Configuring Logging of SOAP Messages and HTTP Headers” on page 109.

7.2.2 Using Touch Files

Table 7-4 describes the touch files that control configuration options for Linux Access Gateway that aren't available from the Administration Console. Filenames are case-sensitive.

Table 7-4 *Touch Files*

Filename	Description
`.~newInstall`	Located in the `/var/novell` directory. The Linux Access Gateway creates this file by default during every start. If you want the Linux Access Gateway to come up without the contents cached in the previous run, or to purge all cache, remove this file before you restart the Linux Access Gateway.
`.modVia`	Located in the `/var/novell` directory. Adds the device ID in the Via header that is sent by the Linux Access Gateway to the Web server. The Linux Access Gateway sends the Via header in the following format: `Via: 1.0 www.mylag.com (Access Gateway 3.0.1-72-D06FBFA8CF21AF45)`
`.enableInPlaceSilentFill`	Located in the `/var/novell` directory. To be used for the Linux Access Gateway Form Fill. When this touch file is used, the login page is not modified.This enables single sign-on to certain Web sites that require the login page to remain as is without any modifications to its structure. When this touch file is used, the Linux Access Gateway does not generate a new page if autosubmit is enabled, but fills the page received from the Web server and hides the text/password/unspecified type fields. Form-Fill issues for CRM applications and teaming and conferencing applications are resolved with this touch file. However, when this touch file is used, the *Debug Submit* and *JS Functions to Keep* options of the Form Fill policy do not work.

Filename	Description
.enableInPlaceSilentFillNew	Located in the /var/novell/ directory. This touch file is to be used to fill forms with complex JavaScript or VBScripts. You must use this touch file along with the .enableInPlaceSilentFill file. To use this file, the Form Fill policy must have the *Statements To Execute on Submit* option enabled and the policy must contain a function to execute as shown in the following example: `function anynamefunction() { ...some statements... } function executeJavaScript() { <<... any functions you want to be called..>> document.forms[0].submit(); }`
lagDisableAuthIPCheck	Located in the /etc directory. Enabling this touch file switches off the proxy authentication cookie binding to client IP. Use this in a setup where two L4 switches are configured in parallel and the browser requests get bounced between the these L4 switches.
.alwaysUseJSFor302	Located in the /var/novell directory. Uses JavaScript for redirection. A 200 OK response is sent back with the redirect metatag instead of the 302 redirect, when this touch file is used.
.useJSFor302withIE7	Located in the /var/novell directory. When Internet Explorer 7 browser is used, 200 OK response is sent back with the redirect metatag instead of the 302 redirect.
.useRelativeUrlInJS	Located in the /var/novell directory. Sends back the 200 OK response with the metatag redirect header referencing a relative URL rather than full URL (scheme, host, path). This touch file should be used when .useJSFor302withIE7 and alwaysUseJSFor302 files are used.
.useHTMLBodyIn302	Located in the /var/novell directory. The Linux Access Gateway sends 302 redirects without any content by default. When this file is present, the following content is sent for any 302 redirects: `<html><head><title>Redirection</title></head><body>Your browser should support redirection.</body></html>`

Filename	Description
.forceUTF8CharSet	Located in the /var/novell directory. When this file is enabled, the Linux Access Gateway serves the Form Fill page to the browser in the UTF-8 character set.
.ignoreDnsServerHealth	Located in the /var/novell directory. Ignores the DNS server health status while reporting health to the Administration Console.
.EnableSecureCookie	Located in the /var/novell/ directory. Adds the word secure at the end of set-cookie so that only HTTPS sites can access it. This file works when the Force Secure Cookie option is disabled in the Administration Console.
.noURLNormalize	Located in the /var/novell/ directory. Disables the URL normalization protection for back-end Web servers. This touch file resolves issues in serving Web content from Web servers which had double byte characters such as Japanese language characters.
.AllowUnknownHTTPMethods	Located in the /var/novell/ directory. When this file is present, the Linux Access Gateway forwards any unknown HTTP methods to the Web server.
.noGzipSupport	Located in the /var/novell directory. Disables GZIP functionality in the Linux Access Gateway. This ensures that the Linux Access Gateway does not send Accept-Encoding: gzip deflate headers to the Web server.
.useAlternate	Located in the /opt/novell/conf/keys directory. This file can be used when you have problems with the SSL listeners in the Linux Access Gateway. The following error message is displayed in the ics_dyn.log: NiciStore unprotect data failed When you use this file, re-push the certificates used by the Linux Access Gateway listeners, apply the changes, then restart the Linux Access Gateway.
.doNotUseTLS	Located in the /var/novell directory. Use this touch file if there is a problem in accelerating Oracle application servers. After creating the touch file, restart the Linux Access Gateway. When this file is enabled, it prevents the Linux Access Gateway from using TLS to communicate with the back-end Web servers.

Filename	Description
.ForceHTTPSSchemeInESPRedirection	Located in the /var/novell directory. Forces the Linux Access Gateway to always return the URL in the HTTPS schema. Use this if the Linux Access Gateway is located behind an SSL terminator. In this case, the original URL accessed by the browser is rewritten with the HTTPS scheme. This ensures that the traffic is sent back to the browser after the authentication contains the right protocol (SSL/TLS).
.overwrite_AuthHeader_With_IIData	Located in the /var/novell directory. This touch file ensures that when a browser sends an authentication header, the Linux Access Gateway Appliance overwrites it with the authentication header configured in the Identity Injection policy.
.PasswordMgmt	Located in the /var/novell directory. Use this touch file to refresh the user's credentials to match password changes. You must use this touch file if you have configured resources to use Identity Injection policies to inject the user's password and the Identity Server is configured to use a password management service. If this touch file is not enabled, when users authenticate and change their credentials, the Access Gateway uses the old password for identity injection.
.enableichaincompatibility	Located in the /var/novell directory. Does protected resource matching, similar to iChain.
.matchLagIchainCookieName	Located in the /var/novell directory. Forwards a proxy session cookie to a back-end application. Cookie without a touch file looks like: IPCZQX03a36c6c0a=*xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx* Cookie with a touch file looks like: IPCZQX01a36c6c0a=*xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx*

Filename	Description
.spnetworkplaces	Located in the /var/novell directory. Helps user to use Microsoft Network Places client to connect to WebDAV folders such as SharePoint that is accelerated by Linux Access Gateway as path-based multi-homing service. For this touch file to function as specified, you should add the following lines to the file, and restart Linux Access Gateway. SHAREPOINTPATH=/<accelerated path> HOSTNAME=<accelerated host name>
.AllowMSWebMiniRedir	Located in the /var/novell directory. Helps the user to disable the following functionality which is enabled by default: If a client (Windows network place) sends an OPTIONS request with MS-WebDAV-MiniRedir user-agent to Linux 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.
.reqPostSize	Located in the /var/novell directory. Helps user to specify POST size up to 50 MB. POST size defaults to 1 MB without this touch file. In this touch file, configure the POST size as: REQPOSTSIZE=*<value in terms of MB>* Even if you specify a value greater than 50 MB, the value limits to 50 MB.
.rewriteAlwaysHTTPS	Located in the /tmp directory. If this touch file is enabled, Linux Access Gateway rewrites all HTTP links to HTTPS while serving to the browser.

Filename	Description

`.modifyRequestURI`	Located in the `/var/novell` directory.
	
	When clients use Internet Explorer and MS office 2007 to access SharePoint resources protected by Access Gateway, some requested URLs are not sent to the correct path-based proxy service.
	
	For example, assume that the SharePoint server is accelerated by a reverse proxy service https://`sharepoint.CompanyA.com/share1`. The browser, instead of sending the request URL as `https://sharepoint.CompanyA.com/share1/_vti_bin/webs.asmx`., sends the URL as `https://sharepoint.CompanyA.com/_vti_bin/webs.asmx`, without the path `/share1`.This causes Access Gateway to serve request to the wrong service.
	
	To workaround this problem, configure the `.modifyRequestURI` file with the following information:
	
	◆ Published DNS name of the proxy service accelerating the SharePoint server.
	◆ URLs that require path injection in the request URL.
	◆ Path (or paths) of the SharePoint service under this proxy service, that must be prepended to the listed URLs.
	
	An example file looks similar to the following:
	
	`HOSTNAME=sharepoint.CompanyA.com`
	`PATH1=/share1`
	`PATH2=/share2`
	`URL1=/_vti_bin/webs.asmx`
	`URL2=/_vti_bin/lists.asmx`
	`URL3=/_vti_bin/Copy.asmx`
	`URL4=/_vti_inf.html`
	
	NOTE: If you are adding multiple paths, make sure that these path-based services belong to the same domain.
	
	When this file is present with the required configuration, the incoming request URL is compared with the URLs in the touch file. If a match is found and the host name of request URL matches the HOSTNAME value, then the following occurs:
	
	◆ If only one path is configured, the path is injected to the request, and the request is sent to this path-based service.
	◆ If multiple paths are configured, Access Gateway looks for the last path-based service accessed by this user. This path is injected to the request, and the request is sent to this path-based service.

Filename	Description
	For example, if the last resource accessed by User A is `https://sharepoint.CompanyA.com/share2/` and the next URL request is `https://sharepoint.CompanyA.com/_vti_bin/webs.asmx`, the request URL is changed to `https://sharepoint.CompanyA.com/share2/_vti_bin/webs.asmx` and the request is sent to `/share2`.
`.setsecureESP`	Located in the `/var/novell` directory. When this touch file is used, the JSESSIONID cookie of the Embedded Service provider is marked as secure. To enable this touch file, you need one of the following: ◆ All services that need authentication must use the secure communication channel or HTTPS. ◆ Access Gateway device must be behind an SSL terminator. For more information, see Section 2.5.1, “Securing the Embedded Service Provider Session Cookie,” on page 71
`.releaseclosewait`	Located in the `/var/novell` directory. This touch file is useful if Linux Access Gateway goes into a non-responsive mode and there are a large number of connections in the close_wait state in the public listener port. You can use this touch file to enable the forced cleanup of connections in a close_wait state.

The Linux Access Gateway must be restarted in order to get the desired functionality. Use the following command to restart when a touch file is created or removed:

```
/etc/init.d/novell-vmc stop

/etc/init.d/novell-vmc start
```

Creating a File

To create a file, use the following command as a root user:

```
touch <pathname>/<filename>
```

For Example, `touch /var/novell/.modVia`

Removing a File

To remove a file, use the following command as a root user:

```
rm <pathname>/<filename>
```

For example, `rm /var/novell/.modVia`

7.3 Protected Resource Issues

7.3.1 HTML Frames Are Lost

When a protected resource on an Access Gateway includes pages with multiple frames, the page displays incorrectly under the following conditions:

- The user’s session times out, the user is redirected to the login page, and the user successfully reauthenticates.
- The user logs out, and the logout page redirects the user to a page with multiple frames.

Under these conditions, only the top frame of the page is displayed. To correct this problem:

1. Create a custom login page for the protected resource.

   This can be as simple as creating a copy of the `nipd.jsp` file and renaming it. For more information on customizing the login page, see “Customizing the Identity Server Login Page” in the *Novell Access Manager 3.1 SP1 Identity Server Guide*.

2. Copy the custom login page to the JSP directory of the Identity Server.

   Linux: `/var/opt/novell/tomcat5/webapps/nidp/jsp`

   Windows: `C:\Program Files\Novell\Tomcat\webapps\nidp\jsp`

3. Modify the `top.jsp` file in the JSP directory.

   3a Locate the following lines in the top.jsp file:

   ```
   <!--
        top.location.href='<%=url%>';
   -->
   ```

   3b Replace these lines with the following:

   ```
   <!--
       location.href='<%=url%>';
   -->
   ```

4. (Conditional) If the Identity Server belongs to a cluster, copy the modified `top.jsp` file and the custom login page to each Identity Server in the cluster.

5. Add two property values to the method that creates the contract for the protected resource.

   If multiple protected resources are using the contract, you can create a custom method and contract rather than modifying the existing method. For information on this process, see “Configuring Authentication Methods” and “Configuring Authentication Contracts” in the *Novell Access Manager 3.1 SP1 Identity Server Guide*.

   5a In the Administration Console, click *Devices* > *Identity Servers* > *Edit* > *Methods*.

5b Click the name of the method that is used by the contract for the protected resource.

5c In the Properties section, click *New,* then specify the following values:

Property Name: MainJSP

Property Value: true

5d Click *OK*.

5e In the Properties section, click *New,* then specify the following values:

Property Name: JSP

Property Value: <*custom_login_page*>

Replace <*custom_login_page*> with the name of your page, without the JSP extension. (see Step 1). Property values are case sensitive.

5f Click *OK* twice.

6 Click *Devices* > *Identity Servers*, then update the Identity Server.

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

8 (Conditional) If you created a new contract for the protected resource, assign the new contract to the protected resource, then update the Access Gateway.

9 To verify that the modifications have solved the problem:

9a Access the page and log in.

9b Wait for the session to timeout.

9c Access the page again.

9d Authenticate as prompted and make sure all the frames are displayed.

7.3.2 Troubleshooting HTTP 1.1 and GZIP

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

To use GZIP, you enable your Web servers to send GZIP-compressed data. Be aware that some Web servers do not respond with compressed (GZIP) data when the Access Gateway sends the Via header to the Web server. Check you Web server documentation.

When the Web server sends compressed data and the rewriter needs to process the data, the data is decompressed, rewritten, and then recompressed. When Form Fill needs to process the data, the data is decompressed and then processed. If the Access Gateway does not need to perform any rewriting of the data or if Form Fill does not need to process the data, the compressed data is sent unchanged from the Web server to the browser. This is the default behavior.

To turn off the GZIP feature:

1 Add the following touch file

`/var/novell/.noGzipSupport`

Use the `touch` utility to create this blank file.

2 Restart the Linux Access Gateway.

In the presence of this touch file, Linux Access Gateway does not forward the ACCEPT-ENCODING header to the Web server. Without this header, the Web server does not send any data with GZIP or Deflate encoding to the Linux Access Gateway.

To allow the Linux Access Gateway to receive GZIP or Deflate encoded data, remove the touch file and restart the Linux Access Gateway.

7.3.3 Protected Resources Referencing Non-Existent Policies

If your protected resources contain references to policies that do not exist, use the following procedures to remove them.

1. Click *Auditing* > *Troubleshooting*.
2. In the *Access Gateways with Protected Resources Referencing Nonexistent Policies* section, click *Repair*.

   This removes the link between the protected resource and the policy.
3. Verify that correct policies are enabled on the protected resources. Click *Devices* > *Access Gateways* > *Edit* > *[Name of Reverse Proxy]* > *[Name of Proxy Service]* > *Protected Resources*.
4. Change to the *Policy View*.
5. (Optional) Click the *Used By* link to modify existing assignments.
6. Click *OK*, then click the *Access Gateways* link.
7. Click *Update* > *OK*.

7.3.4 Protected Resource Configuration Changes Are Not Applied

If you modify the configuration for a protected resource by modifying its *URL Path List* or its Authorization, Identity Injection, or Form Fill policies, save these changes and apply them by clicking *Update*, then return to the resource and the changes have not been applied, the protected resource has a corrupted configuration. To repair the configuration:

1. Click *Auditing* > *Troubleshooting*.
2. In the *Access Gateways with Corrupted Protected Resource Data* list, select the resource with the problem, then click *Repair*.

   This repairs the configuration for the selected protected resource.
3. Reconfigure the protected resource with the changes that weren't applied.

7.3.5 Error AM#300101010 and Missing Resources

Image display problems can arise when an unprotected page references multiple protected resources. The best practices for HTML is to avoid situations where an unprotected page contains references to multiple, automatically loaded protected resources. For example, the unprotected page `index.html` might contain references to two GIF image files. Both GIF files are protected resources. The browser automatically attempts to load the GIF files during the initial load of `index.html`. Because of multiple requests happening at the same time, one or more of the GIFs might be denied access. To avoid this, you should add the page and the `index.html` page as a protected resource. Doing this avoids the possibility of missing GIFs.

7.3.6 Unable to View Contents of Mail When Outlook Web Access is Protected by Access Gateway

If you are not able to view contents of mail when Outlook Web Access is protected by Linux Access Gateway and you see a login page instead of content, configure /exchweb/* as a public resource.

7.3.7 Redirection Issue with Some IE7 Versions

With Internet Explorer 7, if the Linux Access Gateway redirects the first request after authentication to a secure site and if the certificates are not present in the browser, the browser is not redirected to the proper site.

To workaround this problem use the `/var/novell/.useJSFor302withIE7` touch file.

When this touch file is used and Internet Explorer 7 browser is used, 200 OK response is sent back with the redirect metatag instead of the 302 redirect.

7.4 Hardware and Machine Resource Issues

7.4.1 Error: novell-vmc-chroot Failed to Start

You might see the following error message displayed:

```
novell-vmc-chroot Failed to Start.Please refer to the online guide for
troubleshooting.
```

This error usually occurs when the disk is full, which prevents `novell-vmc` from starting. To work around this problem, free some disk space before proceeding with any other configuration changes.

7.4.2 Mismatched SSL Certificates in a Cluster of Access Gateways

Sometimes a newly added server in a cluster does not receive the certificate that the rest of the cluster is using for SSL. To fix this problem:

1. Click *Devices* > *Access Gateways* > *Edit* > *[Name of Reverse Proxy]*.
2. For the Server Certificate, click the *Select Certificate* icon, then select a different certificate, such as the test-connector certificate.
3. Click *OK* to ignore the warnings that the certificate CN does not match the reverse proxy.

   This is what you want.

4 Click *OK*.

5 Click *[Name of Reverse Proxy]*.

This needs to be the same reverse proxy that you selected in Step 1.

6 For the Server Certificate, click the *Select Certificate* icon, select the certificate whose CN matches the published DNS name of the parent proxy service, then click *OK*.

7 Click *OK*.

When you click OK, the correct certificate is added to the keystore.

8 Repeat Step 1 through Step 7 for each reverse proxy that uses a unique certificate. If all the reverse proxies use the same certificate, continue with Step 9.

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

The configuration changes are pushed to the Access Gateway, and the Access Gateway loads and uses the new certificate.

7.4.3 Recovering from a Hardware Failure on an Access Gateway Machine

If an Access Gateway machine experiences a hardware failure, such as a failed hard disk, you can preserve its configuration and have it applied to the replacement machine. For information about this procedure, see "Restoring an Access Gateway" in the *Novell Access Manager 3.1 SP1 Administration Console Guide*.

7.4.4 Reinstalling a Failed Access Gateway

If the hardware of your Access Gateway fails and the Access Gateway is not a member of a cluster, you might receive the following message when you reinstall it:

```
Start unsuccessful. Reason: Unable to read keystore: /opt/novell/devman/jcc/
certs/esp/signing.keystore.
```

If you receive this message, use the following process to solve the problem:

1 Add the failed Access Gateway to a cluster.

Ignore the pending status of this command.

2 Reinstall the Access Gateway with a new IP address.

3 Add the new Access Gateway to the cluster and make it the primary cluster server.

4 Delete the failed Access Gateway from the cluster and from the Administration Console.

5 (Optional) If you want the Access Gateway to use the old IP address:

5a Reinstall the Access Gateway by using the old IP address.

5b Add it to the cluster.

5c Make it the primary cluster server.

5d Delete the Access Gateway that is using the new IP address from the cluster and from the Administration Console.

7.4.5 COS Related Issues

The following sections explain how to troubleshoot COS (cache object store) partition issues:

- “Viewing COS Partition Details” on page 189
- “Checking if the COS Partition Is Mounted” on page 189

Viewing COS Partition Details

You can view COS partition details either through YaST or through the nash prompt.

Using YaST

1. Log in as the `root` user.
2. At command prompt, enter the following command:

   ```
   fdisk -l
   ```

   The partition details are displayed. Check for COS partition details. Make sure that a partition is created with a partition ID of `68` and that the file system is created as type `unknown`.

Using nash

1. At the command prompt, enter the following command:

   ```
   nash
   ```

2. At the `nash` shell prompt, enter the following command:

   ```
   configure .current
   ```

3. Enter the following command:

   ```
   vm scan
   ```

   If the COS partition is already created, the details are displayed.

Checking if the COS Partition Is Mounted

1. Access the Linux Access Gateway main screen.

   For more information on how to access the Linux Access Gateway main screen, see Section 7.1.2, “The Linux Access Gateway Console,” on page 173.

2. Enter the *Proxy Console* option number at the *Pick a Screen* prompt.

   The Linux Access Gateway Console screen is displayed.

3. Enter the *Display Cache Statistics* option number at the *Enter option* prompt.

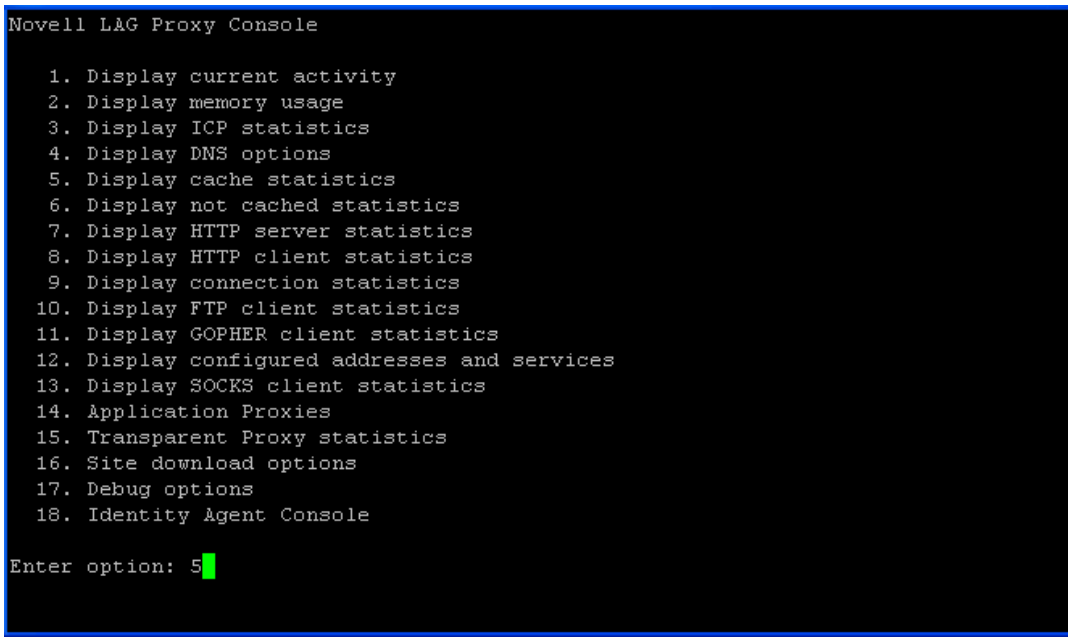

4 Enter the *Display COS Global Statistics* option number at the *Enter option* prompt.

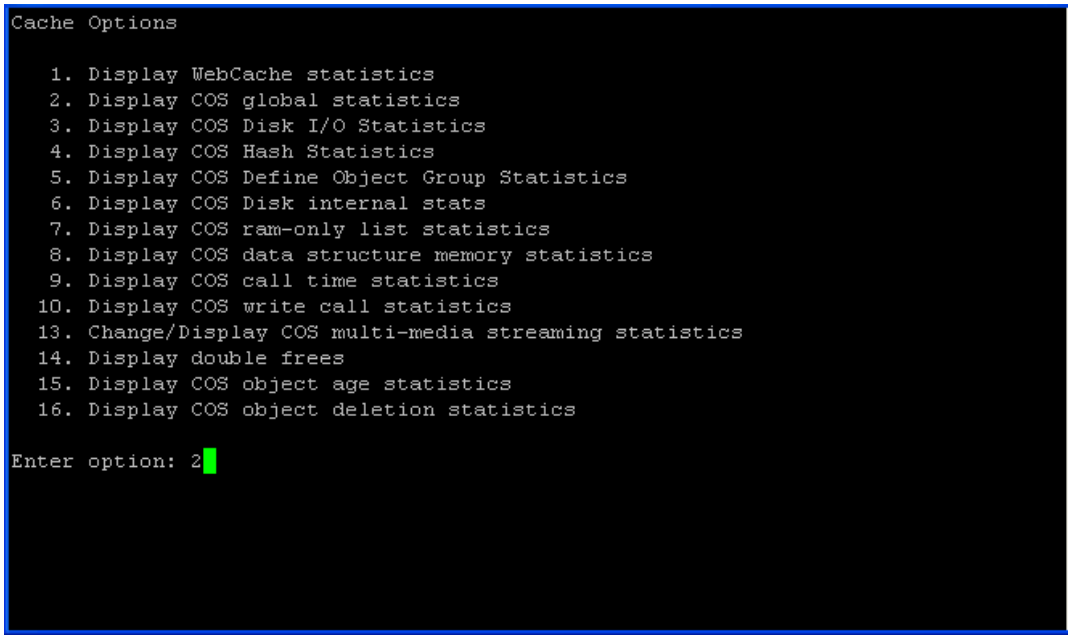

The following details are displayed if the COS partition is mounted:

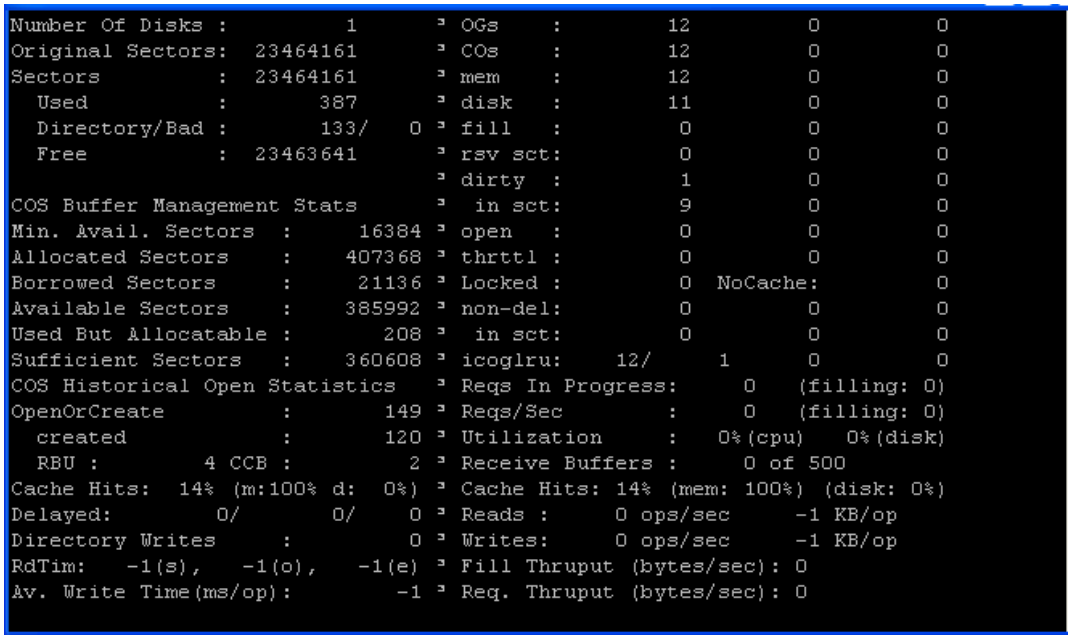

7.4.6 Memory Issues

The following sections explain how to troubleshoot memory issues:

- “Checking Memory Details and Related Information” on page 191
- “Checking Available Memory” on page 191

Checking Memory Details and Related Information

Most of the information, including the memory details, can be accessed by entering the following command at the `bash` prompt:

```
top
```

Ensure that the Linux Access Gateway does not occupy more than the percentage of the memory requirements you set.The ics_dyn process occupies approximately 20 to 25 percent of the total memory by default.

Levels	Requirement
Lower Limit	5 Percent
Requirement for Access Gateway	500 MB
Upper Limit	80 percent
Default	20 percent

Checking Available Memory

As the `root` user, enter the following command at the bash prompt:

```
cat /proc/meminfo | grep MemTotal
```

7.5 Rewriter Issues

7.5.1 Discovering the Issue

To isolate a rewriter issue:

1. Go to the Web server, access the page that is causing the rewriter problem, use view source option of the browser, then copy the source to a text file.
2. Access the page from Access Manager, view the source, and copy it to a text file.
3. Use a diff tool to compare the differences between the two files.

   This should help you identify the URLs that need to be rewritten but aren’t being rewritten.

7.5.2 Rewriting Fails on a Page with Numerous HREFs

Although the rewriting failure occurs when downloading large amounts of data from a protected Web server, it is not the size or the timeout of the page that is the issue. It is the number of links to be rewritten. The Access Gateway has a data size limit for the number of references that the rewriter can rewrite on a page.

The solution is to reduce the number of HREFs on the page that need to be rewritten. If the problem is occurring because the rewriter is rewriting HTTP to HTTPS, you can solve this problem by disabling multi-homing for the Web server and by rewriting the Web page to use relative links. This reduces the number of links that need to be rewritten.

7.5.3 Links Are Broken Because the Rewriter Sends the Request to the Wrong Proxy Service

When links on the Web server are rewritten to the wrong proxy service, the reverse proxy and Web servers might have the following configuration:

- The initial request from the browser is to a path-based multi-homing proxy service.
- The reverse proxy is configured to service one or more path-based proxy services.
- The path-based proxy services are configured to *Forward Received Host Name* and to *Remove Path on Fill*.
- The Web servers protected by these path-based proxy services have links to each other.

With this configuration, the rewriter cannot determine whether the link is to the current proxy service, one of the other path-based proxy services, or the parent proxy service. With the path removed, all the path-based proxy services have the same name. For example if one proxy service has the published name of mycompany.provo.novell.com/sales and a second path-based proxy service has a name of mycompany.provo.novell.com/app, the names are the same as the parent proxy service when the path is removed. The HTTP header does not help, because the proxy services are forwarding the same host name: mycompany.provo.novell.com.

There are a number of ways to solve this problem. One of the easiest ways is to set up DNS names for the Web servers, then configure the proxy services so that the *Host Header* option is set to *Web Server Host Name* and the DNS name of the Web server is specified in the *Web Server Host Name* field. This places the DNS name of the Web Server name in the HTTP Host header, allowing the rewriter to distinguish it from the other Web servers protected by the reverse proxy.

7.5.4 Reading Configuration Files

If the rewriter is successful in reading the configuration files, and you have enabled the log level to LOG_INFO, the following message is displayed in the `/var/log/ics_dyn.log` file:

```
Reading Config File

Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:Configuration information read
successfully
```

For more information on configuring log levels, see "Configuring Log Levels" on page 107.

If the rewriter fails to read the configuration files, the following message is displayed:

```
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:Reading configuration failed for
ssTypeName=www.mynovell.com
```

If this happens, re-create the corresponding proxy service and restart the Linux Access Gateway service.

7.5.5 Rewriter Does Not Rewrite Content in Files with a Non-Default Extension

If the Web server sends data, whose file extensions do not match with any of the default rewriter profiles, then rewriter does not rewrite the content. The following content-type extensions that are rewritten by default are html, htm, shtml, jhtml, asp, jsp, js, php, and css. In order to work around this problem, do the following:

1. In the Administration Console, click *Devices* > *Access Gateways* > *Edit* > *[Name of Reverse Proxy]* > *[Name of Proxy Service]* > *HTML Rewriting*.
2. Do one of the following:
   - If the Web server sends a different content type for a non-default file extension, then configure the new content type in the *Content-Type Header*.
   - If the Web Server does not send any content type for a non-default extension, then configure `extension/<file_extension>` as the *Content-Type Header*. For example, if the data sent is http://www.myproxy.com/test.mytxt, then you must configure the *Content-Type Header* as `extension/mytxt`.

7.5.6 Additional DNS Name Without a Scheme Is Not Rewritten

Rewriter rewrites URLS based on the port configured for *Connect Port*, when domain name without scheme is added to the additional URL list. For example, if the Connect port is configured as 80, Web server rewrites only HTTP URLs and not HTTPs URLs. To work around this problem,

1. In the Administration Console, click *Devices* > *Access Gateways* > *Edit* > *[Name of Reverse Proxy]* > *[Name of Proxy Service]* > *HTML Rewriting*
2. Add the additional DNS name with the scheme in the *Additional DNS Names List* in the following format:

   ```
   scheme://DNS_name
   ```

   For example, https://example.com

7.5.7 Rewriting a URL

Set the log level to LOG_DEBUG to view rewriter log messages in the `/var/log/ics_dyn.log` file. (See "Configuring Log Levels" on page 107.)

For example, if the Rewriter successfully rewrites the URL, the following messages are displayed:

```
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:URL:'http://
www.mynovell.com:9090/common/inc/nav/main.js' Content type match, Will Rewrite
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:URL:'http://
www.mynovell.com:9090/common/inc/nav/main.js' Unknown Content-Type - automatic
match - Will Rewrite
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0::'http://www.mynovell.com:9090/
common/inc/nav/main.js' NULL Content-Type - automatic match - Will Rewrite
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:In
RewriterOption::shouldRewriteUrl, returning TRUE.
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:URL:'http://
www.mynovell.com:9090/common/inc/nav/main.js' Unknown extension - automatic
match - Will Rewrite
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:URL:'http://
www.mynovell.com:9090/common/inc/nav/main.js' NULL extension - automatic match
- Will Rewrite
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:URL:'http://
www.mynovell.com:9090/common/inc/nav/main.js' Extension type match - Will
Rewrite
```

If the conditions for rewriting a URL fail, the following messages are displayed:

```
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:URL:'http://
www.mynovell.com:9090/favicon.ico' - Did not match INCLUDE list, Content-Type
and Extension type
Aug 16 04:16:51 proxy140 LINUX_AG:REWRITER:0:In
RewriterOption::shouldRewriteUrl, returning FALSE.
```

Check the rewriter configuration. Ensure that your content type, extension type, and include URL list are valid.

7.6 Troubleshooting Crashes and Hangs

- Section 7.6.1, "The Access Gateway Hangs When the Audit Server Comes Back Online," on page 195

7.6.1 The Access Gateway Hangs When the Audit Server Comes Back Online

When the Platform Agent loses its connection to the audit server, it enters caching mode. The default size of the audit cache file is unlimited. This means that if the connection is broken for a long time and traffic is high, the cache file can become quite large. When the connection to the audit server is re-established, the Platform Agent becomes very busy while it tries to upload the cached events to the audit server and still process new events. When coming out of caching mode, the Platform Agent appears unresponsive because it is so busy and because it holds application threads that are logging new events for a long period of time. If it holds too many threads, the system can appear to hang. You can minimize the effects of this scenario by configuring the following two parameters in the `logevent` file.

Table 7-5 *Parameters for the logevent File*

Parameter	Description
LogMaxCacheSize	Sets a limit to the amount of cache the Platform Agent can consume to log events when the audit server is unreachable. The default is unlimited.
LogCacheLimitAction	Specifies what the Platform Agent should do with incoming events when the maximum cache size limit is reached. You can select one of the following actions: ◆ Delete the current cache file and start logging events in a new cache file. ◆ Stop logging which preserves all entries in cache and stop collecting new events.

When you set a finite cache file size, it limits the number of events that must be uploaded to the audit server when caching mode is terminated and keeps the Platform Agent responsive to new audit events that are registered.

For more information about the `logevent` file and these parameters, see Logevent (http://www.novell.com/documentation/nsureaudit/nsureaudit/data/al36zjk.html#alibmyw).

7.6.2 Access Gateway Crashes When the Log Files Are Removed

If you have enabled the debug level of logging for the `laghttpheaders` and the `lagsoapmessages` log files and these files grow to be over 200 MB, manual deletion of these files can cause the Access Gateway to crash.

To solve the problem, restart the Access Gateway after manually deleting the files.

7.6.3 Troubleshooting a Failed Linux Access Gateway Configuration

If the IP address and other network configurations are not reflected in the installed Linux Access Gateway, log in as a `root` user and run the following commands:

```
rm /opt/novell/legacy/etc/proxy/.novell_lag_lock
/etc/init.d/novell-vmc stop
/etc/init.d/novell-vmc start
```

7.6.4 Troubleshooting a Linux Access Gateway Crash

The Linux Access Gateway might have crashed because of the following reasons:

- SIGSEGV
- ASSERT (for a debug build only)

The following sections explain how to gather the files that need to be sent to Novell for a resolution of the problem.

- "Linux Access Gateway Logs" on page 196
- "Event Log" on page 196
- "Core Dump" on page 198
- "Proxy Hang Core" on page 199
- "Packet Capture" on page 199

Linux Access Gateway Logs

1. Enter the following command from the bash shell to collect the debug log files that are generated:

   ```
   /chroot/lag/opt/novell/bin/getlaglogs.sh
   ```

2. The `laglogs.tgz` tar file is located in the `/var/log` directory.
3. Send this tar file to Novell® Support.

Event Log

By default the event log size is 15 MB. The size of event log can be controlled by configuring the required event log size in the `eventlogsize.cfg` file, located at the `/chroot/lag/etc/opt/novell` directory. For example, if you specify 350 in the file, you can configure an event log of size 350 MB. This file should contain only the file size information. This file should not contain any other characters or new lines.

The procedure for obtaining the event log depends upon the build type:

- "Event Log for a Production Build" on page 197
- "Event Log for a Debug Build" on page 197

Event Log for a Production Build

To get the event log for the production build:

1. Log in as the `root` user.
2. To disconnect all instances of Linux Access Gateway, enter the following command:

   ```
   /etc/init.d/novell-vmc stop
   ```
3. Enter the following command to change the root environment:

   ```
   chroot /chroot/lag
   ```
4. To start the process, enter the following command:

   ```
   gdb /opt/novell/bin/ics_dyn 2>/var/log/ics_dyn.log
   ```
5. At the GDB prompt, run the following command:

   ```
   run -m <memory>
   ```

   Where <*memory*> is the percentage of total memory to be used for ics_dyn process. It is recommended to set this value in the range of 20-30 percent.
6. Repeat the scenarios to reproduce the issue.

   6a If you are trying to reproduce the proxy crash, you see the GDB prompt as soon as the crash is reproduced.

   6b If you are trying to reproduce a functionality issue, press Crtl+C to enter the GDB prompt as soon as the issue is reproduced.

   For a list of commands that can be entered in the debugger, see "Useful Debugger Commands" on page 198.
7. To save event logs to a file, enter the following command:

   ```
   d ,save 1
   ```

   This stores all the events in the `/chroot/lag/opt/novell/debug/<pid>all_events.0.txt` file.
8. Tar or Zip this file and send it to Novell Support.

Event Log for a Debug Build

To get the event log:

1. Log in as the `root` user.
2. To stop all instances of Linux Access Gateway, enter the following command:

   ```
   /etc/init.d/novell-vmc stop
   ```
3. To start the Novell Linux Access Gateway in debugging mode, enter the following command:

   ```
   /etc/init.d/novell-vmc gdb
   ```
4. To run the Linux Access Gateway process, enter the following command at the GDB prompt:

   ```
   run -m <memory> 2>/var/log/ics_dyn.log
   ```

   Where <*memory*> is the percentage of total memory to be used for ics_dyn process. You should set this value with a range of 20-30 per cent.
5. Repeat the scenarios to reproduce the issue.

   5a If you are trying to reproduce the proxy crash, you will enter the GDB prompt as soon as the crash is reproduced.

5b If you are trying to reproduce a functionality issue, enter the following command to enter the GDB prompt as soon as the issue is reproduced:

```
Crtl+C
```

NOTE: For a list of commands that can be entered in the debugger, see "Useful Debugger Commands" on page 198.

6 To save all event logs to a file, enter the following command:

```
d ,save 1
```

This stores all the events in the `/chroot/lag-debug/opt/novell/debug/<pid>all_events.0.txt` file.

7 Tar or zip this file and send it to Novell Support.

Useful Debugger Commands

Table 7-6 *GDB Commands*

Command	Function
`gcore`	Generate core file
`k`	Kill process
`q`	Quit GDB prompt
bt	Print the back trace

Core Dump

Before you begin, make sure there is free space in `root` to hold the core file and that the space is at least equal to the RAM size

To collect a core dump:

1 Log in as the `root` user.

2 To disconnect all instances of the Linux Access Gateway, enter the following command:

```
/etc/init.d/novell-vmc stop
```

3 At the bash prompt, specify the following command:

```
touch /tmp/.dumpcore
```

4 Enter the following command to start the Linux Access Gateway:

```
/etc/init.d/novell-vmc start
```

5 Repeat the scenarios to reproduce the issue.

The core is dumped to the `/chroot/lag core.<pid>` file.

`<pid>` is the process ID of ics_dyn process.

After the core is dumped, the Linux Access Gateway restarts.

6 Tar or zip the core dump send it to Novell Support.

Proxy Hang Core

To analyze the proxy hang and create a core file:

1. Enter the following command to change the root environment:

   ```
   chroot /chroot/lag
   ```

2. Enter the following command to attach the ics_dyn process to the debugger:

   ```
   gdb /opt/novell/bin/ics_dyn <pid>
   ```

   Where <*pid*> refers to the Process ID of the ics_dyn process.

3. At the GDB prompt, enter the following command:

   ```
   set logging on <filename>
   ```

   Where <*filename*> specifies the name of the file that will store the output of the executed debugger commands.

4. Enter the following command to collect a stack trace of all threads:

   ```
   thread apply all bt
   ```

5. Enter the following command to turn off logging:

   ```
   set logging off
   ```

6. Enter the following command to save the core dump in the `/chroot/lag` directory.

   ```
   gcore
   ```

   The core dump is saved as `core.<pid>`.

7. Tar or zip this file and send it to Novell Support.

Packet Capture

The `tcpdump` utility allows you to capture network trace packets.

1. Log in as the `root` user.
2. Enter the following command:

   ```
   tcpdump -s0 -n -t -p -i 'any' -w filename.cap
   ```

3. Tar or zip this file and send it to Novell Support.

7.6.5 Linux Access Gateway Not Responding

If the Linux Access Gateway is not responding, do the following:

1. Enter the following command to change the root environment:

   ```
   chroot /chroot/lag
   ```

2. Enter the following command to attach the ics_dyn process to the debugger:

   ```
   gdb /opt/novell/bin/ics_dyn <pid>
   ```

   Where <*pid*> refers to the process ID of the ics_dyn process. You can get the process ID by entering the following command:

   ```
   pgrep ics_dyn
   ```

3. At the GDB prompt, enter the following command:

   ```
   set logging file <filename>
   ```

Where <*filename*> specifies the name of the file that will store the output of the executed debugger commands.

4 Enter the following command to start logging:

```
set logging on
```

5 Enter the following command to collect a stack trace of all threads:

```
thread apply all bt full
```

6 Enter the following command to turn off logging:

```
set logging off
```

7 Enter the following command to save the core dump in the `/chroot/lag` directory.

```
gcore
```

The core dump is saved as `core.<pid>`.

8 Tar or zip this file and send it to Novell Support.

7.7 Connection and Authentication Issues

This section provides various troubleshooting scenarios and frequently asked questions that you might encounter while using the Linux Access Gateway, and suggests appropriate actions.

7.7.1 Connection Details

To obtain connection information:

1 Log in as the `root` user.

2 At the bash prompt, enter one of the following `netstat` commands:

Command	Details
`netstat -anp`	Provides the connection information
`netstat -s -t`	Provides the connection statistics

7.7.2 Network Socket Issues

This section lists various issues related to network sockets and provides information on how to verify bind and connection issues:

Socket Listener Bind

To verify whether the socket listener is bound to the required port:

1. Log in as the `root` user.
2. At the bash prompt, enter the following command:

   ```
   netstat -anp | grep LISTEN
   ```

   All ports are displayed.
3. Search for the desired port.

   If the required port is not visible in the list, a bind failure has occurred.

Issues with Outgoing Connections

To verify that the Access Gateway is able to make outbound connections:

1. Log in as the `root` user.
2. At the bash prompt, view the following log file:

   ```
   /var/log/ics_dyn.log
   ```
3. Search for a connection message. If the service is unavailable, the file contains messages similar to the following:

   ```
   ERROR Connection FAILED with peer
   ```

7.7.3 Authentication Issues

This section provides information related to authentication:

- “User Details” on page 201
- “Error Codes” on page 203

User Details

To check the details about the users logged in to the Linux Access Gateway:

1. To access the console, enter the following command:

   ```
   netcat localhost 2300
   ```
2. Press Enter at the `Please enter terminal type` prompt.

   This displays the Linux Access Gateway console screens.

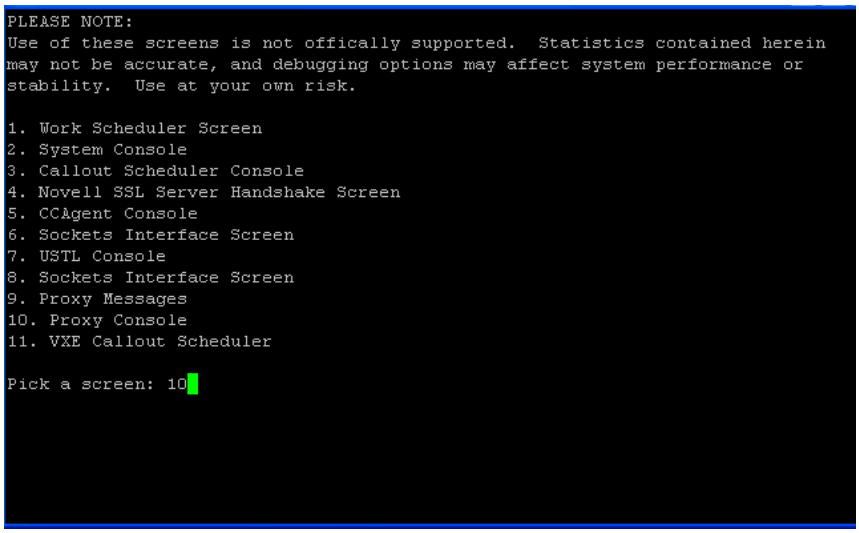

3 Enter the *Proxy Console* option number at the *Pick a Screen* prompt.

The Linux Access Gateway Console screen is displayed.

4 To select the *Identity Agent Console* option, enter the option number at *Enter Option*.

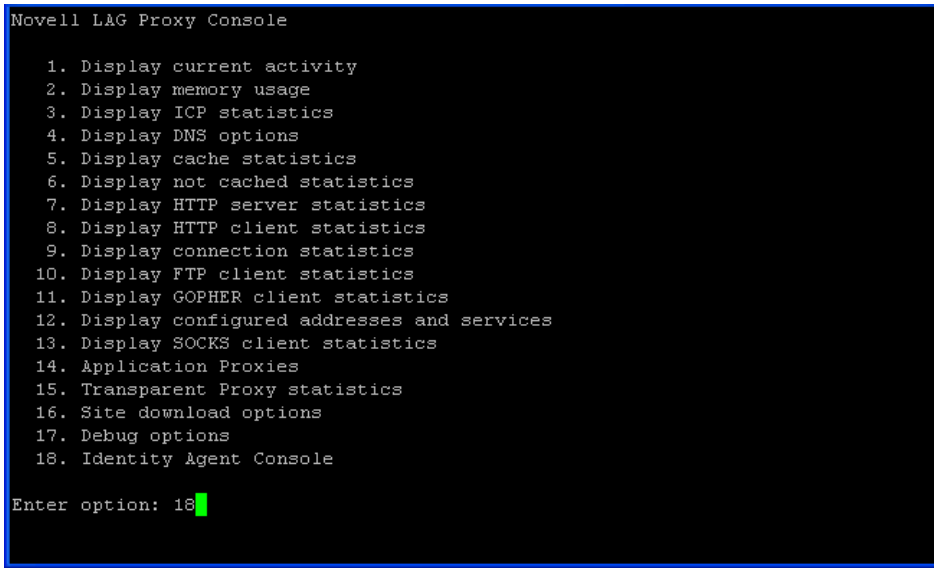

The Identity Agent Console screen is displayed.

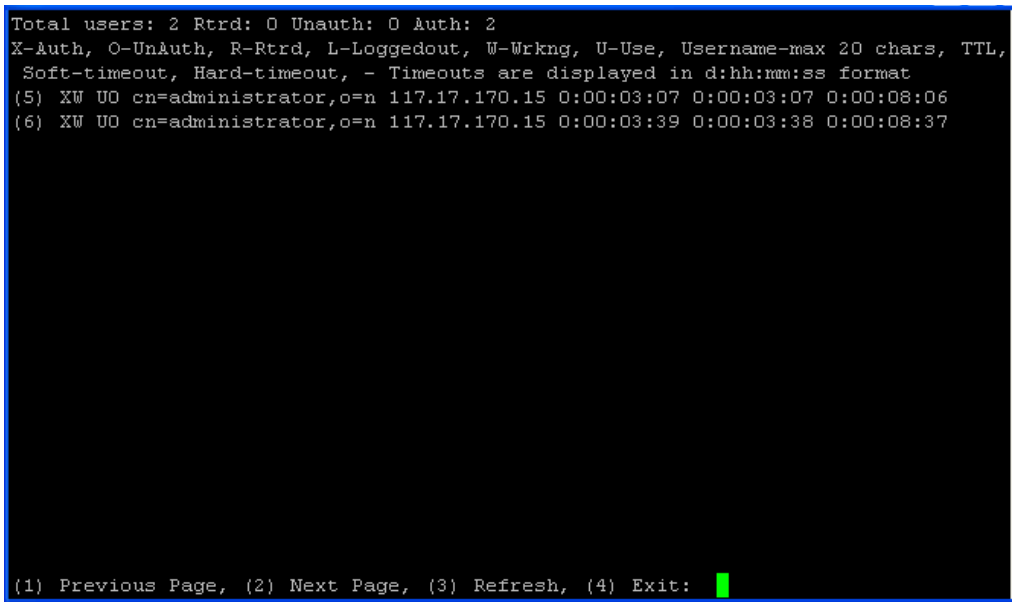

The user information contains the following items:

- **X:** An authenticated user.
- **O:** An unauthenticated user.
- **R:** A retired user; the user session has timed out. The default time-out is 3 minutes. In this state, the user session is deleted. If the user makes another request from the browser session, the Linux Access Gateway requires the user to authenticate.
- **L:** The user has logged out of the session.
- **W:** The user session is functional.
- **U:** The use count is more than zero.
- **Username:** The full distinguished name of the user. The username can contain a maximum of 20 characters.
- **TTL:** The time remaining before the user session goes to the retired state if the user session remains idle.
- **Timeout:** The session timeout is displayed in d:hh:mm:ss format.

The screen displays 20 users at a time. The screen also displays the browser IP address. The following options are available at the bottom of the screen:

- **Previous Page:** Lets you go to the previous page.
- **Next Page:** Lets you go to the next page (to view the next set of users).
- **Refresh:** Refreshes the page to reflect the latest user status.
- **Exit:** Exits the console.

Error Codes

The following error codes indicate authentication problems:

- "500 Internal Server Error" on page 204
- "504 Gateway Timed Out" on page 204

500 Internal Server Error

Possible Cause: Authentication failed because of a system error.

Action: Contact Novell Support.

504 Gateway Timed Out

Possible Cause: The authentication back-end channel is not working.

Action: Check to see if the Embedded Service Provider is listening on the loopback address 127.0.0.1 at port 8080: Use the following command:

```
netstat -na | grep 8080
```

If the Embedded Service Provider is down, restart the service provider from the Administration Console.

If the issue persists, contact Novell Support.

7.8 Form Fill Issues

Form Fill error messages are logged only if you set the log level to LOG_DEBUG. The entries are logged in the `ics_dyn.log` file. Search for entries with a correlation tag of `AM#504507`. For more information, see "Form Fill Traces" in the *Novell Access Manager 3.1 SP1 Policy Management Guide*.

This section contains the following information about form fill issues:

- Section 7.8.1, "Form Fill Does Not Process Forms with Complicated JavaScript Functions when Data is Auto-Submitted," on page 204
- Section 7.8.2, "Form Fill Error Messages," on page 205
- Section 7.8.3, "Alert: SSO (Form Fill) Failed Due to Malformed HTML," on page 205
- Section 7.8.4, "Form Fill Failure Because of Incorrect Policy Configuration," on page 205
- Section 7.8.5, "Browser Spinning Issues," on page 205

7.8.1 Form Fill Does Not Process Forms with Complicated JavaScript Functions when Data is Auto-Submitted

If Form Fill fails to process forms with complicated JavaScript or VBScript functions when data is auto-submitted even if the *Statements to Execute on Submit* option is selected, then do the following:

1. Log in as `root`.
2. Specify the following command to create the `.enableInPlaceSilentFill` file:

   ```
   touch /var/novell/.enableInPlaceSilentFill
   ```
3. Specify the following command to create the `enableInPlaceSilentFillNew` file:

   ```
   touch /var/novell/.enableInPlaceSilentFillNew
   ```
4. Specify the following command to restart Linux Access Gateway:

   ```
   /etc/init.d/novell-vmc stop
   ```

   ```
   /etc/init.d/novell-vmc start
   ```

For more information on the touch files, see Section 7.2.2, “Using Touch Files,” on page 177

7.8.2 Form Fill Error Messages

You might get the following errors when sending a browser request:

- `DataStore Error`
- `The service provider is not running at the moment. Please retry after a few seconds.`

These errors indicate that the Access Gateway cannot retrieve the information that is essential to process the browser request, or is unable to save the information provided by the user because the Embedded Service Provider is down. Retry the action after a few seconds. If the error persists, restart the Embedded Service Provider from the Administration Console.

7.8.3 Alert: SSO (Form Fill) Failed Due to Malformed HTML

Sometimes you might get the following error message:

```
Alert: SSO (Form Fill) Failed Due to Malformed HTML
```

The cause and action for that error could be the following:

Possible Cause: If this message appears on the login page which was to be filled by the Linux Access Gateway Form Fill, then the HTML page is malformed.

Action: You have to manually fill the form.

Possible Cause: If this message is displayed in any page other than the login page that was to be filled by the Linux Access Gateway, then this implies that the CGI or the page matching criteria configured for the Linux Access Gateway Form Fill policy matched the other pages and that there was a failed attempt to fill those pages.

Action: Check and modify the CGI and the Page Matching Criteria in the policy in such a way that the policy is applied only to the login page that you want the Linux Access Gateway to fill.

7.8.4 Form Fill Failure Because of Incorrect Policy Configuration

Form fill fails if the policy is not configured correctly. For configuration information, see “Creating Form Fill Policies” in the *Novell Access Manager 3.1 SP1 Policy Management Guide*.

7.8.5 Browser Spinning Issues

Browser spinning can occur if inappropriate data is filled in the form because of one of the following reasons:

- Shared secrets are configured, the user provided incorrect data to the Linux Access Gateway, and there are no appropriate actions configured to handle login failure.
- A Credential Profile with LDAP attributes has been configured, and there is a mismatch between the username used to authenticate to the Linux Access Gateway and the username used to authenticate to the accelerated Web server.

When a Form Fill policy succeeds and the authentication to the Web server fails, the Web server redirects the browser to its authentication page again and again, if auto-submit is enabled. In such a situation, if there is no appropriate login-failure action configured in the policy, the browser "spins" endlessly.

If this happens, do the following:

- Kill the browser session. If you are unable to do this, run the following commands to restart the Linux Access Gateway:

  ```
  /etc/init.d/novell-vmc stop
  /etc/init.d/novell-vmc start
  ```

- If the issue is with a Credential Profile with LDAP attributes, verify which LDAP attributes are required by the Web server, and create the appropriate entries in the Form Fill policy.
- If the issue is with shared secrets, delete the corresponding values from the Secret Store. If it is not possible to delete the value, modify the corresponding policy to use a different or a new custom attribute or shared secret attribute. For more information on modifying the policy, see "Implementing Form Fill Policies" in the *Novell Access Manager 3.1 SP1 Policy Management Guide*.

7.9 Authorization and Identity Injection Issues

7.9.1 Authorization and Identity Injection Error Messages

If you have already configured the Identity Injection policies, you might receive the following errors while trying to send a browser request:

- `Service provider is in halted state. Please contact your administrator to restart Service Provider from Administrator Console.`
- `Policy engine is sending invalid response. Please contact your administrator to restart Service Provider from Administrator Console.`
- `Unable to process your request.`
- `Unable to process your request due to parseXML failure.`

These errors indicate that the Embedded Service Provider is down. Every Identity Injection policy has a policy ID, which is sent to the Access Gateway by the Embedded Service Provider. If the Embedded Service Provider is down, the Access Gateway does not get the policy ID, and an error is thrown. Restart the Embedded Service Provider from the Administration Console as follows:

1. In the Administration Console, click *Devices* > *Access Gateways*.
2. Select the server, then click *Actions*.
3. Click *Service Provider* > *Restart Service Provider*.
4. Click *OK*.

7.9.2 Identity Injection Failures

Identity injection might fail while trying to inject authentication headers because of improper policy configuration or because the Identity Server is not sending values to the Access Gateway.

Check the `/var/log/ics_dyn.log` file for the following error messages:

- `Customer Header Injection Failed.`
- `Query String Injection Failed.`
- `Authentication Header Injection Failed.`

To receive help resolving identity injection failures, send the following information to Novell Support:

- Linux Access Gateway logs. For more information on how to get Linux Access Gateway log files, see "Gateway Appliance Logs" on page 107.
- Packet Capture. For more information on how to get packet captures, see "Packet Capture" on page 199.

7.9.3 Identity Injection Problems When Using a Password Management Service

If you have configured the Identity Server to use a password management service and you have also configured resources to use Identity Injection policies that inject the user's password, you need to enable the following touch file:

`/var/novell/.PasswordMgmt`

This file causes the Access Gateway to refresh the user's credentials so that they match password changes. Without it enabled, if users authenticate, then change their passwords, the Access Gateway uses the old password in Identity Injection policies.

7.10 YaST Goes into a Non-Responsive Mode When a Partition Is Deleted or Created

YaST goes into a non-responsive mode if you click *Finish* after adding, deleting, or modifying a partition. To work around this problem, click *Apply*, then click *Quit* instead of clicking *Finish*.

7.11 Upgrading the Linux Access Gateway Randomly Halts the Embedded Service Provider

After upgrading, the embedded service provider sometimes halts at the end of the upgrade process. When this happens, restart the Linux Access Gateway. In the Administration Console, click *Access Manager* > *Access Gateways*, select the Access Gateway, then click *Reboot*.

7.12 Using Curl to Download Large Files

If you use the `curl` utility to download large files, sometimes, the files might get corrupted. If this happens, download the file by using the `wget` utility.