

Novell iChain 2.1 Administration Guide

www.novell.com



Novell®

Legal Notices

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

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

You may not export or re-export this product in violation of any applicable laws or regulations including, without limitation, U.S. export regulations or the laws of the country in which you reside. This product may require export authorization from the U.S. Department of Commerce prior to exporting from the U.S. or Canada.

Copyright © 1993-2002 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

U.S. Patent No. 5,349,642; 5,608,903; 5,671,414; 5,677,851; 5,758,344; 5,784,560; 5,818,936; 5,828,882; 5,832,275; 5,832,483; 5,832,487; 5,870,561; 5,870,739; 5,873,079; 5,878,415; 5,884,304; 5,913,025; 5,933,503; 5,933,826; 5,946,467; 5,956,718; 6,047,289; 6,065,017; 6,081,900; 6,105,132; 6,167,393. Patents Pending.

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

www.novell.com

Novell iChain 2.1 Administration Guide

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

Novell Trademarks

Novell iChain is a registered trademark of Novell, Inc.

Novell BorderManager is a registered trademark of Novell, Inc.

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

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

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

ConsoleOne is a registered trademark of Novell, Inc.

DirXML is a registered trademark of Novell, Inc.

eDirectory is a trademark of Novell, Inc.

NLM is a trademark of Novell, Inc.

Novell Certificate Server is a trademark of Novell, Inc.

Novell Client is a trademark of Novell, Inc.

Novell SecretStore is a trademark of Novell, Inc.

Third-Party Trademarks

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

About This Guide

Introduction

The purpose of this documentation is to help you install, configure, and administer the Novell® iChain® infrastructure.

The audience for this documentation is network administrators.

This guide is divided into the following sections.

- ◆ **Chapter 1, “Overview,” on page 7** — An explanation of the benefits of iChain and an overview of the components that make up iChain.
- ◆ **Chapter 2, “Installing iChain Services,” on page 13** — Instructions for how to install iChain, including an installation scenario, system requirements, software installation instructions and instructions for how to install iChain ConsoleOne® Snap-ins.
- ◆ **Chapter 3, “Setting Up a Basic Configuration,” on page 27** — An explanation of the tasks you need for setting up a basic iChain infrastructure.
- ◆ **Chapter 4, “Setting Up Authentication Services,” on page 67** — An explanation of how to set up Authentication Services.
- ◆ **Chapter 5, “Setting Up Web Single Sign-on Services,” on page 95** — An explanation of how to set up Web Single Sign-on Services.
- ◆ **Chapter 6, “Setting Up an Advanced Configuration,” on page 111** — An explanation of how to augment or alter the basic iChain implementation and how to use more advanced iChain features.
- ◆ **Chapter 7, “Using and Tuning iChain Features,” on page 123** — An explanation of how to fine tune your appliance.

- ◆ **Appendix A, “iChain Management,” on page 217** — An explanation of how to use the iChain 2.1 Setup Wizard.
- ◆ **Appendix B, “iChain Proxy Server Management,” on page 275** — An explanation of how to configure and manage the iChain Proxy Server.
- ◆ **Appendix C, “Rewriter Support,” on page 369** — An explanation of how to configure and use iChain’s rewrite filter.
- ◆ **Appendix D, “Upgrading Your iChain System,” on page 379** — An explanation of how to upgrade your current iChain installation to iChain 2.1 Proxy Server and Authorization Server software

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.

Documentation Updates

For the latest iChain documentation, including updates to this administration guide, see the online documentation at the [Novell documentation Web site \(http://www.novell.com/documentation\)](http://www.novell.com/documentation).

1 Overview

Novell® iChain® is an integrated security and access-management infrastructure that protects your network and safeguards sensitive eBusiness and identity data. iChain facilitates your eBusiness and remote-access initiatives by providing secure authentication and access to portals, Web-based content, and Web applications. As an unparalleled foundation for your commercial success, iChain incorporates Novell eDirectory™, the world's most scalable and widely used directory. iChain also offers personalization, simple installation, Web Single Sign-on, and the ability to secure access to data and applications across the Internet. With iChain, you can simplify, secure, and accelerate your eBusiness transformation.

This section examines the need for iChain, lists and describes the components that comprise iChain, and lists and describes the features and services provided by iChain.

Understanding the Value of iChain

Novell iChain is an integrated security solution that offers identity-management and access-management services within a powerful eBusiness infrastructure. iChain provides secure authentication and access to portals, Web-based content and Web applications. This means that all types of eBusiness and remote-access initiatives are more securely available than ever before. To help you successfully meet your eBusiness goals, iChain encompasses the following core components:

- ◆ iChain Proxy Server
- ◆ iChain Authorization Server
- ◆ Novell eDirectory

Through these services, iChain provides Web Single Sign-on, multi-factor authentication, fine-grained access control, and confidential data delivery across the Internet. With these features, iChain enables you to bring all aspects of your business together and move them online. In fact, iChain is the most secure foundation for identity and access management for your eBusiness.

iChain Benefits

With iChain, you can accomplish the following:

- ◆ **Easily Create Public, Restricted, and Secure iChain Resources**
- ◆ **Define Identity-based Security**
- ◆ **Multiple Factor Authentication**
- ◆ **Give Your Users the Convenience of Web Single Sign-on**
- ◆ **Dynamically Encrypt Data**
- ◆ **Simplify Your Management and Administrative Duties**
- ◆ **Provide Enhanced Installation and Configuration Options**
- ◆ **Securely Integrate Various eBusiness Applications**

Easily Create Public, Restricted, and Secure iChain Resources

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

- ◆ **Public** — access is granted with no security checks
- ◆ **Restricted** — requires user authentication only
- ◆ **Secure** — requires authentication and association to an access control object (see **“Define Identity-based Security” on page 9** for more information)

These protected resources are essentially a listing of URLs. To add flexibility when defining these protected URLs, iChain offers wildcarding (*) and entire folder (?) options.

Define Identity-based Security

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

In addition to a user being granted access based on his or her association to an iChain Access Control object, iChain 2.1 also introduces a new dynamic access control process which looks at specified details of a user's identity (for example, jobTitle=manager) and grants access based on that information.

Multiple Factor Authentication

Although iChain seeks to connect your business with the rest of the world, it would hardly be worth the effort if this connection caused your corporate security to be compromised. Identity-based access control is a very secure method of protecting your data, but it must rely on some form of user recognition to ensure that the person attempting to access the protected resources is who he or she claims to be. To guard against unauthorized users, iChain supports a number of authentication methods, including user identifiers (name, e-mail address, and other LDAP attributes), passwords, token-based authentication, and X.509 digital certificates.

iChain also supports multi-factor authentication, which combines several authentication methods to produce an even higher level of security. For instance, a company can require that a user present a valid userID and password as well as an X.509 certificate before granting access to a user. Different levels of access can be required for different Web servers.

To accommodate secure, token-based authentication, iChain uses the Remote Authentication Dial-in User Service (RADIUS) protocol. RADIUS enables communication between remote access servers and a central server. Secure token authentication through RADIUS is possible because iChain includes Novell Modular Authentication Service RADIUS software that you can run on your existing NetWare® server.

In addition, iChain's Proxy Server offloads Secure Sockets Layer (SSL) encryption tasks from the Web server. As well as freeing the server from the processor-intensive task of encoding and decoding SSL protocols, iChain adds another layer of security between the Internet and your network. Before users

authenticate to eDirectory, iChain's Proxy Server establishes an SSL session with the users' browsers. This prevents unauthorized users from intercepting passwords and other authentication credentials. iChain also leverages Novell Certificate Server™, a security product that ships with NDS® eDirectory 8.5. Novell Certificate Server enables you to create and manage digital certificates of your own or import them from third-party vendors.

Give Your Users the Convenience of Web Single Sign-on

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

Single Sign-on is possible because iChain authenticates the user from a centralized eDirectory profile. When the user requests access to a specific server, iChain retrieves the appropriate user credentials and transparently submits them to the Web server, usually in the form of username and password. The user sees no login request; but sees only the end result as access is either granted or denied.

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

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

Increase the Overall Security of a Web Server

Having the iChain Proxy Server as the gatekeeper (the only point of access) to Web applications increases the overall security of the Web server and identity information. Users never get direct access to either Web server or directory

information. When platforms are being used to host Web applications that potentially have a higher risk of being subject to hacking attempts, iChain will ensure that only HTTP requests are serviced, and that those requests are specific to DNS names rather than to an IP address. The iChain proxy immediately stops any other request.

Reduce the Required Number of Public IP Addresses

Generally each Web server available on the Internet requires its own IP address, which can increase the cost of a solution and can require more firewall configurations. iChain offers a multi-homing facility whereby a single IP address can be used to access multiple backend Web servers. As long as these services are represented as a single domain (for example, xxx.novell.com), all services can operate over the standard port 80 (HTTP) or 443 (HTTPS - encrypted).

Dynamically Encrypt Data

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

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

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

Simplify Your Management and Administrative Duties

Today, many companies manage user access to internal Web-based material on a server-by-server basis. These servers often run on different platforms, especially in large enterprises that have many divisions spread across a wide geographic area. A good example is a government agency with many separate departments. Each department employs its own set of standalone servers and Web applications. Something as common as modifying a user's access rights

would require the IT staff to manually change all the involved systems, a time-consuming process that could necessitate a physical visit to each network server. If those servers are scattered across the entire country, the situation becomes expensive and impractical — either a single IT staff member is constantly traveling, or it becomes necessary to maintain a separate IT staff for each part of the network.

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

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

Provide Enhanced Installation and Configuration Options

Novell iChain includes an installation wizard that is a time-saving, cost-effective configuration solution. The wizard enables you to customize how iChain's features will complement your network structure and eliminates several steps required by traditional configuration and installation processes. By presenting you with several questions about your configuration preferences, the wizard helps you create a configuration file that has all the necessary parameters to configure iChain.

Securely Integrate Various eBusiness Applications

Trying to run eBusiness applications from different vendors is usually an exercise in frustration. Because applications cannot automatically share data across the enterprise, your business infrastructure becomes fragmented. With iChain you can provide new avenues of data protection while securely consolidating all the elements of your computing environment into one Net.

To facilitate the transformation from traditional business to eBusiness, iChain integrates with Novell DirXML®. DirXML is an Extensible Markup Language (XML) solution that reliably synchronizes databases and directories from various applications and vendors.

2

Installing iChain Services

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

- ◆ “Product Components” on page 13
- ◆ “Installation Scenario” on page 13
- ◆ “System Requirements” on page 14
- ◆ “Installing iChain Services Software” on page 16
- ◆ “Installing the iChain ConsoleOne Snap-ins” on page 19
- ◆ “Activating iChain” on page 20

Product Components

Your iChain installation includes the following components:

- ◆ iChain Proxy Services
- ◆ iChain Authorization Server

Installation Scenario

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

- ◆ One iChain Proxy Server
- ◆ One iChain Authorization Server

- ◆ An iChain extended schema on the NDS eDirectory[®] tree where the iChain Authorization Server resides
- ◆ An iChain Service Object (ISO)
- ◆ An administrator workstation with ConsoleOne[®] iChain snap-ins installed

NOTE: For increased security, we recommend installing the iChain Authorization Server in a tree separate from your corporate file/print tree. DirXML[®] can be used to synchronize user account information between trees.

System Requirements

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

- ◆ “iChain Proxy Server Requirements” on page 14
- ◆ “iChain Authorization Server Requirements” on page 15
- ◆ “Administrator Workstation Requirements” on page 15

iChain Proxy Server Requirements

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

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

Resources

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

iChain Authorization Server Requirements

The iChain Authorization Server can be installed on the following Novell® eDirectory™ version 8.6.x platforms:

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

NOTE: The iChain Authorization Server will run on eDirectory 8.5 and above. However, because of key fixes in eDirectory, it is recommended that users upgrade to at least eDirectory 8.6.1.

For additional information on the supported platforms and full system requirements for Novell eDirectory 8.6.1, refer to the NDS eDirectory 8.6.1 Quick Start available at the [Novell Documentation site \(http://www.novell.com/documentation/lg/ndsedir86/index.html\)](http://www.novell.com/documentation/lg/ndsedir86/index.html).

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

NOTE: For increased security, it is recommended that you install the iChain Authorization Server in a tree that is separate from your corporate file/print tree. DirXML can be used to synchronize user account information between trees if needed. It is also recommended that the authorization server be the first server in the tree so that it contains the master replica of the tree.

Administrator Workstation Requirements

The administrator workstation requirements are as follows:

- ◆ Pentium* 233 MHz processor or higher
- ◆ Minimum 45 MB of free disk space
- ◆ Minimum 128 MB of RAM
- ◆ One LAN card
- ◆ Windows 98, Windows NT, Windows 2000, or Windows XP
- ◆ Current Service Pack for Windows
- ◆ Current Novell Client™
- ◆ ConsoleOne® 1.3.2 or later

- ♦ IP Connectivity between the client, the iChain Authorization Server, and the iChain Proxy Server

The latest Novell Client and ConsoleOne can be downloaded at: (<http://download.novell.com>).

Installing iChain Services Software

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

- ♦ “Installing the iChain Proxy Services Software” on page 16
- ♦ “Installing iChain Services Schema Extensions on the iChain Authorization Server” on page 17
- ♦ “Installing the iChain ConsoleOne Snap-ins” on page 19

Installing the iChain Proxy Services Software

The iChain Proxy Server should only be installed on compatible hardware (see “iChain Proxy Server Requirements” on page 14). To install the proxy server software:

- 1** Insert the *iChain 2.1 Proxy Server* CD in the CD drive of the appliance or machine.
- 2** At the license screen, type **YES** and press Enter.

The disk image is copied. When the copy is complete, the system reboots itself.

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

After installation, the first LAN adapter on the iChain Proxy Server is preconfigured with the IP address 10.1.1.1 and subnet mask 255.255.255.0. In order to administrate the server using the browser-based administration utility, you will either need to have a client workstation with an IP address on the same subnet (such as 10.1.1.2) or you will need to use the command line interface to set the IP address on the iChain Proxy Server.

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

```
>unlock
```


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

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

You will need to restart the server after resetting the eth0 address.

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

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

NOTE: After installation, your iChain Proxy Server will require some basic setup to support your iChain implementation. The basic steps are detailed in [“Setting Up the iChain Proxy Server” on page 28](#).

If you will be using the iChain Wizard to assist with configuration, you will need to enable FTP on at least one IP address for your proxy server. Once you have configured a LAN adapter as described above, enable the FTP server with the following commands:

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

NOTE: Because FTP is an insecure protocol, enabling FTP can be a security risk on your network. We recommend that you enable the FTP server on an IP address which is only accessible from a private network such as an isolated hub or crossover cable. See [“Using the iChain Wizard to Create a Basic Configuration” on page 46](#) for details on using the iChain Wizard.

Installing iChain Services Schema Extensions on the iChain Authorization Server

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

To install iChain schema extensions on the iChain Authorization Server:

- 1 If you have not already done so, install NDS eDirectory 8.5 on the machine that will be your iChain Authorization Server.

- 2 Insert the iChain authorization CD into the CD drive of a Windows client machine with IP connectivity to the iChain Authorization Server.

If this is a Windows 2000 or Windows NT machine, you will need administrator-level access to the client. The installation program launches automatically.

- 3 Click Install iChain Schema.
- 4 At the Welcome screen, click Next.
- 5 Read the license agreement. If you accept the terms of the agreement, click Yes.
- 6 Enter the administrator user name in comma-delimited LDAP format (for example, cn=admin, o=novell).
- 7 Enter the administrator password.
- 8 Enter the IP address (and port, if necessary) for the server where you want to extend the schema.
- 9 Click Next.

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

Common Bind Errors

Some of the most common bind errors are:

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

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

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

ldap_simple_bind failed 81(Can't contact LDAP server), dn: cn=admin,o=novell: Either the IP address/port combination is incorrect or the

LDAP server is not running. Verify the IP address and LDAP port, make sure the server is running, and try again.

NOTE: Contact Novell Technical Support if you are unable to resolve an error or if you have trouble creating or modifying iChain objects after extending the schema.

Common Log File Errors

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

The LBURP extension is not available on the server. Using standard LDAP calls: This generally means the LDAP server is out of date. You should verify that the latest LDAP server (included with NDS eDirectory 8.5 or later) is installed on the server to ensure that the schema is completely extended.

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

Record1: LBURP operation failed: 20(Type or value exists), dn:cn=schema: This error is expected if the server has already been extended with a previous version of iChain. The new iChain 2.1 schema attributes will still be added.

NOTE: Contact Novell Technical Support if you are unable to resolve an error or if you have trouble creating or modifying iChain objects after extending the schema.

Installing the iChain ConsoleOne Snap-ins

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

NOTE: iChain 2.1 requires ConsoleOne 1.3.2 or later for all of the snap-ins to function correctly.

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

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

NOTE: After ConsoleOne is installed, make sure you close it before starting to install the snap-ins.

- 2 Insert the iChain authorization CD into the CD drive of the server or the administrator workstation.

The installation program launches automatically.

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

After completing the full installation, you will need to use ConsoleOne to create the iChain service group object, along with the access control list (ACL) rule objects, and make any other configuration adjustments. See [Chapter 3, "Setting Up a Basic Configuration," on page 27](#) for more details.

Activating iChain

The evaluation version of iChain 2.1 will expire and not work after 90 days. When the evaluation version expires, the iChain Proxy Server will not function. You must purchase an iChain 2.1 license and apply a Product Activation Credential in order for iChain to continue working after 90 days.

Activating iChain involves these tasks:

- ♦ [Generating a Product Activation Request](#)
- ♦ [Submitting an Activation Request](#)
- ♦ [Activating iChain by Applying the Product Activation Credential Received from Novell](#)

This section also contains the following topics:

- ♦ ["Viewing Product Activations for iChain" on page 24](#)
- ♦ ["Troubleshooting iChain Activation Problems" on page 24](#)

After you purchase an iChain 2.1 license, Novell will send you a Customer ID via e-mail. If you do not remember or do not receive your Customer ID, please call the Novell Activation Center at 1-800-418-8373 in the U.S. In all other

locations, call 1-801-861-8373. (You will be charged for calls made using the 801 area code.)

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

Generating a Product Activation Request

You use your Customer ID to generate a Product Activation Request in ConsoleOne. You must have an iChain Service Object (ISO) already set up to generate a Product Activation Request to activate iChain. See [“Creating an iChain Service Object” on page 27](#) and [“Setting Up the iChain Proxy Server” on page 28](#) for information about setting up the ISO. The iChain Proxy Server must be able to access the ISO via LDAP in order to generate the activation request.

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

1 Open ConsoleOne.

You must use the version of ConsoleOne that you installed with iChain 2.1.

2 Choose Wizards > Create an iChain Activation Request.

3 Select the iChain Service Object (ISO) where you want iChain to be activated > click Next.

4 Enter your Novell Customer ID > click Next.

5 A dialog will appear, asking whether you want to notify the proxy server of the activation request automatically or manually. If you have FTP enabled on the proxy (see [“Installing the iChain Proxy Services Software” on page 16](#)), you can click Yes and enter the IP address and config password of any iChain Proxy Server that has been configured to read this ISO.

Otherwise, click No and enter the command:

```
getactivationrequest
```

from the iChain Proxy Server console.

6 The wizard will then wait for the iChain Proxy Server to generate a Product Activation Request. Every 30 seconds the user will be prompted to ask whether they would like to continue to wait for a response. If no response occurs, ensure that there is LDAP communication between the iChain Proxy Server and the ISO.

7 Do one of these steps:

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

Or,

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

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

8 Click Launch to launch the Novell Product Activator Web site.

NOTE: You will need to submit this Product Activation Request to Novell at the Novell Product Activator Web site (<http://www.novell.com/activator>). At that site, you can either paste the text of the Product Activation Request in a text area or specify the path to the file. If you are working on a machine without a browser, save the file to a diskette to upload at a machine that has a browser.

Submitting an Activation Request

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

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

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

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

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

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

3 Click Submit.

4 Mark the product you are activating.

You need to activate each line item.

5 Click Submit.

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

Activating iChain by Applying the Product Activation Credential Received from Novell

You activate iChain 2.1 by installing the Product Activation Credential you received from Novell. You install this file via ConsoleOne. If a tree has multiple iChain Service Objects (ISOs), the Product Activation Credential must be installed for each ISO.

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

1 Open the e-mail from Novell that contains the Product Activation Credential.**2** Do one of these steps:

Save the Product Activation Credential file.

Or,

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

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

3 Open ConsoleOne.**4** Choose Wizards > Install an iChain Activation.**5** Select the iChain Service Object (ISO) > click Next.**6** Do one of these steps:

Specify where you saved the iChain Activation Credential > click Next.

Or,

Paste the contents of the iChain Activation Credential in the text area > click Next.

7 A dialog will appear, asking whether you want to notify the proxy server of the new credential automatically or manually. If you have FTP enabled on the proxy (see “Installing the iChain Proxy Services Software” on page 16), you can click Yes and enter the IP address and config password of any iChain Proxy Server that has been configured to read this ISO.

Otherwise, click No and enter the command:

refreshcredentials

from the iChain Proxy Server console.

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

8 Click Finish.

Viewing Product Activations for iChain

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

Troubleshooting iChain Activation Problems

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

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

In addition, make sure that:

- ◆ The eDirectory schema has been correctly extended for iChain 2.1.
- ◆ You are not trying to install the Activation Request file as an Activation Credential.

- ◆ The credential you are using was created from an activation request generated on the same tree where your ISO is located. The credential will only function on a specific tree.

3

Setting Up a Basic Configuration

This chapter explains the tasks you need to complete to set up a basic Novell® iChain® infrastructure. For details on the basic scenario, refer to “[Installation Scenario](#)” on page 13. This section contains the following topics:

- ♦ “[Creating an iChain Service Object](#)” on page 27
- ♦ “[Setting Up the iChain Proxy Server](#)” on page 28
- ♦ “[Setting Up Protected Resources](#)” on page 31
- ♦ “[Defining iChain Access Control Rules](#)” on page 34
- ♦ “[Using the iChain Wizard to Create a Basic Configuration](#)” on page 46
- ♦ “[Installing the iChain Proxy Server on Your Network](#)” on page 52

Creating an iChain Service Object

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

To create an iChain Service object (ISO):

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

or

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

- 2** Select iChain Service and define a name for the service or domain (for example, Test) > click OK.

Setting Up the iChain Proxy Server

The iChain Proxy Server functions as the primary access point into your iChain infrastructure. The iChain Proxy Server is implemented on approved hardware.

This section provides a brief introduction to the basic steps needed to set up the iChain Proxy Server. For more details, see "Installing the iChain Proxy Server on Your Network" and [Chapter 7, "Using and Tuning iChain Features," on page 123](#).

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

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

For example, `http://xx.xx.xx.xx:1959/appliance/config.html`

where `xx.xx.xx.xx` is the IP address for the proxy server. You should have configured an IP address during the installation of the iChain Proxy Services software.

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

- 2** Accept the default username (do not enter a password) > click OK.
- 3** Click System > Actions > Password > set a password for the proxy server.
- 4** Click Home > Introduction > verify that iChain Proxy Server 2.1 is installed and running on the server. (This is shown as a bitmap that lets you know if you're running version 2.1.)
- 5** Click Network > IP Addresses > configure the Ethernet ports as follows:
 - ◆ Accept the Eth0 adapter existing setting.
 - ◆ Set the Eth1 adapter with the private IP address for the network.
 - ◆ Set the Eth2 adapter with the public IP address for the network.
- 6** Click Gateway-Firewall > set the iChain Proxy Services default gateway to the gateway necessary to access your public IP address.
- 7** Click Network > DNS > specify the DNS domain name (for example, `www.novell.com`) and the IP address of the DNS server.

- 8** Click Apply to have the new settings take effect.
- 9** Click System > Actions > verify the internal connection to your network by pinging a server within your internal network.

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

- 1** Click Configure > Access Control.
- 2** Specify the fully distinguished name of the ISO object name for the iChain service. You must use commas as delimiters, for example, **cn=myISO,o=novell**.
- 3** Specify the following LDAP profile settings:
 - ◆ LDAP server addresses for the iChain LDAP access control servers
 - ◆ LDAP port on the iChain LDAP access control servers
 - ◆ LDAP proxy user
 - ◆ Password
 - ◆ Enable secure access to LDAP server (only if you are using secure LDAP)
 - ◆ LDAP server trusted root file (for secure LDAP)

NOTE: The LDAP user name and password must have supervisor rights to the container you are searching.

- 4** Click Apply.
- 5** Click Refresh ACLCHECK.

To set up access to the iChain Authorization Server for authentication functions, you will need to create one or more authentication profiles. The following steps will create an LDAP authentication profile to authenticate users to your iChain Authorization Server. (You can also create SSL mutual authentication and RADIUS profiles if you want to use these authentication methods.)

- 1** Click Configure > Authentication.
- 2** Insert a new profile > name the profile > select LDAP Authentication and click the LDAP Options button.
- 3** Set the server IP address to the iChain Authorization Server address.
- 4** Select port 389 for non-secured LDAP, or port 636 for secure LDAP (or another port as configured).

- 5** Enable secure access to LDAP server (only if you are using secure LDAP).
- 6** Specify LDAP server trusted root file (only if you are using secure LDAP).
- 7** Specify a username and password for LDAP access (leave the field blank for anonymous bind).
- 8** Set Use Distinguished Name.
- 9** Click Insert > enter an LDAP context (for example, ou=test,o=mycompany). Repeat for each context users will authenticate from.
- 10** Click OK > OK > Apply.

To set up a Web Server accelerator:

- 1** Click Configure > Web Server Accelerator > Insert.
- 2** Specify a name for the accelerator, using a maximum of 8 characters. This must be unique for each Web Server accelerator.
- 3** Specify a DNS name for the accelerator (for example, www.novell.com).
This is the DNS name by which users will access the resource and should resolve to the public IP address of the iChain Proxy Server.
- 4** In Web Server address, click Insert > specify the IP address of the origin Web Server that contains the desired content.

Either the IP address or the DNS name resolving to the origin Web Server can be used. This will usually be on your private network. Clients should not be able to access this server directly, or the iChain infrastructure will be bypassed.
- 5** For the Accelerator IP address, check the public IP address or addresses that the DNS name specified in Step 3 resolves to.
- 6** Check Enable Authentication.
- 7** Click Authentication Options > select an existing profile from the list > click Add to set the profile as the Service Profile.
- 8** Click OK > OK > Apply.

Setting Up Protected Resources

To integrate and allow access to Web-based application resources, you must set the appropriate parameters in the iChain Service object. Sometimes Web-based resources, called Protected Resources in iChain, need additional information about the user to be passed into the application to protect the resource or customize the user interface. This additional information is usually stored in Novell eDirectory or some other database. iChain uses special object-level access control plug-ins to access the database and retrieve the additional information. The iChain Proxy Services then passes the information to the application by adding it to the URL query string or as a header variable.

To set up a protected resource for an iChain service:

- 1** From ConsoleOne, click the Protected Resources tab on the ISO object you created for this configuration.
- 2** Click Add (the icon with the plus [+] sign).
- 3** Specify a name for the resource and the URL for the resource in this format: `http://www.resource.com`, where `www.resource.com` is the DNS name specified when you created the Web Server Accelerator.

If you enter only the DNS name of your iChain Proxy Server, the iChain snap-in will attach the `http://` prefix automatically.

You can enter the URL with subfolders. See [“Differentiating Among Protected Resources” on page 33](#) for more information.

NOTE: If the URL starts with `https://` (that is, this is a secure Web site), you will still specify `http://` in this field. iChain uses this field for matching purposes. It does not affect the URL in the query string.

- 4** Choose whether this protected resource will be Public, Restricted, or Secure (see [“Differentiating Among Protected Resources” on page 33](#) for more information). Also select whether any associated Object-Level Access Control (OLAC) parameters should be sent in the query string or as a header variable (see [“Setting Up Object-Level Access Control” on page 95](#) for more information).
- 5** Click OK > Apply to save the resource.

You can add as many protected resources as you want.

- 6** If you are using the iChain Wizard to administer the iChain Proxy Server (that is, you have configured the FTP server on the iChain proxy), you can refresh the iChain Proxy Server from the snap-in when prompted.

Otherwise, you will need to go to the Web-based administration utility and click Configure > Access Control > Refresh ACLCheck and Refresh OLAC to read the new protected resource.

If you will be specifying OLAC parameters to pass to the origin Web Server, you will need to define how these are stored in eDirectory.

NOTE: For more information on setting up object-level access control, including information on the plug-ins provided with iChain refer to [“Setting Up Object-Level Access Control” on page 95](#).

To specify these parameters:

- 1** Select the resource you just added > click OLAC (the OLAC button is located below the Delete button, which is the icon with the minus [-] sign).
- 2** For each object-level access control attribute to be passed to the application, define the Name, Data Source, and Value.

NOTE: Object-level access control attributes should be created only for protected resources that are configured to process query strings or custom headers. Typically, access to these protected resources is handled by CGI scripts, Web Server plug-ins, dynamic pages, or similar methods. If this is not the case, users may encounter errors when attempting to access URLs within the protected resource.

The Name field contains an identifying name for the attribute that is passed to the application.

The Data Source field contains the name of the database from which to retrieve the information. With the shipping plug-in for OLAC, this field could be LDAP or CONSTANT.

The Value field contains a key for retrieving the attribute value in the data source. The attribute is added to the URL query string in a Name = Value format. With the shipping LDAP plug-in for this OLAC, this will be the LDAP attribute name for the value where the attribute is stored. This can be different from the eDirectory attribute name.

For example, if the application requires the last name of the user and the data is stored in an LDAP-accessible directory under Surname, the entries would be LastName, LDAP, and Surname. If the user's last name is Smith, the attribute LastName=Smith would be added to the URL query string whenever the user accesses the protected resource.

See [“Setting Up Object-Level Access Control” on page 95](#) for a table of appropriate values for the Data Source and Value fields.

NOTE: If you have already defined the necessary object-level access control attributes for this accelerator on another protected resource on the same iChain service object, you can import them directly into this protected resource. At the OLAC Parameters dialog, select the Import button (located under the Modify button). This will display a list of other protected resources that have had OLAC parameters specified. Select the desired resource and click OK. The OLAC parameters from that resource will appear in the OLAC Parameters dialog box.

- 3 Access the URL of the iChain Proxy Server where you installed the iChain Proxy Services software to launch the iChain Proxy Services browser-based administration tool.

For example, `http://xx.xx.xx.xx:1959/appliance/config.html`, where `xx.xx.xx.xx` is the IP address.

- 4 Click **Configure > Web Server Accelerator > Modify**.

- 5 Click **Authentication Options**.

If you want to pass the basic authentication header to the Web Server, check the **Forward Authentication Information to Web Server** check box. If you have changed the DNS name for this accelerator, ensure that the cookie domain is still accurate.

- 6 Click **OK > OK > Apply**.

- 7 To start the OLAC processor, click **Configure > Access Control**, then check the **Enable Object Level Access Control** check box > click **Apply**.

NOTE: OLAC reads the information in the access control page/tab to find out where to retrieve the parameters from. If this information changes, you will need to stop and start OLAC again. To stop OLAC service, click **Configure > Access Control page/tab > uncheck the Object Level Access Control check box > click Apply**. To start OLAC again, perform **Step 7** (above this note).

IMPORTANT: If the OLAC setup is changed from non-SSL to SSL or vice versa, OLAC will need to be restarted.

Differentiating Among Protected Resources

iChain 2.1 provides a feature that allows administrators to differentiate among different protected resources.

iChain 2.1 provides three levels of security for protected resources:

- ♦ **Public** — No authentication or access control will exist for the pages under this protected resource.
- ♦ **Restricted** — Only authentication will exist for the pages under this protected resource.

- ◆ Secure — This is the most secured type of protected resource. To access the pages under this protected resource, the user needs to be authenticated through the proper authentication mechanism and will also need to pass through the access control. This is the default selection for the new protected resource.

NOTE: Protected resources created with earlier versions of iChain snap-ins are treated as secure.

iChain 2.1 allows you to use wildcard characters when specifying the URL for a protected resource. If the protected resource's URL is absolute, ending with a trailing slash (/), iChain will match just the URL. However, if the protected resource's URL ends with a question mark (?), iChain will match all files in the specified folder. For example, `http://novell.ichain.com/dir1/?` will match all the files under the `dir1` folder. If the protected resource's URL ends with an asterisk (*), iChain will match all the files under the specified folder and all the subfolders and their contents. For example, `http://novell.ichain.com/dir1/*` will match all the files under the `dir1` folder and any subfolders below `dir1`.

While matching the protected resource URLs, iChain looks for the most specific match to decide the URL access. For example, if `http://ichain.novell.com` has Public access, but `http://ichain.novell.com/secure/` has Secure access, all of the pages in the secure folder would still require authentication and access control, while all other pages would be considered public pages and would not require authentication or access control.

IMPORTANT: Authentication *must* be enabled for the Public, Restricted, or Secure levels of security to function.

If you have no authentication on the accelerator, you are using the iChain box as a caching appliance only. In this case, even if you set up a Public resource, such as a home page, not only is the home page available to anyone, but all pages below the home page are accessible to anyone.

Defining iChain Access Control Rules

After a user has logged in, Access Control List (ACL) rules control what resources the user can access. By default, the user has access to nothing. Only those resources explicitly listed in your ACL rules (specified by the URL) can be accessed by the organizations, organizational units, groups, or users listed in the Apply To list for the rule. Whenever possible, it is recommended that you use the highest-level object in the list of allowed users, making it easier and faster to configure an ACL rule.

iChain access control checks the ACL rules in the following sequence:

- ◆ Rules for all the user's containers, starting with the highest level container
- ◆ Rules for all communities of all the user's containers (if ACLCHECK was started using the /m option)
- ◆ Rules for the user's groups
- ◆ Rules for all communities of the user's groups (if ACLCHECK was started using the /m option)
- ◆ Rules for the user
- ◆ Rules for all communities of the user (if ACLCHECK was started using the /m option)

When a user tries to access a protected resource that has been defined as Public, the user is immediately granted access. If the resource is defined as Restricted, the iChain system checks the user's browser cookie address to see if he or she is a currently authenticated user and either lets the user access the resource if the user is authenticated or prompts the user for authentication. A current authenticated connection is all that is required. However, when a user attempts to access a URL that has been defined as Secure, the user must log in to eDirectory and provide a password. When the user is authenticated, the ACL rules are checked to see if the user is allowed to access the site.

ACL rules allow the use of an asterisk (*) or question mark (?) as wildcard characters when specifying URLs. The asterisk indicates that the user can have access to the folder contents and all subfolders. The question mark indicates the user can have access to the folder contents, but not the subfolders. Also, each ACL rule can be individually disabled or enabled, allowing you to turn on or off a particular rule for a time without losing its parameter settings.

ACL rules are stored in a cache that is updated periodically at a configurable interval. For performance reasons, the recommended cache refresh interval is three to six hours. If you make changes or additions to the ACL rules and want the cache to be updated immediately, use the manual Refresh option available in the Configure > Access Control tabs of the iChain Proxy Services Browser-based administration tool. If you have FTP enabled on the proxy, you can automatically refresh the iChain proxy when prompted by the snap-in.

When you create an entry in the URL list of an ACL rule, at least one of the two fields (Resource Name and URL) is required.

If only the URL is specified, it must be given as an absolute URL (for example, <http://www.novell.com/index.html>, not [/index.html](#)). The URL may contain wildcards. The ACL rule will match any request for the URL (including wildcards).

If only the Resource Name is specified, the ACL rule will match any request for the exact path of the Resource Name. For example, if the protected resource myserver has been defined as `http://www.novell.com`, and a URL list entry is created with myserver as the Resource Name and with no URL, then the ACL rule will apply to the `http://www.novell.com` URL only.

If both the Resource Name and the URL are specified, the URL must be given as a relative URL (`/index.html`, not `http://www.novell.com/index.html`) and may include wildcards. The ACL rule will match requests for the combined Resource Name and URL, including wildcards. For example, if the Resource Name is myserver and the URL is `/documentation/*`, then the ACL rule will apply to `http://www.novell.com/documentation/*`.

To create a new ACL rule for iChain:

- 1** From ConsoleOne, select File > New > New Object.
or
Click the New iChain Object icon.
- 2** Select ACL Rule > click OK.
- 3** Define a name for the rule > click OK.
- 4** Select the rule you just created and click Properties > Access Control.
- 5** Under the list of Allowed URLs, click Add > define a name and URL for a resource that this rule will control access to.

You can use an asterisk (*) or question mark (?) as a wildcard character when specifying URLs. The asterisk indicates that the user can have access to the folder contents and all subfolders. The question mark indicates the user can have access to the folder contents, but not the subfolders.
- 6** Under the Apply To List, click Add to browse to and select the Os, OUs, groups, and users to which this rule applies.

The Os, OUs, groups, and users in the Apply to List are allowed access to the listed URLs.
- 7** Under the Exception List, click Add to browse to and select the Os, OUs, groups, and users that are exceptions to this rule.

The Os, OUs, groups, and users in the Exceptions List are a subset of the Apply to List and are objects that are denied access to the listed URLs.
- 8** To enable the ACL rule, check the Enable Access Control check box at the General tab.

- 9 To disable the ACL rule and save it for later use, uncheck the Enable Access Control check box.

ACL Exceptions

You can exclude certain users or group members listed in the Apply To List that you do not want to have access to the specified URLs. However, these exceptions are made on a per rule basis. So, although a user may be excluded from one rule, he may still have access to the URL through other ACL rules. Double-check all ACL rules for the resource to be sure exceptions are as you expect.

You can also define a subset of the destination URL as an exception for an ACL rule. For example, an ACL rule could be set on `http://ichain.novell.com/*` for the users in the `o=novell` container. By using the URL exception feature, an administrator could define `http://ichain.novell.com/private/*` as a URL exception. iChain access control would then allow the users in the `o=novell` container to go to all the pages under `http://ichain.novell.com/`, except `http://ichain.novell.com/private/`.

ACL Theory of Operations

The following is an example that explains the process of ACLs:

Table 1 ISO Protected Resource

Name	URL Prefix	Type
A	<code>http://www.novell.com/</code>	Public
B	<code>http://www.novell.com/*</code>	Restricted
C	<code>http://www.novell.com/*.html</code>	Secure
D	<code>http://www.novell.com/*.gif</code>	Secure
E	<code>http://www.novell.com/secret/</code>	Secure

Table 2 ACL Access Control

Resource Name	URL PostFix	Apply To	Exceptions
A	<code>Secret/*</code>	Jack	
B	<code>Secret/*</code>	Jack	

Resource Name	URL PostFix	Apply To	Exceptions
E		Jack	
	http://www.novell.com/ secret	Jack	

1. At the browser, Jack enters: http://www.novell.com/secret.
2. DNS points Jack's browser to the iChain box.
3. The iChain box has a Web Accelerator with www.novell.com defined and the Enable Authentication switch is turned on. (If the switch was not turned on, iChain would simply cache the resource and would provide no security).
4. The ISO entries are compared to the URL request to determine if authentication is required.
5. Jack is asked to enter his name and password for authentication. It is valid.
6. Completion of the initial URL compare takes place.
 - ◆ A — Serves up only the index.html or the default.html. *No match.*
 - ◆ B — The asterisk (*) indicates everything in the doc root directory and everything in every subordinate directory. *Match.*
 - ◆ C — The asterisk (*) in this protected resource has two meanings: any HTML extension-based file in the doc root directory and all subordinate directories. *Match.*
 - ◆ D — Similar to C, but any .GIF-based file in the document root directory and below. *No match.*
 - ◆ E — Similar to A, but this one *matches* the URL Jack entered.

Notice the types of protected resources: Public, Restricted, and Secure. these are labels to help in the URL comparison process.

- ◆ Restricted — This label indicates that once the user has authenticated and the requested resource matches the Restricted protected resource entry, the URL will be displayed. No further security processing is required.
- ◆ Secure — This label indicates that a second (ACL) compare of the URL entered will follow. Although the URL prefix that is labeled as secure can be used/linked to create the pattern that will be (ACL) compared against, it is provided only for semantic convenience. The

field names URL prefix (in the ISO) and URL postfix (in the ACL) suggest a direct link. Although these fields are commonly concatenated, there is no direct link between the Secure protected resource and the ACL compare. (This will become clearer in the next step when the ACL compare is processed.)

NOTE: Two of the protected resource entries, B and C, match up with the URL Jack entered. The most specific match, based on the URL, takes precedence.

7. ACL check takes place next. In our example, there are four entries in the ACL. These are provided to show various methods that can be used to define the entry for the ACL.

- ◆ The first entry shows both fields populated. `http://www.novell.com/` is basically being cut and pasted from the ISO entry and then concatenated with the URL postfix of `Secret/*` to form `http://www.novell.com/Secret/*`. Since Jack is listed in the Apply To list, he is granted access to the resource.
- ◆ The second entry also shows both fields populated. `http://www.novell.com/*` is basically being cut and pasted from the ISO entry to be concatenated with the URL postfix of `Secret/*` to form `www.novell.com/Secret/*`. Notice in this case `iChain` provides the intelligence to strip the trailing asterisk (*) during the cut and paste. Otherwise, the resource would look like: `http://www.novell.com/*Secret/*` and would produce a false negative match.

Since Jack is listed in the Apply To list, he is granted access to the resource.

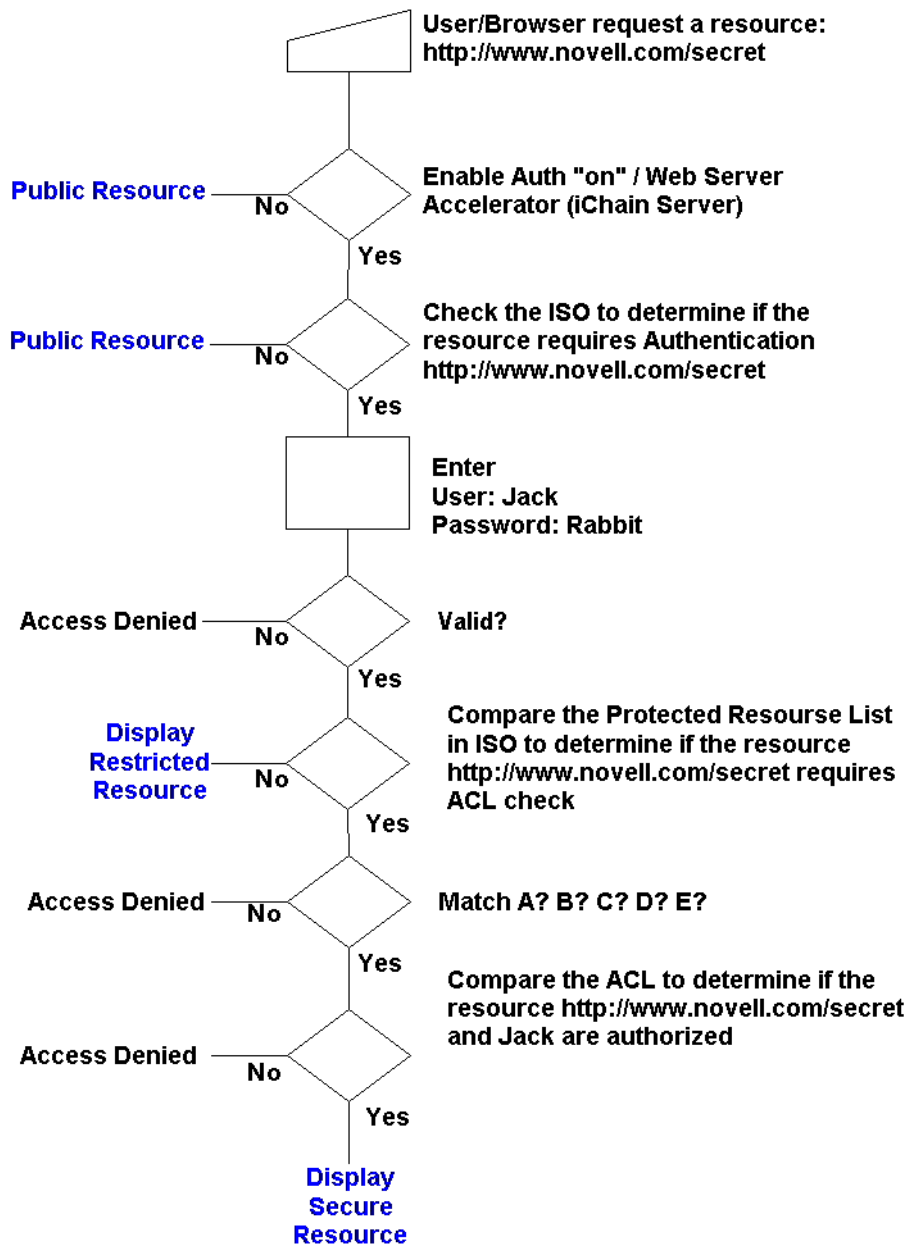
- ◆ The third entry shows only the Name field populated. It simply provides a cut and paste of `http://www.novell.com/secret` to compare to the originally entered URL: `http://www.novell.com/secret` to produce a match and grant access. The ISO object denotes this URL prefix labeled as Secure. At this point in time — the processing of the ACL — this label has no meaning. It was used in the previous authentication process involving the ISO to indicate further security processing (ACL).
- ◆ The fourth entry shows only the URL postfix field populated with `http://www.novell.com/secret` to compare to the originally entered URL: `http://www.novell.com/secret` to produce a match and grant access.

Had Jack entered `www.novell.com/secret/stuff/confid.xml`, the only ISO protected resource that would allow access would be "B". With "B" eliminated, access would not be granted. The ISO protected resources A, C, D, and E do not match up. A new ISO entry could be created similar to C or D, such as `www.novell.com/secret/*.xml`. But

rather than duplicate new entries, a new entry, `www.novell.com/*` (secure), could replace C, D, and the case for XML. During the processing at the ISO protected resource portion, `www.novell.com/secret/stuff/confid.xml` matches with `www.novell.com/*`. This match does not immediately grant access. It simply allows further processing at the ACL portion. Notice the `/*` at the ISO protected resource is appended to provide a match. iChain strips this off in the ACL processing portion.

The following flow chart shows the basic operation of the ACL. (It is not meant to be an all-inclusive code translation.)

Figure 1 Basic ACL Operation



Defining Dynamic Access Control Rules

Dynamic Access Control Rules allow an administrator to set up an access control rule based on a query of user attributes. (If a user's attribute value satisfies a predefined value, the rule would be applied to that particular user.) This type of query can be based on a user's attributes (the user's location, salary, hobby, etc.). For example, an administrator could configure a rule that says "Apply this rule to all the users who are in San Jose and have a salary greater than (>) \$50,000."

Dynamic Access Control Rules are based on the following principles:

- ◆ A rule is applied to the users who satisfy a condition (query) provided by the administrator.
- ◆ Static applicability of the rules (object-based ACLs) will still exist in its current form. The DN exception and URL exception are also supported with the new type of rules. An ACL rule can contain both a dynamic query and a static "Apply To" list. Thus, one ACL rule can act as a dynamic rule and also as a traditional (static) rule.
- ◆ Because user attribute values can be dynamic, the administrator may want to limit the time that a rule can be cached in memory. This cache time, called "Time to Live," can be set in number of minutes and the cache for that rule will be expired after the specified duration elapses. The Time to Live value can only be set for dynamic ACLs. If no value is entered in the Time to Live field, the cache will not expire until the overall cache refresh takes place for ACLCHECK. If the Time to Live value is set at zero, the rule will not be cached.
- ◆ The administrator who creates a dynamic ACL must have write privileges on the ISO. (This is not required for non-dynamic ACLs.)

Setting Up Dynamic Access Control Rules

The dynamic ACL setup GUI allows you to create and test the search filter, then convert it to standard LDAP format to be stored in eDirectory. This query is later used to allow or disallow dynamic access control in iChain.

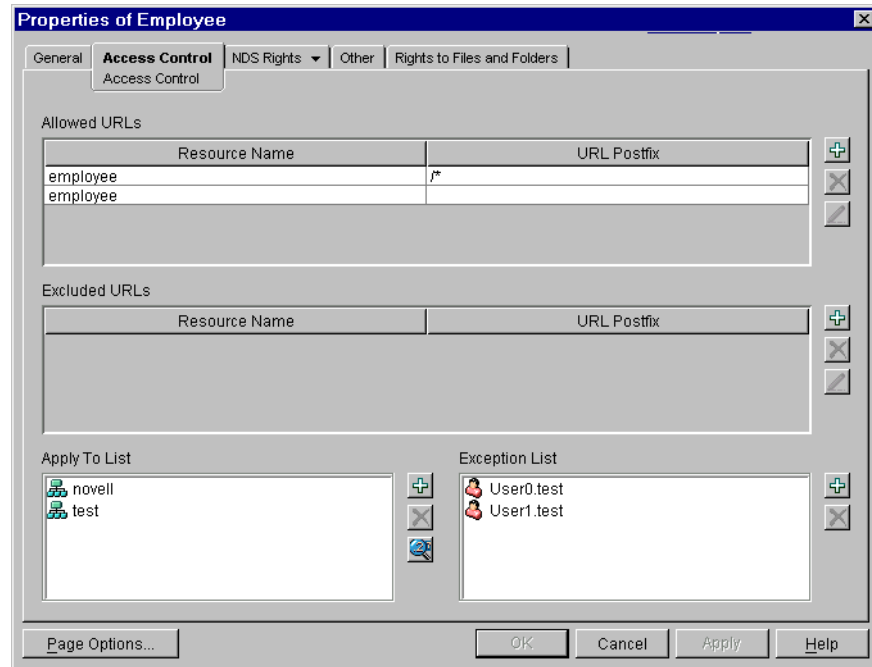
Dynamic ACLs can be defined when using the iChain Web Accelerator Wizard or the iChain Access Control snap-ins.

To create a dynamic ACL:

- 1** From ConsoleOne, open the ACL Rule object.

- 2 At the Access Control tab click the Dynamic ACL button in the Apply To dialog box (see [Figure 2](#)).

Figure 2 Access Control Tab in ACL Rule Object Snap-in



This opens the Dynamic ACL Query Setup dialog box (see [Figure 3](#)).

Figure 3 Dynamic ACL Query Setup

Dynamic ACL Query Setup

iChain service object
CN=iso-atul.O=novell

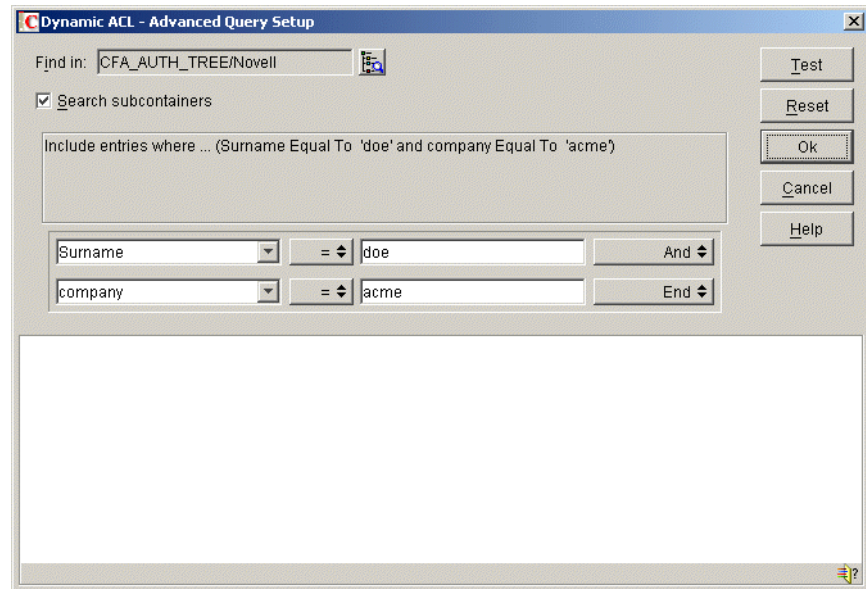
NCP server object
CN=NSD43.O=novell

Search filter
(&(sn=doe)(company=acme))

Specify the time to live for this acl
500 minutes

OK Cancel Help

- 3 Browse and select the active ISO and NCP Server object, if not pre-populated, which points to the appropriate LDAP group object having the updated NDS to LDAP attribute and class mappings.
- 4 Either update the existing LDAP Search filter field to form the query or click on the Advanced Dynamic ACL query setup button (see [Figure 4](#)).

Figure 4 Dynamic ACL Advanced Query Setup

- 5** Specify the time-to-live (in minutes) in cache for the dynamic ACL.
- 6** If you choose to build and test the query in the Advanced Dynamic ACL panel, the Find in field allows you to specify the container where to start searching for the objects.
- 7** The first query field (see [Figure 5](#)) allows you to select an object property to use as a search criterion, or select [Object Type] to search for the objects of a specific class.

NOTE: The [Object Type] is required only to test this query in this panel. Once the query is formed and tested, the administrator should remove this search criteria before saving it since iChain currently allows dynamic access control on user objects only.

Figure 5 Query Field

Click the comparison operator to select a logical operator to use when comparing the value of the specified attribute with the actual attribute value in the Authorization Server (eDirectory).

The second query field allows you to specify the attribute value to compare against the actual attribute value in the Authorization Server (eDirectory). The syntax is the type of data contained in the attribute, such as character string, integer, etc.

Click on a statement connector keyword to select the keyword that specifies how this statement connects with the next statement or group of statements in the query. "And" specifies that both this and the next statement (or statement group) must be true for a match to occur. If there is no next statement, selecting this keyword adds one. "Or" specifies that either this or the next statement (or statement group) must be true for a match to occur. If there is no next statement, selecting this keyword adds one. "Insert Row" adds a new statement below this one. "Delete Row" deletes the statement from the query. "New Group" adds a new statement group below this one. "End" specifies that this statement ends the query.

NOTE: When creating a query there are some limitations with the query tool. It does not allow a mixture of "and/or" conditions in the same group or between groups.

- 8 Click the Test button to test the query in the eDirectory namespace and display the results at the bottom of the dialog box.

You can right-click objects in the result list to perform actions as you do in the ConsoleOne right pane.

- 9 Click the Cancel button to discard the setup and exit the panel.
- 10 Click OK to convert the query (search criteria) to standard LDAP search filter format to be stored in the ACL object. Make sure your eDirectory to LDAP mappings are present in the right LDAP group object as pointed by the NCP Server object. This will close the Advanced window and allow you to edit the formed LDAP search filter before saving it along with its time-to-live-in-cache attribute.

IMPORTANT: It is important to understand the differences and associated limitations while testing your query. Query testing is done in the eDirectory namespace and stored in LDAP format. Thus, there may be situations where you need to specify the value in the second query field in eDirectory format while testing, but change it to LDAP format before saving it.

Using the iChain Wizard to Create a Basic Configuration

The iChain Wizard is an interface included with ConsoleOne that helps to integrate the configuration of the iChain Proxy Server and the iChain Authorization Server. It can dramatically decrease the amount of time needed

to accelerate a resource on the network. While most normal configurations can be done using the wizard, there are some advanced configurations which still require the use of the Web-based administration utility or the command line.

The following walkthrough assumes you have configured a LAN adapter and enabled the FTP server as described in [“Installing the iChain Proxy Services Software” on page 16](#). Only the steps necessary to create a working configuration are described in this section. A detailed listing of all the wizard options is available in [Appendix A, “iChain Management,” on page 217](#).

Launching the Wizard

To launch the iChain Wizard:

- 1** In ConsoleOne, select the container where you want to create your iChain objects in ConsoleOne.
- 2** Select Wizards > iChain Web Server Accelerator.

Several wizard pages will follow, as described in this section.

Wizard Page 1

- 1** If you have already created an iChain Service object, select it by using the browser button next to the iChain Service Object field. Otherwise, you can create an iChain Service object by entering the common name without the context in the iChain Service Object field and clicking Enter. You will then be prompted to create the object in the container you selected when launching the wizard.
- 2** Enter the IP address of the LAN adapter for the iChain Proxy Server that has the FTP server enabled in the Proxy Server IP address field.
- 3** Click Next. You will be prompted for a password. Enter the password you have selected for the Config user on the iChain Proxy Server, or leave this field blank if no password was set.
- 4** Click OK to log in and read the configuration from the iChain Proxy Server. If the login is successful, click Next to proceed to Page 2 of the wizard. Otherwise, you will need to determine why the wizard failed and try to log in again.

Wizard Page 2

- 1** Click Add to create a new iChain Access Control Server. The information on this page corresponds to the Configure > Access Control page in the Web-based administration utility.
- 2** Enter the IP address, port, administrator name, and password for the iChain Authorization Server in the ensuing dialog box. The administrator name should be in comma-delimited LDAP format; if you use the Browse button to select the user, the name will automatically be formatted correctly. If you will be using secure LDAP for access control, check the LDAP over SSL box and enter the name of the Trusted Root file. You will need to import the Trusted Root certificate for that LDAP server through the Web utility. Click OK to save this information.
- 3** Enter optional information as described in [Appendix A, "iChain Management," on page 217](#).
- 4** Click Next to continue to Wizard Page 3.

Wizard Page 3

- 1** If this is an initial configuration, enter the Web Server Accelerator name and DNS name.

The accelerator name must be 8 characters or less and must be unique.

The DNS name is the DNS name by which users will access the Web server and should resolve to the public IP address of the iChain Proxy Server.

When additional accelerators have been configured, you can modify, delete, and set up multi-homing from this page as well.
- 2** Click Next to continue to Wizard Page 4.

Wizard Page 4

- 1** Click Add (located next to the Web Server Addresses field). Enter the private IP addresses or DNS names of the origin Web server in this field. Clients should not be able to access this address directly, or iChain can be bypassed.
- 2** Enter the port the origin Web server is running on in the Web server port field.
- 3** Select the check box next to the public IP addresses through which clients can access this Web server.

NOTE: You cannot currently define new proxy server IP addresses through the wizard. This must be done from the command line or the Web-based administration utility.

- 4** Enter the port through which the content will be delivered from the iChain Proxy Server to the browser in the Accelerator proxy port field.
- 5** Click Next to continue to Wizard Page 5.

Wizard Page 5

- 1** Select the Enable Authentication check box to enable the authentication option fields.
- 2** Select the Enable Secure Exchange check box if you want content delivered from the iChain Proxy Server to the client over a secure channel.
- 3** In the SSL Listening Port field, enter the secure port for authentication and Secure Exchange to use. This port must be different for each accelerator using a given IP address on the iChain Proxy Server.
- 4** If you have created a custom certificate using the iChain certificate creation utility, you can enter the name in the SSL Certificate Name field. Otherwise, leave the setting on Auto and iChain will generate its own certificates.
- 5** In the Session Timeout Interval field, enter the idle time before re-authentication is required.
- 6** Select the Forward Authentication Information to Web Server check box if you want user credentials or OLAC parameters passed to the origin Web server.
- 7** Select the Authenticate over HTTP check box if you want LDAP authentication to occur over an HTTP (clear text) connection.

NOTE: The Authenticate over HTTP option cannot be used simultaneously with Secure Exchange or SSL Certificate Authentication.

- 8** Select the Add button (located next to the Authentication Profiles field).
This area is very similar to the Configure > Authentication tab in the Web-based administration utility, and can be used to add, delete, or modify authentication profiles.
- 9** Enter a unique name, 8 characters or less, in the name field.

- 10** Select the radio button for the type of profile you want to create — either mutual SSL certificate authentication, LDAP authentication, or RADIUS authentication.

This walkthrough assumes LDAP authentication is used. Click the Authentication Options button.

- 11** Click Add (located next to the LDAP Servers field). Enter the IP address, port, and security settings (if applicable) for the iChain Authorization Server, as done in [Wizard Page 2](#).
- 12** Click OK. Repeat Step 11 for any additional servers used for authentication failover.
- 13** Select whether users will authenticate using their distinguished username, e-mail address, or another LDAP field by selecting the appropriate radio button.
- 14** Add the contexts where users can be found by selecting the Add button (located next to LDAP User Contexts for DN logins or LDAP Search Base for e-mail and LDAP field name logins).

You can either manually enter the container name in comma-delimited LDAP format or use the Browser button to select the container.
- 15** For e-mail and LDAP field logins, select whether an anonymous bind or an LDAP proxy user is used to search the tree. If applicable, enter the proxy username and password.
- 16** For LDAP field logins, enter the LDAP attribute name used to log in with.
- 17** Click OK > OK to save the new authentication profile.
- 18** Click the check box next to the desired profiles to use them for this accelerator.
- 19** If multiple authentication profiles are used together (such as an LDAP and a mutual SSL profile), use the Multiple Profile Rule to determine whether only one method (Or) or all the specified methods (And) are required to log in.
- 20** Add any advanced authentication options as described in [Appendix A, “iChain Management,” on page 217](#).
- 21** Click Next to continue to Wizard Page 6.

Wizard Page 6

Page 6 of the wizard is identical to the Protected Resource page for the iChain service object selected on [Wizard Page 1](#).

To create a protected resource:

- 1 Click Add.
- 2 Specify a name for the resource and the URL for the resource. This should be in the form `http://www.resource.com`, where `www.resource.com` is the DNS name specified when you created the Web Server Accelerator. The iChain snap-in will automatically attach the `http://` prefix, so all that needs to be entered is the DNS name.

NOTE: If the URL starts with `https://` (that is, this is a secure Web site), you will still need to specify `http://` in this field. iChain uses this field for matching purposes. It does not affect the URL in the query string.

- 3 Choose whether this protected resource will be Public, Restricted, or Secure. See [“Differentiating Among Protected Resources” on page 33](#) for more information.

This walkthrough assumes a secure resource, requiring both authentication (login) and authorization (ensuring the authenticated user has permission) to access the resource.

- 4 Select whether any associated OLAC parameters should be sent in the query string or as header variables by selecting the appropriate radio button.
- 5 Click OK to save the new protected resource.
- 6 If you want to specify Object-Level Access Control parameters, they can be entered using the OLAC button under the Modify button. See [“Setting Up Protected Resources” on page 31](#) and [“Setting Up Object-Level Access Control” on page 95](#) for more details.
- 7 Click Next to save the information and continue with the wizard.

Wizard Page 7

- 1 Use the Browse button (located next to the Access Control Rule field) to select the desired ACL rule.

You can also create the new rule by typing the context-free common name in the Access Control Rule field and selecting Enter.

2 Once an ACL rule is selected or created, the screen will display the parameters for the selected ACL rule. This is identical to the information in the Access Control properties tab for the ACL rule.

3 Enter the Allowed URLs as defined in [“Defining iChain Access Control Rules” on page 34](#).

A simple configuration to apply this rule to the entire Web site can be created by entering the resource name defined in [Wizard Page 6](#) and adding /* in the URL Postfix. The forward slash corresponds to the root directory of the Web server, and the asterisk is a wildcard meaning all files in this directory and all subdirectories.

4 If desired, add Excluded URLs as defined in [“ACL Exceptions” on page 37](#).

5 Click Add (located next to the Apply To List field) to add the users to which this rule will grant access. You can also select containers and groups as desired.

6 If desired, in the Exception List, select individual users to deny access to.

7 Click Next to continue to the Summary page.

At any time in the wizard, you can activate the changes you have made by clicking the Finish button. This will display a summary of the configuration that you can view before the changes will take effect. Once you complete the wizard, the changes you have made will be visible both in the Web-based administration utility, and in subsequent uses of the wizard. See [Appendix A, “iChain Management,” on page 217](#) for a detailed description of the options available in the iChain Wizard.

Installing the iChain Proxy Server on Your Network

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

This section contains the following topics:

- ◆ [“Managing the iChain Proxy Server” on page 53](#)
- ◆ [“Preparing the Network” on page 54](#)
- ◆ [“Troubleshooting iChain Proxy Server Problems” on page 57](#)
- ◆ [“Accelerating Web Servers” on page 59](#)
- ◆ [“Web Server Accelerator Setup” on page 61](#)

Managing the iChain Proxy Server

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

- ◆ Through the browser-based management tool.
- ◆ From the iChain Wizard included with the ConsoleOne snap-ins.
NOTE: The iChain Wizard in iChain 2.1 includes most of the commands necessary to administrate a standard iChain configuration. However, some advanced features are currently only available from the iChain Proxy Server console or the browser-based management tool.
- ◆ From the command line through a Telnet or null-modem connection.
(You can also use an attached keyboard and monitor if your proxy server has the required connections.)

The Web-Based Administration Utility

The web-based administration utility is unlike other management utilities because its interface appears in your browser as an HTML page originating from the proxy server. The only programs running on your workstation are a Java-compatible Web browser and the Java* components required by the HTML page.

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

For more information about using the web-based administration utility, see [“Proxy Server Browser-Based Administration Tool” on page 275](#).

The iChain Wizard

The iChain Wizard automates many of the tasks of setting up and configuring the iChain system. It uses an FTP interface to send a list of configuration commands to the iChain Proxy Server, which then automatically configures itself. [“Using the iChain Wizard to Create a Basic Configuration” on page 46](#) describes using the iChain Wizard as an alternative way to set up a basic configuration.

The Command Line

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

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

For more information about using the command line, see [“Command Line Reference” on page 355](#).

Preparing the Network

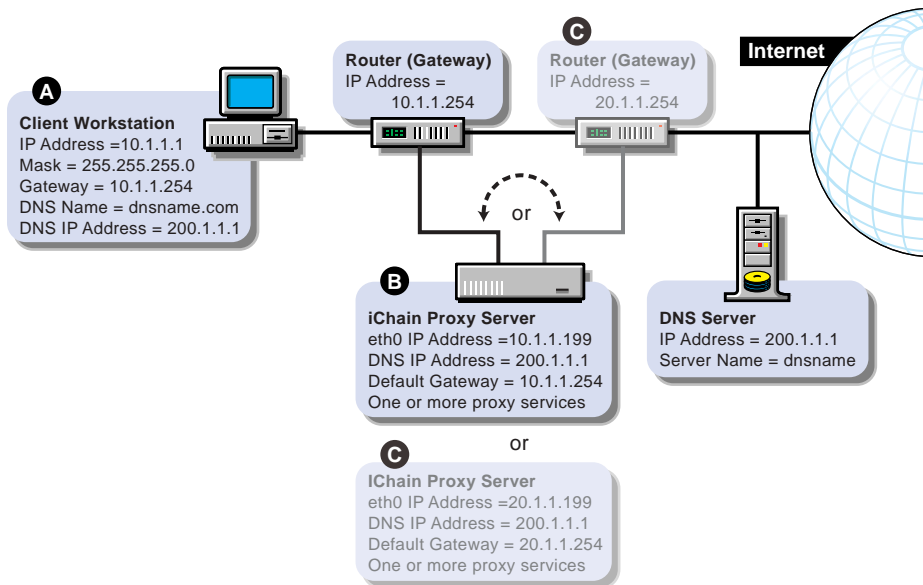
After you complete the initial proxy server setup and test, review this section to ensure that all your network components are properly configured.

Basic Network Configuration Setup

[Figure 6 on page 55](#) provides a visual map for the information in this section.

NOTE: The letters in [Figure 6 on page 55](#) are referenced in the tables that follow. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 6 Basic Network Configuration Setup



Configuring the Client Workstation

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

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

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

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

Configuration Requirements	Do This	Notes
A numeric IP address on the subnet	Refer to setup instructions for the system.	See A in Figure 6 on page 55 .
The subnet mask	The procedure is different for each platform. On a Windows* 95/98 or Windows NT* workstation, for example, right-click the Network Neighborhood icon on the desktop.	The IP address, subnet mask, and gateway address must all be on the same subnet.
The numeric IP address of the default gateway for the subnet		
The numeric IP address of the DNS server the browser will use to resolve DNS names		
The domain name for the DNS server the client will use (optional)		

Configuring the iChain Proxy Server

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

Configure the iChain Proxy Server with the following information:

To Configure	Do This	Notes
IP addresses and subnet masks for the network connections (eth0, eth1, etc.) that will handle proxy services	<ol style="list-style-type: none"> 1. In the browser-based tool, click Network > IP Addresses > Add Address. 2. Enter the addresses in the appropriate fields > click the Assign to Adapter drop-down list > select the appropriate adapter. 3. Click Apply. 	<p>See B and C in Figure 6 on page 55.</p> <p>The proxy server need not be on the same subnet as the browser. If not, its IP address will reflect a different subnet.</p> <p>Also, eth0, eth1, etc., can be on different subnets.</p>
At least one DNS server IP address	<ol style="list-style-type: none"> 1. In the browser-based tool, click Network > DNS. 2. Enter the addresses in the appropriate fields. 3. Click Apply. 	

To Configure	Do This	Notes
The numeric IP address for a gateway (router) on the same subnet as the proxy server	<ol style="list-style-type: none"> 1. In the browser-based tool, click Network > Gateway/Firewall. 2. Enter the address in the Default Gateway IP Address field. 3. Click Apply. 	<p>See <i>B</i> in Figure 6 on page 55.</p> <p>If the iChain Proxy Server is on the same subnet as the client workstation, the proxy server and the workstation will have the same gateway address.</p> <p>If the proxy server is on a different subnet than the browser, its gateway address will be the IP address of the router on the other subnet. See <i>C</i> in Figure 6 on page 55.</p>
Passwords for Config and View users	<ol style="list-style-type: none"> 1. In the browser-based tool, click System > Actions > Password. 2. Click the User drop-down list > select Config. 3. Enter the password information > click Change. 4. Repeat Step 2 and Step 3 for the View user. 5. Click Apply. 	<p>For information about Config and View users, see “Password Dialog Box” on page 287</p> <p>NOTE: Telnet is not secure unless a password is set.</p> <p>We strongly recommend you set system passwords as part of the initialization process. For more information, see “Password Dialog Box” on page 287.</p>
One or more proxy services	See the sections that follow.	

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

Troubleshooting iChain Proxy Server Problems

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

- ♦ [“My proxy server isn’t working” on page 58](#)
- ♦ [“I can’t ping the proxy server from my client” on page 58](#)
- ♦ [“All the numbers are correct and it still won’t ping” on page 58](#)

- ♦ “My browser can’t find the application” on page 58
- ♦ “Nothing ever comes up on my browser” on page 59
- ♦ “None of the changes I made in the browser application are taking effect” on page 59

My proxy server isn’t working

Most problems are caused by invalid IP address configurations. Four things are critical:

- ♦ A numeric IP address with a subnet mask
- ♦ A valid gateway address on the same subnet as the IP address
- ♦ A valid DNS server IP address
- ♦ A valid DNS domain name

I can’t ping the proxy server from my client

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

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

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

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

My browser can’t find the application

- ♦ The correct URL is <http://10.1.1.1:1959/appliance/config.html>.

- ◆ Make sure you specify `http://` in the URL window. Typing the address of the application without `http://` doesn't work.

Nothing ever comes up on my browser

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

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

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

Accelerating Web Servers

This section covers the following topics:

- ◆ [“Overview of Web Server Acceleration” on page 60](#)
- ◆ [“How Origin Web Server Acceleration Works” on page 60](#)
- ◆ [“Benefits of Origin Web Server Acceleration” on page 61](#)
- ◆ [“Working with DNS” on page 64](#)

Overview of Web Server Acceleration

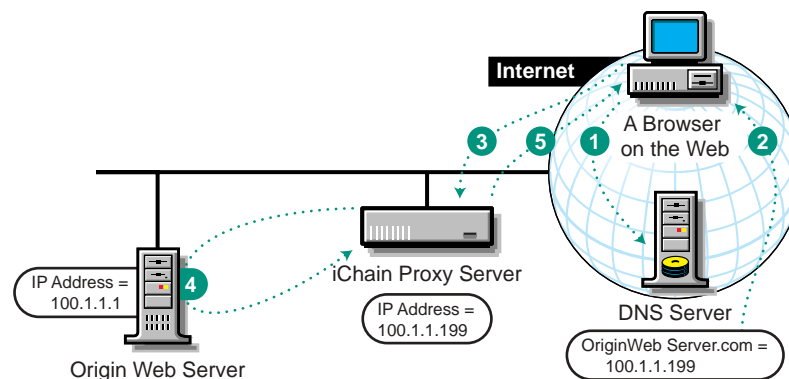
The proxy server's origin Web server accelerator relies on DNS causing the accelerator to receive requests originally targeted at the origin Web server. The Web server accelerator handles the requests, accessing the origin Web server only when needed objects are not cached.

How Origin Web Server Acceleration Works

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

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

Figure 7 Web Server Acceleration



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

Benefits of Origin Web Server Acceleration

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

For tips and guidelines on setting up origin Web server accelerators, see “[Web Server Accelerator Setup](#)” on page 61.

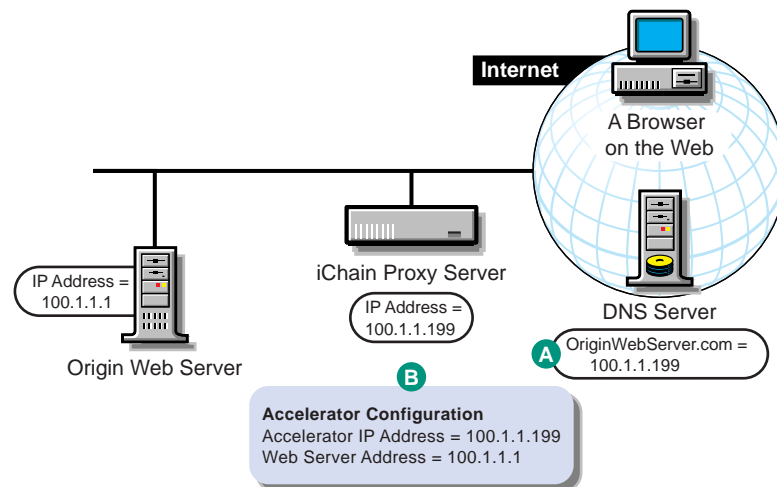
The procedure for configuring DNS to work with Web server accelerators is explained in “[Working with DNS](#)” on page 64.

Web Server Accelerator Setup

[Figure 8](#) provides a visual map for the information in this section.

NOTE: The letters in [Figure 8](#) are referenced in the table that follows. The addresses shown are for illustration purposes only. You will need to substitute actual addresses for your network.

Figure 8 Web Server Accelerator Setup



To	Do This	Notes
Ensure that your basic network configuration is complete for each proxy server	1. See “Configuring the iChain Proxy Server” on page 56.	
Ensure that DNS resolves browser requests to the proxy server IP addresses configured for the Web server accelerator services	1. See “Working with DNS” on page 64.	See A in Figure 8.

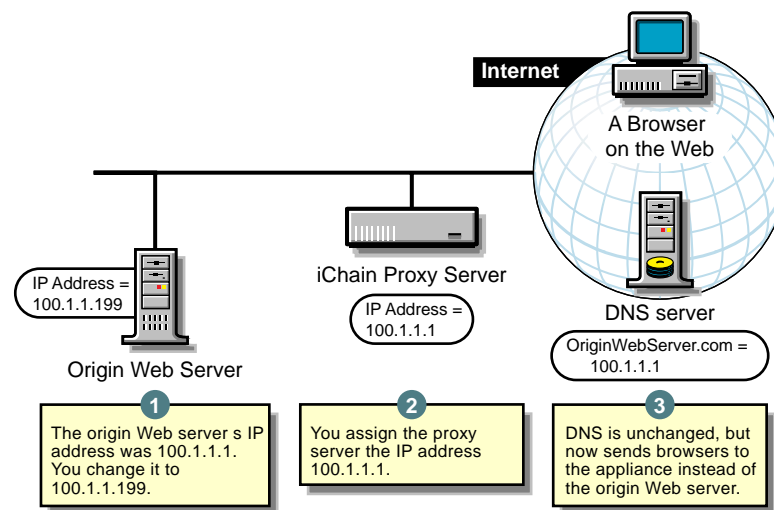
To	Do This	Notes
Set up one or more Web server accelerator services	<ol style="list-style-type: none"> 1. In the browser-based tool, click Configure > Web Server Accelerator > Insert. 2. Enter a name for the Web server accelerator for your tracking purposes. 3. Enter a DNS name. 4. In the Accelerator Proxy Port field, enter the port that the Web server accelerator will receive requests and vend data on. 5. In the Accelerator IP Addresses list, check one or more addresses that the Web server accelerator will receive requests and vend data on. (DNS resolves requests to these addresses.) 6. In the Web Server Port field, enter the port that the proxy server and origin Web server will communicate on. 7. In the Web Server Addresses list, insert one or more IP addresses (or DNS names) that the Web server accelerator will fill its cache from. (The proxy server must be able to fill all requests through any of these names or addresses.) 8. To activate the Web server accelerator, check Enable This Accelerator. 9. Click OK > Apply. 	<p>See <i>B</i> in Figure 8 on page 61.</p> <p>If server persistence is enabled in the Web Server Accelerator tab, the proxy server will use the same Web server to fill browser requests during a session. This setting affects all accelerators on the proxy server and saves e-business users from having to log in multiple times. See "Web Server Accelerator Tab" on page 313.</p> <p>If logging is enabled, accelerator log files for the Web server accelerator will have the same name as the Web server accelerator.</p> <p>The DNS name is required when:</p> <ul style="list-style-type: none"> ◆ You are accelerating multiple Web servers on the same IP address. (When multiple accelerator services use the same IP address.) ◆ You are accelerating a single Web site using path-based multi-homing. <p>If you enter DNS names in the Web Server Addresses list, make sure they are not the names that now resolve to proxy server numeric IP addresses. That would create an endless loop.</p>

Working with DNS

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

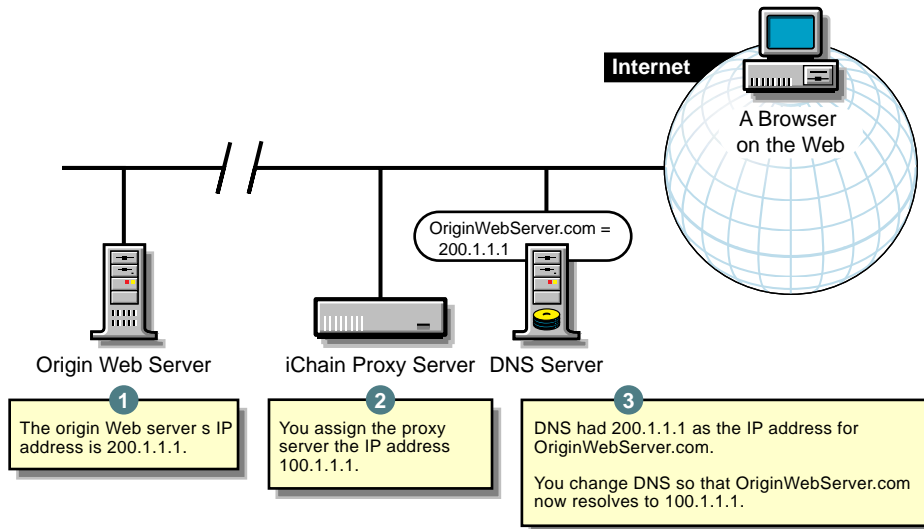
If the proxy server and the origin Web server are on the same subnet, you can swap IP addresses as shown in [Figure 9](#).

Figure 9 Working with DNS



If the origin Web server is on a remote network, you need to alter DNS as shown in [Figure 10 on page 65](#).

Figure 10 Altering DNS



4

Setting Up Authentication Services

This chapter describes the following Novell® iChain® Authentication Services:

- ◆ “Setting Up Mutual SSL” on page 67
- ◆ “Using Third-Party Certificates” on page 79
- ◆ “Using Multi Certificate Authorities” on page 81
- ◆ “Using Cross-Domain Authentication” on page 83
- ◆ “Using Token Authentication with iChain” on page 87
- ◆ “Using the Authentication Session Broker” on page 92

Setting Up Mutual SSL

To review the basics of SSL, see *Web Security and Commerce* by Simson Garfinkel and Gene Spafford, and/or the Netscapes Developer site at <http://developer.netscape.com/tech/security/ssl/howitworks.html>.

SSL provides:

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

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

Mutual SSL Configuration

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

iChain Server Certificate Setup — Certificate Signing Request

- 1** In the browser-based management tool, click Home > Certificate Maintenance > Create.
- 2** Enter an appropriate name for the certificate and subject name.
- 3** Click the Signature Algorithm drop-down list > select the algorithm you want to use (SHA-1 or MD-5).
- 4** Click the RSA Key Size drop-down list > select the RSA key size that you want to use.

You cannot select a key size larger than the maximum key size on the appliance.

- 5** Click Use External Certificate Authority.

If you are requesting a VeriSign* certificate, check the VeriSign CA check box. Otherwise, leave this box unchecked.

- 6** If desired, enter a name for your organization or division.

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

- 7** Enter the city or town where your organization does business.

- 8** Enter the unabbreviated name of the state or province where the organization does business.

This is commonly referred to as the state.

- 9** Enter the International Standards Organization country code for the country where the organization does business.

This is commonly referred to as the country and must be a valid, two-character country code.

- 10** Click OK.

Examine the Action and Status fields. The Action field should have red arrows on the left and the word Request displayed on a green background. The Status should be Building. The red arrows and green background indicate that you need to click Apply.

11 Click Apply.

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

If an error occurs:

11a Click Modify.

11b In the Modify Certificate dialog box, make the changes necessary to resolve the errors > click OK.

11c Click Apply. Repeat the modification process until the Status field displays the words CSR in Progress on a yellow background.

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

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

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

Internet Explorer and other browsers combine them with the CSR text that is in between. Clicking the browser refresh/reload button will often fix the problem. If it doesn't, simply insert the appropriate carriage returns during the next step. After you have copied the text, you can close the browser window. If you don't fix the defect, you can view the source of the HTML file and copy and paste from the source file.

3 Paste the CSR text from the clipboard to the e-mail message or HTML form as required by your CA. The method for sending the CSR will vary depending on the authority. VeriSign, for example, uses a Web page interface.

IMPORTANT: The header and trailer must be on lines separate from the body of the CSR. The header line will be similar to the following:

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

The trailer line will be similar to the following:

----- END NEW CERTIFICATE REQUEST -----

If required, you must use hard returns to separate these two lines from the body of the CSR.

Using Novell as the External Certificate Authority

- 1** In ConsoleOne click on the Issue Certificate on the Tools Menu and past the CSR and follow the default settings or change them if you prefer.
- 2** Click Finish and save the file in a .b64 format.
- 3** At the Certificate tab, export the Self Signed Certificate from the Organization CA object's properties page (found in the Security Container in .b64 format).

Storing the Certificate in the iChain Proxy Server

After the external CA responds with the certificate:

- 1** In the browser-based management tool, click Home > Certificate Maintenance > click the name of the certificate you want to store > click Store Certificate.
- 2** In the Store Certificates dialog box, paste the CA certificate into the CA Certificate Contents box. If you are using Novell CA, this is where the Self Signed Certificate should be placed.

NOTE: If you requested a VeriSign certificate and you checked the VeriSign box in [Step 5 on page 68](#), the CA Certificates box is grayed out. You do not need to paste the VeriSign CA certificate because VeriSign certificates are already stored on the appliance.

- 3** Paste your newly issued certificate in the Server Certificate Contents box.
- 4** Click Create.

Examine the Action and Status fields. The Action field should have red arrows on the left and the word Create displayed on a green background. The Status should be CSR in Process. The red arrows and green background indicate that you need to click Apply.

- 5** Click Apply.

If any errors occur during the certificate creation process, they will be displayed in the Error field on a red background.

If an error occurs:

- 5a** Click Store Certificate.
- 5b** In the Store Certificate dialog box, verify that the correct certificates are pasted in the boxes > click OK.
- 5c** Click Apply. Repeat the modification process until the Status field displays the words Active on a green background.

Create a New Authentication Profile via an iChain Browser GUI

- 1** Select Cache > Authentication > Insert > enter a name for the profile (for example, "SSL").
- 2** Select SSL Certificate mutual authentication > click OK > Apply.
- 3** Add the new profile to the Web Server Accelerator by selecting Cache > Web Server Accelerator.
- 4** Click Authentication Options > select the newly created profile.
- 5** Click Add > And Profiles > OK.
- 6** At the Accelerator tab, highlight the accelerator > click Modify.
- 7** At the SSL Key ID, select the name of the certificate you created > click OK > Apply.

You can also select Mutual Options if you need to configure certificate mapping.

iChain Server — SSL Between the Browser and iChain

NOTE: SSL can also be set up between the iChain appliance and the Web server.

- 1** Using the iChain browser GUI, select Cache > Web Server Accelerator.
- 2** Choose the appropriate Web accelerator > click Modify > Enable SSLizer.

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

Create a User Certificate From a Novell Certificate Authority

- 1** In ConsoleOne as an administrator go to the user object's properties on the Security tab > click Certificates > Create.
- 2** Select the default options (for example, with the private key, etc.). Change only what you need to change (for example, the expiration, etc.).

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

- 3** Save the file in .PFX format with a password.

iChain Server — SSL Between iChain and the Browser (I.E. 6.0)

- 1** At the browser, select Tools > Internet Options.
- 2** Select Content > Certificates.

- 3** Click the Import button to launch the wizard.
- 4** Import a Personal Certificate that has been signed by the same Certificate Authority.

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

Using Certificate Mapping

When using SSL Mutual Authentication, there must be a user in the iChain LDAP Authentication tree that corresponds with the user certificate. Certificate Mapping gives four different ways to map the user certificate to a user in the iChain LDAP Authentication tree. The four mapping types are directory name, mail, serial number & issuer name, and subject name. The proxy server can be configured to use any combination of the four mapping types. Note that when searching for a user with the configured mappings, the first user found will be the user that is used for authentication and access control, even if the other users will map to the same certificate.

Configuring a Certificate Mapping Search Base

At least one search base needs to be configured for Certificate Mapping. The search base is the location in the iChain LDAP Authentication tree to search for user objects that the certificate may map to. More than one search base can be configured. The search will look for matches starting at the search base. All containers below the search base are included in the search.

To add a search base, go to the iChain Proxy Server console and enter the following:

```
add authentication aclcheck ldap searchbase =  
<context>
```

where *<context>* is the LDAP context where you want the search to start. For example, *ou=users;o=novell*. Note that the syntax for the context uses a semi-colon (;) for the delimiter between containers for the full context.

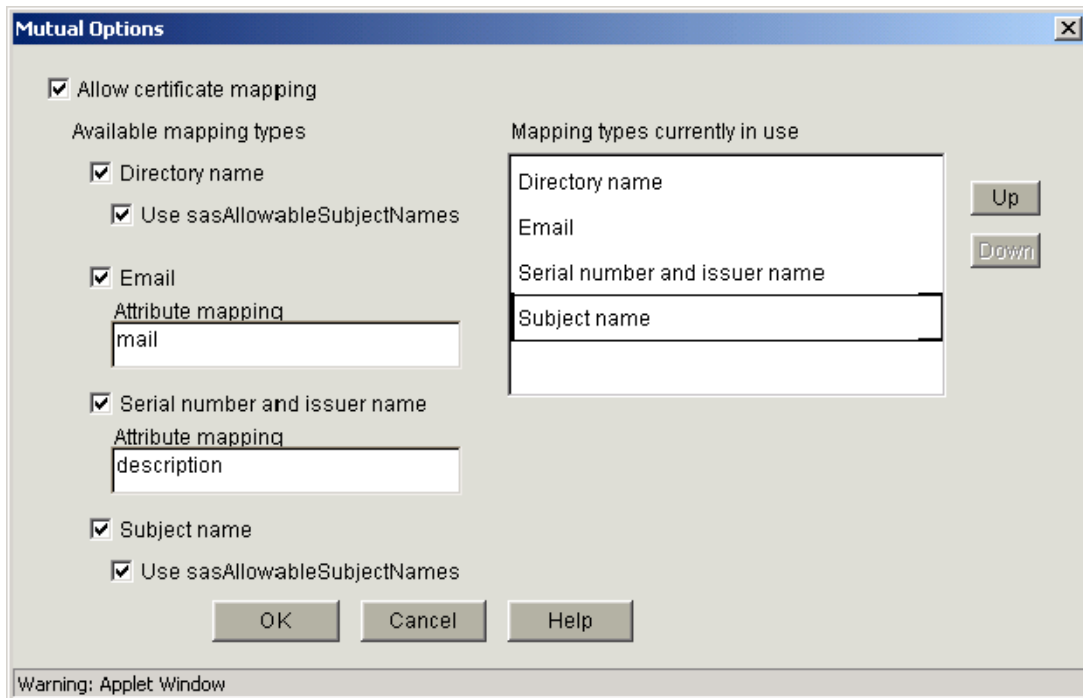
Configuring Certificate Mapping Types

The certificate mapping types are configured from the iChain Proxy Server utility.

To configure certificate mapping types:

- 1 At the iChain Proxy Server utility, choose Configure > Authentication.
- 2 Highlight an authentication profile of type Mutual.
- 3 Click Modify > Mutual Options (see [Figure 11](#)).

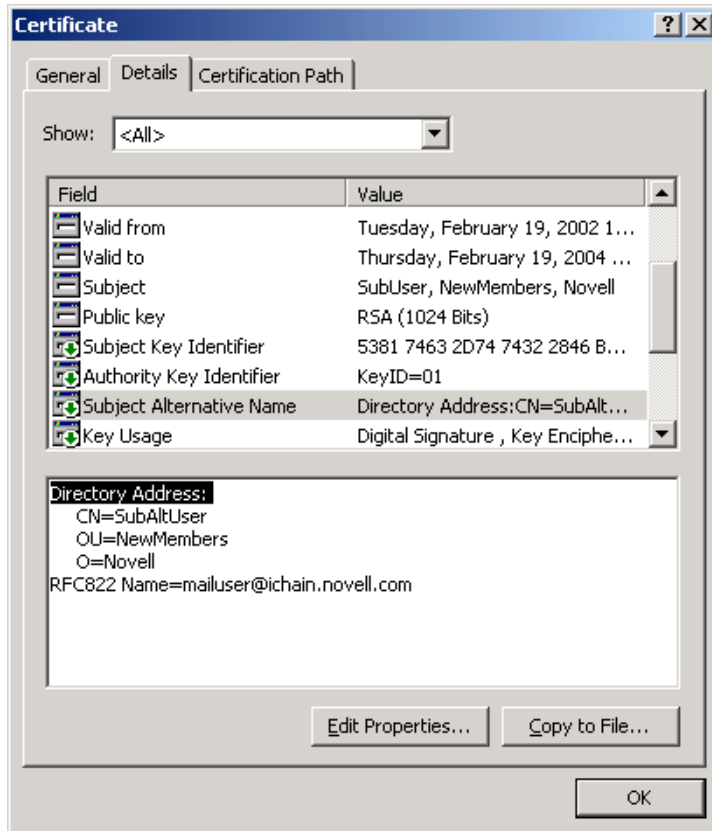
Figure 11 Mapping Types



Directory Name Mapping

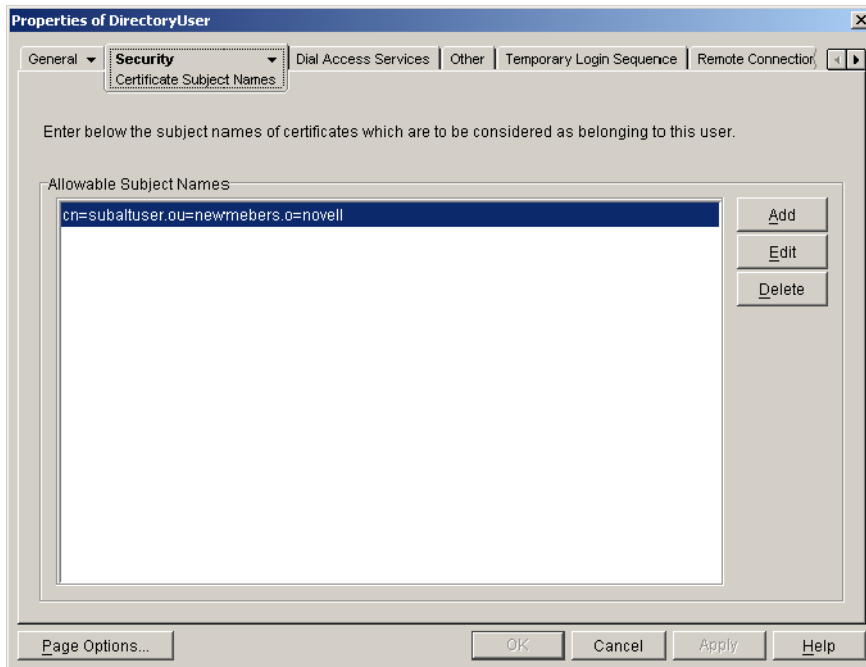
With directory name mapping, the Subject Alternative Name field in the user certificate, with a name type of Directory Name, will be used to identify the certificate portion of the user (see [Figure 12](#)). The name in the certificate can be from root to leaf or from leaf to root.

Figure 12 Subject Alternative Name



A user in the LDAP Authentication Tree matching the Directory Name in the Subject Alternative Name field of the certificate will be checked first. If a user is not found and Use sasAllowableSubjectName is also enabled for directory mapping (see Figure 11), the LDAP Authentication Tree will be searched for a user containing an sasAllowableSubjectName attribute matching the Directory Name in the Subject Alternative Name field of the certificate. If sasAllowableSubjectName is enabled, the LDAP Authentication tree should be configured so that there is no duplication of allowed names between users in the sasAllowableSubjectName attribute.

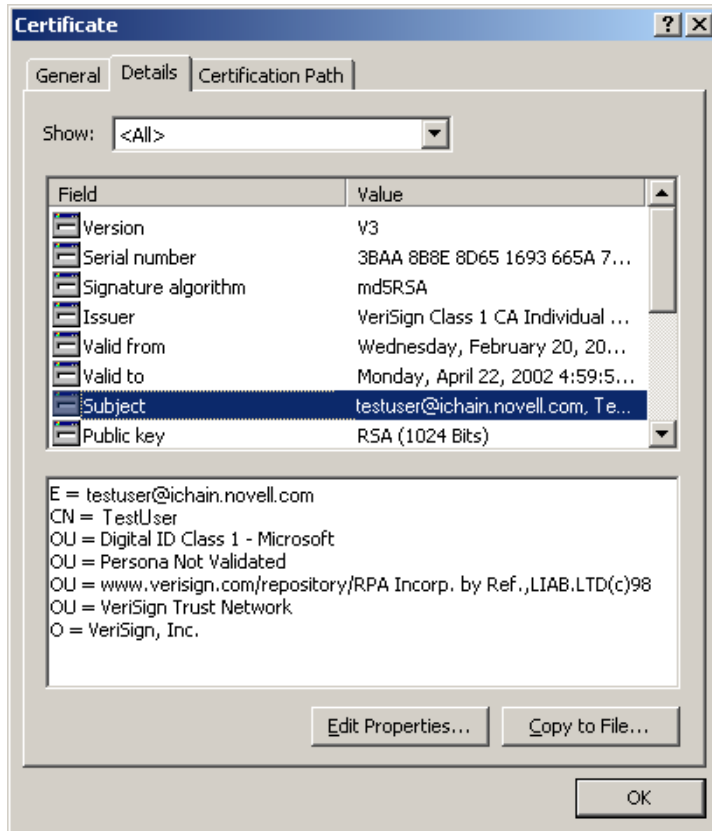
The sasAllowableSubjectName attribute is the same attribute currently used by NMAS for certificate mapping. The ConsoleOne snap-ins and schema updates are part of the NMAS installation on the Authorization Server CD. Figure 13 shows the sasAllowableSubjectName attribute in ConsoleOne.

Figure 13 sasAllowableSubjectName Attribute

Email Mapping

With Email mapping, there are two possible fields in the user certificate that can be used to identify the certificate portion of the user. The first is the Subject Alternative Name field in the user certificate, with a name type of RFC822 (see [Figure 12](#)). The second is when an e-mail name is embedded in the Subject field of the certificate (see [Figure 14](#)). If both the Subject field and the Subject Alternative Name field contain an e-mail address, the Subject Alternative Name will be the only field used.

Figure 14 Email Name Embedded in Certificate Subject Field

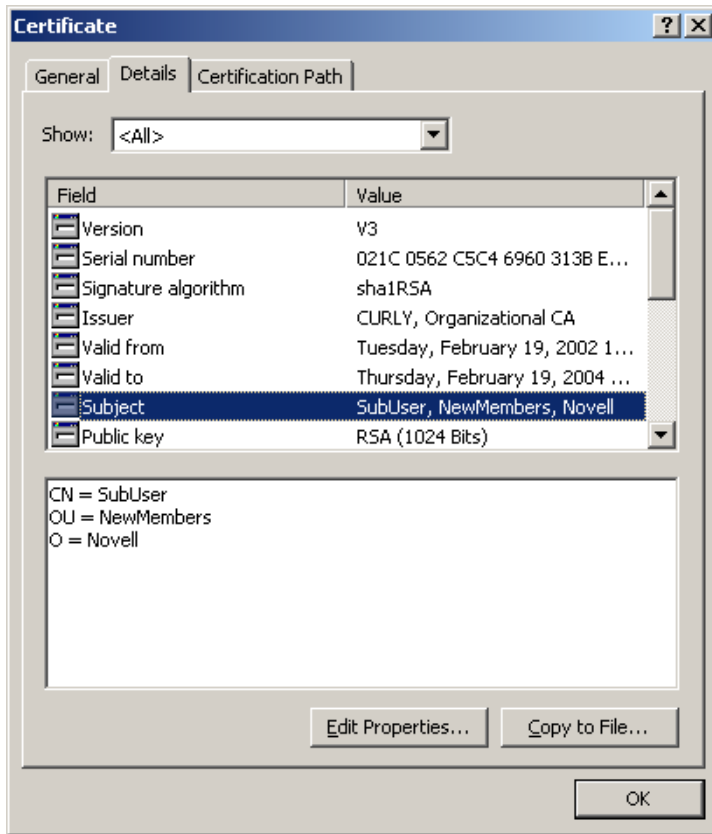


The LDAP attribute configured in the Email attribute mapping (see [Figure 11](#)) will be used to match the Email address from the certificate when searching for a user in the LDAP Authentication tree. The default LDAP attribute is mail, which is the attribute currently used by GroupWise and Novell Certificate Server. The LDAP Authentication tree should be configured so that there is no duplication of Email addresses between users in the configured e-mail attribute mapping.

Subject Name Mapping

With directory name mapping, the Subject field in the user certificate will be used to identify the certificate portion of the user (see [Figure 15](#)). The Subject name in the certificate can be from root to leaf or from leaf to root.

Figure 15 Subject Field in the User Certificate



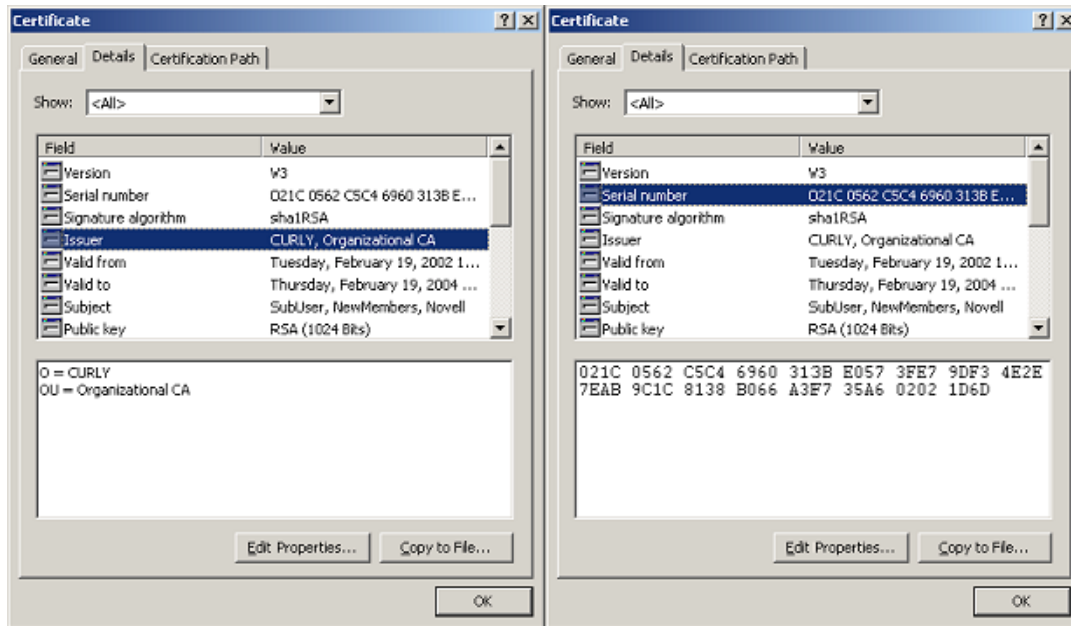
A user in the LDAP Authentication tree matching the Subject Name field of the certificate will be checked first. If a user is not found and the Use sasAllowableSubjectName is also enabled for directory name mapping (see [Figure 11](#)), the LDAP Authentication tree will be searched for a user containing a sasAllowableSubjectName attribute matching the Subject Name field of the certificate. If the sasAllowableSubjectName attribute is enabled, the LDAP Authentication tree should be configured so that there is no duplication of allowed names between users in the sasAllowableSubjectName attribute.

The sasAllowableSubjectName is the same attribute currently used by NMAS for certificate mapping. The ConsoleOne snap-ins and schema updates are part of the NMAS installation on the Authorization Server CD. [Figure 13](#) shows the sasAllowableSubjectName in ConsoleOne.

Serial Number & Issuer Name Mapping

With Serial Number & Issuer Name mapping, both the serial number and the issuer name fields from the certificate will be used together to identify the certificate portion of the user (see [Figure 16](#)).

Figure 16 Serial Number & Issue Name Mapping



Both the issuer name and the serial number need to be put into the same LDAP attribute of the user. The LDAP attribute that is used is specified in the Serial number and issuer name Attribute mapping field (see [Figure 11](#)) of the iChain Proxy Server utility. The LDAP attribute can be any Case Ignore List or Case Ignore String attribute of the user. If you are configuring your own attribute, make sure the attribute is added to the Person class.

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

When using a Case Ignore String attribute, both the issuer name and the serial number need to be in the same attribute separated by a dollar sign (\$) character. The issuer name needs to be in front of the \$ character, with the serial number following the \$ character. Do not use any spaces in front of or

behind the \$ character. (For example, O=CURLY.OU=Organization CA\$021C0562C5C4... could be used for the certificate displayed in [Figure 16](#)).

The issuer name can be from root to leaf or from leaf to root. The issuer name is dot-delimited without a preceding dot. (For example, O=CURLY.OU=Organization CA or OU=Organization CA.O=CURLY could be used for the certificate displayed in [Figure 16](#)).

NOTE: The certificate number is displayed in Internet Explorer with a space after every fourth digit. The certificate number needs to be entered without spaces. For example, the certificate number displayed in [Figure 16](#) is shown with spaces, but should be entered as:

021C0562C5C46960313BE0573FE79DF34E2E7EAB9C1C8138B066A3F735A602021D6D.

Using Third-Party Certificates

Novell iChain includes Novell Public Key Infrastructure Services (PKIS 2.0) to provide cryptography and enable certificate services in your iChain infrastructure. A Novell server certificate is installed and configured automatically when you install Novell iChain; however, you may want to use other third-party certificates, such as Baltimore* certificates, in your infrastructure. In order to use third-party certificates in your iChain infrastructure, you must request a certificate from a Certificate Authority (CA), have the CA sign the certificate, collect and export the certificate and its trusted root, and then import the certificate and its trusted root to the iChain Proxy Server. The following procedure describes the process for a Baltimore certificate.

To create a Certificate Signing Request (CSR) for a server certificate for the iChain Proxy Server:

- 1** From ConsoleOne, access the tree containing the iChain Proxy Server.
- 2** From within the server's OU, click New Object > NDSPKI:KeyMaterial.
- 3** Specify a name for the certificate, such as proxy server certificate.
- 4** Select Custom > click Next.
- 5** Select External Certificate Authority.
- 6** Select the key size > click Next.
- 7** Click Next to accept the default subject name of the certificate.
- 8** Click Finish.

- 9 Click Save to save the CSR.

To sign the CSR, perform the following steps:

- 1 Copy the CSR onto a diskette.
- 2 Insert the diskette with the CSR into the drive of the Baltimore Certificate Authority.
- 3 From the Registration Authority Operatory (RAO) menu of the Baltimore CA, click Face to Face Requests > Register a New User.
- 4 Select the Baltimore policy you have previously created for the proxy server certificate.
- 5 Locate the CSR file on the diskette > select the file > click Open.
- 6 Click Accept to process the CSR.

To collect response to the CSR and export the trusted root, perform the following steps:

- 1 Click Collect Reply from Last Request.
- 2 Click File.
- 3 Click DER Encoded Certificate.
- 4 Save the response file to the diskette as a .DER file.
- 5 Click OK to acknowledge.
- 6 Click OK on the yellow back arrow.
- 7 From the Certificate Authority Operator (CAO) menu, click Open/Create PKI.
- 8 Right-click the CA object > click Export Certificate.
- 9 Click DER Encoded Certificate.
- 10 Save the Trusted Root file to the diskette as a .DER file.
- 11 Select PKI > Done.

To import the new Trusted Root into the server certificate, perform the following steps:

- 1 From ConsoleOne, right-click on the server certificate object you created when you made the certificate request and click Properties.
- 2 Click Certificates > Trusted Root > Import > Read From File.

- 3** Select the Trusted Root file on the diskette > click Open > Next.
- 4** Click Read From File.
- 5** Select the certificate response file from the diskette > click Open.
- 6** Click Finish.

Using Multi Certificate Authorities

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

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

Configuring Multi CAs

To configure Multi CAs, you need to place the alternate CA certificates into your LDAP tree, configure the iChain proxy to use a specified trusted root container, and create new iChain server certificates, as described below in [“Task One: Place Alternate CA Certificates Into Your LDAP Tree” on page 81](#), [“Task Two: Configure the iChain Proxy Server to Use a Specified Trusted Root” on page 82](#), and [“Task Three: Create New iChain Server Certificates” on page 83](#).

Task One: Place Alternate CA Certificates Into Your LDAP Tree

To place the alternate CA certificates into your LDAP tree:

- 1** From ConsoleOne, select the Security object located at the root of your LDAP tree.
- 2** Select File > New > New Object
or
Click the New Object icon.
- 3** Select NDSPKI:Trusted Root > click OK.

- 4** Define a name for the trusted root container (for example, iChain Roots) > click OK.
- 5** Select the object you just created (for example, the iChain Roots object).
- 6** Select File > New > New Object
or
Click the New Object icon.
- 7** Select NDSPKI:Trusted Root Object > click OK.
- 8** Define a name for the trusted root object (for example, Baltimore CA) > click OK.
- 9** Select the Read from File button > browse your system for the trusted root certificate > import it into the dialog box
or
Paste your trusted root certificate into the dialog box. To use this option, you must first open the trusted root certificate in a text editor or some other program and copy the contents to the clipboard. Then right-click inside the box and select Paste. Your certificate will be inserted into the dialog box.
- 10** Click Finish.
If you want to add more trusted root certificates, repeat **Step 5** through **Step 10** for each certificate.

Task Two: Configure the iChain Proxy Server to Use a Specified Trusted Root

To configure the iChain proxy to use a specified trusted root container:

- 1** From ConsoleOne, click the Other tab on the iChain Security object (ISO) you previously created for this configuration.
- 2** Click Add.
- 3** Click iChainTrustedRootContainer > click OK.
- 4** Using the Browse button, browse to the trusted root container previously created (see “**Task One: Place Alternate CA Certificates Into Your LDAP Tree**” on page 81) > click OK.
or
Enter the complete name of the previously created trusted root container (for example, iChain Roots.Security).

5 Click OK.

Task Three: Create New iChain Server Certificates

Only iChain server certificates created or restored after “[Task One: Place Alternate CA Certificates Into Your LDAP Tree](#)” on page 81 and “[Task Two: Configure the iChain Proxy Server to Use a Specified Trusted Root](#)” on page 82 will have the Multi CA feature enabled. Therefore, you need to perform one of the following tasks to enable Multi CA support:

- ♦ Create a new iChain server certificate
- or
- ♦ Back up and restore existing iChain server certificates

Using Cross-Domain Authentication

In iChain 2.1, Cross-Domain Authentication (CDA) provides a graded authentication feature that allows an iChain server administrator to build a trust relationship (CDA) among different domain types (for example, www.c.com, www.l.com, and www.lc.com). Inside this CDA, same types of authentication method will only login once.

CDA Scenario and Examples

The following scenario and examples can be used to clarify how CDA works:

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

www.l.com — LDAP authentication

www.c.com — Certificate (Mutual) authentication

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

Single Sign-on Example (same grade with same authentication methods):

If authentication methods of all accelerators are the same (for example, all LDAP/certificate, or all certificate and LDAP), once a user logs in to a domain, he or she can access any other domains without having to log in to them.

Graded Authentication Example 1 (from grade high to low): If a user accesses www.lc.com first, he or she will be asked to log in twice; once for

certificate and again for LDAP. After the user accesses www.lc.com, he or she can access www.l.com (or www.c.com) without any login.

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

Graded Authentication Example 3 (same grade but with different authentication methods): If a user accesses www.l.com first and then later accesses www.c.com, he or she will be asked to log in twice; once for www.l.com with LDAP and again for www.c.com with certificate (then the user can access www.lc.com without logging in).

These four examples show that with CDA-enabled accelerators, the same type of authentication will require logging in only one time. By doing so, CDA provides the user-friendly features of Single Sign-on.

Selecting Accelerators as Members of CDA and the Cross-Domain Broker

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

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

1. Security (trust and single cookie)
2. Graded authentication (with Single Sign-on)
3. Performance

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

one of the accelerators, he or she will be logged out from all CDA-enabled accelerators.

CDA provides a Single Sign-on feature by allowing accelerators with the same type of authentication to require log in only once.

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

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

The selection of a DB is critical in CDA. Selection of a member of CDA as the DB should be based on the following criteria:

1. You must never disable a DB-enabled accelerator. (You must never disable its authentication, either.)
2. Because there is no extra redirection when accessing a DB-enabled accelerator, we recommend that you select the most frequently accessed accelerator to be the DB.

NOTE: Criterion 1 is required. If Criterion 2 is difficult to meet, you can select any CDA member to be the DB.

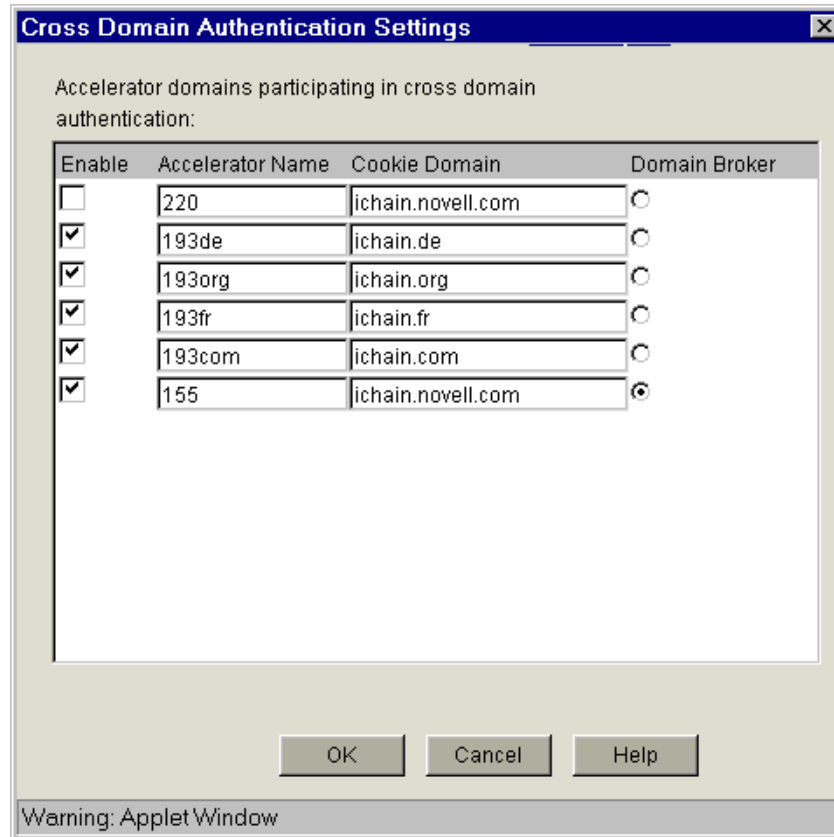
Configuring CDA

You can configure the CDA by following these steps:

- 1** At the iChain Proxy Server GUI, click Configure > Web Server Accelerator > Domain.

The Cross Domain Authentication Settings screen will appear (see [Figure 17](#)).

Figure 17 Cross Domain Authentication Settings



2 At the Cross Domain Authentication Settings screen, select the boxes of accelerators you want to be CDA members.

3 Select Domain Broker from the enabled CDA members by checking the radio button among these members.

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

CDA and Session Broker

When using Session Broker, all groups of accelerators participating in CDA should be configured identically. Session Broker relied on the fact that groups of accelerators are a single entity. Therefore, if an accelerator is participating

in CDA on one machine, then it should be participating in an identical CDA on all other machines with that accelerator. For security reasons, when using CDA with Session Broker, the browser must be closed in order to enforce logouts.

Using Token Authentication with iChain

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

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

1. “Installing NMAS, Novell RADIUS, and a Token Method” on page 87.
2. “Configuring Novell RADIUS Components” on page 88.
3. “Setting Up the iChain RADIUS Client” on page 91.

Installing NMAS, Novell RADIUS, and a Token Method

You must install NMAS into your eDirectory tree. NMAS is included on the *iChain Authorization Server* CD under the \NMASSERVER directory. Change to this directory and run INSTALL.EXE.

As part of the NMAS installation, you can select and install the Novell RADIUS server components.

IMPORTANT: The NMAS installation screen lists NICI 1.5.7 or later on the NMAS server. You will need to install NICI 2.02 on the NMAS server. This version of NICI is included on the *iChain Authorization Server* CD under the \NICI directory.

NOTE: You should also install the NICI on the workstation where you are running ConsoleOne. The location of the NICI files for the workstation is located on the *iChain Authorization Server* CD at /NICI/WINCLIENT. You should run the WCNICIU0 application found in this folder to install the NICI snap-ins on the workstation.

If you run ConsoleOne from a different location than the NMAS server, you will want to install the NMAS ConsoleOne snap-ins to that location. To do this, change to the \NMASSERVER directory on the *iChain Authorization Server* CD and run SNAPININSTALL.EXE. This will allow you to install the NMAS ConsoleOne snap-ins to any location you choose.

IMPORTANT: Your NMAS server must reside in the same eDirectory tree as your iChain LDAP server that holds the Access Control List (ACL). This allows the ACL to recognize users who are authenticating using NMAS and to allow the users access to the information they need.

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

For more information on installing NMAS, see the NMAS 2.0 Installation Quick Start at the Novell Documentation Web site (<http://www.novell.com/documentation/lg/nmas20/index.html>).

For more information on installing and configuring Novell RADIUS, see the Novell RADIUS Administration Guide at the Novell Documentation Web site (<http://www.novell.com/documentation/lg/nmas20/index.html>).

For general information on installing and setting up a login method, see the NMAS 2.0 Administration Guide at the Novell Documentation Web site (<http://www.novell.com/documentation/lg/nmas20/index.html>).

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

Configuring Novell RADIUS Components

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

Perform the following procedures in order:

“Creating a Dial Access System (DAS) Object” on page 89

“Configuring the Login Policy Rules” on page 89

“Adding the iChain Proxy Server As a Client of the Dial Access System (DAS) Object” on page 89

“Creating a RADIUS Dial Access Profile (DAP) Object” on page 90

“Adding an Attribute to the RADIUS Dial Access Profile (DAP) Object” on page 90

“Assigning the Token Method to Each User Object” on page 90

“Assigning the DAS Object to Each User Object” on page 91

“Starting Novell RADIUS Services on Your NMAS Server” on page 91

Creating a Dial Access System (DAS) Object

- 1** Start ConsoleOne.
- 2** Right-click an Organizational Unit container object > click New > Object > RADIUS:Dial Access System.
- 3** Specify the object name.
- 4** Click OK.
- 5** Specify the password.
- 6** Click OK.

Configuring the Login Policy Rules

- 1** Start ConsoleOne.
- 2** From the Security Container, double-click the Login Policy object.
- 3** Click the Rules tab.
- 4** Click + to add a login rule.
- 5** On the User list tab, click + and select the user or container that you want the rule to apply to.
- 6** On the Sequences tab, click + > select the token method > select Mandatory.
- 7** Click OK > OK > OK.

Adding the iChain Proxy Server As a Client of the Dial Access System (DAS) Object

- 1** Start ConsoleOne.
- 2** Double-click the DAS object.
- 3** On the Clients tab, click Add.
- 4** For Address, type the IP address of your iChain proxy server.
- 5** For Vendor Type, use the drop-down list > select Novell.
- 6** Type and confirm a secret for this client.
- 7** Click OK.
- 8** On the User Resolution tab, click the Use Lookup Contexts List to Resolve User Name radio button if the users are not in the same context as the DAS object.

- 9** Click Add.
- 10** Browse and highlight the container where the User objects reside.
- 11** In the Object Name field, type a name for the object.
- 12** Click OK > OK.

Creating a RADIUS Dial Access Profile (DAP) Object

- 1** Start ConsoleOne.
- 2** Right-click an Organizational Unit container object and click New > Object > RADIUS:Profile.
- 3** Click OK.
- 4** Specify the object name.
- 5** Click OK.

Adding an Attribute to the RADIUS Dial Access Profile (DAP) Object

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

IMPORTANT: It is critical that you select FDN so that name resolution will work properly. Otherwise, the users who use this profile will get 403 Forbidden errors when they try to access web pages.

- 7** Click OK > OK.

Assigning the Token Method to Each User Object

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

Assigning the DAS Object to Each User Object

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

Starting Novell RADIUS Services on Your NMAS Server

From the NMAS server console, type **RADIUS**. This will start the RADIUS services.

Setting Up the iChain RADIUS Client

Adding a RADIUS Authentication Profile

- 1 In the iChain Proxy Server Administration tool, click Configure > Authentication > Insert.
- 2 Enter a name for the Radius profile.
- 3 Click RADIUS authentication > RADIUS Options.
- 4 Enter the RADIUS server's IP address.
- 5 Enter 1645 for the Novell NMAS RADIUS server's port number.
- 6 Enter the shared secret set up in [“Adding the iChain Proxy Server As a Client of the Dial Access System \(DAS\) Object”](#) on page 89.
- 7 Click OK > OK > Apply.

Adding RADIUS Authentication to an Accelerator

- 1 In the iChain Proxy Server Administration tool, click Configure > Web Server Accelerator.
- 2 Select the accelerator you want to add the RADIUS authentication profile to.

- 3** Click Modify > Enable Authentication > Authentication Options.
- 4** Highlight the Radius profile created in [“Adding a RADIUS Authentication Profile” on page 91](#).
- 5** Click Add > OK > OK > Apply.

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

Using the Authentication Session Broker

The Session Broker feature is useful when you have more than one iChain server running at your site. Session Broker allows "sessions" (user authentication data) to be shared between multiple iChain servers, which, in turn, allows a user to authenticate only once when browsing across all of them.

For example, if a user browses to a page on your site protected by iChain Server A, iChain will ask the user to authenticate before granting access to the page. Suppose that during the course of browsing, the user is directed to a page protected by iChain Server B. Without Session Broker, the user would be required to authenticate again because iChain Server B has no way of knowing that the user was authenticated on Server A.

When Session Broker is running, iChain servers relay all their authenticated users. When a user who isn't authenticated attempts to access a protected page, iChain will ask the Session Broker if this user is authenticated to a different iChain server. If so, the user is granted access without the need to authenticate again.

NOTE: Only Authentication Profiles with the same name on each iChain server are shared. If the second iChain server doesn't have an Authentication Profile with the same name as the Authentication Profile the user authenticated to on the first server, the user will be required to authenticate again.

Session Broker Configuration

Use the iChain Web Accelerator Wizard in ConsoleOne to configure Session Broker on an iChain server:

- 1** In ConsoleOne, select Wizards > iChain Web Server Accelerator.
- 2** Select an iChain Service object (or create one).
- 3** Enter the two requested IP addresses: the Primary Session Broker IP Address and the Secondary Session Broker IP Address.

The Secondary Session Broker IP Address is optional and only becomes active if the Primary Session Broker is down; otherwise it remains idle.

- 4** Enter the IP address of the iChain servers you have designated to run the Session Broker.

For performance reasons, you may want to put the primary Session Broker on an iChain server that has no other responsibilities. The secondary Session Broker can be configured on an iChain server with other duties since it will only be used for short periods of time.

- 5** Do the following steps on the first machine:

- 5a** Establish a shared secret between your iChain servers and the Session Brokers that can be used to encrypt data passed between them. To do this, enter the following command at any iChain console:

```
createsessionbrokerkey
```

This command will create a floppy with the encryption key on it and this installs the key on the server that generates it.

NOTE: It is possible to disable encryption of data passed between iChain and Session Broker. Do this only if you are certain that the messages passed between them are secure. To disable encryption of data, instead of the command `createsessionbrokerkey`, enter *the command* `createnullsessionbrokerkey`. This creates a null key, telling iChain and Session Broker that no encryption is desired.

- 5b** When prompted, insert a floppy disk into the floppy drive and enter a password to encrypt the shared secret with. The password you type must be at least 6 characters in length.
 - 5c** When prompted, confirm the password.
- 6** Do the following step on all other machines (but not the first one):
 - 6a** Insert the floppy disk with the encryption key on it into the floppy drives of each of your iChain servers, including those designated to run Session Broker. At the console of each server, enter the following command:

```
installsessionbrokerkey
```

- 7** When prompted for the password, enter the password you gave when you created the encryption key (see [Step 5b](#)).

- 8** After creating or installing the encryption key, you must restart your proxy server in order for the server to read in the key and begin encrypting the session broker data.
- 9** At the iChain console of the server(s) you have designated to run Session Broker (including the Session Broker itself), enter the following command:

```
set authentication sessionbrokerenable = yes
```

Session Broker should now be running.

5

Setting Up Web Single Sign-on Services

This chapter describes the following Novell® iChain® Web Single Sign-on services:

- ♦ [“Setting Up Object-Level Access Control” on page 95](#)
- ♦ [“Setting Up Form Fill” on page 100](#)

Setting Up Object-Level Access Control

The iChain service enables you to integrate and allow access to Web-based applications. Sometimes these resources or objects need additional access control or application information about the user to be passed into the application. This additional information about the user can be stored in Novell eDirectory® or some other database. Within iChain, these resources are called Protected Resources and access to them is set up through the Protected Resources tab of the iChain Service object. Refer to [“Creating an iChain Service Object” on page 27](#) for basic setup information.

To implement this feature, a special iChain Object-Level Access Control (OLAC) plug-in (an LDAP plug-in) is available to access the database and retrieve the additional information. By default this plug-in allows you to define attributes in LDAP datastore that are embedded and passed within the HTTP request header or as a query string. You can assign a name as the tag to the data.

iChain also supports an additional plug-in called CONSTANT. This plug-in allows you to pass the same constant literal with every OLAC request. This is particularly valuable when an application requires a constant to be passed and the administrator does not want to include the constant in each user object (for easier setup and maintenance).

The following table lists the LDAP and CONSTANT plug-ins' corresponding entries for the Data Source and Value fields in ConsoleOne.

Plug-in	Description	Data Source	Value
LDAP	Adds user attributes from a directory with LDAP support.	ldap (case-insensitive)	Any LDAP user attribute (for example, surname, givenName).
CONSTANT	Adds the constant literal for every OLAC request, where defined	constant (case-insensitive)	Constant Literal (for example, "string123")

Because the LDAP plug-in is based on iChain APIs, you may customize iChain and create OLAC plug-ins to integrate your applications as needed. The APIs for customizing your iChain infrastructure are available in the iChain Novell Developer Kit (NDK). Novell NDKs are available for download at [Novell Developer site \(http://developer.novell.com/ndk/\)](http://developer.novell.com/ndk/).

NOTE: Only administrators familiar with programming principles and Java* programming syntax should attempt to customize OLAC plug-ins.

The settings for the OLAC Frameworks and its plug-ins are partly stored in the iChain Access Control profile and partly from the OAC.PROPERTIES file, which is typically found in the SYS:/ICHAIN/OAC directory on the iChain Proxy Server. The configuration file contains a section for the framework as well as one for the plug-in. The following table lists the valid OLAC options for each section:

Name	Description	Required?	Default Value
Object-Level Access ControlOptions [OAC] section			
Security Authentication	The method to use when authenticating to the LDAP server. Currently, only "simple" is supported.	No	simple

Name	Description	Required?	Default Value
Server Port	The port on which the OLAC framework will listen for lookup requests from the proxy server.	No	4444
Worker Count	The number of worker threads to create.	No	10
Refresh Time	The number of minutes after which the OLAC configuration will be re-read from the ISO.	No	80
Value Delimiter	The delimiter used to separate multiple values assigned to the same name in the URL query string. For example, if Value Delimiter is specified as ";", the resulting query string might look like COLORS=blue;green;yellow;orange&SHAPE S=circle;square;triangle.	No	, (comma character)
LDAP Plug-In Options [LDAP Processor] section			
Security Authentication	The method to use when authenticating to the LDAP server. Currently, only "simple" is supported.	No	simple
Class Name	The name of the class implementing the LDAP plug-in. Must be com.novell.ichain.oac.Ildap.ParamListBuilder.	Yes	None

OLAC Custom Header Variables

You can pass the OLAC parameters defined for a protected resource as a part of the HTTP header itself. Thus, if needed, you can forward more of the OLAC parameters/values to the Web or application servers than would be possible by sending them as a part of the query string.

You can specify if the OLAC parameters needed to be passed either as a part of the query string or header variables while defining the protected resource in ConsoleOne (as shown by radio buttons at the Add New Protected Resource dialog).

OLAC parameters in the header variables must be taken care of while parsing and are specified in the following format:

```
x-<olac-param name>: <olac-param value>
```

NOTE: The letter x- is prefixed for all of the OLAC parameters (custom variables) in the header. This is done to negate or minimize name collisions between OLAC and non-OLAC parameters residing in the header.

When sending an attribute value from eDirectory through OLAC using LDAP, the total length of the query string should not exceed 255 characters for some Web browsers/servers or the authentication header will fail to appear (or it will appear with the username and password instead of the attribute that was mapped by the ICHAIN_UID parameter). The HTTP headers may also fail to appear.- trying to send an attribute value from eDirectory through OLAC via LDAP.

Customizing the Authentication Header

Using iChain 2.1's OLAC feature, you can customize the authentication header as described below.

By default, iChain puts the fully qualified distinguished name and user-entered password in the authentication header. However, administrators might want to change the content of this header by changing the value of the username or password. This customization may be required because some Web servers (Microsoft IIS*) require the common name (CN) instead of DN as the username. Using OLAC, you can customize the authentication header.

For a particular protected resource, you can define special OLAC parameters, such as ICHAIN_UID and ICHAIN_PWD, to change the values of the authentication header. The values returned by OLAC will be placed in the authentication header as username and password, respectively.

For example, if you define the ICHAIN_UID=CN and ICHAIN_PWD=SSN OLAC parameters for a protected resource, OLAC will return the values of CN and SSN attributes of the logged-in user. iChain will use these values as the username and password to construct the authentication header and will send it to the Web server.

Both of these parameters are optional. If only one parameter is defined, such as ICHAIN_UID=CN, the other parameter value will be filled with default behavior, such as a password, provided by the user.

IMPORTANT: If you have defined these special parameters and OLAC is not enabled or the value of the given attribute is NULL, iChain will pass NULL in the authentication header.

OLAC Caching

OLAC has two levels of caching: The first level of OLAC caching occurs in OLAC server (OACJAVA), which caches the protected resource name and its associated OLAC parameters, data source, and value (the name of the attribute in case of LDAP) initially when the OLAC server is started and/or when it is issued an OACREFRESH command.

Administrators can refresh OLAC from the Web GUI (at the Access Control tab) or they can run OACREFRESH from the NetWare prompt (not the iChain prompt) on the iChain console. OACREFRESH is a java program that opens a socket to the OLAC server and sends the REFRESH command. After receiving the command, the OLAC server refreshes the cache. OLAC is also refreshed when the administrator opts to refresh the iChain Proxy configuration in the ISO snap-in in ConsoleOne after making changes. OLAC also auto-refreshes every "n" minutes, where "n" is specified in the OAC.PROPERTIES file. This value currently defaults to 180 minutes.

The second level of OLAC caching is for "OLAC values" cached in proxy for that user session and cannot be refreshed by any command. The life of this cache is only for that authenticated user session. OLAC refresh gets the changes in the protected resources and associated parameter definitions. It does not refresh any LDAP connections. The administrator must restart OLAC for those changes to take effect. Restarting OLAC can be done easily from the Access Control tab in the Web GUI.

Setting Up Form Fill

This section discusses the format and functionality of the iChain Form Fill Policy file, which controls which forms are identified and filled. The form fill file contains any number of URL policies encoded in XML. The iChain proxy user must have capability to add and modify attributes on all user objects in order for form fill to work.

The following is an example of what a simple policy looks like:

Figure 18 Simple Policy

```
<urlPolicy>
  <name>test</name>
  <url>www.test.com/login.htm</url>
  <actions>
    <fill>
      <input name="userid"      value="~dn">
      <input name="password"   value="~password">
    </fill>
    <post/>
  </actions>
</urlPolicy>
```

The bounds of a policy are delimited by the start tag `<urlPolicy>` and the end tag `</urlPolicy>`. This sample has a `<name>` of Test, which can be anything the user chooses, as long as it is unique within the policy file. The `<name>` is optional, except in a few cases, which will be discussed later in this section. The policy's `<url>` value of `www.test.com/login.htm` matches an exact page so that whenever this URL is requested by a browser, it will match this policy.

When a policy matches on a given page, certain actions are performed. These actions are specified by the `<actions>` tag. The example above designates two actions: `<fill>` and `<post>`.

The `<fill>` action identifies the page as a form that needs to be filled, and the `<input>` tags that follow describe the fields in the form that have values to insert. The first `<input>` tag has a name of `userid` and a value of `~dn`, which instructs the policy to find the `<input>` tag in the form that has the name of `userid` and insert the user's dn (LDAP distinguished name) as the value of the field. The `"~"` in `~dn` is a flag to the parser that this input field has a value that needs to be filled in.

Note that the name that follows the `~` is the name of the attribute in LDAP, not in eDirectory.

The second `<input>` tag is a special case because password is not an LDAP attribute. It instructs iChain to use the same password that the user gave when he or she logged in to iChain.

Since it is possible to have more than one `<input>` field in a form by the same name (or `<input>` fields without a name, the `<input>` tags given in the `<fill>` action are order-dependent. They must be specified in the order they will be seen in the form. Only those `<input>` fields that need to be filled in need to be specified.

The second action, `<post/>`, instructs iChain to process the form and to send it back to the browser. the browser then submits/posts the form automatically without the user's confirmation. This logs the user in automatically and gives the illusion of single sign-on. Care should be taken when using this action so that all login failures are handled properly. Handling failures is explained later in this section.

The following is an example of what a more complex policy looks like:

Figure 19 Complex Policy

```
<urlPolicy>
  <name>test2</name>
  <url>www.test.com/abc/*</url>
  <formCriteria>
    <title>Login</title>
  </formCriteria>
  <actions>
    <fill>
      <input name="userid"      value="~">
      <input name="password"    value="~">
      <input name="useJava"     value="~" type="checkbox">
    </fill>
    <post/>
  </actions>
</urlPolicy>
```

In the above example, the `<url>` value, `www.test.com/abc/*`, does not match a single URL, but rather all URLs below `abc` in the tree. There are some Web applications that don't have a specific login page; they simply return a login form from any of their URLs anytime they determine the user is not logged in.

In this case, there is no specific URL to match; any one of a set of URLs may be (or may not be) a login form. The only way to know is to look in the HTML being returned.

<formCriteria> Tag

When deciding on a policy, iChain first matches the URL of the page with the <url> value or pattern. If that matches, iChain then searches for whatever is placed between the form criteria on and off tags. If more than one line of text is given as criteria, iChain searches for each line individually. Criteria must match exactly, so it is a good idea to cut and paste a line directly from the login form. It should also be unique to the form, and as small as possible, because the larger the criteria, the more time-consuming the searches will be.

NOTE: By defining <formCriteria>, performance may be increased, as it specifically looks for the string that you specified rather than parsing every packet that matches the URL. Not defining <formCriteria> may cause Form Fill to not work properly with certain applications when the <url> that you specified is not specific enough. It is recommended that you define both the <url> and <formCriteria> tags as specifically as possible.

Figure 19 shows some new uses of <input> tags.

Notice the values of "~". This instructs iChain to remember this value the first time the user enters it and to supply it automatically thereafter. The last <input> tag refers to a check box. Since check boxes and radio buttons behave differently than ordinary input fields, their types must be explicitly specified.

This policy also specifies the <post/> action so each user will be prompted once for a user ID, password, and with a check box titled useJava. Once the user has supplied these values, the values will be remembered and single sign-on will be simulated from then on (or until the login fails).

<cgiCriteria> Tag

Currently a Java class (usually a servlet) in the Web server generates some login and logout forms that only have a difference in the CGI parameter-value pair(s).

For example, a login is:

```
webaccess.provo.novell.com/servlet/  
webacc?action=User.Login&.....
```

and a logout is:

```
webaccess.provo.novell.com/servlet/  
webacc?action=User.Logout&.....
```

In order to distinguish between the login and the logout, the Form Fill Policy includes a `<cgiCriteria>` tag beside the `<url>`. The lines between the tags will be used for the string comparison (case-sensitive).

Figure 20 `<cgiCriteria>` Tag

```
<url>webaccess.provo.novell.com/servlet/webacc</url>
<cgiCriteria>
  action=User.Login
  .....=.....
</cgiCriteria>
```

`<cgiCriteria>` is a complementing tag for `<url>`. A policy should have `<url>`, but `<cgiCriteria>` is an additional option for `<url>`. When the URL is matched, form fill will try to match the `<cgiCriteria>` tag if there is one. Users can put `<cgiCriteria>` before `<url>` but it is recommended that the `<cgiCriteria>` tag follow the `<url>` tag.

Because form fill uses a sequential search to match a policy, the most specific policy should always be placed before the generic one. For example, the URL of login and logout forms are the same, but the specific one includes the `<cgiCriteria>` tag (the generic one does not). Otherwise, the specific one will never be triggered (see [Figure 21](#) and [Figure 22](#)).

Figure 21 Specific Example

```
<!-- This is the normal WebAccess exit. This is a specific one -->
<urlPolicy>
  <name>GWEexit</name>
  <url>webaccess.provo.novell.com/servlet/webacc</url>
  <cgiCriteria>
    action=User.Logout
  </cgiCriteria>
  <actions>
    <redirect>http://webaccess.provo.novell.com:1959/data/gwlogout.htm</redirect>
  </actions>
</urlPolicy>
```

Figure 22 Generic Example

```

<!-- This is the normal WebAccess login. This is a generic one. -->
<urlPolicy>
  <name>GWLlogin</name>
  <url>webaccess.provo.novell.com/servlet/webacc</url>
  <formCriteria>
    loginForm
  </formCriteria>
  <actions>
    <fill>
      <input name="User.id" value="~">
      <input name="User.password" value="~">
    </fill>
    <post/>
  </actions>
</urlPolicy>

```

List Box Support

Form fill supports the list box. The syntax for list box in the policy is:

```

<select name="list box name" type="listbox"
value="~"

```

The tag name is "select". Type must be "listbox" and value must be "~".

Form fill only supports list box of single selection.

For example:

```

<select name="MyChoice">
<option>Apple
<option selected>Banana
<option>Pear
</select>

```

is a portion of the origin server's login page.

The corresponding Form Fill Policy should be:

```

<select name="MyChoice" type="listbox" value="~">

```

Login Failure Policies

The following is an example of a simple login failure policy:

Figure 23 Simple Login Failure Policy

```
<urlPolicy>
  <name>testLoginFailure</name>
  <url>www.test.com/abc/loginerr.htm</url>
  <actions>
    <deleteRemembered>test2</deleteRemembered>
    <redirect>www.test.com/abc/login.htm</redirect>
  </actions>
</urlPolicy>
```

This example shows the normal way to handle a login failure and introduces two new actions: `<deleteRemembered>` and `<redirect>`. When a login fails, it usually means that iChain has remembered values that are incorrect. `<deleteRemembered>` gives a way to remove those values from iChain's memory. It is necessary to tell iChain the name of the form that remembered the values to be deleted.

The above example instructs iChain to delete the remembered values of a policy called test2. This is a case where the `<name>` value is not optional. A form-fill policy that has login failure policies needs to be named so those policies can reference it.

The other case is when more than one policy matches a given URL. Suppose you have a policy with a `<url>` of `www.a.com/b/*` and another policy with a `<url>` of `www.a.com/*`. The actual URL of `www.a.com/b/c.html` matches both policies. In this case, both policies must be named to avoid ambiguity. It is recommended that you name all policies to avoid potential problems later on.

When deciding on a policy to use for filling a form, iChain uses the first matching policy in the policy file. For this reason, it is a good idea to put more specific policies first in the file, and more general ones after.

Redirect Action

The `<redirect>` action takes a URL that instructs iChain where to redirect the user's browser. Usually this is the original login page where the user will have the chance to try logging in again.

Login Policy with Form Fill

The following example is a login failure policy that also does form fill:

Figure 24 Login Failure Policy with Form Fill

```

<urlPolicy>
  <name>testLoginFailure2</name>
  <url>www.test.com/abc/*</url>
  <formCriteria>
    Login failed.
  </formCriteria>
  <actions>
    <deleteRemembered>test2</deleteRemembered>
    <fill>
      <input name="userid" value="~">
      <input name="password" value="~">
      <input name="useJava" value="~" type="checkbox">
    </fill>
  </actions>
</urlPolicy>

```

The <url> tag specifies a range of URLs to match so it has <formCriteria> to note when the failure actually occurs. In this case, it searches for the text Login failed. When this policy matches, it deletes the remembered fields of the policy called test2, but instead of redirecting to a login form, it fills the current form. Some Web applications return an error and another login form as part of the same response. This case is handled as shown in [Figure 24](#). Note that the fields remembered as the result of this policy are actually remembered in behalf of test2, the normal form-fill policy for which this policy is the login-failure policy. As previously mentioned, this association is made via the <deleteRemembered> action.

Most applications will present a logout or exit link to allow the user to exit the application. In most cases, an HTML page is returned, indicating the session status and also offering the user the chance to log in again via a link to the original login page. In some rare cases, a page is sent back, redirecting the user's browser to retrieve the original login page. Both scenarios would cause confusion to the user when automatic form fill is enabled. A false sense of security is implied in the first case: anybody who takes over the workstation can hit the re-login button and log in (automatically done by form fill) as the previous user because the browser session is still valid.

The second case would create the impression that the user will be re-logged in (automatically done by form fill) again when the logout/exit button is selected.

A logout mechanism resolves the above issues. A form fill policy needs to be configured to intercept the logout/exit page and redirect the user to a customized page which allows him or her to log out of iChain.

The following is an example of a policy with a logout mechanism in place:

Figure 25 Policy without Logout Mechanism

```
<urlPolicy>
  <name>OracleExit</name>
  <url>accelar.provo.novell.com/apps/plsql/icx_admin_sig.Startover</url>
  <formCriteria>
    ICXINDEX2.htm
  </formCriteria>
  <actions>
    <redirect>http://(IP address of iChain/ICS)/data/ilogout.htm</redirect>
  </actions>
</urlPolicy>
```

The page (ilogout.htm) and its associated image files that the logout/exit action is redirected to should be deposited in the SYS:\ETC\PROXY\DATA directory on the iChain proxy server. Remember to use the IP address of the iChain proxy server to complete the path to the page.

A sample page for the above policy might look like the following:

Figure 26 Sample Page

```
<html>
<head>
  <meta http-equiv="Content-Type" content="text/html; charset=iso-8859-1">
  <meta name="Author" content="someone">
  <meta name="GENERATOR" content="Mozilla/4.7 [en] (WinNT; U) [Netscape]">
  <meta http-equiv="Pragma" content="no-cache">
  <meta http-equiv="Cache-Control" content="no-cache">
  <title>ilogout</title>
</head>
<body>
  &nbsp;
  <br>&nbsp;
  <br>&nbsp;
  <center>
  <p><img SRC="novell.gif" height=200 width=300>
  <p>You have just logged out of Oracle application,
  <br>press the button below to log out from <b><font
  color="#FF0000">iChain</font></b>
  as well.
  <p><a href="http://accelar.provo.novell.com/cmd/BM-Logout"><img
  SRC="logout.gif" BORDER=0 height=44 width=56</a></center>
  <p><br>
  </body>
</html>
```

The key to the actual iChain Proxy Server logout action lies on the reference link associated with the logout button: <http://accelar.provo.novell.com/cmd/BM-Logout>, replacing the URL with the proper domain name associated with the link (accelar.provo.novell.com in this case).

Storing User Credentials with SecretStore

User-entered data in form fields will be stored in the user object's `ichainFormFillCrib` attribute by default. To make the storage of these credentials more secure, you can use Novell SecretStore™.

SecretStore is a service in eDirectory that allows applications to store sensitive data securely. This service guarantees protection for integrity and confidentiality of application data in storage and during transmission between client and SecretStore.

In order to use SecretStore, you must upgrade to NCI version 2.0.2, install SecretStore, then enable SecretStore. (These procedures are explained in the following sections.)

NOTE: Verify that all the LDAP servers listed in the Access Control tab have SecretStore installed on them or Form Fill will not work properly.

Installing SecretStore on Windows NT

- 1 Insert the *iChain Authorization Server* CD.
- 2 Run **NICI/WCNICIU0.EXE**.
- 3 Follow the screen prompts to complete the install.
- 4 Run **SSO/NT.SETUP.EXE**.

During the SecretStore install, you will need to log in to your tree as Administrator in order for the schema to be modified.

- 5 Start the NDS Services Console.
- 6 Configure each of the following services to start up automatically:

ssldp.dlm
ssncp.dlm
sss.dlm

- 7 Continue with [“Enabling SecretStore” on page 109](#).

Installing SecretStore on NetWare

- 1 Insert the *iChain Authorization Server* CD into the CD drive on the server.
- 2 From a remote console prompt, enter **CDROM** to mount the *iChain Authorization Server* CD.

The server mounts the CD as ICHAIN_AUTH_SR.

- 3** To upgrade NICE, enter **NWCONFIG**. Choose Product Options > Install a Product Not Listed.
- 4** Press Esc and then F3, > enter the path to the NICE upgrade:
ICCHAIN_AUTH_SR:/NICE/NWSERVER > press enter.
- 5** Accept the license agreement to complete the NICE upgrade.
- 6** To install SecretStore, enter **NWCONFIG**. Choose Product Options > Install a Product Not Listed.
- 7** Press Esc and then F3, > enter the path to the SecretStore installation:
ICCHAIN_AUTH_SR:/SSO/NETWARE > press enter to complete the SecretStore installation.
- 8** Run **\SSO\UTILS\SSINIT.EXE** from the *Novell iChain Authorization Server* CD, using the NetWare server as the default server. (This is done from the client.)
- 9** Continue with **“Enabling SecretStore” on page 109**.

Enabling SecretStore

After SecretStore has been installed, you must enable it to work with Form Fill as explained in the following procedure.

- 1** Log on to the LDAP server.
- 2** Export your trusted root for the LDAP server.
 - 2a** From ConsoleOne[®], view the properties of the key material object (usually named SSLCertificate*).
 - 2b** Select the Certificates tab > select Trusted Root > click Export.
 - 2c** Save the trusted root in Base64 format to the local drive.
- 3** View the properties of the ISO object.
- 4** From ConsoleOne, select the Form Fill Policy tab > mark Use Novell Secret Store for Form Fill > click OK.
- 5** From the iChain Proxy Services Admin GUI, import the trusted root.
 - 5a** Select Configure (this icon is in the bottom left corner) > select the Access Control tab > mark Enable Secure Access to LDAP Server.
 - 5b** Select Import Trusted Root and enter a name for the imported file.
The file name must adhere to 8.3 naming conventions.

5c Using a text editor, open the file that was exported from the LDAP server to a local drive > copy all the contents of the file to the text field in the Trusted Root dialog > click OK.

5d Mark Enable Form Fill Authentication and click Apply.

NOTE: If you already had Form Fill enabled, you need to refresh Form Fill from the Access Control tab.

6

Setting Up an Advanced Configuration

The previous chapters in this document described how to install and set up the basic implementation described in “[Installation Scenario](#)” on page 13. To meet your company’s networking needs, you may need to augment or alter this implementation and use some of the more advanced features of Novell® iChain® services. This chapter describes the following Novell iChain configuration procedures:

- ♦ “[Setting Up Secure Exchange](#)” on page 111
- ♦ “[Advanced Access Control Configuration](#)” on page 112
- ♦ “[Multi-homed Configurations](#)” on page 115

Setting Up Secure Exchange

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

To set up Secure Exchange:

- 1** Export a base-64 trusted root file for the Web server where you want to enable secure access.
- 2** Access the URL of the proxy server where you installed the iChain Proxy Services software to launch the proxy server browser-based administration tool.

For example, `http://xx.xx.xx.xx:1959/appliance/config.html`
where `xx.xx.xx.xx` is the IP address.

- 3** Click `Configure > Web Server Accelerator > Modify`.
- 4** Check `Enable Secure Exchange`.
- 5** Specify the SSL listening port to use for Secure Exchange traffic.
- 6** Specify the certificate to use to establish the Secure Exchange session.
- 7** Click `Secure Exchange Options`.
- 8** Click `Enable Secure Access Between Secure Exchange and Web Server`.
- 9** Click `Insert > Import Trusted Root`.
- 10** Enter a file name with a `.DER` extension.
- 11** Paste the contents of the previously exported base-64 trusted root file in the lower text box.
- 12** Click `OK > OK > Apply`.

Advanced Access Control Configuration

This section contains information about the following topics:

- ◆ [“Enabling ACL Rule Checking for Community Objects” on page 112](#)
- ◆ [“Enabling Debugging Messages for Access Control” on page 113](#)
- ◆ [“Using ACLCHECK options” on page 113](#)

Enabling ACL Rule Checking for Community Objects

iChain 2.1, by default, will not check community objects for ACL rules. Community objects existed in previous versions of iChain but are no longer provided in version 2.1; however, the functionality is provided to allow the use of pre-existing community objects.

To enable ACL rule checking for community objects when upgrading from iChain 1.5 to iChain 2.1, administrators should do the following:

- 1** Unlock the console.
- 2** Edit the `appstart.ncf`.
- 3** Change the load `aclcheck` entry to `load aclcheck /m`.

4 Restart the machine.

After specifying changes in the configuration, ACL rules will be checked in the following sequence:

- OU's
- OU's communities*
- groups
- groups' communities*
- user
- user's communities*

If a specified option is not provided, checking for the italicized portions of the above list will not be performed for checking the ACL rules.

Enabling Debugging Messages for Access Control

The module that provides iChain's Access Control (ACLCHECK.NLM) can be configured to output debug information. The administrator can choose one of two levels of increasingly more detailed information. This information can be helpful to developers and consultants. By default, no debug information is output.

To enable these debugging options, an administrator should:

- 1** Edit the APPSTART.NCF file on the iChain Proxy Server.
- 2** Find the line containing the **LOAD ACLCHECK** command and add a debug level switch at the end of that line, for example, **LOAD ACLCHECK /D2**.
- 3** Shut down and restart the proxy server.

Using ACLCHECK options

The ACLCHECK utility can be used with a number of options to refine rule checking. These options are not case sensitive. When you change an ACLCHECK option, the update is stored in the appstart.ncf file.

Table 10 ACLCHECK command line options

Option	Syntax and Example	Explanation
Check dynamic ACLs	<pre>ACLCHECK /Q ACLCHECK /Q</pre>	By default, dynamic ACLs are checked after checking all traditional (static) ACLs. If this option is specified, ACLCHECK first checks for dynamic ACLs. This option should be used when you have mainly dynamic ACLs.
Cache refresh interval	<pre>ACLCHECK /F<i>number_of_minutes</i> ACLCHECK /F300</pre>	<p>Default: 180 minutes</p> <p>Keep this number higher if you are not likely to change DS information quickly. This can improve performance since ACLCHECK does not need to throw away the already built-up cache.</p>
Maximum log file size	<pre>ACLCHECK /S<i>max_file_size_in_KB</i> ACLCHECK /S2000</pre>	Default: 1 MB
Number of connection handles for the LDAP server	<pre>ACLCHECK /C<i>number_of_connections</i> ACLCHECK /C70</pre>	<p>Default: 10</p> <p>If you see an error message stating that ACLCHECK was unable to obtain any LDAP handles, increase this number to avoid that problem. The maximum recommended number of connections is 70.</p>
Debug level	<pre>ACLCHECK /D<i>level</i> ACLCHECK /D2</pre>	<p>Default: 0</p> <p>Debug information can be helpful to developers and consultants. Set the level at 1 or 2 for more detailed information.</p>
Utility help	<pre>ACLCHECK /H ACLCHECK /H</pre>	Gives you information about ACLCHECK.

Multi-homed Configurations

This section contains information about the following topics:

- ◆ “What is Multi-homing?” on page 115
- ◆ “Multi-homing Web Server” on page 116
- ◆ “Host-based Multi-homing” on page 116
- ◆ “Domain-based Multi-homing” on page 117
- ◆ “Path-based Multi-homing” on page 118
- ◆ “Path-Based Multi-homing Examples” on page 120

What is Multi-homing?

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

Within a large company or corporation, IP address (for example, 10.1.1.1) are valued resources that are managed with great care. Opening port numbers can be especially time-consuming and political because of firewall and security issues.

For example, imagine using iChain to accelerate 100 backend origin Web servers. Each of these web servers will have their own IP address. In order to grant access to these Web servers through iChain, a DNS server must map each Web server name to an IP address that iChain is listening on. A simple mapping would mean that 100 IP addresses would be needed by iChain to map to 100 backend Web servers.

It is possible for iChain to multi-home all 100 backend origin Web servers through one IP address and port. Most installations, however, will organize the origin Web servers into a handful of groups where each group is a group of multi-homed accelerators listening on one IP address.

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

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

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

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

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

Multi-homing Web Server

A multi-homing Web server is a Web server that can listen on a single IP address and multi-home the host name into multiple logical Web servers on the same machine. This is an easy way of having multiple alias names representing one host name. At the Web Server Accelerator screen, the user should select the radio button, "Use host name sent by browser (multi-homing web server)". If a request is made on the accelerator IP address and port, the host name sent by the browser will not be checked against the DNS name of the accelerator.

This kind of multi-homing has some potential problems with iChain. The first issue is that any host name can come through that is unknown to iChain's Protected Resource List or Access Control List. iChain could restrict more URL references than the administrator's name.

Host-based Multi-homing

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

Table 11

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

Perhaps this configuration looks ready for customer use. A problem will arise, however, when either authentication is enabled or Secure Exchange is enabled for two or more of the above accelerators. When authentication or Secure Exchange is enabled, an SSL listening port will be issued to the accelerator. Two or more accelerators cannot listen on the same SSL port unless multi-homing is enabled. This is because a certificate must be used to establish an SSL connection. The certificate has the host name or a wild card host name (*.novell.com). A single certificate cannot be served to secure both www.a.com and www.b.com. Usually each accelerator will have its own unique SSL listening port for an IP address. See [Table 12](#).

Table 12

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

Domain-based Multi-homing

Domain-based multi-homing is selectable when "Enable multi-homing" is checked. To be more specific, this feature could be named "Cookie-domain-based multi-homing" because all of the DNS names for the accelerators in the multi-homing have the same cookie domain (see acme.com in [Table 13](#)).

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

2.1 version, the master also defines the authentication and Secure Exchange options.

Table 13

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

The same SSL port can be used for a group of domain-based accelerators if the SSL certificate supports wild card host names. In [Table 13](#), the SSL certificate would be "*acme.com".

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

Path-based Multi-homing

Path-based multi-homing is selectable when "Enable multi-homing" is checked. Path-based multi-homing could be considered the most secure option because an SSL certificate containing a host name is used. The belief exists that wild card certificates are less secure than a certificate with just a host name. All accelerators that participate in a path-based group have the same DNS name. All children must have a unique starting or ending path.

The same rules apply as they do for a domain-based multi-homing master: There must be a "master" accelerator already defined. A master accelerator is not part of a host-based multi-homing group and is not a child or a multi-homing group. The master defines the cookie domain, accelerator IP address, accelerator proxy port, SSL listening port, and certificate. In the initial iChain 2.1 version, the master also defines the authentication and Secure Exchange options.

Table 14

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

There are three types of options for a path-based multi-homing child. [Table 15](#), [Table 16](#), and [Table 17](#) outline the URL requested from the browser and the URL that will be submitted to the origin Web server.

Table 15 "Starts with" option, sub-path match string = /Bstuff, remove sub-path from URL checked

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

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

Table 16 "Starts with" option, sub-path match string = /Bstuff

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

This second example expects the sub-path match string to be present as a valid path on the origin Web server. The rewriter doesn't need to update the absolute references as it did in [Table 15](#).

Table 17 "Ends with" option, sub-path match string = jpg

URL from browser	URL to web server
http://a.acme.com/logos/companyLogos.jpg	http://b.internal.com/logos/companyLogos.jpg

Path-Based Multi-homing Examples

Example One

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

The administrators sets up two accelerators for www.zxy.org on the same IP address and port number, and configures them with path-based multi-homing rules.

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

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

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

Example Two

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

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

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

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

Browsers can now access the single www.cab.org Web site and have requests for www.cab.org/main.asp get directed to the ASP Web server farm, requests for www.cab.org/logo.gif and www.cab.com/photo.jpg get directed to the graphics Web server farm, and requests for www.cab.com/directory.html get directed to the third Web server farm.

122 Place Book Title Here

7

Using and Tuning iChain Features

Novell® iChain® includes a variety of logging tools that allow you to view information about access rules, news servers, mail servers, and Web server activity to help you manage your infrastructure.

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

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

Managing Appliance Certificates

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

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

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

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

To create appliance-specific certificates, see the instructions in [“Creating Certificates Using the Appliance CA”](#) on page 125.

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

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

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

To generate a CSR and store the issued certificate, complete all the instructions in [“Obtaining a Certificate from an External CA” on page 126](#)

Because the creation process is different for internal and external certificates, they are described separately in [“Creating Certificates Using the Appliance CA” on page 125](#) and [“Obtaining a Certificate from an External CA” on page 126](#).

Naming Certificates

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

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

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

3. Choose the subject name that the browser expects to find in the certificate.

- ◆ For accelerator services, the Subject Name field must contain the DNS name, with the fields separated by periods (.).

For example, the www.foo.gov Web server accelerator certificate must have a Subject Name of www.foo.gov.

- ◆ For client accelerator and transparent services, the Subject Name field must contain the IP address on which the request is sent, for example, 110.1.1.199.

Creating Certificates Using the Appliance CA

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

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

To create an appliance CA certificate:

- 1** In the browser-based management tool, click Home > Certificate Maintenance > Create.
- 2** Type an appropriate name for the certificate as explained in ["Naming Certificates" on page 124](#).
- 3** Type an appropriate subject name as explained in ["Naming Certificates" on page 124](#).
- 4** Click the Signature Algorithm drop-down list > select the algorithm you want to use (SHA-1 or MD-5).
- 5** Click the RSA Key Size drop-down list > select the RSA key size that you want to use.

You cannot select a key size larger than the maximum key size on the appliance.
- 6** Check Use Local Certificate Authority.
- 7** Click the Validity Period drop-down list > select the length of time that you want the certificate to be valid.
- 8** Click OK.
- 9** Look at the Action and Status fields.

The Action field should have red arrows on the left and the word Create displayed on a green background. The Status should be Building.

The red arrows and green background indicate that you need to click Apply.

- 10** Click Apply.

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

- 11** If an error occurs, click Modify

- 12** In the Modify Certificate dialog box, make the changes necessary to resolve the errors > click OK.
- 13** Click Apply and repeat the modification process until the Status field displays the word Active.

Obtaining a Certificate from an External CA

Requesting the CSR

- 1** In the browser-based management tool, click Home > Certificate Maintenance > Create.
- 2** Type an appropriate name for the certificate as explained in [“Naming Certificates” on page 124](#).
- 3** Type an appropriate subject name as explained in [“Naming Certificates” on page 124](#).
- 4** Click the Signature Algorithm drop-down list > select the algorithm you want to use (SHA-1 or MD-5).
- 5** Click the RSA Key Size drop-down list > select the RSA key size that you want to use.

You cannot select a key size larger than the maximum key size on the appliance.
- 6** Click Use External Certificate Authority.
- 7** If you are requesting a VeriSign* certificate, check the VeriSign CA check box. Otherwise, leave the box unchecked.
- 8** If desired, type a name for your organization or division.

This is commonly referred to as the Organizational Unit and is used to differentiate organizational divisions or to describe departments or divisions.
- 9** Type the city or town where your organization does business.

This is commonly referred to as the Locality.
- 10** Type the unabbreviated name of the state or province where the organization does business.

This is commonly referred to as the State.
- 11** Type the International Standards Organization (ISO) country code for the country where the organization does business.

This is commonly referred to as the Country and must be a valid, two-character ISO country code.

12 Click OK.

13 Look at the Action and Status fields.

The Action field should have red arrows on the left and the word Request displayed on a green background. The Status should be Building.

The red arrows and green background indicate that you need to click Apply.

14 Click Apply.

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

15 If an error occurs, click Modify

16 In the Modify Certificate dialog box, make the changes necessary to resolve the errors > click OK.

17 Click Apply and repeat the modification process until the Status field displays the words CSR in Progress on a yellow background.

NOTE: As an added precaution, "update clone" can be used to help safeguard the private key of the certificate until the certificate is returned and stored. After the certificate is returned and stored, it can then be backed up. "Update clone" is found in the iChain Proxy Server browser-based administration tool under System > Actions.

Sending the CSR

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

2 Select and copy the complete CSR text into your computer's clipboard. Internet Explorer and other browsers sometimes combine them with the CSR text that is in between. Clicking the browser refresh/reload button will often fix the problem. If it doesn't, simply insert appropriate carriage returns during the next step. After you have copied the text you can close that browser window.

If you don't fix the defect, you can view the source of the HTML file and copy and paste from the source file.

3 Paste the CSR text from the clipboard to the e-mail message or HTML form as required by your CA.

The method for sending the CSR will vary depending on the authority. VeriSign, for example, uses a Web page interface.

IMPORTANT: The header and trailer must be on lines separate from the body of the CSR.

The header line will be similar to the following:

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

The trailer line will be similar to the following:

```
-----END NEW CERTIFICATE REQUEST-----
```

If required, you must use hard returns to separate these two lines from the body of the CSR.

- 4 Wait for the certificate to be returned from the external CA.

Storing the Certificate

After the external CA responds with the certificate:

- 1 In the browser-based tool, click Home > Certificate Maintenance > the name of the certificate you want to store > Store Certificate.
- 2 In the Store Certificates dialog box, paste the CA certificate into the CA Certificate Contents box.

NOTE: If you requested a VeriSign certificate and you checked the VeriSign box in [Step 7 on page 126](#), the CA Certificate Contents box is dimmed. You do not need to paste the VeriSign CA certificate because VeriSign certificates are already stored on the appliance.

- 3 Paste your newly issued certificate in the Server Certificate Contents box.
- 4 Click Create.
- 5 Look at the Action and Status fields.

The Action field should have red arrows on the left and the word Create displayed on a green background. The Status should be CSR in Process.

The red arrows and green background indicate that you need to click Apply.

- 6 Click Apply.

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

- 7 If an error occurs, click Store Certificate

- 8** In the Store Certificate dialog box, make sure the correct certificates are pasted in the boxes > click OK.
- 9** Click Apply and repeat the modification process until the Status field displays the words Active on a green background.

Viewing (Exporting) a Certificate's CA

To view (export) a certificate's Certificate of Authority (CA):

- 1** In the browser-based management tool, click Home > Certificate Maintenance > the certificate you want to export > Export CA Certificate > View Source of HTML.

The contents of the CA certificate are displayed in a new browser window.

Modifying a Certificate

Only certificates that have an error or the status Building can be modified.

- 1** In the browser-based management tool, click Home > Certificate Maintenance > the certificate you want to modify > Modify.
- 2** After making the necessary changes, click OK to accept the changed values.
- 3** In the Modify Certificate dialog box, make the desired changes.
- 4** If the Action field displays the word Request or Create on a red background, you must click Apply to make the changes.

Deleting a Certificate

If a certificate has expired or you are unable to resolve an error, you might want to delete a certificate.

IMPORTANT: Use caution when deleting certificates. You should never delete system-generated certificates.

- 1** In the browser-based management tool, click Home > Certificate Maintenance > a certificate you have generated that has expired or has an unresolvable error.
- 2** Click Delete.
- 3** In the Delete Certificate dialog box, click Yes.

- 4** The certificate is removed from the certificates list.
If you have deleted the certificate in error, click Cancel.
- 5** Click Apply to remove the certificate from the appliance.
After clicking Apply, the certificate cannot be restored unless you have created a backup copy.

Backing Up a Certificate

Only active certificates can be backed up.

- 1** In the browser-based management click Home > Certificate Maintenance > the certificate you want to back up.
- 2** Click Backup.
- 3** In the Backup Certificate dialog box, type a password to use when restoring the certificate.
- 4** In the Confirm Password field, retype the same password.

IMPORTANT: Although the password is optional, we strongly suggest you use one. If you don't enter a password, the backed-up certificate can be used by anyone who has access to the file.

- 5** Check either Disk or Floppy to indicate where the backup file should be placed.
- 6** Click OK.

The Action field should display red arrows and either Backup (Disk) or Backup (Floppy) on a green background.

If you want to cancel the backup action, click Cancel Backup by the Action field.

- 7** If the Action field is green, click Apply.

The Backed Up status field for each certificate indicates whether a certificate has been backed up and where the backup file was placed (disk, floppy, or both).

If any errors occur during the backup process, they are displayed on the Error line and the background turns red.

You can then click Backup and repeat the process taking care to avoid the errors indicated.

Backed-up certificates are stored in a file named *CERTIFICATE.PFX*, where *CERTIFICATE* is the name of the certificate that was backed up.

IMPORTANT: If the certificate was backed up to the appliance hard disk, you should transfer the file from the appliance to another secure location, or the backup copy will be lost if the appliance fails and has to be reimaged.

Certificate backup files are stored in ETC/PROXY/APPLIANCE/CONFIG/USER/CERT/BACKUP. See [“Using FTP” on page 166](#) for help using appliance FTP services.

If the certificate was backed up to a floppy disk, the file is in the root directory of the disk and the floppy should be stored in a safe place in case the certificate must be restored.

Restoring a Certificate

Only certificates that were previously backed up can be restored.

Prior to completing the following steps, make sure the backup file is in one of the following locations:

- ◆ On a floppy disk in the appliance’s floppy drive
- ◆ In ETC/PROXY/APPLIANCE/CONFIG/USER/CERT/BACKUP on the appliance’s hard disk.

Unless the appliance has been damaged or reimaged, the backup file will be in the expected location.

If the file is not on the appliance, you must retrieve a copy from your secure location and either copy it to a floppy disk or to the appliance using FTP.

- 1** In the browser-based management tool, click Home > Certificate Maintenance > Restore.
- 2** In the Restore Certificate dialog box, type the certificate name, which is the PFX filename.
- 3** Type the same password you used when creating the backup file.
- 4** Click OK.
- 5** Click Disk or Floppy to indicate where the backup file is.
- 6** Click OK.

The Action field should display red arrows and either Restore (Disk) or Restore (Floppy) on a green background. The Status field should display Building.

If you want to cancel the restore action, click Cancel Restore by the Action field.

7 Click Apply.

If any errors occur during the restore process, they are displayed on the Error line and the background for the text will turn red.

The only way to fix a restore error is to delete the certificate and try the restore process again.

A restoration failure might mean that the backup file didn't exist or you had the wrong password.

Certificate Error Handling

Currently if accelerators have mutual authentication (this may include mutual and other authentication) enabled, when users present bad certificates (expired or revoked) to access these accelerators, the browsers display a "page not found" error. The certificate error handling feature enables administrators to configure the error messages so that they can define what the problem is with a particular certificate. For example, if a certificate is expired, the administrator can configure an error message to let the user know that the certificate has expired. This feature helps administrators more effectively troubleshoot certificate problems. It also helps the user better understand why his or her certificate may not be working.

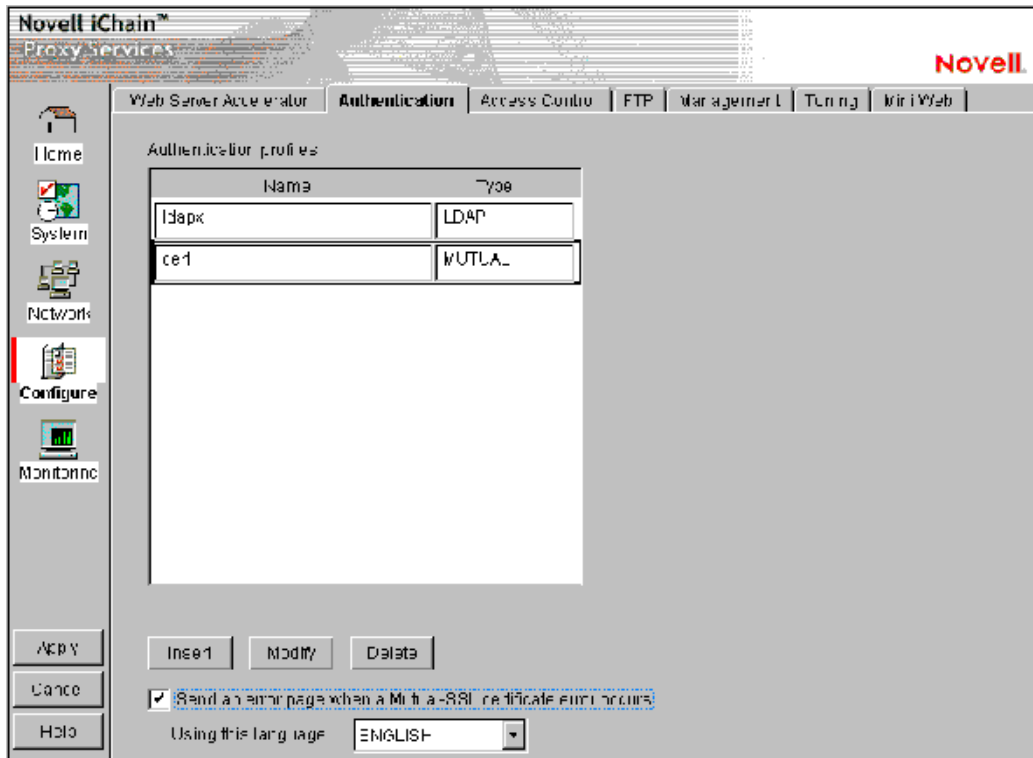
NOTE: In iChain 2.1, if users use mutual authentication or other authentication and cancel a certificate or have a bad certificate, the authentication will be failed at mutual authentication and they will not be prompted for other authentication. By turning on the certificate error handling feature, users will be prompted for other authentication. There will be no error page for failure of mutual authentication.

Using Certificate Error Handling

To use the certificate error handling feature, at the Authentication tab, a user must check the "Send an error page when a Mutual-SSL certificate error occurs" option, then select the language from the drop-down list (see [Figure 27](#)).

NOTE: The certificate error handling feature applies to all accelerators. After this feature is enabled or disabled, the user must restart the iChain server. Also, if users change the error messages and/or error pages, the iChain server must be restarted.

Figure 27 Certificate Error Handling



Customizing Error Messages

There are two files, CRTERRP.GCFG (message file) and CRFTERRPG.HTM (error page). These files are located at:
 SYS:\ETC\PROXY\DATA\ERRPAGE\NLS\ENGLISH.

The CRTERRP.GCFG file is for error status and description in the CERERRPG.HTM.

There are 58 messages in the CRTERRP.GCFG file. Users can change the content after "Translated Message=". For messages 9 and 13, users must keep "%d" for error code.

The messages are paired for status and description fields in the error page. For example, message 11 (status) and message 12 (description) are shown here:

Figure 28 Original Message 11 (Status) and Message 12 (Description)

```
[Message 11]
Message ID=LOC_MSG_CERT_NOT_YET_STATUS
English Message=Certificate Not Valid Yet
Translated Message=Certificate Not Valid Yet
```

```
[Message 12]
Message ID=LOC_MSG_CERT_NOT_YET_DESC
English Message=Your certificate's start date is in the future. Please try later or contact your system administrator.
Translated Message=Your certificate's start date is in the future. Please try later or contact your system administrator.
```

When customized, message 11 (status) and 12 (description) may look like this:

Figure 29 Customized Message 11 (Status) and Message 12 (Description)

```
[Message 11]
Message ID=LOC_MSG_CERT_NOT_YET_STATUS
English Message=Certificate Not Valid Yet
Translated Message=Certificate Not For Now
```

```
[Message 12]
Message ID=LOC_MSG_CERT_NOT_YET_DESC
English Message=Your certificate's start date is in the future. Please try later or contact your system administrator.
Translated Message=Your certificate is not for now. Please try later.
```

NOTE: Only the content following "Translated Message=" is changed.

Customizing the Error Page

To customize an error page, change the static messages in the CRTERRPG.HTM. Users can redesign their own pages as long as they use the CRTERRPG.HTM as the file name and they have <ERROR_STATUS> and <ERROR_DESCRIPTION> fields in the HTML page.

Localizing Error Messages

The current default language is English, however, users can translate error messages and the error page in other languages.

To localize error messages:

- 1 In the CRTERRPG.CFG file, change the character set from ISO-8859-1 to ISO-10646-1.

- 2** Translate the content of the Translated Message for every message (Message 1 to 14). The *_STATUS and *_DESC messages in the CRTERRPG.CFG file are used to replace the <ERROR_STATUS> and <ERROR_DESCRIPTION> fields in the CRTERRPG.HTM file.

Localizing Error Pages

To localize an error page, translate the static messages in the CRTERRPG.HTM. Users can redesign their own pages as long as they use the CRTERRPG.HTM as the file name and they have <ERROR_STATUS> and <ERROR_DESCRIPTION> fields in the HTML page.

User Management Servlets

iChain User Services is a collection of four Java* servlets that provide a lightweight and easy-to-manage user self-provisioning environment that is based on open standards (LDAP). No Novell client is needed. The user self-provisioning services include:

- ◆ User self-registration
- ◆ User account modification
- ◆ Change password
- ◆ Password manager and redirect

Java Servlets

The services in iChain User Services are provided by the following four Java servlets:

- ◆ iChainAddUser
- ◆ iChainModifyUser
- ◆ iChainPasswordChange
- ◆ iChainPasswordMgr

These servlets can be found on the Authorization Services CD under the \servlets directory. This directory contains both a compiled file (.class) and the source file (.java) for each servlet.

Servlet Requirements

The following components are required for using iChain User Services:

- ♦ Java servlet engine
- ♦ JVM 1.2.x or higher
- ♦ LDAP Classes for Java, available from the Novell NDK site (<http://developer.novell.com/ndk>). Look for "LDAP Class Libraries for Java" on the NDK. There is one download file for NetWare[®] and Windows* and a separate download file for Linux* and UNIX*.

iChainAddUser

This servlet enables Web-based user self-registration, as shown in [Figure 30](#) and [Figure 31](#).

Figure 30 Web-Based User Self-Registration

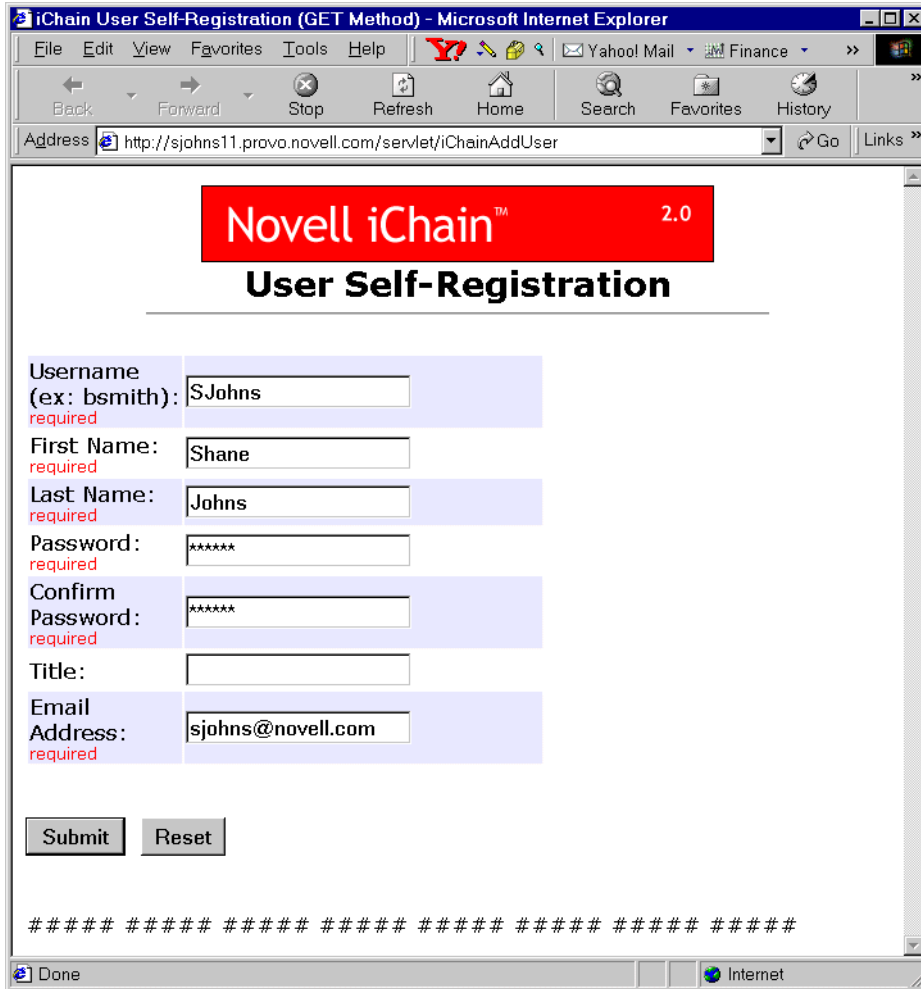
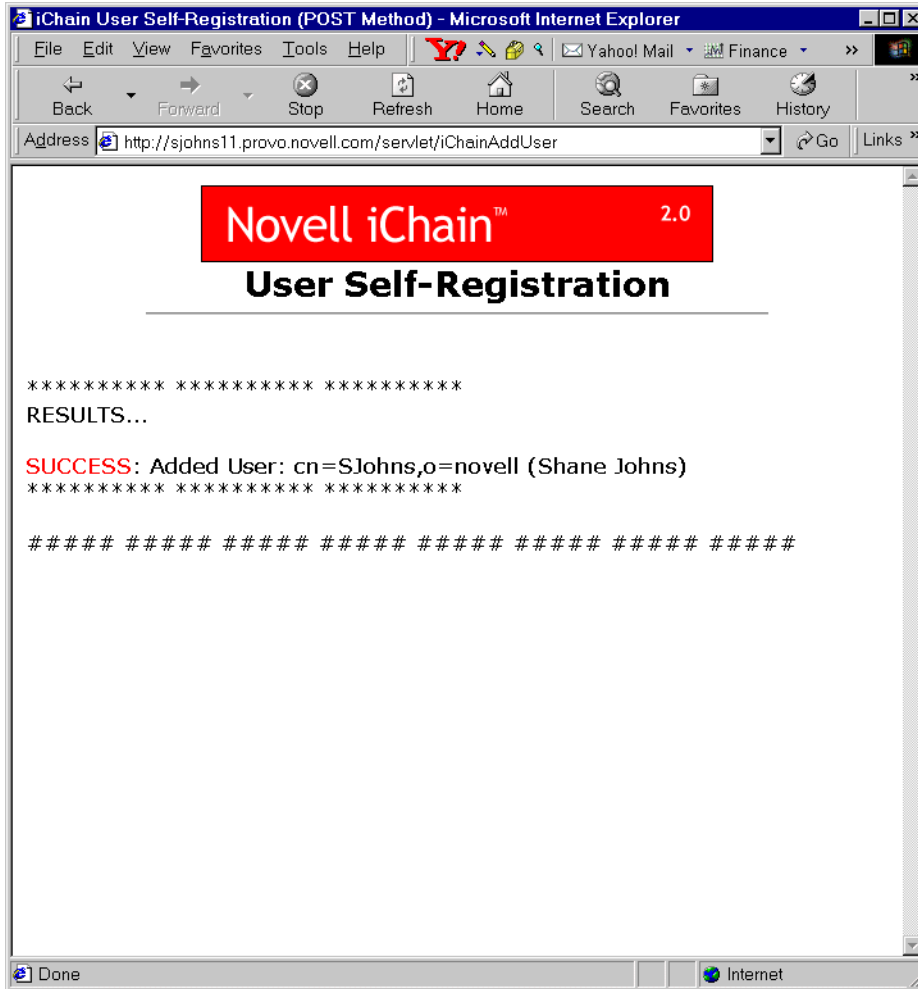


Figure 31 User Self-Registration Results



iChainModifyUser

This servlet enables Web-based user account modification, as shown in [Figure 32](#) and [Figure 33](#):

Figure 32 Web-Based User Account Modification

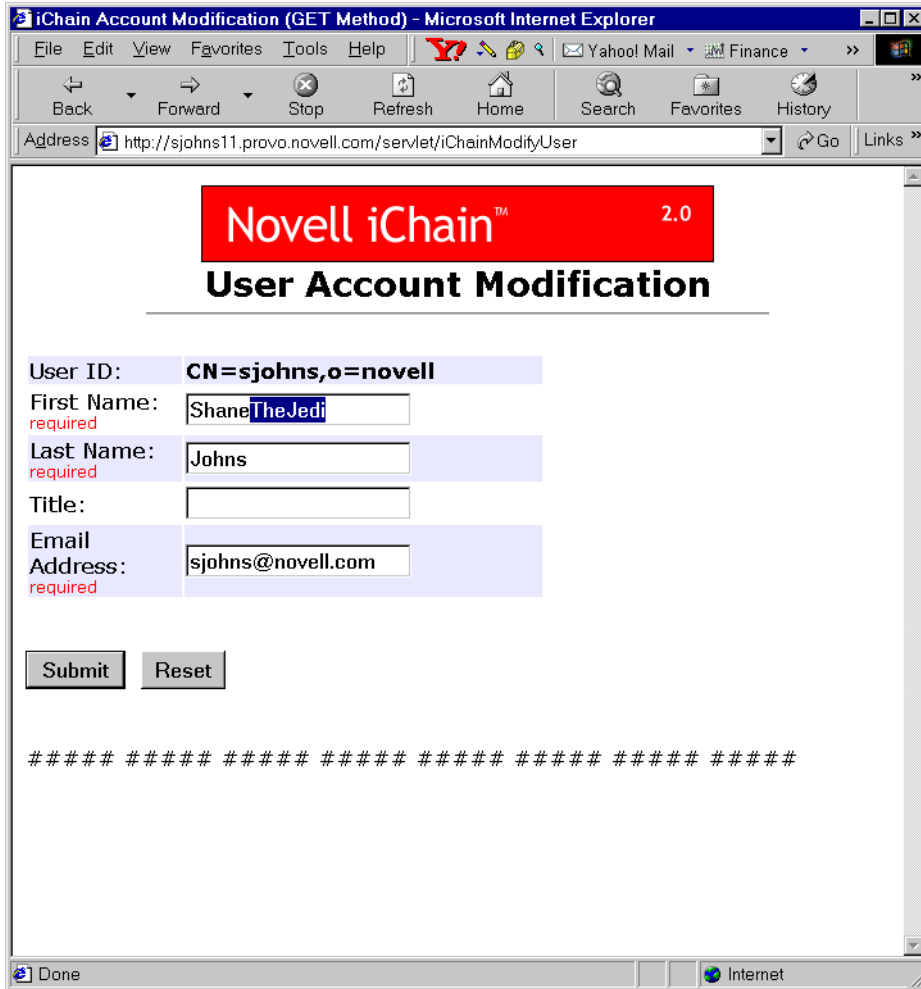
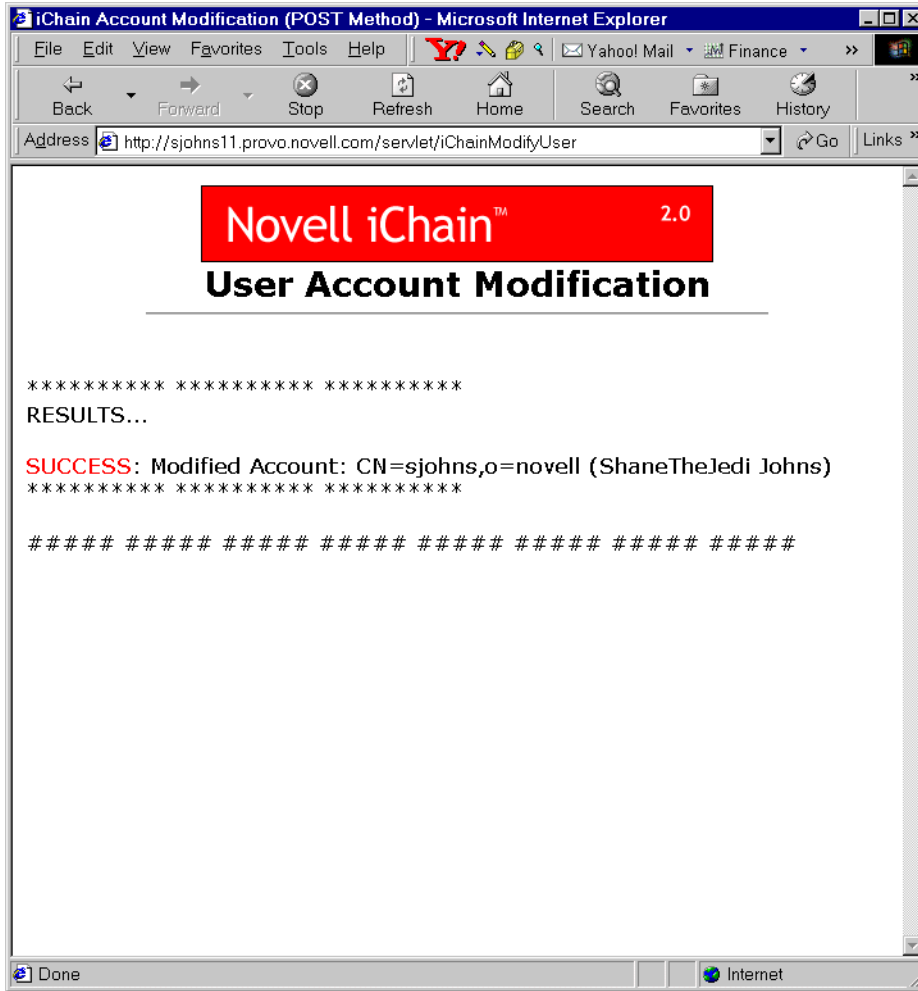


Figure 33 User Account Modification Results



iChainPasswordChange

This servlet enables users to change their own passwords, as shown in [Figure 34](#):

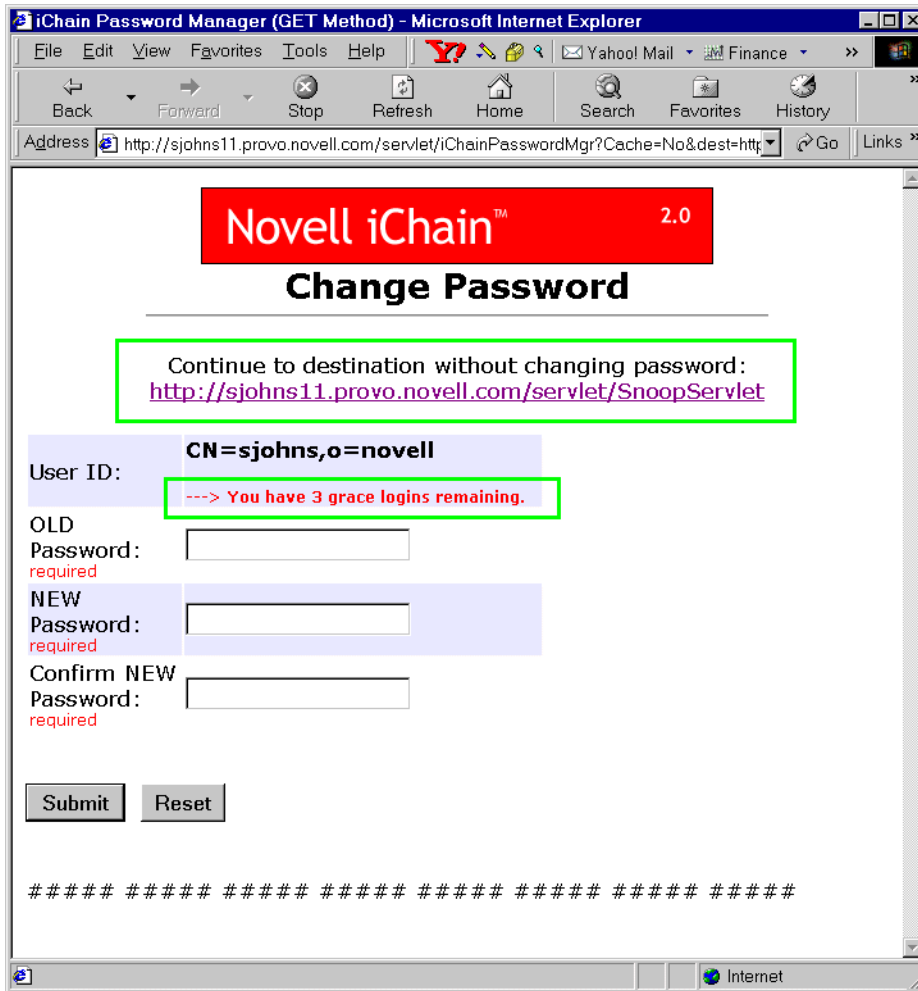
Figure 34 Password Change

The screenshot shows a web browser window titled "iChain Change Password (GET Method) - Microsoft Internet Explorer". The address bar contains the URL "http://sjohns11.provo.novell.com/servlet/iChainPasswordChange". The main content area features a red header with "Novell iChain™ 2.0" and the title "Change Password". Below the title, the "User ID" is set to "CN=sjohns,o=novell". There are three password input fields: "OLD Password:" (required), "NEW Password:" (required), and "Confirm NEW Password:" (required). Each field contains six asterisks. At the bottom of the form are "Submit" and "Reset" buttons. A line of ten hash symbols is displayed below the buttons. The browser's status bar at the bottom shows "Done" and "Internet".

iChainPasswordMgr

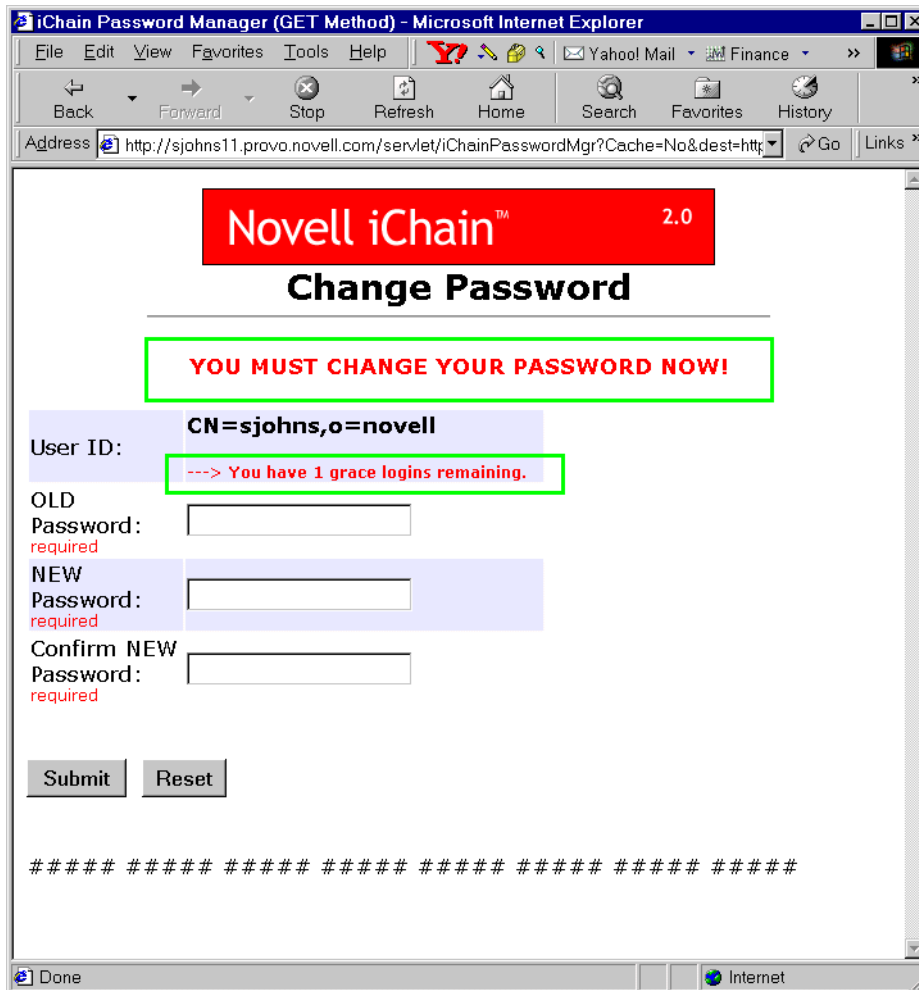
This servlet enables a Password Manager that will give users a way of changing their passwords after their passwords expire, while they still have grace logins remaining. The Password Manager will check (and display to the user) the number of grace logins remaining and will provide an automatic destination redirection once a password has been successfully changed. See [Figure 35](#).

Figure 35 Password Manager



If the grace logins are greater than 1, the user will have the option to bypass the password change screen. If the grace login equals 1, the user will be forced to change his or her password before continuing on. See [Figure 36](#).

Figure 36 Password Grace Logins



To set up the iChainPasswordMgr servlet:

- 1** From the proxy server Web administration GUI, select Configure > Access Control tab.
- 2** Enter the servlet information in the Password Management Servlet URL field.
- 3** Use the full http URL (for example, <http://ichain.provo.novell.com/servlet/iChainPasswordMgr>).

Servlet Configuration File

The servlets use a configuration file that they read upon initialization to customize their functionality to your environment and to easily provide localization (or modification) of all message and error strings. This configuration file also allows you to make any LDAP directory attributes available to the user for user account creation or modification. The name of this required configuration file is ICHAINAPPCONFIG.TXT.

NOTE: The servlets expect the ICHAINAPPCONFIG.TXT file to reside in the same directory as the servlets themselves (where you place the *.class servlet files).

The settings that can be modified/customized via this configuration are:

- ◆ HostName (IP address of LDAP directory)
- ◆ AdminName
- ◆ AdminPassword
- ◆ UsersContainer (container where users are created during self-registration)
- ◆ ISO (iChain ISO object)
- ◆ LDAP directory attributes for user creation/modification
- ◆ Message and error strings

The following is an example of the format and syntax of how to make the LDAP attributes of "title" (Job Title) and "mail" (E-mail Address) available to the user during account creation or modification:

```
-----
# ATTRIBUTES for User Creation or Modification
# Format: <LDAP name>, <"HTML display">, <Required: yes/no>
# Example: givenName, "First Name", yes
# NOTE: "cn", "givenName", "surname", & "UserPassword" are
#       automatically provided and required by default
#       (do NOT list these 4 attributes in this file)
NumberOfAttributes=2
attr1=title, "JobTitle", no
attr2=mail, "Email Address", yes
-----
```

The attributes can also be designed as "required" or "not required" when they are presented to the user during account creation or modification.

Auxiliary Class Support

The ICHAINAPPCONFIG.TXT configuration file does support auxiliary class attributes. If the attributes that you want to list in the above attributes section are attributes in an auxiliary class, you will need to add an extra entry at the end of the attrX line as shown in the following example:

```
attr3=commerceAccountID, "Commerce ID", no, commerceAcct
```

In this example, commerceAccountID is the LDAP attribute name and commerceAcct is the auxiliary class that this attribute is a member of.

Servlet Installation and Setup

Installation

There are seven files that need to be placed in the proper locations for the iChain User Services to work correctly:

- 1** Place the four servlets (iChainAddUser.class, iChainModifyUser.class, iChainPasswordChange.class, and iChainPasswordMgr.class) in the servlet directory of your servlet engine.
- 2** Place the iChainAddUser\$ISOPasswordTemplate.class class file in the same directory as the servlets (the servlet directory of your servlet engine).
- 3** Place the ICHAINAPPCONFIG.TXT configuration file in the same directory where you placed the servlets.
- 4** Place the TOP.GIF image file at the root documents directory of your Web server. (For example: for IIS, place the image file in C:\INETPUB\WWWROOT. For Netscape/Novonyx, place the image file in NOVONYX\SUITESPOT\DOCS.)

Setup

To set up the iChain User Services, modify the iChainAppConfig.txt configuration file to match your environment. Make sure this configuration file resides in the same file system directory as the servlets.

In order for the servlets to get the user identity and credential information (via the Authorization section of the HTTP header), you must enable the Forward Authentication Information to Web Server authentication option on the iChain Proxy Server for the Web server running these servlets at Configuration > Web Server Accelerator > Modify > Authentication Options.

The servlets in the iChain User Services are optimized for use with iChain 2.1. This means that the work of user authentication is offloaded to the iChain Proxy Server rather than having the servlets themselves perform user authentication. Of the four servlets, iChainAddUser is designed to be configured as an iChain public resource, while the other three servlets (iChainModifyUser, iChainPasswordChange, and iChainPasswordMgr) are designed to be configured as private (restricted or secure) resources.

The three restricted or secure iChain servlets extract the user's identity from the base64-encoded Authorization section of the HTTP header (which is populated by the iChain Proxy Server after the user authenticates).

IMPORTANT: Make sure that the directory path where these *.class files reside on your server are in the runtime CLASSPATH of your servlet engine's JVM. If they are not, the servlets will fail upon initialization because of the `getClass().getResourceAsStream(configFileName)` method call that they make when attempting to read the configuration text file (that should be in the same directory as the servlet *.class files) when the servlets initialize.

Configuring iChainAddUser As a Public Resource

To configure the iChainAddUser servlet as a Public resource while also configuring the other servlets in the same directory as Restricted resources, you will need two Protected Resource entries on the ISO object.

For example, on the ISO object's Protected Resource tab, enter the following:

URL Prefix	Access
<code>http://ichain.novell.com/servlet/iChainAddUser</code>	Public
<code>http://ichain.novell.com/servlet/*</code>	Restricted

This makes the iChainAddUser servlet a Public resource, and all other servlets in that directory are restricted (login/authentication is required before access is given).

Enabling a Password Dictionary

To enable support for a Password Dictionary file (which lists passwords that users are not allowed to use):

- 1** Select the ISO object > Password Policy tab.
- 2** Select the Check Password Dictionary check box.

- 3** In the Password Dictionary File field, list the full HTTP URL where the Dictionary file resides. For example, `http://137.65.215.225/Dictionary.txt`.

The Dictionary file must be a text file, and each word must be on its own line. For example, the contents of your text file might look like the following:

```
A
Aaa
Abc
Ansi
Az
etc.
```

OS/Servlet Engine Environments

The iChain User Services servlets require LDAP class support that is commonly provided by an LDAP.JAR file. You can get the latest copy of this required LDAP.JAR file from Novell's Java LDAP SDK, titled "LDAP Classes for Java". This kit is available on the [Novell Developer Web site \(http://developer.novell.com/ndk/jldap.htm\)](http://developer.novell.com/ndk/jldap.htm).

For convenience, Novell has included this file (renamed to LDAP-NOVELL.JAR) on the *iChain Authorization Server* CD, in the \SERVLETS directory. You can use this LDAP-NOVELL.JAR file without having to install any additional library kit. Simply configure your servlet engine to recognize the LDAP-NOVELL.JAR file, typically by either adding this file to your server's CLASSPATH or by copying it to your servlet engine's LIB directory.

NOTE: The LDAP-NOVELL.JAR file does NOT have to be renamed to LDAP.JAR to be used (filenames do not matter in this case).

Because there are many different OS and servlet engine environments, it would be virtually impossible to document the installation for every environment. Here are a few examples:

NetWare 5.1 Server with Novell Servlet Gateway and NetWare Enterpriser Web Server

In this example, the Novell Servlet Gateway is previously installed. See the Novell Developer site (<http://developer.novell.com>) to download this gateway.

- 1** Verify the functionality of the Novell Servlet Gateway by running a servlet such as SnoopServlet.

- 2** Verify the server has the current support pack.
- 3** Check the version of java.nlm at the Novell NDK site (<http://developer.novell.com/ndk/jvm12.htm>). If it is less than 1.22, it will need to be updated to Novell JVM for NetWare v1.X.X (1.2.2. or higher).
- 4** Install the LDAP Class Libraries for Java on the NetWare server. This will place the ldap.jar file in the sys:/java/lib directory.
- 5** Follow the installation instructions to add a classpath to the sys:\etc\java.cfg file:


```
SERVLETCLASSPATH=$SERVLETCLASSPATH;\java\lib\rt.jar;sys:\java\lib\i18n.jar

SERVLETCLASSPATH=$SERVLETCLASSPATH;\java\lib\ldap.jar

SERVLETCLASSPATH=$
SERVLETCLASSPATH;\java\lib\servgate.jar

SERVLETCLASSPATH=$SERVLETCLASSPATH;\java\lib\jsdk.jar

SERVLETCLASSPATH=$SERVLETCLASSPATH;\java\lib\njgwap.jar
```
- 6** Place the four servlets (iChainAddUser.class, iChainModifyUser.class, iChainPasswordChange.class, iChainPasswordMgr.class) in the SYS:\JAVA\SERVLETS directory.
- 7** Place the ICHAINAPPCONFIG.TXT configuration file in the same directory. Modify it as described in “[Servlet Configuration File](#)” on page 144.
- 8** Place the TOP.GIF image file at the root documents directory (SYS:NOVONYX\SUITESPOT\DOCS).

NetWare 6 Server with Tomcat and Apache

- 1** Verify that Tomcat is functioning properly. This can be done by accessing a servlet in the SYS:\TOMCAT\33\WEBAPPS\ROOT\WEB-INF\CLASSES directory, such as SnoopServlet.
- 2** Install the LDAP Classes for Java on the NetWare server. This will place the ldap.jar file in the sys:/java/lib directory. Follow the install instructions to add a classpath to the SYS:TOMCAT\33\BIN\TOMCAT33.NCF file:

```
envset TOMCAT_CLASSPATH=sys:\java\lib\classes.zip
```

```
envset
TOMCAT_CLASSPATH=$TOMCAT_CLASSPATH;$TOMCAT_HOME\lib\tomcat.jar
```

```
envset
TOMCAT_CLASSPATH=$TOMCAT_CLASSPATH;SYS:\java\njclv2\lib\jndi.jar
```

```
envset
TOMCAT_CLASSPATH=$TOMCAT_CLASSPATH;SYS:\java\njclv2\lib\njclv2.jar
```

```
envset
TOMCAT_CLASSPATH=$TOMCAT_CLASSPATH;SYS:\java\lib\ldap.jar
```

- 3** Place the four servlets (iChainAddUser.class, iChainModifyUser.class, iChainPasswordChange.class, iChainPasswordMgr.class) in the SYS:\TOMCAT\33\WEBAPPS\ROOT\WEB-INF\CLASSES directory.
- 4** Place the ICHAINAPPCONFIG.TXT configuration file in the same directory. Modify it as described above.
- 5** Place the TOP.GIF image file at the root documents directory (SYS:\NOVONYX\SUITESPOT\DOCS).

5.1 Server with Enterprise Web Server: WebSphere 3.53 (running in servlet engine only mode)

- 1** Copy servlets and .txt files to OnDemand/WAS server
SYS:\WEBSHERE\APPSERVER\SERVLETS
- 2** Edit ICHAINAPPCONFIG.TXT to add proper LDAP server address, contexts, usernames, and passwords.
- 3** Copy TOP.GIF to SYS:\NOVONYX\SUITESPOT\DOCS
- 4** Expand LDAP class libraries (this puts the ldap.jar file in SYS:/JAVA/LIB directory).
- 5** Edit STARTWEBSHERESERVLETENGINE.NCF. Add the line:

```
envset WAS_CP=WAS_CP;$JAVA_HOME\lib\ldap.jar
```
- 6** Access the servlets with a URL similar to `http://<hostname>/servlet/iChainAddUser`.

Open Source

Although the iChain User Servlets are functional and usable right out of the box, the full source code to these servlets is included for those who require more functionality or more customization than is normally provided with these core user services.

Custom Login Pages

iChain mandates that the same login page be presented to all users using Novell's iChain/iChain Proxy Server as a front-end proxy with authentication enabled. In order to understand the means by which customer requests to modify the login page can be accommodated, it is beneficial to examine the current implementation.

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

- ♦ user's name
- ♦ distinguished name
- ♦ field name

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

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

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

```
SYS:\ETC\PROXY\DATA\CALOGLDP.HTM
```

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

```
SYS:\ETC\PROXY\DATA\CALOGMA.HTM
```

In order to be able to allow an administrator to specify a custom login page, in the iChain 2.1/iChain Proxy Server Web Administration tool, a field in the setup window exists to specify a subdirectory where the login page for that accelerator can be found. The actual file name will continue to be predetermined by the profile and login name format.

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

```
SYS:\ETC\PROXY\DATA\NIKE\CALOGMA.HTM
```

Currently the designated login pages are:

- ◆ SYS:\ETC\PROXY\DATA\CALOGLDP.HTM — Login page for LDAP profile based on a login page with a login name format of distinguished name or field name.
- ◆ SYS:\ETC\PROXY\DATA\CALOGMA.HTM — Login page for LDAP profile based on a login name format of user's e-mail.
- ◆ SYS:\ETC\PROXY\DATA\CALOGRAD.HTM — Login page for RADIUS authentication profile.

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

Path/Host-Based Multi-homing

The same mechanism that is described in [“Custom Login Pages” on page 150](#) is provided for child accelerators that are part of either path-based or host-based multi-homing.

Limitations

Because the login pages are serviced from memory, this limits the types of graphics that can be supported. All streaming types of graphical/sound widgets (that is, avi, wav, mpeg, Quicktime, etc.) are not supported; however, BMP, JPG, GIF, and other types of clipart graphics are supported.

Coding of Login Pages

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

- ◆ SYS:\ETC\PROXY\DATA\CALOGLDP.HTM — Used for accelerators that have an LDAP profile with login name formats of distinguished with LDAP contexts.
- ◆ SYS:\ETC\PROXY\DATA\CALOGLNC.HTM — Used for accelerators that have an LDAP profile with login name formats of distinguished without LDAP contexts (fully distinguished).
- ◆ SYS:\ETC\PROXY\DATA\CALOGLFN.HTM — Used for accelerators that have an LDAP profile with login name formats of field name.
- ◆ SYS:\ETC\PROXY\DATA\CALOGMA.HTM — Used for accelerators that have an LDAP profile with login name formats of e-mail.
- ◆ SYS:\ETC\PROXY\DATA\CALOGRAD.HTM — Used for RADIUS-based authentication profiles.

SYS:\ETC\PROXY\DATA\CALOGLDP.HTM

The requirements for an LDAP profile login page for login name formats of distinguished with LDAP contexts can be found in CALOGLDP.HTM:

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

SYS:\ETC\PROXY\DATA\CALOGLNC.HTM

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

SYS:ETC\PROXY\DATA\CALOGLFN.HTM

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

SYS:ETC\PROXY\DATA\CALOGMA.HTM

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

SYS:ETC\PROXY\DATA\CALOGRAD.HTM

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

Custom Logout Page

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

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

or

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

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

```
SYS:ETC\PROXY\DATA\CALOGOUT.HTM
```

Custom logout pages are similar to custom login pages.

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

```
set accelerator <acc name> loginpage=<subdir from etc proxy data>
```

For example:

```
set accelerator nike loginpage=nike
```

the proxy will then look for:

SYS:ETC\PROXY\DATA\NIKE\CALOGOUT.HTM

Strong Cryptography

The strong cryptography settings allow the server to be configured to force strong encryption to be used in SSL sessions (as in https). Client mode (when the proxy server initiates the SSL session) and server mode (when the proxy server accepts an SSL session from another machine) can be configured separately. The default is to not force the use of strong cryptography in either mode.

The configuration can be done from the iChain Proxy Server system console using the following commands:

```
set authentication strongserverenable = (yes/no)
```

No — Clients can initiate an SSL session with the proxy server using null, weak, or strong cryptography.

Yes — Clients must initiate an SSL session with the proxy server using strong cryptography, or the session will fail.

```
set authentication strongclientenable = (yes/no)
```

No — The proxy server will initiate an SSL session with another server using any cryptography that server supports (null, strong, or weak).

Yes — The proxy server will only initiate an SSL session with another server using strong crypto; if unsupported by the other server, it will fail.

Applying these settings will store them in the ISO object and create a NILE.CFG file. This file is read by NILE.NLM at startup, so the server must be restarted for these settings to take effect.

Cryptography Settings

- ◆ Null cryptography — No encryption is used, even though SSL has been enabled.
- ◆ Weak cryptography — Encryption is done with key sizes less than 128 bits.
- ◆ Strong cryptography — Encryption is done with key sizes of 128 bits or larger.

Cache Freshness

When first introduced to Web content caching, many network administrators assume that the object cache on an proxy server is basically the same as a browser's cache, which all users access when they click the Back button. The logical extension from this assumption is the fear that iChain Proxy Services will serve stale content that doesn't accurately reflect the fresh content on the origin Web server.

Actually, most time-sensitive Web content is flagged by Webmasters in such a way that it cannot become stale unless a caching system ignores the Webmaster's settings. The proxy server honors all flags that affect cache freshness, including Time to Expire, Don't Cache, and Must Revalidate directives.

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

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

This is administered at Configure > Tuning > Cache Freshness. For more information on configuring the iChain Proxy Services for cache freshness, see [“Managing Cache Freshness” on page 155](#) and [“Cache Freshness Dialog Box” on page 343](#).

Managing Cache Freshness

Cache freshness is a primary concern of most appliance administrators. The following sections briefly explain how your appliance ensures fresh content for network users and the options you have for adjusting this appliance feature.

How the iChain Proxy Server Checks for Object Freshness

Although the following explanation is an over-simplification, it lays the foundation for the specific examples that follow this section.

An iChain Proxy Server has timers that it applies to every cached object.

Each time an object is cached or revalidated, the appliance starts a timer for that object. As long as the timer is running, the appliance will vend the object from cache. After the time has expired and when the appliance receives a request for the object, it will issue an IF-MODIFIED-SINCE request to the origin Web server.

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

If the object has not changed, the iChain Proxy Server vends the object from cache and resets the timer, and the countdown for vending the object from cache begins again.

If a browser forces a refresh of the object, the iChain Proxy Server honors the browser request, retrieves and caches the object regardless of whether it has changed, and restarts the timer.

How a Proxy Server Keeps the Oldest Cached Objects Fresh

More than 80% of all Web objects have either no Time to Expire directives or they are set to stay cached for as long as weeks or even months.

Since many of these objects actually change fairly frequently, the appliance has two timers for ensuring their freshness. You can configure these timers in the **Cache Freshness Dialog Box**, administered at Configure > Tuning > Cache Freshness.

HTTP Maximum: This timer overrides an object's Time to Expire settings if it is longer than the timer's value.

The default timer value is six hours. This means that iChain Proxy Services will not vend an object that has been in cache longer than six hours without first checking whether it should be refreshed.

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

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

How the iChain Proxy Server Handles the Freshest Objects in Cache

Most Webmasters ensure that their time-sensitive objects have appropriate Time to Expire directives. Late-breaking news stories and photographs, for example, might stay in cache for only a few minutes before expiring.

By default, the proxy server simply honors the Webmasters' instructions and revalidates the objects in cache as directed.

However, some appliance installations, such as those connected through a modem, might need to limit how often these objects are refreshed. The appliance has a third timer for this purpose, also accessible in the **Cache Freshness Dialog Box**.

HTTP Minimum: This timer sets the minimum number of hours or minutes the proxy server will serve HTTP data from cache before revalidating it against content on the origin Web server. No requested object will be revalidated sooner than specified by this value.

The default value for this timer is 0, meaning that the proxy server honors the Time to Expire directive for each object (assuming, of course, it is not longer than the HTTP Maximum timer).

If the timer is set to a value other than 0, it then overrides any object's Time to Expire directive that is shorter than the value set.

Fine-Tuning Cache Freshness on Your Appliance

The default timer settings explained in the previous sections are tuned for most appliance installations. However, you might have special requirements you have that need the default settings to be adjusted.

Perhaps you are accelerating content that doesn't contain Time to Expire directives but changes frequently and needs to be refreshed more often than every two hours. You can adjust the HTTP Default timer in the **"Cache Freshness Dialog Box"** on page 343 so that iChain Proxy Services refreshes the objects more frequently.

Perhaps you have severe Internet bandwidth restrictions and an environment with users who don't require object freshness checks every six hours. You can adjust the HTTP Maximum timer in the **"Cache Freshness Dialog Box"** on page 343 to a different setting that meets your requirements and conserves bandwidth.

If you choose to adjust the timer values, avoid settings that result in objects refreshing more often than is necessary. Otherwise you could negate the bandwidth and response-time benefits of having the appliance on your network.

Using Custom Cache Control Headers

In addition to fine-tuning cache freshness with the system's global HTTP timers as explained in [“Fine-Tuning Cache Freshness on Your Appliance” on page 157](#), you can configure each proxy service to recognize custom headers in HTTP packets. Your Web server can then use these headers for transmitting caching instructions that only the configured appliance services will recognize and follow.

How Headers Work

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

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

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

- ◆ You can specify that browsers and external caches cannot cache the objects, but the accelerator can.
This lets you offload request processing from the origin Web server while still requiring that users return to the site each time they request an object.
- ◆ You can also specify separate cache times for browsers, external caches, and the accelerator you are defining.

Implementing Custom Cache Control Headers

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

- ◆ Enter a header string in the Custom Cache Control Header dialog box, for example, MYCACHE.

- ◆ Configure the Web server to send an HTTP header containing the defined string and the time in seconds that the object should be retained in cache, for example, MYCACHE: 60.

If the number is non-zero, the proxy server treats the reply as if it had the following headers:

```
Cache-Control: public
Cache-Control: max-age= number
```

If the number is zero (0), iChain Proxy Services treats the reply as if it had the following header:

```
Cache-Control: no-cache
```

- ◆ Ensure that the Web server continues to send standard HTTP cache-control headers so that browsers and external caches follow the caching policies you intend them to.

For example, you could do the following:

- ◆ Use an Expires or Cache-Control: Max-Age header to specify that browsers should cache an object for two minutes
- ◆ Use a Cache-Control: Private header to prevent external caches from caching the object at all
- ◆ Use a custom cache control header, such as MYCACHE: 1800, to indicate that the accelerator should cache the object for 30 minutes.

Custom Cache Control Headers override the following standard HTTP cache-control headers on the appliance but 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
```

For example, you might do the following:

1. While configuring a Web server accelerator service, you insert a string in the Custom Cache Control Header list with the value of FOOTTL.

The appliance will now recognize FOOTTL as a custom cache control header on objects requested through the service you are configuring.

2. You then configure the accelerated Web server to send FOOTTL: 600 in the headers of objects you want to be cached at the appliance.

The appliance will recognize this header as overriding the standard HTTP cache-control headers listed above when objects are requested through the accelerator service you are configuring.

3. Finally, you ensure that the Web server continues to send the following standard HTTP cache-control headers:
 - ◆ Cache-Control: Max-Age headers that cause browsers to cache objects for no longer than two minutes
 - ◆ Cache-Control: Private headers that cause external caches to not cache the objects.

When your Web server sends an object with the FOOTTL header in response to an appliance request made through the accelerator service, your appliance recognizes the custom header and caches the object for 10 minutes. Requesting browsers cache the object for only two minutes. External caches do not cache the object.

Thus, the appliance offloads a processing burden from the Web server by caching the frequently requested objects for 10 minutes (the value you specified in Step 2). Browsers, on the other hand, must always access the appliance to get the objects if their previous requests are older than two minutes. And the objects in the appliance's cache are kept fresh due to their relatively brief time-to-live value.

Managing Appliance Security Features

This section contains the following topics:

- ◆ [“Using the Console Lock Feature” on page 160](#)
- ◆ [“Accessing Proxy Internals” on page 161](#)

Using the Console Lock Feature

The iChain Proxy Services console is locked by default to prevent unauthorized access. The password to unlock the console is the Config user password you specified during the initial configuration.

To use the command line interface, you must unlock the console by entering the following command:

```
unlock
config_user_password
```

NOTE: If a config_user_password is not set, the password is null.

Once the console is unlocked, it remains unlocked until you lock it using the lock command.

Accessing Proxy Internals

A few iChain features require the administrator to access the internals of the iChain Proxy Server.

WARNING: Changes to the internal proxy server configuration should be limited to those items specified in authorized iChain documentation or as directed by Novell support personnel. Undocumented changes may result in proxy server malfunction and may require Novell support personnel to request a software re-image for that server.

To access the internals of the proxy server, enter the following command at the proxy server console:

```
debug
proxy_debug_password
```

NOTE: The proxy_debug_password is "proxydebug".

Automatic Configuration Mechanisms

To	See
Learn about appliance configuration files	“About Appliance Configuration Files” on page 162
Learn about the three methods of managing configuration files	“Managing Configuration Files” on page 165
Save appliance configurations	“Using Customized Configuration Files to Change the System Configuration” on page 164
Change the current appliance configuration	“Using Customized Configuration Files to Change the System Configuration” on page 164
Back up the appliance configuration	“Backing Up the Appliance Configuration” on page 166

To	See
Configure multiple appliances	“Creating Appliance Configuration Shortcuts” on page 167
Restore the original factory settings	“Restoring Factory Settings” on page 167
Change the clone image. The appliance uses this to restore the system if it senses the system has become unstable.	“Restoring the Appliance to the Clone Image” on page 168
Reimage the appliance	“Reimaging and Restoring the Appliance System” on page 169

About Appliance Configuration Files

Configuration files are ASCII text files that store the command line syntax used to configure the appliance. Each line in the file represents a single configuration command. When you use the browser-based management tool, the system generates multiple commands in the correct order to cause the configuration changes you specify. These commands are then recorded, in the correct sequence, in configuration files on the appliance.

The following is a sample from a configuration file with the missing portion indicated by the ellipsis (...).

```
set eth1 name=eth1
set eth1 speed=default
set eth1 duplex=default
clear eth1 address
add eth1 address=10.1.1.2,mask=255.255.255.0
set eth0 name=eth0
set eth0 speed=default
set eth0 duplex=default
clear eth0 address
add eth0 address=10.1.1.1,mask=255.255.255.0
set floppy poll=no
set floppy interval=120
```

```
set floppy saveonapply=no
. . .
apply
```

System-Generated Configuration Files

The iChain Proxy Server employs four configuration files: FACTORY.NAS, CURRENT.NAS, AUTOLOAD.NAS, and iChain.NAS.

FACTORY.NAS

This file contains the appliance configuration as it came from the factory. This is a system file that is never modified.

CURRENT.NAS

This file contains the appliance's current configuration settings since the last apply command was issued.

You can view this file in the browser-based management tool if you are interested in seeing all of the commands used to create the current appliance configuration. To view the file, click System > click Import/Export > select Current under Configuration Files on Appliance > click Download.

AUTOLOAD.NAS

This file is saved by iChain Proxy Services whenever a floppy disk is in the appliance's floppy drive and automatic polling is enabled.

NOTE: System monitoring of the AUTOLOAD.NAS file is enabled by default with a polling interval of 30 seconds. You can change these settings in the browser-based management tool > Import/Export tab. See ["Import/Export Tab" on page 291](#).

By default, the AUTOLOAD.NAS file contains the appliance's configuration settings since the last apply command was issued.

After the system is reimaged and after the clone image is applied, iChain Proxy Services checks the floppy disk for an AUTOLOAD.NAS file. If the file is found, iChain Proxy Services immediately applies the commands it contains.

AUTOLOAD.NAS is useful in two situations:

- ♦ If you have reimaged the appliance, you can quickly configure it by inserting a floppy disk containing an AUTOLOAD.NAS file you have previously saved. For more information, see [“Reimaging and Restoring the Appliance System” on page 169](#).
- ♦ If you want to apply specific configuration settings, such as a filtering configuration, you can save only these settings in an AUTOLOAD.NAS file and have the system automatically apply the settings when you insert the diskette.

IMPORTANT: Remember that the last command in AUTOLOAD.NAS is always the apply command, which causes the system to immediately update AUTOLOAD.NAS.

You can prevent automatic updating of AUTOLOAD.NAS by opening the write-protect tab on the floppy disk after the desired AUTOLOAD.NAS file is in place, or by setting the file properties of AUTOLOAD.NAS to read-only on a different workstation.

Using Customized Configuration Files to Change the System Configuration

In addition to using system-level configuration files, iChain Proxy Services lets you save the appliance’s current configuration to arbitrarily named .NAS files and apply them to the system later.

The Import feature lets you save backup copies of the appliance configurations you have created, and the Export feature lets you quickly apply any previously backed-up configuration to the appliance.

For more information about importing and exporting configuration files, see [“Managing Configuration Files” on page 165](#), [“Backing Up the Appliance Configuration” on page 166](#), and [“Import/Export Tab” on page 291](#).

Configuration files have an 8.3 DOS-style filename, the last three characters of which must be NAS.

You can save the configuration settings on the appliance or to a floppy disk on the appliance through the browser-based management tool, Telnet, and the command line interface. You can then quickly reconfigure the appliance using the configuration files.

IMPORTANT: We recommend storing copies of your customized configuration files on a floppy disk. This ensures that you have the files if the clone image is ever applied or the appliance is ever reimaged.

Two situations where having customized configuration files is an advantage are summarized in [“Backing Up the Appliance Configuration” on page 166](#) and [“Creating Appliance Configuration Shortcuts” on page 167](#).

Managing Configuration Files

You can manage appliance configuration files using the browser-based management tool, Telnet, or the command line interface, and using the appliance's FTP functionality. The next three sections briefly explain how to use each of these management options.

Using the Browser-Based Management Tool

You can export and import configurations and manage the creation of the autoload configuration from the browser-based management tool. For more information, see [“Import/Export Tab” on page 291](#).

Using Telnet or the Command Line

From Telnet or a command line, you can import and export configuration files. Do *not* specify the three-digit NAS extension when using either of these methods.

If You Want To	Then Enter	Notes
Apply an autoload file from a floppy	<code>import floppy</code>	First verify that the disk containing AUTOLOAD.NAS is inserted into the appliance.
Export a named configuration file to the appliance's hard drive	<code>export filename</code>	<i>Filename</i> is the name of the configuration file without the .NAS extension specified.
Export a named configuration file to a floppy	<code>export filename, floppy</code>	The file will be saved on the DOS-formatted floppy disk inserted into the appliance.
Apply a named configuration file from the appliance's hard drive	<code>import filename</code>	<i>Filename</i> is the name of the configuration file without the .NAS extension.
Apply a named configuration file from a floppy	<code>import filename, floppy</code>	The file will be loaded from the DOS-formatted floppy disk inserted into the appliance.

Using FTP

You can use FTP to move the configuration files to and from the appliance using the get and put commands. You can also apply a configuration file you are moving by using the execute option specified after a comma on the command line.

After starting the FTP client and pointing it to an IP address for the appliance (see [“Starting an FTP Session with the Appliance” on page 199](#)), use one of the following commands, where *filename* is the name of your configuration file:

Command	Description
<code>get filename.nas</code>	Downloads the specified configuration file to your FTP local directory on your client workstation
<code>put filename.nas</code>	Uploads the specified configuration file from the FTP local directory to the appliance
<code>put filename.nas,execute</code>	Uploads the configuration file and applies it to the proxy server

Backing Up the Appliance Configuration

Because the backup configuration file, AUTOLOAD.NAS, is automatically updated by the caching system each time any configuration change is applied (see [“AUTOLOAD.NAS” on page 163](#)), we recommend saving your appliance configuration to a .NAS file with another name.

If you ever need to reimage your proxy server for some reason, having an alternatively-named .NAS file will provide a configuration backup in case the AUTOLOAD.NAS file is overwritten with factory settings before your configuration is restored.

If you want to automatically load a configuration using the settings in an alternatively-named file, you can use a workstation and replace the system-created AUTOLOAD.NAS file with a copy of your file on the floppy disk before inserting it into the appliance's floppy disk drive.

Creating Appliance Configuration Shortcuts

You might want to have more than one configuration for an appliance, depending on business or other conditions. An alternate method to manually reconfiguring the appliance is to save various configurations in separate configuration files and use these to turn services on and off through FTP services. For example, you could use two files named FORWARD.NAS and REVERSE.NAS to quickly configure the appliance to provide the services indicated by the filenames.

Restoring Factory Settings

You can quickly return the appliance to its original factory settings from the browser-based management tool, a Telnet session, or the command line. After restoring factory settings, you must either reinitialize the appliance as described in the Initial Installation Guide or use a previously created AUTOLOAD.NAS file on a floppy diskette to restore the appliance's configuration settings.

An appliance's original factory settings include the following:

- ◆ The eth0 network adapter is bound to IP address 10.1.1.1 on subnet 10.1.1.0 with a subnet mask of 255.255.255.0.
- ◆ Other network adapters have no addresses bound.
- ◆ No caching, proxy cache, caching hierarchy, filtering, or other appliance services are configured.

WARNING: Restoring factory settings removes all the settings you have configured except passwords. This includes network addresses and all appliance cache services.

In most cases, you can automatically restore the settings if you have prepared an AUTOLOAD file on a floppy disk. See ["AUTOLOAD.NAS" on page 163](#) and ["Import/Export Tab" on page 291](#).

You should also prepare an alternatively-named backup configuration file as a precaution. For further details, see ["Backing Up the Appliance Configuration" on page 166](#).

From the Browser-Based Management Tool:

- 1** Click System > Actions > Factory Settings.
- 2** Restore factory settings by clicking Restore.

or

Cancel the action by clicking Do Not Restore.

From a Telnet Session or the Command Line:

- 1** At the system prompt, enter
factorysettings
- 2** Do one of the following:
Restore factory settings by entering **apply**.
or
Cancel the action by entering **cancel**.

After restoring factory settings, you must either reinitialize the appliance as described in the Initial Installation Guide or use a previously created AUTOLOAD.NAS file on a floppy diskette to restore the appliance's configuration settings. See [“AUTOLOAD.NAS” on page 163](#) and [“Import/Export Tab” on page 291](#).

Restoring the Appliance to the Clone Image

Each appliance stores a clone image that, initially, is the same as the factory image. If the appliance experiences an abnormal shutdown four times within a half hour period, or if it is restarted six times within a half-hour period, the appliance assumes the current configuration is faulty and automatically replaces it with the clone image.

If the default factory image is restored, you must either reinitialize the appliance using the instructions in the Initial Installation Guide, or, if you have saved the appliance configuration, you can use an AUTOLOAD.NAS file to restore the configuration. See [“AUTOLOAD.NAS” on page 163](#).

To prevent automatic restoration to the default factory settings in the event of system problems, you can overwrite the default clone image after you have applied an alternate configuration to the appliance. You can also apply the clone image as an alternate method for reconfiguring the appliance. For more information, see [“Actions Tab” on page 285](#).

IMPORTANT: You should update the clone image whenever you perform an upgrade. Be aware, however, that this process causes the appliance to reboot, resulting in a temporary interruption of services. See [“Upgrade Tab” on page 294](#) and [“Actions Tab” on page 285](#).

Reimaging and Restoring the Appliance System

The appliance comes with a CD that can be used to reimage the system. This reformats the hard disks and reinstalls the appliance system. After reimaging an appliance, you must either reinitialize it as described in the Initial Installation Guide or use a previously created AUTOLOAD.NAS file to restore your configuration settings.

WARNING: Reimaging the system removes all the settings you have configured, including passwords, network addresses, and all cache services.

In most cases, you can automatically restore the settings if you have prepared an AUTOLOAD file on a floppy disk. See [“System-Generated Configuration Files” on page 163](#) and [“Import/Export Tab” on page 291](#).

You should also prepare an alternatively-named backup configuration file as a precaution. For further details, see [“Backing Up the Appliance Configuration” on page 166](#).

To reimage and restore an appliance using an AUTOLOAD.NAS file:

- 1 Locate the appliance system CD.

IMPORTANT: If the system CD is in the appliance, remove the CD, shut down the appliance, recycle the appliance's power switch and allow the appliance to restart.

- 2 If the appliance configuration has not been previously saved, insert a formatted, blank floppy disk into the appliance's floppy disk drive.

If you have previously saved the appliance configuration, skip to [Step 7](#).

- 3 If you have access to the appliance through the browser-based management tool, click System > Import/Export; otherwise, continue to [Step 5](#).

If there is an AUTOLOAD file on the floppy disk, skip to [Step 6](#).

- 4 If you need to create an AUTOLOAD file on the floppy disk, type **autoload** in the Export Configuration File to Floppy field > click Export To > skip to [Step 6](#).

- 5 If you do not have appliance access through the browser-based management tool, establish a Telnet or null-modem session with the appliance. (You can also use an attached keyboard and monitor if your appliance has the required connections.)

At the appliance command line, enter the following:

```
export autoload floppy
```

An AUTOLOAD.NAS file is created on the floppy disk.

- 6** Remove the floppy disk from the appliance.
- 7** Insert the appliance system CD into the CD drive.
- 8** Turn the appliance's power switch off, wait a few seconds, and then turn the appliance back on.
The CD automatically launches and the appliance reinitializes.
- 9** After the initialization process starts, insert the configuration diskette with the AUTOLOAD.NAS file into the appliance.
- 10** After all disk activity ceases and the system prompt appears, remove the appliance system CD and the floppy disk.
- 11** Shut down the appliance, recycle the appliance power switch, and wait for the system prompt to appear or for the start-up beep sequence to sound.

The appliance should now be restored to its previous operating configuration.

DNS Name Resolution

As iChain Proxy Services processes browser requests, it uses the DNS system to obtain the IP addresses of origin Web servers.

Since the DNS names in browser requests are not always straightforward, the proxy server tries various permutations to try and locate the Web server. As a result, DNS names ending with domain extensions other than .com, .org, and so on, are sometimes resolved in unexpected ways.

If users of your appliance are experiencing this problem, you can customize how the appliance resolves DNS names.

How the Appliance Resolves DNS Names

When the appliance receives a browser request, it creates a DNS query based on the URL in the request and sends the query to one of the DNS name servers defined for the appliance.

How the Appliance Formulates Subsequent DNS Queries

If the DNS name server can't resolve the query, the appliance formulates subsequent DNS queries based on the following:

170 Place Book Title Here

- ◆ The appliance's domain name
- ◆ The appliance's R_APPEND.CFG file

For example, assume the following:

- ◆ The browser request URL is webserver.
- ◆ The appliance's domain name is acme.com.
- ◆ The appliance's R_APPEND.CFG file has the following content:

```
www.%s.com
www.%s.ed
www.%s.org
www.%s.gov
www.%s.net
%s.com
%s.edu
%s.org
%s.gov
%s.net
www.%s
```

After the initial request fails, the appliance formulates subsequent requests as follows:

1. The appliance formulates a second query by appending the appliance's domain name to the URL as follows:
webserver.acme.com
2. If this query fails, the appliance appends the appliance's subdomain name to the URL as follows:
webserver.com
3. If this query fails, the appliance appends each entry in the R_APPEND.CFG file in the order listed until one of the following occurs:
 - ◆ The DNS server returns an IP address for the name.
 - ◆ The appliance's query options are exhausted and it returns a DNS error to the browser.
4. If a DNS name has already been tried, the appliance skips the query and moves to the next item in the list.

Continuing with the example, the appliance would submit the following queries, substituting webserver for the %s variable in the lines of the R_APPEND.CFG file.

www.webserver.com
www.webserver.edu
www.webserver.org
www.webserver.gov
www.webserver.net
webserver.edu
webserver.org
webserver.gov
webserver.net
www.webserver

Because webserver.com was tried previously, the appliance skips the sixth line (%s.com) in the R_APPEND.CFG file.

Modifying the R_APPEND.CFG File

To modify the R_APPEND.CFG file:

- 1 Start an FTP client on a workstation with access to the appliance.
For help, see [“Starting an FTP Session with the Appliance” on page 199](#).
- 2 Point the FTP client to one of the appliance’s IP addresses.
- 3 Enter the following command:

```
get /etc/proxy/appliance/config/user/r_append.cfg
```

The file is transferred to the FTP client's default directory.
- 4 Referring to the example in [“How the Appliance Formulates Subsequent DNS Queries” on page 170](#), modify the R_APPEND.CFG file using an ASCII editor.
Ensure that the lines in your file reflect the query order and content you want the appliance to use when attempting DNS name resolution. For example, you might want to reorder the domains listed or include two-letter country codes in the list.
- 5 Use the put command to place the modified R_APPEND.CFG file back in \ETC\PROXY\APPLIANCE\CONFIG\USER on the appliance.
- 6 Restart the appliance.

Dynamic Bypass

The Dynamic Bypass feature lets you configure the appliance so that specific errors from Web sites are explicitly not cached and subsequent requests to the Web sites are simply passed through for a specific time period.

For help configuring the Dynamic Bypass feature, see [“Enable Dynamic Bypass” on page 341](#).

Why Dynamic Bypass Is Needed

The appliance follows the HTTP 1.1 specification, which stipulates that no 400- or 500-series errors are cached (except HTTP 410: Data Moved Permanently). If these errors occur, the proxy server passes them through to the requesting browsers.

Some Web servers erroneously pass cache control public headers along with the errors they generate. These headers override the default error handling by the caching system and cause the error to be cached on the appliance.

For example, if a Web server receives an unauthorized request for data and responds with an HTTP 403: Forbidden error accompanied by a cache control public header, the error will be cached. The appliance will then respond to all subsequent requests, including those from authorized users, with the 403 error.

How the Proxy Server Handles Dynamic Bypass

If dynamic bypass is enabled for a given error and iChain Proxy Services receives the error from an origin Web site, the site’s URL is dynamically added to the bypass list and retained in the list for the period of time specified. All subsequent requests to that URL during the dynamic bypass time period are passed through to the Web site and do not go through the cache.

To continue with the above example, if the appliance has dynamic bypass enabled and the 403 error checked, it would not cache the error but pass it directly back to the browser. It would also add the site’s URL to its dynamic bypass list. On subsequent requests to the same URL, iChain Proxy Services would bypass cache and passthrough requests to the Web site for the period of time specified. This lets the Web server determine whether to accept or reject access requests on a case-by-case basis.

Appliance Error Messages

The appliance lets you specify a language for the error messages it sends to browsers. This section explains appliance error message support and provides instructions to modify error message text and create support for additional languages.

The appliance provides error messages to browsers through a set of language-specific directories, each of which contains two files.

When you select a language in the browser-based management tool, you are actually selecting one of these language-specific directories and the files it contains.

The two files the appliance uses are:

- ◆ ERRPAGE.CFG, which contains the text of all appliance error messages
- ◆ PXYERR.HTM, which is the HTML template file that applies a format to the applicable error text and other error information and is sent to the receiving browsers

The language-specific directories that contain these two files are located in `\ETC\PROXY\DATA\ERRPAGE\NLS\LANGUAGE`, where LANGUAGE is the English name of the language.

For example, the English files are stored in `\ETC\PROXY\DATA\ERRPAGE\NLS\ENGLISH`.

Other common appliance error message directories include:

```
\ETC\PROXY\DATA\ERRPAGE\NLS\GERMAN
\ETC\PROXY\DATA\ERRPAGE\NLS\SPANISH
\ETC\PROXY\DATA\ERRPAGE\NLS\PORTUGUESE
\ETC\PROXY\DATA\ERRPAGE\NLS\FRENCH
\ETC\PROXY\DATA\ERRPAGE\NLS\JAPANESE
```

You can specify the language for error page vending in `Configure > Mini Web`. For more information, see [“Mini Web Tab” on page 346](#).

Checking the Language Directories on Your Appliance

To see a list of the directories on your appliance:

- 1 In the browser-based management tool, click Configure > Mini Web > the drop-down list.

Customizing the Appliance Error Template and Message Files

You can create error message support for additional languages and customize existing error message text and format.

Creating a New Language

- 1 Start FTP and log in as explained in [“Starting an FTP Session with the Appliance” on page 199](#).

- 2 Enter the following:

```
get /etc/proxy/data/errpage/nls/english/pxyerr.htm
get /etc/proxy/data/errpage/nls/english/errpage.cfg
```

- 3 Modify the ERRPAGE.CFG file.

There are explicit instructions in the file that clearly indicate which parts can be translated.

You cannot delete or add error messages, nor can you change the number or order of the messages.

- 4 Modify the PXYERR.HTM file.

Because this is an HTML file, you can customize it in a variety of ways. To interface with the appliance's message delivery mechanisms, you must ensure the following:

- ♦ The keywords <PROXY_ADDRESS>, <ERROR_STATUS>, and <ERROR_DESCRIPTION> must be retained because they are dynamically replaced with the information that their names imply.
 - ♦ Graphics that you add should be put in the SYS:\ETC\PROXY\DATA directory. References to these graphics must use the <PROXY_ADDRESS> keyword as a starting reference point. See the usage of ALERTBAR.GIF in the PXYERR.HTM file.
- 5 Using FTP, create a new language directory in the path given in [“Appliance Error Messages” on page 174](#).

- 6 Enter the following, where *language* is the new language directory you created:

```
put pxyerr.htm /etc/proxy/data/errpage/nls/language/
errpage.cfg
```

```
put pxyerr.htm /etc/proxy/data/errpage/nls/language/  
pxyerr.htm
```

The new language is dynamically available in the browser-based management tool and from the command line. You do not need to restart the appliance.

Customizing the Error Message Text of an Existing Language

Referring to the procedure in [“Creating a New Language” on page 175](#) for details, complete the following basic steps:

- 1 Open the ERRPAGE.CFG file from an existing language-specific directory.
- 2 Modify the file.

IMPORTANT: There are limitations on what you can change in this file. See [“Creating a New Language” on page 175](#) for details.

- 3 Replace the file when modifications are completed.

Customizing the Error Message Format of an Existing Language

Refer to the procedure in [“Creating a New Language” on page 175](#) for details, then complete the following basic steps:

- 1 Open the PXYERR.HTM file from an existing language-specific directory.
- 2 Modify the file.

IMPORTANT: There are limitations on what you can change in this file. See [“Creating a New Language” on page 175](#) for details.

- 3 Replace the file when modifications are completed.

Logging

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

iChain offers the following logging services:

- ◆ You can turn on logging for reverse proxy as well as for URL filtering.

- ◆ You can have the appliance automatically download files to an FTP server and automatically delete downloaded files.
- ◆ You can control the deletion of old log files based on an older-than-*x* time period or the number of log files in the system.

The appliance can create logs using both the common and extended log formats. A wide variety of tools exists for manipulating and processing these files.

Using Appliance Logging

iChain Proxy Services provides a high performance proxy cache system capable of handling thousands of transactions per second. Although the iChain Proxy Server can log extensive details of each transaction, and although the disk space reserved for log files is quite generous on most appliances, if transaction volume is high and log entries consume a few hundred bytes each, iChain Proxy Services can fill up the available disk space in a matter of minutes.

This section explains how appliance logging works and presents management options you can use to ensure optimal use of the available log file disk space and timely migration (downloading) of log files to other storage devices.

What the Appliance Can Log

The following table shows the transactions the appliance can log and the formats available for each service type.

Service	Common	Extended
Web Server Accelerator	Yes	Yes
Dynamic Bypass	Yes*	No

* These common log formats differ from the industry standard proxy cache common log format.

The Costs of Logging

Performance

Turning on logging for a given service increases system overhead and causes some degradation of performance. Therefore, logging should be used only

when service transactions must be tracked for customer billing purposes or other compelling reasons.

Disk Space

Transaction volume and log entry size can cause available log disk space to fill up quickly. Proxy cache disk space is unaffected by log files.

See “[Planning Step 3: Calculating Log Rollover Requirements](#)” on page 181 for formulas you can use to estimate how quickly your logging disk will fill.

System Constraints for Logging

To plan a logging strategy you must know the capacity and limitations of your appliance.

Preset Disk Space

Logging disk space is not user-configurable; it is preset to 1 gigabyte on the iChain Proxy Server. Plan to download and delete log files before the disk space is filled.

Rolling Over Log Files Before Deletion

iChain Proxy Services will not allow the deletion of active log files (files that are currently in use by the caching system). Only log files that have been rolled over and closed can be deleted.

You can ensure that there are closed files on the system by scheduling regular rolling over of log files. During each rollover, the current log file is closed and a new log file is opened.

You must plan for log files to roll over on a schedule that coincides with your download and deletion schedule. This is to ensure that there is at least one closed log file per service when the download and delete cycle starts.

NOTE: Although you can download active log files, this is normally useful only for periodic administrative checks.

Active files contain only the transaction data up to the moment of the download and are incomplete from customer-billing and other business standpoints.

Logging Ceases When the Logging Disk Is Full

When the appliance encounters a log disk full condition, it stops logging and closes all active log files. Information that would have been logged after that point is lost. Other appliance functions continue without interruption.

Log Filenames

The appliance automatically generates log filenames as follows:

- ◆ Six numbers representing the year, month, and day the file was created
- ◆ A dash separating the date from a single letter identifier
 - NOTE:** The dash is not included after the letters double.
- ◆ Letter identifiers running from A through ZZ

This naming convention accommodates up to 702 log files per service per day. If the rollover options are set so that all the possible filenames are used in one day, the log file with the ZZ letter identifier is not closed until the start of the next day (unless the logging disk becomes full).

Log Rollover Options

Appliance log rollover options let you specify when the appliance closes active log files and opens new files so that the closed files can be downloaded and deleted.

Because of the limitations explained in this section, it is essential that you develop a solid log file rollover plan. This will ensure that your appliance doesn't run out of logging disk space or overwrite log files before they are downloaded and deleted.

You can have the appliance roll log files over according to time or file size as explained in the following table:

Option	Considerations
Roll Over by Time	<p>If you plan to download and delete older log files at a set time, you must configure the appliance so that at least two log files exist per service at the time you've scheduled for downloading and deleting. One file will be active, the other closed and ready for download and deletion.</p> <p>For example, if you determine that your log disk space will fill every 12 hours, then you must configure the appliance to roll the log files over in intervals of less than 6 hours, so at least one log file per service is closed and ready to be downloaded and deleted.</p>

Option	Considerations
Roll Over by Size	<p>If you aren't certain how long it will take to fill your appliance's logging disk space, you can roll the log files over by size.</p> <p>For example, you might be logging transactions for three services and have a log volume size of 6 GB. Because you must have at least two log files per service before the disk space fills, each log file must be smaller than 1 GB when the appliance rolls it over.</p>

Planning Your Logging Strategy

As explained in “[The Costs of Logging](#)” on page 177 and “[System Constraints for Logging](#)” on page 178, logging of caching transactions involves system and maintenance overhead. If your situation requires logging, you should plan carefully so that the information you are tracking aligns with specific requirements. This will ensure optimal use of appliance resources.

Because logging requirements and transaction volume vary widely, it is impossible to make recommendations regarding specific logging strategies.

The following sections step you through the logging strategy planning process. We recommend you record the information you gather on a planning sheet of some kind.

Planning Step 1: Determining Your Logging Requirements

To plan a logging strategy, you should first determine the requirements behind the need for logging. We recommend you complete the following steps:

- 1** Identify the business and other reasons for tracking service transactions.
Possible examples might include customer billing requirements, statistical analysis, growth planning, etc.
- 2** Determine which services you need to track.
- 3** Record this information for further reference.

Planning Step 2: Selecting a Log File Format and Optimizing the Log Entry Size

If you use the common log format, log entry size is fixed. If you use the extended log format, log entry size depends on the number of log fields selected.

Complete the following steps for each service you need to track:

- 1** Referring to the “[Log Options Dialog Box](#)” on page 319, record which log fields must be tracked for each service to be logged.
- 2** Carefully scrutinize the information you plan to track to ensure that the log data collected is essential.

For example, if you select URI, also selecting URI-STEM and URI-QUERY would be redundant since URI = URI-STEM + URI-QUERY. Also, logging cookie information can consume a lot of space and might provide little, if any critical information.

A few bytes can add up quickly when the appliance is tracking thousands of hits every second.

Planning Step 3: Calculating Log Rollover Requirements

As explained in “[Log Rollover Options](#)” on page 179, you can have the appliance roll over log files based on time or on size, but not both.

If you already know which option you want to use, scan this section and then complete the calculations pertinent to your choice.

If you don't know which option best matches your situation, completing the calculations in this section should help you decide.

Variable Definitions

The following variables are used in the formulas:

logvol_size: The total disk capacity reserved for log files on your appliance.

The default size is 1 gigabyte.

logentry_size: The average log entry size.

You can determine this by configuring your appliance to track the required information, generating traffic to the appliance, downloading the log files, determining how large each entry is, and calculating the average.

request_rate: The peak rate of requests per second.

You can estimate this rate or place your appliance in a service and get more accurate data by accessing the browser-based management tool's Monitoring tabs.

num_services: The number of services for which you plan to enable logging.

logs_per_service: The number of log files, both active and closed, that you want the appliance to generate for each service before the disk fills.

You must plan to have at least two logs per service before the disk is filled. See [“Rolling Over Log Files Before Deletion” on page 178](#).

Calculating DISKFULL_TIME

Using the following formula, you can calculate how long it will take the appliance to fill your logging disk space:

```
diskfull_time seconds = logvol_size / (request_rate *  
    logentry_size * num_services)
```

For example, if you assume the following:

- ◆ *logvol_size* = 1 GB
- ◆ *request_rate* = 1000 requests per second
- ◆ *logentry_size* = 1 KB
- ◆ *num_services* = 1

Then *diskfull_time* = (1 GB) / (1000 * 1KB * 1) = 1048 seconds (17.47 minutes).

The logging disk space will fill up every 17.47 minutes.

If this time is too short, you must reduce the log entry size by configuring the appliance to log less information per transaction. This is because you can't increase the disk space or limit the requests being logged.

To calculate the *diskfull_time* for your appliance:

- 1** Determine the values of the four variables listed in the bullet list above.
For more information, refer to [“Variable Definitions” on page 181](#).
- 2** Using the *diskfull_time* formula, calculate how often you can expect your logging disk to fill, then use the result in [Calculating MAX_ROLL_TIME](#).

Calculating *MAX_ROLL_TIME*

Using the following formula, you can calculate the maximum roll-over time value you should specify in the Rollover Every field of the Log Options dialog box.

$$\text{max_roll_time} = \text{diskfull_time} / \text{logs_per_service}$$

For example, if you assume the following:

- ♦ *diskfull_time* = 12 hours
- ♦ *logs_per_service* = 2

Then *max_roll_time* = 12 / 2 = 6 hours.

If you roll your logs over by time intervals, the maximum time should be less than six hours. Otherwise scheduling the download and deletion of log files is much more complicated and the window in which this can be done is narrower.

To calculate the *max_roll_time* for your appliance:

- 1** Determine how many log files you want the appliance to generate per service before log space fills.
The minimum number is two.
- 2** Using the *max_roll_time* formula and the *diskfull_time* value obtained in “Calculating *DISKFULL_TIME*” on page 182, calculate how often you should have the appliance roll over the log files.
- 3** Record the *max_roll_time* result on your planning sheet.

Calculating *MAX_LOG_ROLL_SIZE*

Using the following formula, you can calculate the maximum log file size you should specify in the Rollover When File Size Reaches field of the Log Options dialog box.

$$\text{max_log_roll_size} = \text{logvol_size} / (\text{num_services} * \text{logs_per_service})$$

For example, if you assume the following:

- ♦ *logvol_size* = 600 MB
- ♦ *num_services* = 2
- ♦ *logs_per_service* = 3

Then $\text{max_log_roll_size} = 600 \text{ MB} / (2 * 3) = 100 \text{ MB}$.

If you roll your logs over when they reach a specific size, the file size must be no more than 100 MB. Otherwise, the system will run out of disk space before you have three complete log files and scheduling the download and deletion of log files will be much more complex.

To calculate the *max_log_roll_size* for your appliance:

- 1 Determine the values of the three variables listed in the bullet list above.
- 2 Using the *max_log_roll_size* formula, calculate the maximum size a log file should reach before the appliance rolls it over.
- 3 Record the *max_roll_time* result on your planning sheet.

Configuring Logging Options

Based on the planning you have completed in “[Planning Your Logging Strategy](#)” on page 180, you must now configure the log options for each affected service.

Configuration Step 1: Opening the Appropriate Log Options Dialog Box

- 1 For each service you are logging, open the Log Options dialog box in the browser-based management tool.

Refer to the services you selected in “[Planning Step 1: Determining Your Logging Requirements](#)” on page 180.

The path for the Web Server Accelerator service is Configure > Web Server Accelerator > Insert > Enable Logging for This Accelerator > Log Options.

The following sections discuss each of the areas within the Log Options dialog box.

Configuration Step 2: Selecting a Log Format

- 1 In the Log Options dialog box, specify the log format for the service based on the planning you did in “[Planning Step 2: Selecting a Log File Format and Optimizing the Log Entry Size](#)” on page 181.

Remember that each bit of information you log increases the size of each log entry, and affects the rate at which logging disk space is used.

Configuration Step 3: Specifying Rollover Options

- 1 In the Log Options dialog box, specify how the appliance rolls over the log files for the service based on the planning you did in “[Planning Step 3: Calculating Log Rollover Requirements](#)” on page 181.

Configuration Step 4: Specifying Handling of Older Files

You must schedule regular download and deletion of log files to avoid running out of log disk space.

Whenever possible, we recommend you use the FTP log push feature for this task. However, you can also manage log files manually using the browser-based management tool or the appliance’s Mini FTP Server. See “[Manually Downloading and Deleting Log Files](#)” on page 188.

The appliance also provides three options for dealing with old files as a failover precaution.

Ideally, iChain Proxy Services will never actually use the old file option you select because you will schedule the downloading and deleting of log files so that the system never becomes full.

Two of these options automatically dispose of older files to avoid the disk full condition. The third option is not recommended for most situations.

- ♦ **Limit Number of Files To:** This option lets you limit the total number of log files retained for each service. After the limit for each is reached, the oldest file for the service is deleted each time a new file is created. All logging data in deleted files is lost.
- ♦ **Delete Files Older Than:** This option lets you configure the appliance to delete files when they are older than the time you specify. All logging data in deleted files is lost.
- ♦ **Do Not Delete:** This option is not recommended because it can lead to a disk full condition if files are not manually downloaded and deleted. If, however, the older logging data has more value for some reason, this option will preserve the oldest log files unless you manually delete them or specify their deletion in the FTP Log Push Configuration dialog box.

To specify how the appliance handles older files:

- 1 In the Log Options dialog box, select an old file option that matches your requirements. (To review option specifics, see the bullet list above.)

As with log format and rollover options, you can specify different old file options for each service. We recommend, however, that you avoid potential confusion by using the same old file settings for each service.

- 2** (One time only) Click FTP Log Push > configure the FTP log push options.

For help with setting options in the FTP Log Push Configuration dialog box, refer to [“Using FTP Push to Automatically Download and Delete Log Files” on page 187](#), then return to this procedure.

- 3** In the Log Options dialog box, double-check the Old File Options settings against either your FTP log push configuration or your schedule for manual download and deletion to ensure that log files won't reach the deletion threshold (number or age) prior to a scheduled download and deletion.
- 4** If you need to configure log options for other services, return to [“Configuration Step 1: Opening the Appropriate Log Options Dialog Box” on page 184](#); otherwise, continue with the next section.

Configuration Step 5: Monitoring and Refining Your Logging Strategy

As with all appliance operations, you should monitor what is happening with your logging strategy over time and make adjustments and refinements if necessary.

The following are examples of monitoring you might want to consider.

- ◆ Ensure that all the logging information you are gathering is being used. If not, you might be able to further reduce your logging record size.
- ◆ Ensure that your log file sizes match the estimated averages you used to plan your log file roll-over strategy. If not, you might need to adjust the frequency or even the method used to trigger log file rollover.
- ◆ Ensure that your logging strategy is leaving a buffer of free log disk space adequate for possible surges in appliance traffic.
- ◆ Ensure that the external storage capacity (FTP server or other storage) is adequate.
- ◆ Ensure that all aspects of your logging strategy are keeping pace with increases in traffic through the appliance.

About the FTP Log Push Feature

The FTP Log Push feature lets you configure the appliance to push log files to an FTP server at specified intervals: on the first day of the month and/or on specified days of the week. However, log files cannot be pushed more often than once a day.

The feature operates within the following parameters:

- ◆ iChain Proxy Services will try as many times as necessary to establish one connection with the FTP server during the hour of the scheduled push. When the hour changes, iChain Proxy Services stops trying until the next interval you have specified.
- ◆ When a connection with the FTP server is established, iChain Proxy Services assumes that the pushing of log files is successful. Any errors that prevent the successful pushing of log files are not detected by iChain Proxy Services.

For example, you specify that log files are to be pushed on every day of the week at 12 midnight. When the system clock reaches the target hour, iChain Proxy Services begins trying to establish a connection with the FTP server.

If a connection cannot be established before the hour changes to 1 a.m., iChain Proxy Services stops trying to connect and doesn't try again until 12 midnight the next day.

If a connection is established but an error occurs that prevents a successful push, the error is not detected, and iChain Proxy Services doesn't try to connect again until 12 midnight the next day.

Using FTP Push to Automatically Download and Delete Log Files

To configure your appliance to use the FTP Log Push feature:

- 1 In the browser-based management tool, access the FTP Log Push Configuration dialog box by clicking FTP Log Push on any of the Log Options dialog boxes.

Paths to the dialog boxes are summarized under [“Configuring Logging Options” on page 184](#).

IMPORTANT: Although the FTP Log Push Configuration dialog box is accessed through one of the service-specific Log Options dialog boxes, it is unaffected by the path used to reach it. The settings you specify affect all the log types you check in the box.

This lets you set the FTP push options for all log types in a single place.

2 In conformance with your logging strategy, specify the following information:

- ◆ Which log file types to push (all of the log types to be managed through FTP push must be checked).
- ◆ Your FTP server information.
- ◆ The method the appliance uses for determining when to push log files.

If your FTP server is always available, we recommend using the Push Logs When the Logs Roll Over option rather than setting specific download times. This will protect your appliance from sudden surges in traffic, which can fill the disk sooner than expected.

- ◆ Whether the appliance should delete the files from the log disk once they have been pushed.

We recommend deleting log files after they have been pushed unless there is a compelling reason for manually deleting them. Automatic deletion also protects your appliance from sudden surges in traffic.

3 When you have configured your FTP log push options, click OK to return to the Log Options dialog box of the service you are configuring.

Manually Downloading and Deleting Log Files

Whenever possible, we recommend that you use the FTP Log Push feature for downloading and deleting log files. See [“Using FTP Push to Automatically Download and Delete Log Files” on page 187](#).

If you need to manage your log files manually, we recommend that you establish a regular schedule and ensure that all those responsible for downloading and deleting log files know the following information:

- ◆ When log files are to be downloaded and deleted
- ◆ How to determine the name of each log file to be downloaded and deleted
- ◆ Where to save the log files

You will want to develop specific procedures for your situation. The following sections contain general ideas for accomplishing these tasks.

When to Download and Delete Log Files

The primary consideration is that log files must be downloaded and deleted before the logging disk space fills up.

Getting Log Filenames

Before you can download or delete a log file you must know its exact name.

Appliance log filenames can be listed in the browser-based management tool in Monitoring > Cache Logs. They can also be listed from the command line, or through a Telnet session using the get command.

The appliance automatically generates log filenames as follows:

- ◆ Six numbers representing the year, month, and day the file was created
- ◆ A dash separating the date from a letter identifier
- ◆ Letter identifiers running from A through ZZ

This naming convention accommodates up to 702 log files per day. If the rollover options are set so that all the possible filenames are used in one day, the log file with the ZZ letter identifier is not closed until the start of the next day.

To list log files using FTP, you must know the path to the files. Use the following table to determine the paths to various log files.

File	Location
All log files	LOG:ETC/PROXY/DATA/LOGS/
Transparent and forward proxy log files in common format	LOG:ETC/PROXY/DATA/LOGS/FORWARD/COMMON/
Transparent and forward proxy log files in extended format	LOG:ETC/PROXY/DATA/LOGS/FORWARD/EXTENDED/
Filter log files in appliance filtering common format	LOG:ETC/PROXY/DATA/LOGS/FILTER/COMMON/
Web server accelerator log files in common format	LOG:ETC/PROXY/DATA/LOGS/REVERSE/COMMON/ <i>name</i>
	The variable <i>name</i> is the name of the Web server accelerator.

File	Location
Web server accelerator log files in extended format	LOG:ETC/PROXY/DATA/LOGS/REVERSE/EXTENDED/ <i>name</i> The variable <i>name</i> is the name of the Web server accelerator.
Dynamic Bypass log files in common format	LOG:ETC/PROXY/DATA/LOGS/D_BYPASS/COMMON/

Using the Browser-Based Tool to Get Filenames

You can most easily view log filenames in the browser-based management tool. To do so, click Monitoring > click Cache Logs > select a log format > select a service.

Using FTP to Get Filenames

The Mini FTP Server supports the CWD command for changing to the target log directories. You can also use the LS command in connection with full paths to list log files.

For example, the following command lists transparent and forward proxy log files in common format:

```
ls log:etc/proxy/data/logs/forward/common/
```

For a complete list of log file directory paths, see [“Getting Log Filenames” on page 189](#).

Using the Command Line or Telnet to Get Filenames

You can also see a list of log filenames from the command line. However, you cannot download files from the command line.

The following table presents some command line/Telnet examples.

If You Want To	Then Enter
See a list of available forward/transparent log files in common format	<code>get comlog forward</code>

If You Want To	Then Enter
See a list of available Web server accelerator log files in common format	<code>get comlog reverse:name</code> (The variable <i>name</i> is the name of the Web server accelerator.)
See a list of available filtering log files in proxy server filtering common format	<code>get comlog filter</code>
See a list of available forward/transparent log files in extended format	<code>get extlog forward</code>
See a list of available Web server accelerator log files in extended format	<code>get extlog reverse:name</code> (The variable <i>name</i> is the name of the Web server accelerator.)

Downloading Log Files

Using the Browser-Based Management Tool to Download Log Files

You can download the files in the browser-based management tool as you view them. After you click Download, when the browser asks what you want to do with the file, save it to your designated log file storage location.

Using FTP to Download Log Files

You can use FTP from the storage location to retrieve the files with the `get` command. You must first obtain each filename by using one of the options explained in [“Getting Log Filenames” on page 189](#).

After you have the log filename, you can transfer the file to your workstation. For example, to download a forward proxy common format log file, you would use the following command after starting an FTP session with the appliance:

```
get log:/etc/proxy/data/logs/forward/common/filename.log
```

The *filename* variable is the name of the log file you have previously obtained.

You can also use the `mget` command, but be aware that this command also downloads active log files that are not complete.

The appliance doesn't currently support the FTP server `put` command.

Deleting Downloaded Log Files

After the log files have been downloaded and saved to another location, delete the files using one of the following options:

- ◆ The Delete button in the browser-based management tool
- ◆ The del command in FTP

About Extended Log Field Headers

The following information about field values in extended log files might help you interpret the content of the files.

- ◆ Fields within the file are delimited by the tab character.
- ◆ A field can be one of two types: string or non-string.
- ◆ String fields are enclosed in quotation marks ("").
- ◆ If a string field contains a quotation mark, that character is repeated once for every occurrence to enable unambiguous file parsing.
- ◆ If a string field has no value, it is represented by two quotation marks ("").
- ◆ Non-string fields containing no value are represented by a hyphen (-).
- ◆ Field headers starting with s- are associated with the appliance.
- ◆ Field headers starting with c- are associated with the client/browser.
- ◆ Field headers starting with sc are associated with flow from the appliance to the client/browser.
- ◆ Field headers starting with cs are associated with flow from the client/browser to the appliance.

The information in the following table is supplementary to the W3C Extended Log Format Specification found on the [Extended Log File Format \(http://www.w3.org/TR/WD-logfile\)](http://www.w3.org/TR/WD-logfile) Web site. You might find it useful for interpreting the content of extended log field headers.

Name	Description	Type	Selectable	Comments
date	GMT date in YYYY-MM-DD format	non-string	No	

Name	Description	Type	Selectable	Comments
time	GMT time in HH:MM:SS format	non-string	No	
c-ip	Client (browser) IP address	non-string	No	
cs-authname	Username if applicable	non-string	Yes	
s-ip	The appliance IP address	non-string	Yes	
s-sitename	Reverse proxy or accelerator site name	non-string	Yes	
cs-method	The HTTP method the browser sent to the appliance	non-string	Yes	
cs-uri	The HTTP URL the browser sent to the appliance	non-string	Yes	The URL must not have spaces per the HTTP specification.
cs-uri-stem	The stem portion of the HTTP URL the browser sent to the appliance	non-string	Yes	The URL stem is everything up to the first question mark (?). If the URL has no question mark, the cs-uri-stem is the same as the cs-uri. This field is redundant if cs-uri is selected.
cs-uri-query	The query portion of the HTTP URL the browser sent to the appliance	non-string	Yes	The query portion is the first question mark through the end of the URL. If the URL has no question mark, cs-uri-query has no value. This field is redundant if cs-uri is selected.

Name	Description	Type	Selectable	Comments
c-version	The HTTP version specified in the URL the browser sent to the appliance	non-string	Yes	
sc-status	The HTTP status code the appliance sent to the browser	non-string	Yes	
sc-bytes	The number of bytes of HTTP response data the appliance sent to the browser	non-string	Yes	
cs-bytes	The number of bytes of HTTP request data the appliance received from the browser	non-string	Yes	
time-taken	The time in seconds it took appliance resources to deal with the request	non-string	Yes	
cs(User-Agent)	The User-Agent HTTP request header value the browser sent to the appliance	string	Yes	
cs(Cookie)	The Cookie HTTP request header value the browser sent to the appliance	string	Yes	The appliance doesn't cache cookie information.
cs(Referer)	The Referer HTTP request header value the browser sent to the appliance	string	Yes	The appliance reads the field header as it is.

Name	Description	Type	Selectable	Comments
cs(X-Forwarded-For)	The X-Forwarded-For HTTP request header value the browser sent to the appliance	string	Yes	Do not confuse this with the X-Forwarded-For option that causes the appliance to generate or forward headers to upstream proxies or Web servers.
cached	The value indicating whether the request was filled from cache	non-string	Yes	1 = filled from cache 0 = not filled from cache
x-fill-proxy-ip	The IP address of the upstream proxy	non-string	Yes	Assumes the appliance is configured with an upstream proxy and brought the request from that proxy
x-origin-ip	The IP address of the origin server	non-string	Yes	Assumes the appliance brought the request directly from the origin server

Enabling and Viewing the ACL Rule Log File

The logs for authorized access attempts and unauthorized access attempts can be turned on or off globally. The rules for logging authorized access attempts for an individual access control list (ACL) can also be turned on or off. However, because unauthorized access attempts are usually the result of a user not being defined in any ACL rule, logging of unauthorized access attempts cannot be turned on or off for individual ACL rules.

To enable or disable ACL rule logging on a global level:

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

For example, <http://xx.xx.xx.xx:1959/appliance/config.html> where xx.xx.xx.xx is the IP address.

- 2 Click Configure > Web Server Accelerator > Modify.
- 3 Check the Enable Logging check box.

To enable or disable ACL rule logging for an individual ACL rule:

- 1 From ConsoleOne, right-click an ACL Rule object.

2 Select Properties.

3 Check the Authorized Logging check box.

The ACL log files for each 24-hour period are saved to LOG:\ETC\PROXY\DATA\LOGS\REVERSE\EXTENDED\ACLCHECK\yymmdd-a.log, where yymmdd is today's date represented by two digits for the year, month and date. The default maximum size of the file is 1 MB. The default size can be changed; see [“Using ACLCHECK options” on page 113](#) for more information. If a log exceeds the maximum size, a new file named yymmdd-b.log is created.

Each file contains the following fields:

- ◆ Date
- ◆ Time
- ◆ Source IP address
- ◆ Destination IP address
- ◆ Protocol
- ◆ Source port
- ◆ Destination port
- ◆ TCP flag
- ◆ Access (Allow=1, Deny=0)
- ◆ IP headers
- ◆ IP payload
- ◆ Username
- ◆ Destination host name
- ◆ URL (the user requested)
- ◆ Rule Object name (if access was allowed; if denied, the field displays a —)

FTP Services

You can manage several aspects of the appliance using an FTP client on a workstation connected to a network where the appliance is visible

The appliance's FTP services let you get and put configuration files, log files, and the optional splash screen (which you can customize with your own HTML).

You can also configure the appliance to provide FTP forward proxy (client acceleration) and FTP reverse proxy (server acceleration) services. For more information, see ["FTP Tab" on page 337](#).

Tips for Using FTP Services

The following sections contain important information about using iChain Proxy Services FTP services.

Firewalls and Passive FTP

The appliance's system supports access from both active and passive FTP clients.

The fact that passive FTP is often required to traverse a firewall has certain implications for using FTP with the appliance.

FTP Access Through a DOS Window

Because DOS uses only active FTP, you cannot access an appliance that is outside a firewall through a DOS window on a client inside the firewall.

The reverse is also generally true. If the appliance is inside a firewall, you will not usually be able to access it through a DOS window on a client outside the firewall.

FTP Access Through a Browser Outside the Firewall Is Also Limited

You cannot generally access an appliance inside a firewall through a browser that is outside the firewall.

FTP Access Through a Browser Inside the Firewall

Since Netscape browsers use passive FTP, you can generally access an appliance outside the firewall from a Netscape* browser inside the firewall.

To access an appliance outside a firewall using Internet Explorer 5 inside the firewall, you must configure the browser to use passive FTP.

- 1** In the browser, click Tools > Internet Options > Advanced.
- 2** Under Browsing, check Use Web-Based FTP.

NOTE: This option name varies according to browser version. In Internet Explorer 5.5, for example, the option is Use Passive FTP for compatibility with some firewalls and DSL modems.

3 Click OK.

Appliance Routing and Transparent Proxy Limitations on FTP

When using the built-in appliance router capabilities in a *transparent* proxy situation, users will not be able to perform browser-based FTP using ftp:// as the protocol. Browser-based FTP works normally with forward proxy.

Setting Up Appliance FTP Services

Before using FTP services to manage the appliance, you must ensure the appliance's Mini FTP Server is properly configured.

- 1** Start the browser-based management tool > click Configure > FTP.
- 2** Ensure that at least one of the IP addresses in the Server IP Addresses list is checked.

You will use the checked address for your FTP session. IP address 10.1.1.1 is checked by default.

Limitations of the Appliance's Mini FTP Server

The appliance's Mini FTP Server was originally designed only for uploading and executing appliance configuration (.NAS) files. This functionality has been expanded slightly and currently supports the following commands:

- ◆ CWD
- ◆ DELE
- ◆ LIST
- ◆ PASS
- ◆ PASV
- ◆ PORT
- ◆ PWD
- ◆ QUIT
- ◆ RETR

- ◆ STOR
- ◆ SYST
- ◆ TYPE
- ◆ USER

The Mini FTP Server has no support for wildcard characters.

The Mini FTP Server will not work with directory or file names that contain spaces.

Mini FTP server access is limited to the following directories and their subdirectories:

- ◆ SYS:ETC\PROXY\MFTP
- ◆ SYS:ETC\PROXY\APPLIANCE
- ◆ SYS:ETC\PROXY\APPLIANC
- ◆ SYS:ETC\PROXY\DATA
- ◆ SYS:ETC\APPLIANCE
- ◆ SYS:ETC\APPLIANC
- ◆ LOG:ETC\PROXY\DATA

When you log in to the FTP server, the SYS:ETC\PROXY\APPLIANCE\CONFIG\USER directory is the default.

To execute a .NAS configuration file, you must be in this default directory and use the following syntax:

```
put local_filename remote_filename,execute
```

The *local_filename* variable is the name of the .NAS file on your local machine and the *remote_filename* variable is the name after the file is uploaded to the appliance.

Starting an FTP Session with the Appliance

- 1** Launch your FTP application and enter a valid appliance IP address, for example:

```
ftp 10.1.1.1
```

- 2** Log in to the appliance using the Config username and the password you have set.

Changing the FTP Working Directory

The default working volume and directory for FTP sessions is SYS:\ETC\PROXY\APPLIANCE\CONFIG\USER. This means that the SYS: volume is implied for all FTP commands unless the LOG: volume is specifically included.

You can use the cd (change directory) command to change the current volume and directory path. For example, cd log:/etc/proxy/data changes the working path to the LOG: volume and the directory path to \ETC\PROXY\DATA.

FTP commands that include only a directory path and that are issued subsequent to the cd command will use the newly specified volume.

You can specify a full path (volume and directory) when using the get and put commands to copy files to or from any FTP-accessible location. However, the default volume and directory are unaffected by these commands. Only the cd command changes the FTP working path.

Managing Configuration Files with FTP

All appliance settings are contained in configuration text files with the extension .NAS. These files can be edited and sent through FTP back to the appliance where they can be executed as a means of instant configuration.

By using several configuration files, you can quickly apply different scenarios, such as turning various proxy services on or off. The appliance updates the default file CURRENT.NAS every time you apply a setting. Additionally, if you have created a file named AUTOLOAD.NAS on a floppy disk, it is updated with each change. From the browser-based management tool or the Telnet/command line interface, you can export other configuration files.

Using FTP to Download a Configuration File to Your Workstation

- 1 Start FTP and log in as explained in [“Starting an FTP Session with the Appliance” on page 199](#).
- 2 Enter the following, where *filename* is the name of your configuration file:

```
get filename.nas
```

The file is transferred to your FTP client's default directory.

Using FTP to Move a Configuration File to the appliance from a Workstation

- 1 Start FTP and log in as explained in “Starting an FTP Session with the Appliance” on page 199.
- 2 Enter the following, where *filename* is the name of your configuration file:

```
put filename.nas
```

The file is transferred to the appliance.

Using FTP to Move a Configuration File to the Appliance and Execute It

- 1 Start FTP and log in as explained in “Starting an FTP Session with the Appliance” on page 199.
- 2 Enter the following, where *filesrc* is the name of the source configuration file and *filedst* is the name used at the destination:

```
put filesrc.nas filedst.nas,execute
```

Using FTP to Customize the Appliance Splash Screen

The appliance has an optional splash screen that displays before pages are vended.

The splash screen has the root filename BMSPLASH. A three-letter extension indicates whether the splash screen is disabled (.OFF) or enabled (.HTM). The screen is disabled by default. For more information, see “Tuning Tab” on page 342.

To alter the appearance of the splash screen:

- 1 Start FTP and log in as explained in “Starting an FTP Session with the Appliance” on page 199.
- 2 To download the splash screen to you can modify it, do one of the following:
 - ♦ If the splash screen is enabled, enter:

```
get /etc/proxy/data/bmsplash.htm
```
 - ♦ If the splash screen is not enabled, enter:

```
get /etc/proxy/data/bmsplash.off
```
- 3 Modify the file.

4 To upload the splash screen after you have modified it, do one of the following:

- ◆ If the splash screen is enabled, enter:

```
put bmsplash.htm /etc/proxy/data/bmsplash.htm
```

- ◆ If the splash screen is not enabled, enter:

```
put bmsplash.off /etc/proxy/data/bmsplash.off
```

Object Pinning

This section contains the following topics:

- ◆ [“The Pin List” on page 202](#)
- ◆ [“URL Mask” on page 202](#)
- ◆ [“Pin Type” on page 203](#)
- ◆ [“Pin Links” on page 204](#)
- ◆ [“Pin Images” on page 204](#)
- ◆ [“Refresh Frequency/Time” on page 205](#)
- ◆ [“Processing URL Masks” on page 205](#)
- ◆ [“Wildcards in Pin Lists” on page 208](#)
- ◆ [“Pin List Examples” on page 209](#)

The Pin List

The pin list contains URL patterns for identifying objects on the Web. You configure each URL pattern in the list with specific handling instructions as explained in the following sections.

Pinned objects remain in the cache indefinitely unless it fills up. This ensures that the lists are available from cache and will not be bumped out by more recently requested objects.

URL Mask

The URL mask can contain complete or partial URL patterns. A single URL mask might apply to a large set of URLs, or it might be so specific that only a single file on the Web matches it. For more information, see [“Pin List Examples” on page 209](#).

The appliance processes the masks in the pin list in order of specificity. A mask containing a host name is more specific than a mask that specifies only a file type. The action taken for an object is the action specified for the first mask that the object matches. For more information, see “[Processing URL Masks](#)” on page 205.

If the mask contains an asterisk, only the pin type can be specified. The Pin Links, Pin Images, and Refresh Frequency/Time options are not available for URLs containing this wildcard. Objects matching a mask with an asterisk are not automatically downloaded, but are pinned in cache only as individually requested.

Pin Type

The pin type specifies whether and how the appliance will cache objects that match the URL mask.

- ◆ **Normal:** iChain Proxy Services handles objects matching the mask in the same way it handles any other requested objects. In other words, objects are cached but not pinned.

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

For example, you could specify a URL mask of `/*.jpg` with a pin type of `bypass` and a second URL mask of `www.foo.gov/graphics/*.jpg` with a pin type of `normal`. The result would be that the `.jpg` files in the `graphics` directory on the `foo.gov` Web site would be cached as requested. They would not, however, be pinned in cache because of the `normal` pin type. Assuming there were no other URL masks in the pin list, all other `.jpg` graphics would not be cached because of the first URL mask.

- ◆ **Cache:** iChain Proxy Services keeps the pinned objects in cache as long as possible, although they might be written to the appliance’s hard disk.
- ◆ **Memory:** iChain Proxy Services keeps the pinned objects in memory as long as possible, writes them to disk when memory gets too full, and places them back in memory as soon as they are requested by a user of the cache.
- ◆ **Bypass:** iChain Proxy Services does not cache the objects. In other words, you can use this option to prevent objects from being cached. Keep in mind that URL masks set with a type of `bypass` operate differently than URL masks set with a type of `cache`. For example, `http://www.foo.com/` with a type set as `bypass` will bypass only the root

directory. To bypass the entire URL, the trailing forward slash must be removed or an asterisk must be added (<http://www.foo.com> or http://www.foo.com/*).

Pin Links

This specifies how many link levels iChain Proxy Services will follow for the pin type rule you've established. Selecting levels 1 or 2 causes all linked objects, including the images on the host, to be downloaded and cached when the pin list is applied to the appliance configuration, and then to be periodically refreshed as specified.

For example, if the requested object is an HTML page and you have specified a pin links level of 1, the HTML page will be downloaded and cached when the pin list is applied along with all the items linked from the page. These cached objects will also be refreshed at the refresh frequency and time specified.

To use levels 1 or 2 you must specify an absolute address, including the scheme, host, and path for the URL mask, for example, <http://www.foo.gov/documents/>. The tool will let you insert masks that do not meet this requirement, but the entries are removed when you click Apply.

Attempting to include an asterisk wildcard immediately hides this option.

Pin Images

This option is used to pin image files that reside on a different host than the page requested. It works in conjunction with the Pin Links option, which specifies how many levels of links iChain Proxy Services will follow when downloading a page.

For example, if the requested HTML page uses images that reside on another host and you have checked this option, the HTML page will be cached along with all the image files associated with the page, including those on the other host. If you have also specified a pin link level, images on the linked pages that reside on another host will also be pinned.

On the other hand, if the Pin Images option is not checked, iChain Proxy Services only pins the images that reside on the same host as the requested page.

Refresh Frequency/Time

This lets you specify a refresh frequency and time for the URL that is different from the default values shown above the pin list.

Processing URL Masks

There are four basic types of URL masks you can enter in the pin list. The following table lists each type, provides a few examples of each, and provides information on how they are processed by iChain Proxy Services.

Type	URL Mask Examples by Specificity	Notes
Hostname	<p>http://www.foo.gov/documents/picture.gif</p> <p>http://www.foo.gov/documents/</p> <p>http://www.foo.gov</p> <p>foo.gov/documents/</p> <p>foo.gov/</p> <p>*.foo.gov/</p>	<p>Although these entries can include the protocol or scheme, the DNS name, the path, and the filename, only the DNS or hostname must be present in the mask. All DNS label portions must be indicated, if only by an asterisk wildcard.</p> <p>iChain Proxy Services processes hostname entries before it processes other mask types. It also processes the most specific URL mask entries first.</p> <p>When an object match occurs, iChain Proxy Services applies the pin type rule, and processing of the object is finished.</p> <p>For example, if the first URL mask in the examples column has a pin type rule of bypass, PICTURE.GIF will not be cached regardless of the pin type rules for the other URL masks.</p> <p>Hostname entries can have a dramatic impact on object pinning and cache bypassing.</p> <p>For example, if the first two URL masks in the examples column were not present, a pin type of Bypass on the third URL mask would prevent caching of all objects delivered through HTTP on the www.foo.gov Web site.</p> <p>If no scheme (HTTP, FTP, etc.) is indicated, the mask applies to all schemes. The last three masks would apply to objects delivered through any Web protocol.</p> <p>Finally, Configure interprets hostnames literally. For example, the sixth entry would cover www.foo.gov, ww1.foo.gov, army.foo.gov, etc., but the fourth and fifth entries would not, because a scheme is assumed to immediately precede the hostname.</p>

Type	URL Mask Examples by Specificity	Notes
Path	/documents/picture.gif /documents/picture.gif/ /documents/	<p>iChain Proxy Services processes path entries after all hostname entries have been considered. It assumes that the first forward slash immediately follows a hostname.</p> <p>A leading forward slash must always be used when specifying a directory. The leading slash always references the root directory of the Web server.</p> <p>For example, the first entry would apply only to a graphics file named PICTURE.GIF that is located in a DOCUMENTS directory at the root of the host.</p> <p>The forward slash in the second entry causes iChain Proxy Services to assume that PICTURE.GIF is a directory. The pin type rules associated with this entry would apply to any matched objects that have a URL directory path that starts with a documents directory followed by a subdirectory named PICTURE.GIF.</p> <p>The third entry would apply to any matched objects that contain a DOCUMENTS directory at the Web server's specified root directory.</p>

Type	URL Mask Examples by Specificity	Notes
Filename	/picture.gif	After the path entries have all been processed, iChain Proxy Services looks for specific filenames.
	/widget.js	
	/default.htm	A leading forward slash must be used and, as opposed to a path-based mask, does not reference the root directory of the Web server. For example, if requested objects named PICTURE.GIF, WIDGET.JS, and DEFAULT.HTM have not been covered by one of the hostname or path entries above, the files will have the pin type rule for their respective filename mask applied to them. If the first entry carries a pin type rule of Bypass, all PICTURE.GIF files that didn't match previously processed hostname or path masks would not be cached.
File Extension	/*.gif	File extension entries are processed last.
	/*.js	These are simply filename entries with the root of the filename replaced by an asterisk, which makes them less specific than complete filenames.
	/*.htm	A leading forward slash must be used and, as opposed to a path-based mask, does not reference the root directory of the Web server. For example, if the examples shown all had pin types of Bypass, then only those .GIF, .JS, and .HTM files that had been cached and pinned because of hostname, path, or filename masks would be stored in cache. All other files with the named extensions would not be cached.

Wildcards in Pin Lists

Only the asterisk (*) wildcard is allowed in pin list entries.

208 Place Book Title Here

iChain Proxy Services interprets everything between an asterisk and the next delimiter to the right (a forward slash [/], a period[.], or a colon [:]) as a wildcard. This effectively allows only one asterisk between delimiters.

Pin List Examples

The following table provides brief examples of sample pin list entries and their effects on appliance caching.

URL Mask	Pin Type	Pin Links	Pin Images	Effect on Cache
http://www.foo.gov/documents/	cache	1	Yes	<p>As a general rule, you should always include fully qualified DNS or hostnames in the pin list. iChain Proxy Services resolves these more quickly than other masks, and you will be able to track the effects on pinning more easily.</p> <p>For this URL mask, iChain Proxy Services downloads, caches, and pins all objects whose URL starts with the mask. In other words, all objects below the documents directory will be downloaded, cached, and pinned. Also, all objects that are linked from one of the pinned objects will be downloaded, cached, and pinned. And finally, images that reside on other hosts will be downloaded, cached, and pinned as well.</p> <p>Objects will be refreshed according to the refresh settings (default or specific) as specified in the pin list entry.</p>
www.foo.gov/groups.html	cache	1	No	<p>iChain Proxy Services downloads, caches, and pins objects (including images) in the GROUPS.HTML page and in pages linked from that page. Any images referenced from other hosts, however, are not included.</p>

URL Mask	Pin Type	Pin Links	Pin Images	Effect on Cache
www.foo.gov/groups.html/	normal	1	Yes	<p>iChain Proxy Services downloads and caches objects in the subdirectory named groups.html and in pages linked from any of those objects.</p> <p>The forward slash at the end of the path tells iChain Proxy Services that this is a directory rather than a file.</p> <p>Objects are cached but not pinned in cache, meaning they might be bumped by more frequently accessed objects or objects that are pinned.</p> <p>Images linked from other hosts are downloaded and cached.</p>
www.foo.*	bypass	n/a	n/a	<p>iChain Proxy Services doesn't cache objects from any URLs whose DNS names begin with www.foo.</p> <p>All domain extensions (.com, .net, .org, etc.) are covered by the asterisk wildcard.</p> <p>Link and image pinning is not available for bypass pin types.</p> <p>If this entry appeared in a pin list with either of the previous two entries, it would not prevent caching of objects covered by them because it is less specific than they are.</p>
w*.f*.com	bypass	n/a	n/a	<p>iChain Proxy Services doesn't cache objects for any URLs whose first domain label begins with w and second domain label begins with f, providing the domain extension is .com.</p> <p>This mask doesn't prevent caching of objects on other domains such as .net, .gov, etc.</p>

URL Mask	Pin Type	Pin Links	Pin Images	Effect on Cache
w*.f*.*	bypass	n/a	n/a	This mask functions like the previous entry, but the domain is not limited to .com.
.foo.	cache	n/a	n/a	<p>This causes all objects on any Web server whose second domain label is foo to be pinned in cache.</p> <p>Link and image pinning are not available because the mask contains asterisks.</p> <p>This mask would not cover DNS names that don't have a domain label before foo. For example, foo.gov would not normally be covered. However, if foo.gov happened to resolve in DNS to the same IP address as www.foo.gov, the iChain Proxy Server would apply the pinning rules specified for www.foo.gov to foo.gov. To understand more about IP addresses and URL masks, see "Using the Proxy Server to Record IP Addresses When Resolving URL Masks" on page 211.</p>

Using the Proxy Server to Record IP Addresses When Resolving URL Masks

As stated earlier, you should include fully qualified DNS or hostnames in URL masks whenever possible.

The iChain Proxy Server resolves DNS names to their respective IP addresses and uses those addresses when pinning objects.

You can use this fact when constructing your pin list entries.

For example, if you use the DNS name www.foo.gov as the URL mask and you know that the DNS name foo.gov resolves to the same IP address, you don't need to include foo.gov in the pin list.

Because both URLs resolve to the same IP address, iChain Proxy Services will treat objects for both DNS names the same.

On the other hand, if `www.foo.gov` and `foo.gov` resolve to different IP addresses, separate pin list entries would be required to cover both sites.

Router Capabilities

Having the appliance double as a router impacts appliance performance, but it is a low-cost router option that delivers acceptable performance in some low-volume networks.

Each appliance is normally configured with a default gateway. If the appliance is not acting as a router, the default gateway is the appliance's next hop.

For more information about appliance routing, see [“Gateway/Firewall Tab” on page 306](#) and [“Routes Dialog Box” on page 311](#).

Using Appliance Routing

If the appliance is acting as a router, it routes requests to IP addresses based on the information in its routing table. If a request could be routed through multiple gateways, the appliance chooses the gateway associated with the most restrictive mask (the smallest range of destination addresses).

Routing table entries fall within the following three basic groups:

- ◆ Host gateways for specific destination addresses
- ◆ Network gateways for destination addresses that fall within specific subnets
- ◆ The default gateway for destination addresses that aren't covered by host or network gateways

The syntax for this gateway is often expressed in router configuration tables as `0.0.0.0 / 0.0.0.0 / iii.iii.iii.iii`, where the *i*'s represent the IP address of the default gateway.

You define these gateways in the browser-based management tool by clicking `Network > Gateway/Firewall > Additional Gateways` or by clicking `Configure > Client Accelerator > Router Options`.

IMPORTANT: If the appliance is acting as a router and you don't specify a default gateway, the appliance routes only those requests whose destination addresses are covered by a host or network gateway. Other requests are not routed.

For more information on routing concepts, see one of the TCP/IP references available at any bookstore carrying computer reference manuals.

Shutting Down and Restarting

If you need to shut down or restart an appliance, you should do it properly to protect the data in memory and ensure it is written to disk. There are several ways to properly shut down or restart the appliance.

Restarting from the Browser-Based Management Tool

- 1 Start the browser-based management tool.
- 2 Click System > Actions.
- 3 Shut down iChain Proxy Services by clicking Shut Down or shut it down and restart it by clicking Restart.

You are given a chance to verify your selection.

If you choose to shut down iChain Proxy Services, you hear a three-beep sequence that repeats until the appliance is turned off or restarted.

If you choose to restart iChain Proxy Services, you hear a two-beep, four-beep sequence that repeats for a short period and then stops, signifying that your appliance has restarted.

If you do not have access to the physical location of the appliance, you can test to see if the appliance has restarted by pinging its address on port 1959. If the ping succeeds, the appliance has restarted.

Shutting Down and Restarting from Telnet or the Command Line

You can shut down or restart an appliance from the command line.

NOTE: Both actions break the connection. If you *restart* the appliance from a remote connection, you will be able to reconnect after the appliance restarts. However, if you *shut down* the appliance, someone will need physical access to the appliance to restart it.

To restart the appliance from the command line, enter

Restart

If you are near the appliance, you will hear a two-beep, four-beep sequence that repeats for a short period and then stops, signifying that your appliance has restarted.

To shut down the appliance from the command line, enter the following:

Shutdown

If you are near the appliance, you will hear a continuous three-beep sequence once the system is down. You can restart the system from an attached keyboard by pressing Ctrl+Alt+Del, or you can shut off the power once you hear the three-beep sequence.

Using the SOCKS Client Service

For sites that have deployed SOCKS firewalls, the appliance provides a SOCKS Client service to redirect all forward proxy traffic through the firewall. Redirecting appliance traffic to the SOCKS firewall might significantly reduce appliance performance.

The appliance currently supports both SOCKS4 and SOCKS5 protocols.

For information about setting up SOCKS client support, see [“Gateway/Firewall Tab” on page 306](#).

Time Synchronization

Time settings offered within the management tool are more than adequate for most system needs. For more information on the specific parameters available, see [“Using the Browser-Based Management Tool to Set the Time” on page 215](#) and [“Date/Time Tab” on page 283](#).

Additional flexibility in setting system time, including changing the GMT offset and daylight saving parameters, is available through the command line interface.

Synchronizing Time

You can either set the time manually or synchronize it using the network time protocol (NTP). The appliance uses NTP by default and comes pre-configured with two servers:

132.163.4.101

132.163.4.103

You can add additional servers or delete them using the browser-based management tool and the command line interface. For more information regarding NTP functionality, see [“Using the Command Line to Set NTP Time Sources” on page 215](#).

Using the Browser-Based Management Tool to Set the Time

To add or delete an NTP Server:

- 1** Start the browser-based management tool > click System > Date/Time.
- 2** Check Use Network Time Protocol.
- 3** Do one of the following:
 - ◆ To add a server, click Insert > type the URL or IP address of the server.
 - ◆ To delete a server, select the server > click Delete.

Changes in the Date/Time tab are immediately effective.

To set the time manually:

- 1** Start the browser-based management tool > click System > Date/Time.
- 2** Check Set Time Manually.
- 3** Click Set Time.
- 4** Using the drop-down lists, select the correct time and date.
- 5** Click OK.

Using the Command Line to Set NTP Time Sources

- 1** Do one of the following:
 - ◆ Add an NTP server address by entering:
`add ntp server=128.115.14.97`
 - ◆ To enable NTP, enter:
`set ntp enable=yes`
 - ◆ To disable NTP, enter:
`set ntp enable=no`

2 To have the changes take effect, enter **apply**.

For more command line options, refer to the appliance's command line help.

NTP Date/Time Synchronization Is Not Immediate

When you specify an NTP server, synchronization between the NTP server clock and the appliance clock might not be immediate.

If the NTP server clock has an earlier date and time setting than the appliance clock, the system will slow down the appliance clock until the two are synchronized. This provides for proper incrementation of log files and other time-sensitive information during the synchronization process.

If the NTP server clock is later than the appliance clock, synchronization between the two will generally be immediate. However, in certain situations you might observe the appliance clock incrementing by 600-minute intervals. This is normal system behavior.

The fact that the Apply button changes from Wait back to Apply indicates only that the NTP configuration change has been made, not that the appliance clock is fully synchronized with the NTP server.

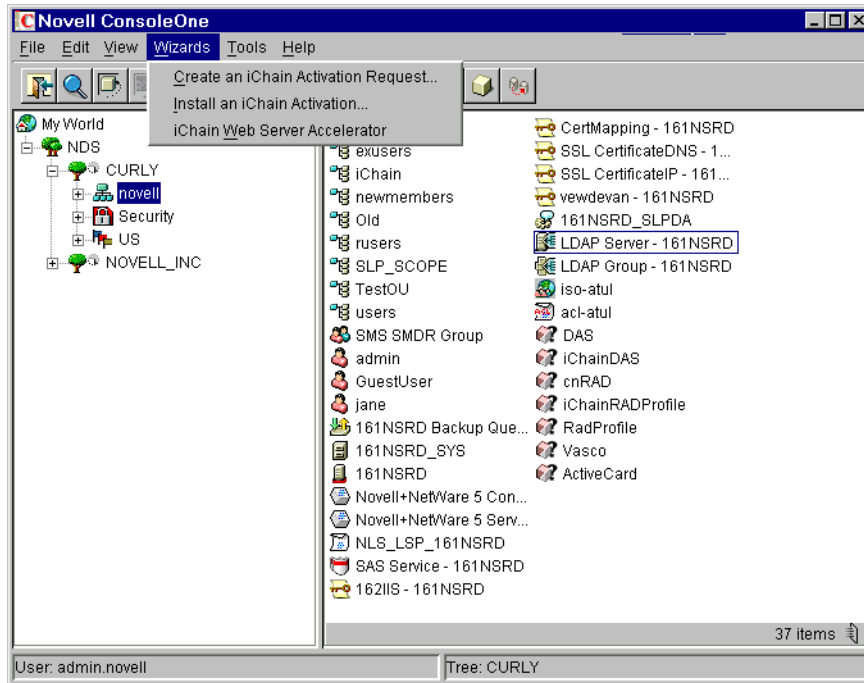
If the above features are problematic in your situation, you can set appliance time manually to the target time and then re-enable the NTP feature.

A iChain Management

This appendix describes the Novell® iChain® Web Server Accelerator menu and the pages of the iChain setup wizard. For each wizard page, the fields and controls on each page are discussed, with explanations of the purpose, content, and functionality included.

The iChain wizard is located at ConsoleOne® > Wizards, as shown in [Figure 37, “Wizards Menu in ConsoleOne,”](#) on page 218.

Figure 37 Wizards Menu in ConsoleOne



The Wizards menu will become selectable when the user has selected a container in a Novell eDirectory[®] tree where iChain objects (iChain Service or ACL) can be created, typically an O (Organization) or OU (Organizational Unit) container.

To create any iChain objects, the user must have object create rights in the container he or she has selected.

Proxy Server Access Configuration Page

The first page of the iChain setup wizard allows the user to specify the Proxy Server Access Configuration settings, as shown in [Figure 38](#).

Figure 38 Proxy Server Access Configuration Page

The screenshot shows a window titled "iChain Web Server Accelerator Wizard" with a close button in the top right corner. The main area is titled "Proxy Server Access Configuration" and contains the following fields:

- iChain service object:** A text box containing "CN=iso-atul.O=novell" and a browse button (represented by a folder icon).
- Proxy Server IP Address:** A text box containing "151.155.164.159".
- Primary session broker server IP address:** An empty text box.
- Secondary session broker server IP address:** An empty text box.

At the bottom of the window, there is a horizontal line and a row of five buttons: "< Back", "Next >", "Cancel", "Finish", and "Help".

To configure the proxy server from ConsoleOne, the user must know the IP address of the proxy server. This address is stored in an iChain Service object in the directory. It is used to log in to the proxy server so that the proxy configuration may be read and written by the wizard.

The following describes the fields on this page:

Field Name	Description	Status
iChain service object	The name of the iChain service object that is associated with the proxy server. The user can type a new name into this field and then press Enter, which will then prompt the user to create the object or browse for an existing object.	Required
Primary session broker server IP address	Displays the IP address for the primary session broker server.	Optional
Secondary session broker server IP address	Displays the IP address for the secondary session broker server.	Optional

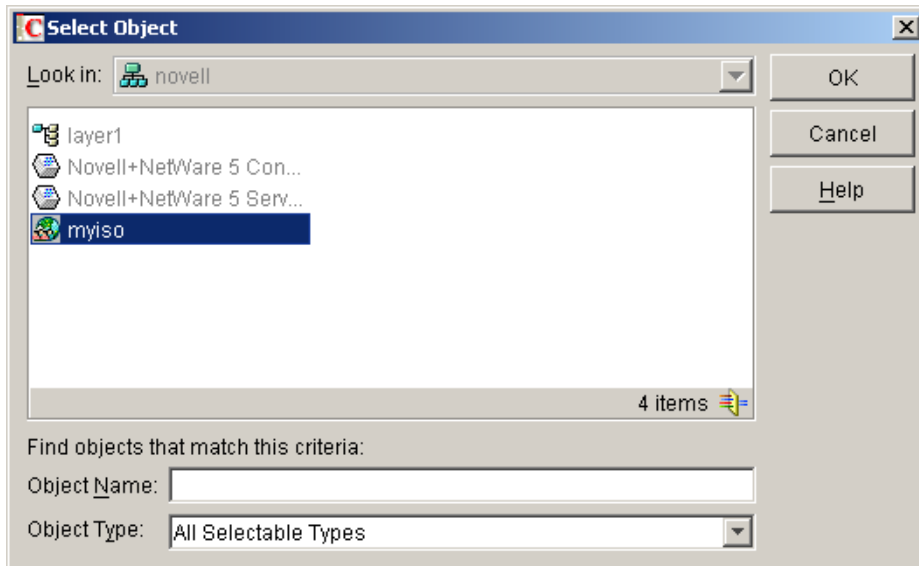
Controls for Proxy Server Access

This section describes the following buttons:

- ◆ [“Browser” on page 220](#)
- ◆ [“Next” on page 221](#)
- ◆ [“Finish” on page 221](#)

Browser

The browser button, located next to the iChain service object field, brings up an object entry selector browser that enables the user to select an existing iChain service object. When an object is selected, the rest of the fields on the page will be populated with data read from the object. The object entry selector browser is shown in [Figure 39](#).

Figure 39 Select Object Dialog Box**Next**

When the user clicks the Next button, he or she is prompted to log in to the proxy server. A dialog box appears and the user can enter the proxy server password. This is necessary to be able to read and write information to and from the proxy server, as shown in [Figure 40](#).

Figure 40 Connect to Proxy Server Dialog Box**Finish**

The Finish button is enabled when all the required information has been entered on the page. This allows the user to quickly enter information about

the proxy configuration, save it, and exit the wizard without needing to continue through to the last page.

Access Control Configuration Page

The second page of the wizard enables the user to specify the access control server parameters used during system access control, as well as form fill, OLAC, password management parameters, and load balance, as shown in [Figure 41](#).

Figure 41 Access Control Configuration Page

The screenshot shows the 'iChain Web Server Accelerator Wizard' window. The title bar reads 'iChain Web Server Accelerator Wizard'. The main content area is titled 'Access Control Configuration'. Underneath, there is a section for 'LDAP servers' which contains a table with three columns: 'IP Address', 'Port', and 'Connection'. The table has one row with the values '151.155.164.161', '389', and 'Normal'. To the right of the table are three buttons: a plus sign (+), a minus sign (-), and a diagonal slash (/). Below the table are three checkboxes: 'Form fill (SSO) enabled', 'OLAC enabled', and 'Load balance at session level only'. Below these checkboxes is a text field labeled 'Password management servlet URL' containing the text 'http://www.novell.com/pswd.html'. At the bottom of the window are five buttons: '< Back', 'Next >', 'Cancel', 'Finish', and 'Help'.

IP Address	Port	Connection
151.155.164.161	389	Normal

Form fill (SSO) enabled

OLAC enabled

Load balance at session level only

Password management servlet URL
http://www.novell.com/pswd.html

< Back Next > Cancel Finish Help

The servers contained in this table are used by the proxy and authorization servers to provide access control to the actual Web servers specified in the Web server accelerators. Each server must be installed in the same directory services tree so that the login information may be used to connect to each one.

Multiple servers are provided for access control failover. If a server fails or is shut down, the system tries the next server in the list, thus maintaining a continuation of service.

The following describes the fields on this dialog:

Table 18

Field Name	Description	Status
LDAP servers	Lists the IP address, port, and connection type of each access control server that will be used. The user can add, delete, or edit table entries.	Required
Form fill (SSO) enabled	Allows user to specify whether to enable Form Fill for this proxy server. If checked, Form Fill is enabled.	Optional
OLAC enabled	Allows user to specify whether OLAC is enabled. If checked, OLAC is enabled.	Optional
Password management servlet URL	Shows location of the password management servlet that will be used by the proxy system when it determines when the user needs to change his or her password. If a user's password expires and he or she tries to log in, the system will use this servlet to prompt the user to change the password. If the user does not populate this field upon prompt, the system will use a default password management function.	Optional

Field Name	Description	Status
Load balance at session level only	The proxy server will use the same Web server for all fills during a session. This prevents eBusiness users from needing to log in multiple times. This setting affects all Web server accelerators configured on the proxy server.	Optional

Controls for Access Control

This section describes the following buttons:

- ◆ “Add” on page 224
- ◆ “Delete” on page 227
- ◆ “Edit” on page 227

Add

The Add button allows the user to enter information that identifies an access control server. When this button is selected, a dialog box appears which specifies the IP address, port, admin name, password, and security connection parameters, as shown in [Figure 42](#).

Figure 42 New LDAP Access Control Server Dialog Box

When you add the first entry in the table, all the fields are active. For the second and subsequent entries, only the IP address field is active. This forces all of the servers to be located within the same eDirectory tree.

The fields in [Figure 42](#) are described in the table below:

Table 19

Field Name	Description	Status
IP Address	The IP address of the server.	Required
Port	The Port the server communicates through.	Required

Field Name	Description	Status
Administrator	The administrator's name used to log in to the server. This field is a non-editable field. To populate the field, the user must click the browser button next to it to use an object entry selector browser to choose the administrator from the directory.	Required
Password	The administrator's password.	Required
Password confirmation	Confirms the administrator's password.	Required
Connection	Specifies whether the connection will be secured (Secure Exchange).	Optional
Trusted root file	The name of the trusted root file that will be used to secure the connection.	Required only if using a secure connection

Figure 43 shows how the dialog box appears during the second and subsequent server entries:

Figure 43 New LDAP Access Control Server Dialog Box

New LDAP Access Control Server

Identity

IP Address
137 . 65 . 159 . 221

Port
389

Administrator
cn=admin,o=novell

Password

Password confirmation

Connection

Use a secure connection (LDAP over SSL)

Trusted root file

OK Cancel Help

Delete

The Delete button will delete a server entry from the list of LDAP servers. If the first server is deleted, the second server in the list becomes the controlling server.

Edit

The Edit button allows the user to edit all entries in a table; however, the user can only change all of the information (rather than just the IP address) when editing the first entry. When the first server entry is changed, all of the other servers will change accordingly to ensure that all servers are in the same tree. See [Figure 44](#).

Figure 44 Modify an LDAP Access Control Server Dialog Box

Modify An LDAP Access Control Server

Identity

IP Address
137 . 65 . 159 . 222

Port
389

Administrator
cn=admin,o=novell

Password

Password confirmation

Connection

Use a secure connection (LDAP over SSL)

Trusted root file

OK Cancel Help

Accelerator Specification Page

The third page of the wizard is the accelerator specification page, where the user can create a new accelerator, edit an existing accelerator, or delete an existing accelerator. See [Figure 45](#).

Figure 45 Accelerator Specification Page

iChain Web Server Accelerator Wizard

Create a new web server accelerator

Name: DNS name:

Edit an existing web server accelerator

Name: DNS name:

Delete an existing web server accelerator

Name:

Enable multi-homing Multi-home master:

At this page, the user will see actual accelerator configuration data read from the proxy server if any data exists.

The table below describes the fields on this page:

Table 20

Field Name	Description	Status
Create a new web server accelerator radio button	Specifies the create operation. When selected, this option will enable the Name and DNS name fields directly under the Create section, while disabling all other fields.	
Name	The name of the new accelerator, limited to 8 characters. This name must be unique. If a name is entered that is already being used, the user is prompted to choose a unique name. This name should reflect the nature of the Web server or site. That is, it should be as descriptive as possible, so as to identify what Web site is being accelerated. For example, Novell, ETN1, Compaq, etc.	Required when creating an accelerator
DNS name	The DNS (Domain Naming Service) name is the name by which the Web server will be accessed. This is the name that is typically entered in a Web browser's location (URL) line. For example, the DNS name for Novell is www. novell.com.	Required when creating an accelerator
Edit an existing web server accelerator radio button	Specifies the edit operation. When selected, this option will enable the Name and DNS name fields directly under the Edit section, while disabling all other fields.	

Field Name	Description	Status
Name	Lists all the existing Web server accelerators currently found on the proxy server. If the user wishes to edit an existing accelerator, he or she must select a name from the list. The corresponding DNS name will appear in the DNS name field.	
DNS name	Shows the DNS (Domain Naming Service) name used by the accelerator shown in the list box to the left on the interface. This field is editable; the user can choose to edit the selected accelerator or change the DNS name. The DNS name is editable unless the selected accelerator is functioning as a path-based multi-homing child accelerator (in which case it inherits the DNS name from its master). NOTE: If you are changing the DNS name on an accelerator which has authentication enabled, the existing cookie domain may not be valid with the new DNS name if the DNS name is not a sub-domain of the cookie domain (resulting in the browser displaying a 403 Forbidden Error message). To check the cookie domain from the browser-based administration utility, select Configure > Web Server Accelerator > Modify > Authentication Options > Cookie domain.	Required when editing an accelerator
Delete an existing web server accelerator radio button	Specifies the delete operation. When selected, this option will enable the Name list box to allow the user to select an accelerator to be deleted.	

Field Name	Description	Status
Name	Shows a list of all the existing Web server accelerators currently found on the proxy server. In order to delete an existing accelerator, the user must select a name from the list. NOTE: Any accelerator that is a path-based multi-homing master will not show up in the list of accelerators that can be deleted until all of its child accelerators have been deleted.	
Enable multi-homing for this accelerator	If selected, this check box allows the user to set up path-based or domain-based multi-homing options for the accelerator.	

Controls for Accelerator Specification

This section describes the following buttons:

- ◆ “Delete” on page 232
- ◆ “Options” on page 232

Delete

The Delete button allows the user to delete an accelerator. This button will only be enabled when the user selects an accelerator from the Name list box. If the user selects an accelerator from the list and clicks Delete, a warning dialog box will appear and prompt him or her to continue to delete it. If the user clicks the No button on the warning dialog box, the accelerator will not be deleted. The deletion takes place only when the user clicks Finish at the end of the wizard session.

Options

The Options button displays the path-based multi-homing options setup dialog. This dialog box provides the necessary information to set up an accelerator to use a master for path-based multi-homing, as shown in [Figure 46](#):

Figure 46 Multi-homing Options Dialog Box

The following table describes the fields on this page:

Table 21

Field Name	Description	Status
Multi-home Master	Specifies the accelerator that will become the current accelerator's model, or master.	Required
Domain-based multi-homing	When selected, multi-homes the accelerator based on a DNS name.	Required if the Domain-based multi-homing button is selected
DNS name	The DNS name ends with the cookie domain of the selected multi-home Master. The user must add names to the front of the cookie domain so that this DNS name is unique for all accelerators.	

Field Name	Description	Status
Path-based multi-homing	When selected, uses the DNS name of the multi-home master. Multi-homing will be based on the sub-path match string.	Required if the Path-based multi-homing button is selected
Sub-path match string	Shows the string that the system uses to route the browser request to the proper Web server.	Required
Starts with	When selected, specifies that the text entered in the sub-path match string will be at the beginning of the path rule.	Required
Remove sub-path	Specifies that the sub-path match string is to be removed during routing to the Web server.	Optional
Ends with	When selected, specifies that the text entered in the sub-path match string will be at the end of the path rule.	Required

Accelerator/Web Server Page

On the fourth page of the wizard, the user can enable the accelerator and specify the Web servers that will be accelerated, as well as which ports they will be accelerated through and which proxy (accelerator) address and port will be used. See [Figure 47](#).

Figure 47 Accelerator/Web Server Page

iChain Web Server Accelerator Wizard - AcmeCorp

DNS name:

Enable this accelerator

Web server addresses

Web server port:

Accelerator IP addresses

137.65.159.217
 137.65.159.218
 137.65.159.221

Accelerator proxy port:

Act as a tunnel
 Tunnel only SSL traffic

< Back Next > Cancel Finish Help

The following table describes the fields on this page:

Table 22

Field Name	Description	Status
DNS name	Displays the DNS name of the accelerator currently under construction or modification. The accelerator name is shown in the caption following the hyphen after the page title. In Figure 47 , the accelerator name is AcmeCorp. The DNS Name field is non-editable and exists as an information source for the user.	Non-editable
Enable this accelerator	Selecting this check box enables the accelerator so that the data entered will be used in accelerating the specified Web servers. When the box is checked, all of the fields will be enabled unless the accelerator is a path-based multi-homing child (in which case only the Web server addresses and Web server port fields are enabled). If deselected, the fields are not editable.	Optional
Web server addresses	Lists all the Web servers being accelerated by this accelerator. The Web server address may be in one of two formats: IP address or DNS name. Entries in the list may be added or deleted by using the buttons to the right of the field on the interface.	Required when the accelerator is enabled
Web server port	Specifies the port on the Web server by which the proxy server will communicate with the Web server. The default value is 80.	Required when the accelerator is enabled.

Field Name	Description	Status
Accelerator IP addresses	Displays which IP addresses will do the actual acceleration. This list is populated by the proxy server. That is, it is not populated by the user, but shows the IP addresses on the proxy server that are available for accelerating. Selecting an entry in the table will toggle the check box on the line on or off. A checked entry will participate in accelerating the Web server.	Required when the accelerator is enabled
Accelerator proxy port	Specifies the port on the proxy server by which the proxy server will communicate with the Web server. The default value is 80.	Required when the accelerator is enabled.
Act as a Tunnel	The Act as a Tunnel option lets you create one or more accelerator services for the specific purpose of tunneling non-HTTP traffic through the appliance to the origin Web server. When this option is checked, the accelerator sets up a tunnel for all incoming traffic.	Optional
Tunnel only SSL traffic	If you decide to have the accelerator act as a tunnel, you can elect to have it tunnel only SSL traffic. If this option is checked the service will then verify that the address and port being accessed are actually an SSL Web site. If verification fails, the service will tear down the connection. NOTE: The SSL port number for the SSL tunnel is specified via the Accelerator Proxy Port and not the SSL listening port.	Optional

Controls for Accelerator/Web Server


This section describes the following buttons:

- ◆ “Add” on page 238
- ◆ “Delete” on page 239

Add

When the Add button is selected, a dialog box appears where the user can enter an IP address or DNS name for a Web server, as shown in [Figure 48](#):

Figure 48 Add Web Server Dialog Box



The following table describes the fields in this dialog box:

Table 23

Field Name	Description
IP address	If selected, an IP address is used to specify the Web server address. The address is entered in the field directly below this button, as shown in Figure 48 .
DNS name	If selected, a DNS name is used to specify the Web server address. The DNS name is entered in the field directly below this button on the interface, as shown in Figure 48 .

Delete

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

Accelerator Authentication Parameter Page

The fifth page of the wizard is where the accelerator authentication parameters are specified. The user enables or disables authentication, enables or disables Secure Exchange, and creates authentication profiles. See [Figure 49](#).

Figure 49 Accelerator Authentication Parameter Page

The screenshot shows the 'iChain Web Server Accelerator Wizard - AcmeCorp' dialog box. It features a blue title bar and a grey background. The main area contains several controls: a checked checkbox for 'Enable Authentication' with an 'Authentication Options' button to its right; an unchecked checkbox for 'Enable Secure Exchange' with a 'Secure Exchange Options' button to its right; a text field for 'SSL Listening Port' containing '443'; a dropdown menu for 'SSL Certificate Name' set to 'Auto'; a text field for 'Session Timeout Interval' containing '3' and a 'Minutes' dropdown; two unchecked checkboxes for 'Forward Authentication Information to Web Server' and 'Authenticate over HTTP'; a list box titled 'Authentication Profiles' containing 'Idap', 'ssl', 'rad', 'Idap2', 'Idaprad', and 'Idapcert', with 'ssl' and 'Idap2' checked; radio buttons for 'Multiple Profile Rule' set to 'AND profiles'; and an unchecked checkbox for 'Create another accelerator'. At the bottom, there are five buttons: '< Back', 'Next >', 'Cancel', 'Finish', and 'Help'.

The following table describes the fields on this page:

Table 24

Field Name	Description	Status
Enable Authentication	Checking this option forces a user to authenticate to access this Web server	Optional
Enable Secure Exchange	Checking this option enables Secure Exchange (formerly known as SSLizer). Advanced options for Secure Exchange are not currently available from the wizard, but can be set from the proxy server administration application.	Optional
SSL Listening Port	The SSL port that the user is redirected to for authentication if Secure Exchange is enabled.	Required if authentication or Secure Exchange is enabled
SSL Certificate Name	The certificate name for this accelerator. If the name does not appear in the drop-down list, it can be entered manually.	Required if Secure Exchange is enabled
Session Timeout Interval	The amount of time a connection can be inactive before re-authentication is required.	Required if authentication is enabled or Secure Exchange is enabled
Forward Authentication Information to Web Server	Sends username and/or password to the Web server	Optional
Authenticate over HTTP	Allows authentication over unencrypted HTTP instead of HTTPS. This feature is not compatible with RADIUS authentication profiles.	Optional

Field Name	Description	Status
Authentication Profiles	Each existing profile is listed, those in use appearing with a check box. At least one profile must be checked when authentication is enabled. When multiple profiles are in the list, more than one may be enabled. Currently, only Mutual SSL profiles may be used with LDAP or RADIUS profiles. LDAP and RADIUS profiles can not be used together.	Required if authentication is enabled.
Multiple Profile Rule	Only valid if multiple Authentication Profiles are checked. Selects whether only one profile is required (OR) or if all selected authentication methods need to be fulfilled before authentication is granted (AND). OR is the default when multiple profiles are checked.	

Controls for Accelerator Authentication Parameters

This section describes the following buttons:

- ◆ “Advanced Options” on page 241
- ◆ “Add” on page 241
- ◆ “Delete” on page 242
- ◆ “Edit” on page 242

Advanced Options

The Advanced Options button launches the Advanced Authentication Options dialog as shown in [Figure 50](#).

Add

The Add button launches the Add Authentication Profile dialog box.

Delete

The Delete button allows the user to delete an existing Authentication Profile.

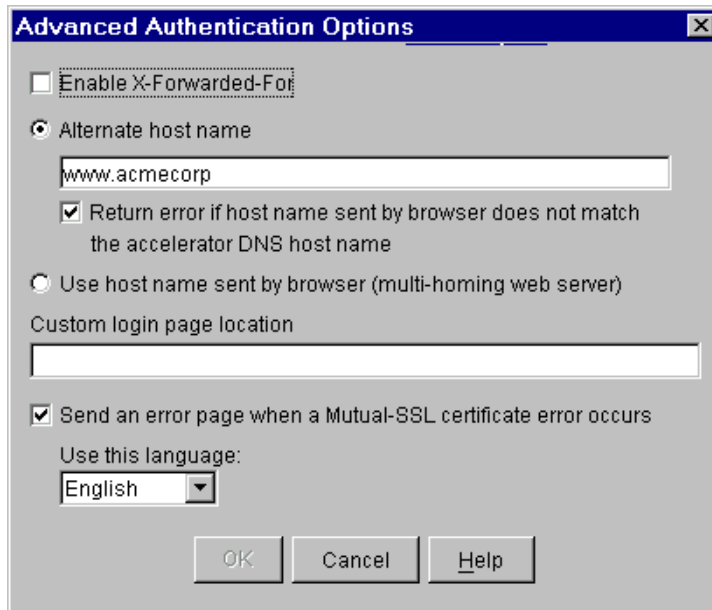
Edit

The Edit button launches the Modify Authentication Profile dialog box.

Advanced Authentication Options Dialog Box

The Advanced Authentication Options dialog box allows you to specify advanced authentication options, including options that are set under special circumstances. See [Figure 50](#).

Figure 50 Advanced Authentication Options Dialog Box



The following table describes the fields in this dialog box:

Table 25

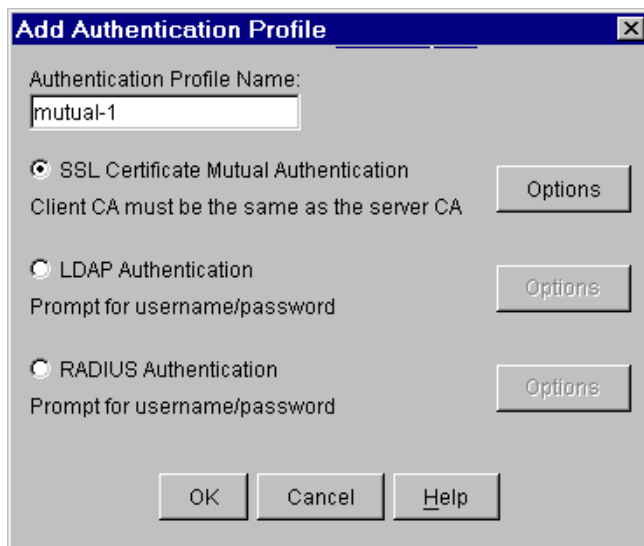
Field Name	Description	Status
Enable X-Forwarded-For	Checking the X-Forwarded-For option causes the appliance to either add information to an existing X-Forwarded-For or Forwarded-For header, or to create a header if one doesn't already exist.	Optional
Alternate host name	Checking this option causes the specified string to be substituted for the host name in the HTTP header before the request is forwarded to the Web server.	Optional
Return error if host name sent by browser does not match the accelerator DNS host name	Checking this option causes iChain Proxy Services to match the host name in the DNS header that came from the browser against the DNS name specified in this accelerator definition. If the names don't match, the request is not forwarded to the Web server. Instead, iChain Proxy Services returns an error to the requesting browser.	Optional
Use host name sent by browser	Checking this option preserves the host name in the HTTP header exactly as it came in the browser request.	Optional
Use custom login page	Checking this option allows for the specification of different login/error pages for this accelerator.	Optional
Location	Specifies the location of the login page for this accelerator. The login page must exist on the iChain Proxy server.	Required only if Use custom login page is checked.

Field Name	Description	Status
Send an error page when a Mutual-SSL certificate error occurs	Check this option to send a specific error page when a Mutual-SSL certificate error occurs. Otherwise a "page not found" message is always given.	Optional
Use this language	Select the language to be used for the error page. The default language is English.	Required only if "Send an error page when a Mutual-SSL certificate error occurs" is checked.

Add (Modify) Authentication Profile Dialog Box

The Add Authentication Profile dialog box allows you to name and create authentication profiles. The Modify Authentication Profile dialog box is exactly the same except for the dialog box title. See [Figure 51](#).

Figure 51 Add Authentication Profile Dialog Box



The following table describes the fields in this dialog box:

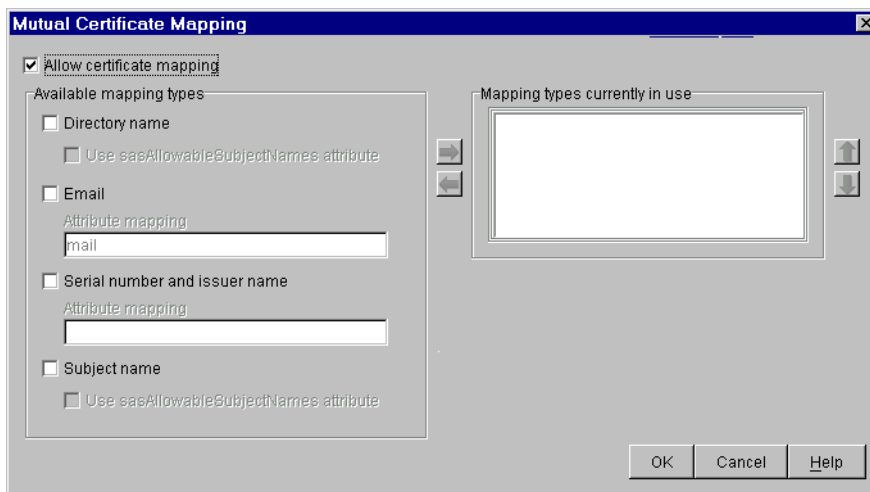
Table 26

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

Mutual Certificate Mapping Dialog Box

The Mutual Certificate Mapping dialog box allows you to configure certificate mapping types. See [Figure 52](#).

Figure 52 Mutual Certificate Mapping Dialog Box



The following table describes the fields in this dialog box:

Table 27

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

Field Name	Description	Status
Attribute mapping	Both the issuer name and the serial number need to be put into the same LDAP attribute of the user. The LDAP attribute that is used is specified in this field. The LDAP attribute can be any Case Ignore List or Cast Ignore String attribute of the user. If you are configuring your own attribute, make sure the attribute is added to the Person class.	
Subject name	A user in the LDAP Authentication tree matching the Subject Name field of the certificate will be checked first.	
Use sasAllowableSubjectNames attribute	If a user is not found with Subject name and Use sasAllowableSubjectNames is also enabled for directory name mapping, the LDAP Authentication tree will be searched for a user containing an sasAllowableSubjectName attribute matching the Directory Name in the Subject Alternative Name field of the certificate. If sasAllowableSubjectName is enabled, the LDAP Authentication tree should be configured so that there is no duplication of allowed names between users in the sasAllowableSubjectName attribute.	
Add button	The iChain Proxy Server can be configured to use any combination of the four mapping types. This button allows type to be added to the Mapping types currently in the use list.	
Remove button	Allows a type to be removed from Mapping types currently in the use list.	
Order up button	Allows for a mapping type within the Mapping types currently in the use list to be moved up. NOTE: When searching for a user with the configured mappings, the first user found will be the user that is used for authentication and access control, even if the other users will map to the same certificate. See “Using Certificate Mapping” on page 72 for more information.	

Field Name	Description	Status
Order down button	<p>Allows for a mapping type within the Mapping types currently in the use list to be moved down.</p> <p>NOTE: When searching for a user with the configured mappings, the first user found will be the user that is used for authentication and access control, even if the other users will map to the same certificate. See “Using Certificate Mapping” on page 72 for more information.</p>	

Controls for Authentication Profiles

This section describes the following buttons:

- ◆ [“LDAP Options” on page 248](#)
- ◆ [“RADIUS Options” on page 248](#)

LDAP Options

The LDAP Options button launches the LDAP options dialog box.

RADIUS Options

The RADIUS Options button launches the RADIUS options dialog box.

LDAP Options Dialog Box

The LDAP Options dialog box allows the user to specify LDAP authentication parameters. It is functionally identical to the corresponding dialog box in the iChain Proxy Server administration application. See [Figure 53](#).

Figure 53 LDAP Authentication Profile Options Dialog Box

LDAP Authentication Profile Options

Authentication Server Configuration

LDAP servers

IP Address	Port	Connection
137.65.159.220	389	Normal

LDAP Login Name Format

Use Distinguished Name
 Use User's Email Address
 Use LDAP Field Name

LDAP Search Base:

O=novell

Use anonymous bind for LDAP search
 Use username/password bind for LDAP search

Username: cn=admin,o=novell
 Password: ****
 Password confirmation: ****

LDAP Field Name: cn

OK Cancel Help

The following table describes the fields in this dialog box:

Table 28

Field Name	Description	Status
LDAP servers	This table lists the IP address, port, and connection type for all the LDAP servers used for this profile. Currently, the port and connection type must be the same for all servers.	Required
Use Distinguished Name	Selecting this option requires users to log in using their DS name.	Optional
Use User's Email Address	Selecting this option requires users to log in using their e-mail address.	Optional
Use LDAP Field Name	Selecting this option requires users to log in using some LDAP field.	Optional
LDAP Search Base (LDAP User Contexts)	This field will display as LDAP Search Base when either Use User's Email Address or Use LDAP Field Name is selected. It allows entry/deletion/modification of LDAP search bases or user contexts.	Required
Use anonymous bind for LDAP search	Bind anonymously to search the LDAP directory.	Optional
Use username/password bind for LDAP search	Bind with a proxy server to search the LDAP directory.	Optional
Username	Proxy username in LDAP format.	Required when Use username/password bind for LDAP search is selected
Password	Proxy user password.	Required when Use username/password bind for LDAP search is selected
Password confirmation	Proxy user password confirmation.	Required when Use username/password bind for LDAP search is selected
LDAP Field Name	LDAP field name to search for (only visible with Field Name).	Required when Use LDAP Field Name is selected.

Controls for Authentication Profile Options

This section describes the following buttons:

- ◆ “Add LDAP Server” on page 251
- ◆ “Delete LDAP Server” on page 251
- ◆ “Edit LDAP Server” on page 251
- ◆ “Add LDAP Context” on page 251
- ◆ “Delete LDAP Context” on page 251
- ◆ “Edit LDAP Context” on page 251

Add LDAP Server

The Add LDAP Server button allows you to launch the New LDAP Authentication Server dialog box.

Delete LDAP Server

The Delete LDAP Server button allows you to delete an authentication server from the list.

Edit LDAP Server

The Edit LDAP Server button allows you to launch the Modify LDAP Authentication server dialog box.

Add LDAP Context

The Add LDAP Context button allows you to launch the dialog box to add an LDAP Search Base/User Context (if DN is selected).

Delete LDAP Context

The Delete LDAP Context button allows you to delete an LDAP Search Base/User Context from the list.

Edit LDAP Context

The Edit LDAP Context button allows you to launch the dialog box to modify an LDAP Search Base/User Context (if DN is selected).

New LDAP Authentication Server Dialog Box

The New LDAP Authentication Server dialog box allows you to specify the parameters for new LDAP authentication servers. The Modify LDAP Authentication Server dialog box is exactly the same except for the dialog box title. See [Figure 54](#).

Figure 54 New LDAP Authentication Server Dialog Box



The following table describes the fields in this dialog box:

Table 29


Field Name	Description	Status
IP Address	The IP address of this LDAP server.	Required
Port	The LDAP port to communicate over. Currently, this is only modifiable for the first LDAP server in the list.	Required

Field Name	Description	Status
Use a secure connection (LDAP over SSL)	If checked, authentication information will be sent over LDAPS (encrypted). This is only modifiable for the first LDAP server.	Optional
Trusted root file	Specifies the trusted root file to be used for secure communications. This is only modifiable for the first LDAP server.	Required when Use a secure connection is selected.

Add LDAP Context Dialog Box

The Add LDAP Context dialog box provides the input of LDAP search bases or user contexts. The Modify LDAP Context dialog box is exactly the same except for the dialog box title. See [Figure 55](#).

Figure 55 Add LDAP Context



The following table describes the field on this dialog:

Table 30

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

Controls for Add LDAP Context

This section describes the following button:

- ◆ [“Object Browser” on page 254](#)

Object Browser

The Object Browser button allows you to launch an object browser to select the desired container.

Radius Options Dialog Box

The Radius Options dialog box allows you to specify the parameters for RADIUS profiles. This dialog box is functionally identical to the corresponding iChain Proxy Server administration application dialog box. See [Figure 56](#).

Figure 56 RADIUS Profile Options Dialog Box

The screenshot shows a dialog box titled "Radius Profile Options". It contains the following fields and values:

- RADIUS server address: 137.65.159.217
- RADIUS server listening port: 1812
- RADIUS server shared secret: SecretStuff
- RADIUS server reply time in seconds: 7
- RADIUS server resend time in seconds: 2
- User search base(s) for all RADIUS profiles: O=Novell

At the bottom of the dialog are three buttons: OK, Cancel, and Help.

The following table describes the fields in this dialog box:

Table 31

Field Name	Description	Status
RADIUS server address	The IP address of the RADIUS server.	Required
RADIUS server listening port	The port number on which the RADIUS server listens for incoming authentication.	Required
RADIUS server shared secret	The string the RADIUS server uses to verify that the appliance can request authentication of users.	Required
RADIUS server reply time in seconds	The total time the appliance will wait for a response from the RADIUS server before authentication fails. The default is 7 seconds.	Required
RADIUS server resend time in seconds	The interval in seconds between appliance requests to the RADIUS server. The default is two seconds. This means that the appliance could send three requests before the 7-second default limit expires and the authentication request fails.	Required
User search base(s) for all RADIUS profiles	Lists the contexts that the proxy server will use when searching for the user being authenticated when using non-Novell RADIUS authentication. This list applies to all RADIUS profiles, not just the current one being created or modified.	Optional

Controls for RADIUS Options

This section describes the following buttons:

- ◆ [“Add Search Base” on page 256](#)
- ◆ [“Delete Search Base” on page 256](#)

Add Search Base

The Add Search Base button allows you to launch an object browser to select the desired container.


Delete Search Base

The Delete Search Base button allows you to delete a search base from the list.

Protected Resource Page

The sixth page of the wizard provides the capability to create protected resources which are stored in the iChain Service Object (ISO) selected on page one of the wizard. Operation of this page is identical to the Protected Resource tab found in the ISO snap-in in ConsoleOne. See [Figure 57](#).

Figure 57 Protected Resource Page



The following table describes the field on this page:

Table 32

Field Name	Description	Status
Protected Resource	Table that displays all the protected resources defined for this ISO.	Required

Controls for Protected Resources

This section describes the following buttons:

- ◆ “Add New Protected Resource” on page 258
- ◆ “Delete Protected Resource” on page 258
- ◆ “Edit Protected Resource” on page 258
- ◆ “OLAC Parameters” on page 258

Add New Protected Resource

The Add New Protected Resource button allows you to launch the Add New Protected Resource dialog box.

Delete Protected Resource

The Delete Protected Resource button allows you to delete a protected resource from the table.

Edit Protected Resource

The Edit Protected Resource button allows you to launch the Modify New Protected Resource dialog box.

OLAC Parameters

The OLAC Parameters button allows you to launch the New OLAC Parameters dialog box.

Add New Protected Resource Dialog Box

The Add New Protected Resource dialog box allows the user to specify a new protected resource. The Modify Protected Resource dialog box is the same except for the dialog box title. See [Figure 58](#).

Figure 58 Add New Protected Resource Dialog Box

The following table describes the fields in this dialog box:

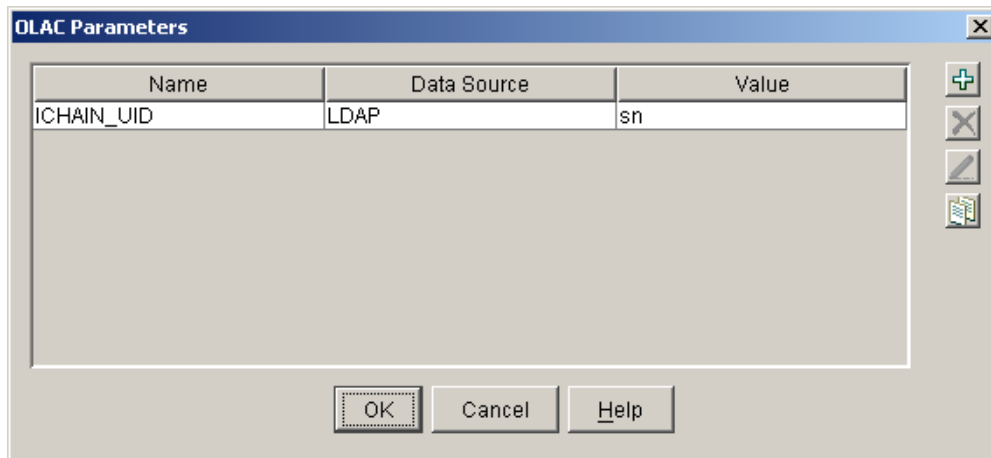
Table 33

Field Name	Description	Status
Resource Name	The name of this protected resource.	Required
URL Prefix	The URL prefix for this resource.	Required
Access	Defines whether the resource is secure (requires both authentication and access control), restricted (requires authentication only), or public (any user can access). This information is also shown in the table on the ISO tab listing all protected resources. Secure is the default.	
Pass Parameters in the Query String	If this option is selected, OLAC parameters are passed in the query string.	
Pass Parameters as Header Variables	If this option is selected, OLAC parameters are passed as header variables.	

OLAC Parameters Dialog Box

The OLAC Parameters dialog box gives the user the means to enter OLAC parameters for a protected resource. See [Figure 59](#).

Figure 59 OLAC Parameters Dialog Box



Controls for OLAC Parameters

This section describes the following buttons:

- ◆ “New OLAC Parameters” on page 260
- ◆ “Delete OLAC Parameters” on page 260
- ◆ “Edit OLAC Parameters” on page 261
- ◆ “Import Parameters” on page 261

New OLAC Parameters

The New OLAC Parameters button allows you to launch the New OLAC Parameters dialog box.

Delete OLAC Parameters

The Delete OLAC Parameters button allows you to delete the selected set of OLAC parameters.

Edit OLAC Parameters

The Edit OLAC Parameters button allows you to launch the Modify OLAC Parameters dialog box.

Import Parameters

The Import Parameters button allows you to launch the Import Parameters dialog box.

New OLAC Parameters Dialog Box

The New OLAC Parameters dialog box allows the user to specify a set of OLAC parameters including the name, data source, and value. The Modify OLAC Parameters dialog box is exactly the same except for the dialog box title. See [Figure 60](#).

Figure 60 New OLAC Parameters

The following table describes the fields in this dialog box:


Table 34

Field Name	Description	Status
Name	The name of the OLAC variable.	Required
Data Source	The data source from where the variable's data will be pulled.	Required
Value	The attribute name of the data as it is stored in the data source.	Required

Import Parameters Dialog

The Import Parameters dialog box provides a list of all protected resources that have OLAC parameters defined for them from which the user can select to import into the current protected resource. See [Figure 61](#).

Figure 61 Import Parameters Dialog Box



Combination Page from the Selected ACL Object

The seventh page of the wizard shows the combination page from the ACL object selected. Like the first page of the wizard, the user can either browse to an existing ACL rule or create a new rule in the container from which the wizard was launched. Once the ACL rule object has been selected, the other tables and lists can be populated as necessary. See [Figure 62](#).

Figure 62 Combination Page from Selected ACL Object

Access Control Rule:

Allowed URLs

Resource Name	URL Postfix
employee	/*
employee	

Excluded URLs

Resource Name	URL Postfix

Apply To List

- novell
- test

Exception List

- User0.test
- User1.test

< Back Next > Cancel Finish Help

The following table describes the fields on this page:

Table 35

Field Name	Description	Status
Access Control Rule	Specifies the name of the ACL rule object. Type an object common name and press Enter. If the object exists, the object's data will be read in. If not, you will be prompted to create a new object.	Optional

Field Name	Description	Status
Allowed URLs	Shows a table that specifies all URLs that this ACL rule grants access to.	Optional
Excluded URLs	Shows a table that specifies a subset of URLs in the Allowed URLs field that this ACL rule does not grant access to.	Optional
Apply To List	The users, groups, or containers that have access granted by this ACL rule.	Optional
Exception List	Users, groups, and containers that are excluded from access by this rule	Optional
Dynamic ACLs button	Dynamic Access Control Rules allow an administrator to set up an Access Control Rule based on a query of user attributes. See "Defining Dynamic Access Control Rules" on page 42 for more information.	Optional

Controls for the Selected ACL Object

This section describes the following buttons:

- ◆ ["ACL Rule Browser"](#) on page 265
- ◆ ["Add Allowed URL"](#) on page 265
- ◆ ["Delete Allowed URL"](#) on page 265
- ◆ ["Edit Allowed URL"](#) on page 265
- ◆ ["Add Excluded URL"](#) on page 265
- ◆ ["Delete Excluded URL"](#) on page 265
- ◆ ["Edit Excluded URL"](#) on page 265
- ◆ ["Add Item To Apply To List"](#) on page 265
- ◆ ["Delete Item From Apply To List"](#) on page 266
- ◆ ["Add Item To Exclusion List"](#) on page 266

- ◆ [“Delete Item From Exclusion List” on page 266](#)

ACL Rule Browser

The ACL Rule Browser button displays the name of the OLAC variable.

Add Allowed URL

The Add Allowed URL button allows you to launch the Add New Resource dialog box.

Delete Allowed URL

The Delete Allowed URL button allows you to delete the selected resource from the Allowed URL table.

Edit Allowed URL

The Edit Allowed URL button allows you to launch the Modify Existing Resource dialog box.

Add Excluded URL

The Add Excluded URL button allows you to launch the Add New Resource dialog.

Delete Excluded URL

The Delete Excluded URL button allows you to delete the selected resource from the Excluded URL table.

Edit Excluded URL

The Edit Excluded URL button allows you to launch the Modify Existing Resource dialog box.

Add Item To Apply To List

The Add Item To Apply To List button allows you to launch an object selection browser to select objects to add to the list.

Delete Item From Apply To List

The Delete Item From Apply To List button allows you to delete the selected object from the list.

Add Item To Exclusion List

The Add Item To Exclusion List button allows you to launch an object selection browser to select objects to add to the list.


Delete Item From Exclusion List

The Delete Item From Exclusion List button allows you to delete the selected object from the list.

Add New Resource Dialog Box

The Add New Resource dialog box allows the user to specify a resource to be added to the Allowed URL or Excluded URL lists. The Modify Existing Resource dialog box is exactly the same except for the dialog title. See [Figure 63](#).

Figure 63 Add New Resource Dialog Box



The following table describes the fields on this page:

Table 36

Field Name	Description	Status
Resource Name	The name of the resource.	Optional unless URL Postfix is left blank
URL Postfix	The string that specifies what directories or files are accessible through this resource.	Optional unless Resource Name is left blank

Controls for Add New Resource

This section describes the following button:

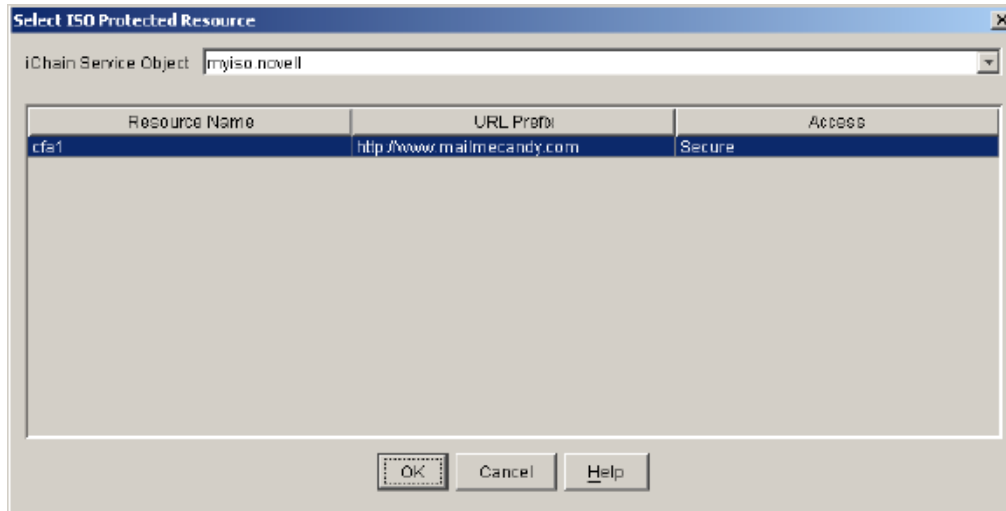
- ◆ [“Resource Browser” on page 267](#)

Resource Browser

The Resource Browser button allows you to launch the Select ISO Protected Resource dialog.

Select ISO Protected Resource Dialog Box

The Select ISO Protected Resource dialog box provides a list of iChain Service Objects and their associated protected resources from which the user can select to add to either the Allowed URL or Excluded URL list. The list is a drop-down list which will contain all iChain Service Objects in the currently selected context. See [Figure 64](#).

Figure 64 Select ISO Protected Resource Dialog Box

The following table describes the field in this dialog box:

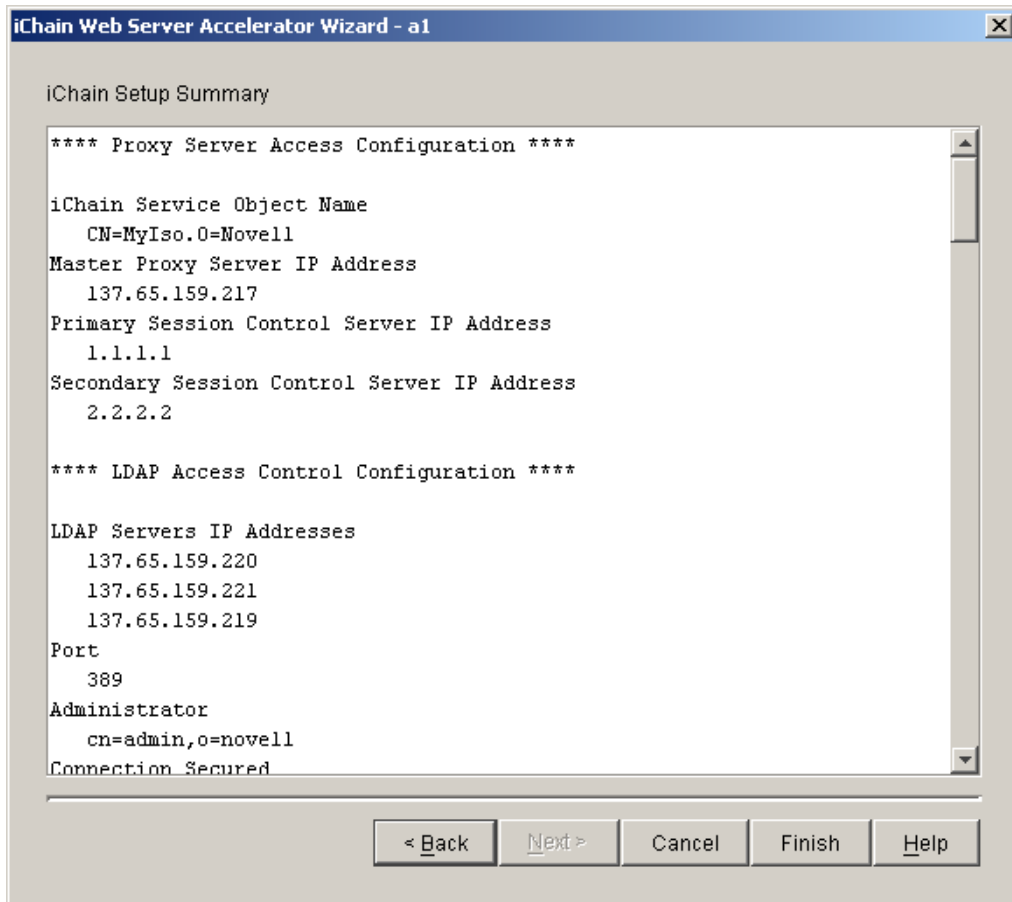
Table 37

Field Name	Description	Status
iChain Service Object	Lists all the existing iChain Service Objects.	Required

Summary Page

The eighth page of the wizard is the summary page. On this page, all of the changes and settings for the ISO and Web server are displayed in a scrollable text box which the user can review prior to submitting the changes to the proxy server. See [Figure 65](#).

Figure 65 Summary Page



The following table describes the field on this page:

Table 38

Field Name	Description
iChain Setup Summary	A scrollable text area that contains all the changes and settings.

Controls for the Summary Page

This section describes the following buttons:

- ◆ “Back” on page 270
- ◆ “Cancel” on page 270
- ◆ “Finish” on page 270

Back

The Back button will take you back to the previous page in the wizard. All changes are preserved.

Cancel

The Cancel button will cancel all changes and return to ConsoleOne. No prompt is given before leaving the wizard. No changes are preserved.


Finish

The Finish button allows you to display the Confirm Proxy Server Configuration Update message box. The user will be prompted to proceed or cancel.

Confirm Proxy Server Configuration Update Dialog Box

The Confirm Proxy Server Configuration Update dialog box prompts the user to answer yes or no before proceeding to update the proxy server. If the user selects the Yes button, the configuration is sent to the proxy server. If the user selects No, he or she is taken back to the Summary Page and no changes are made to the proxy server. See [Figure 66](#).

Figure 66 Confirm Proxy Server Configuration Update Dialog Box



Controls for Confirming the Proxy Server

This section describes the following buttons:

270 Place Book Title Here

- ◆ “Yes” on page 271
- ◆ “No” on page 271

Yes

The Yes button saves the configuration changes to the proxy server and returns you to ConsoleOne. This finishes the wizard session. If the configuration is sent successfully, the Proxy Server Configuration Update Complete dialog box is displayed.


No

The No button will return you to the wizard without saving any changes to the proxy server. All changes made in the wizard are preserved.

Proxy Server Configuration Update Complete Dialog Box

The Proxy Server Configuration Update Complete dialog box confirms that the configuration was sent to the proxy server. When you select OK at this dialog box, you are returned to ConsoleOne. See [Figure 67](#).

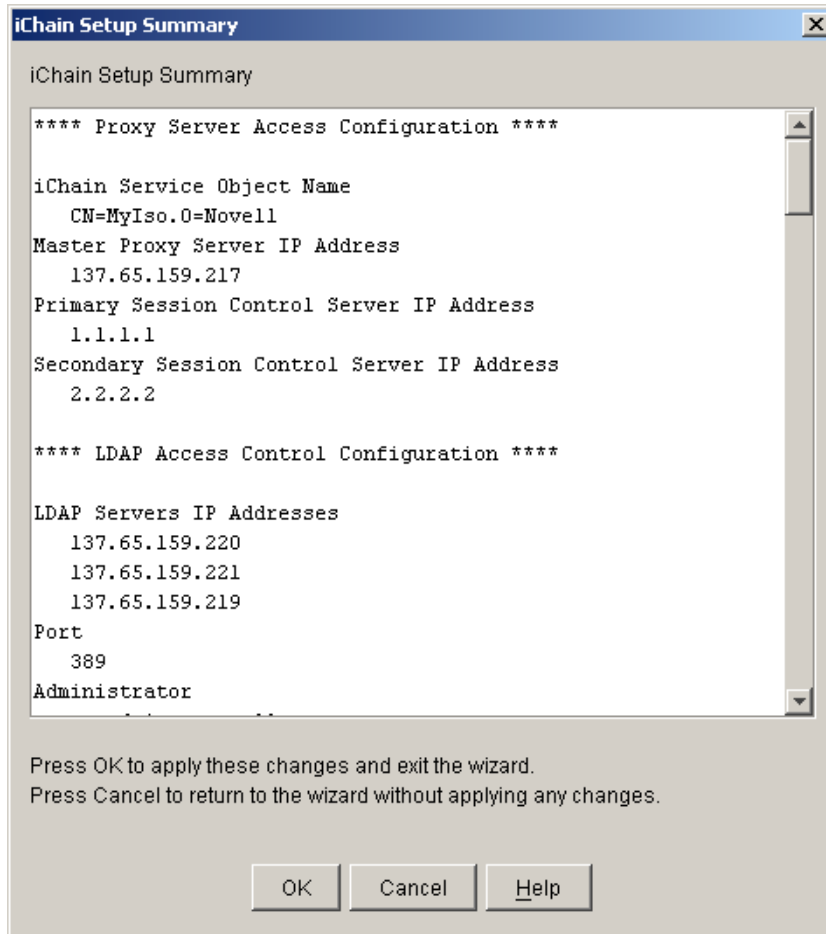
Figure 67 Proxy Server Configuration Update Complete Dialog Box



iChain Setup Summary Dialog Box

The iChain Setup Summary dialog box is displayed if the Finish button is selected on any wizard page before reaching the summary (eighth) page of the wizard. The dialog box is virtually the same as page 8 of the wizard, but it allows you to make changes to settings and then exit without going through the rest of the wizard. See [Figure 68](#).

Figure 68 iChain Setup Summary Dialog Box



The following table describes the field on this dialog box:

Table 39

Field Name	Description
iChain Setup Summary	Displays a scrollable text area that contains all the changes and settings you've made.

272 Place Book Title Here

Controls for iChain Setup Summary

This section describes the following buttons:

- ◆ “OK” on page 273
- ◆ “Cancel” on page 273

OK

The OK button saves the configuration changes you have made to the proxy server and returns you to ConsoleOne. This finishes the wizard session. No confirmation box is displayed when the OK button is selected.

Cancel

The Cancel button will return you to the wizard without saving any changes to the proxy server. All changes made in the wizard are preserved.

274 Place Book Title Here

B iChain Proxy Server Management

The Novell® iChain® Proxy Server is a key component of the iChain infrastructure. The iChain Proxy Server is optimized to perform the functions needed for your iChain infrastructure. This appendix details the configuration and management of the iChain Proxy Server.

Proxy Server Browser-Based Administration Tool

The iChain Proxy Server supports a browser-based administration utility allowing you to manage and administer your iChain Proxy Server from a browser. To launch the utility, you simply access a special management URL on the iChain Proxy Server. The URL must contain either the 10-net management address or the IP address you have already configured for the server, followed by :1959/appliance/config.html. For example,

<http://10.1.1.1:1959/appliance/config.html>

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

Prerequisites for Running the Management Tool

You need the following:

- ◆ An proxy server that has been initialized and is currently running
- ◆ A Java-enabled browser, such as Netscape* Navigator* 4.07 (or higher), Netscape Communicator* 4.5 (or higher), or Internet Explorer 4.01 (or higher) running on your workstation
- ◆ SSL 2.0 and SSL 3.0 (where available) enabled on the browser

- ◆ A network or crossover cable connection to the proxy server
- ◆ The IP address of the proxy server

Once the appliance has been configured with an IP address and mask, a gateway server, and a DNS server, you can administer it over the network via any client that can communicate with it over IP.

Until you have completed that configuration, however, you must use a crossover cable and a client with the following constraints:

- ◆ Client IP address set to 10.1.1.2 (or another available 10-net IP address) with a mask of 255.255.255.255
- ◆ Client gateway address set to 10.1.1.1 (the management address of the appliance)
- ◆ Client DNS server address set to 10.1.1.1

Starting the Management Tool

- 1** Start the browser on your client workstation.
- 2** Point the browser to the URL of the appliance you want to manage.

The URL must contain either the 10-net management address or an IP address you have already configured on the appliance, followed by :1959/appliance/config.html, for example:

`http://10.1.1.1.:1959/appliance/config.html`

- 3** Accept the SSL certificate.

IMPORTANT: You must have SSL 2.0 and SSL 3.0 (where available) enabled in your browser. Otherwise, the browser will display an error indicating that the page cannot be displayed.

- 4** Enter a password if you have previously specified one for the appliance.

Applying and Canceling Changes

As you make changes to appliance parameters in the management tool, these changes are tracked and accumulated in a buffer until you either apply or cancel them. You can make changes in multiple tabs and wait to apply them all at once.

This does not apply to the Actions and Date/Time tabs. Changes in these tabs are immediately effective. If you change the NTP server, the appliance time will change with the next synchronization cycle (normally about 15 minutes).

Except in the cases just mentioned, clicking Apply commits all changes made in any tab since the last time you started the appliance or clicked Cancel. Clicking Cancel cancels all changes made since the last time you started the appliance or clicked Apply. Clicking Cancel is also a quick way of requesting that the appliance reread the currently displayed settings.

Once you click Apply or Cancel, the action cannot be undone.

The Help Button

Click the Help button in the left frame to display the online documentation with a table of contents in the left frame. To navigate through the documentation, click the titles in the table of contents.

Encryption

If you have specified passwords for appliance management purposes, communications regarding the password are transmitted through HTTPS. All other communications with the appliance are not normally encrypted.

The Home Panel

The Home panel provides access to general information regarding the appliance, such as the caching system version currently running and the general health of the current configuration.

Introduction Tab

Path: Home > Introduction

Figure 69 Introduction Tab

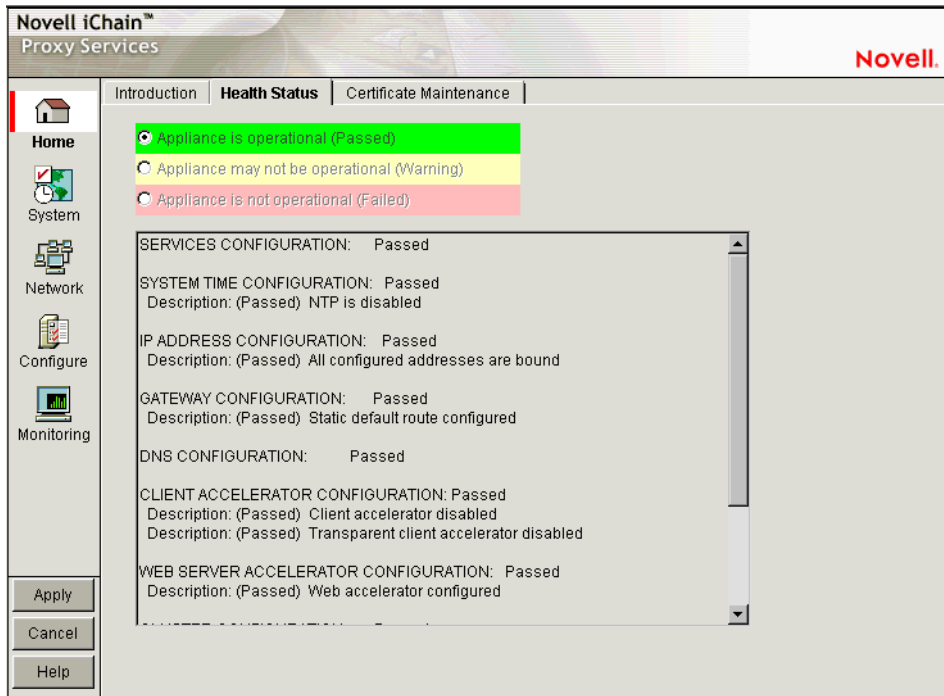


The Introduction tab displays iChain Proxy Server information, such as version, build number, and licensing information.

Health Status Tab

Path: Home > Health Status

Figure 70 Health Status Tab



The Health Status tab indicates general status of appliance configurations, including which services are currently configured and the operational status of selected services.

A green status indicates iChain Proxy Services has not detected any configuration discrepancies.

A yellow status indicates iChain Proxy Services might be functioning sub-optimally due to configuration discrepancies.

A red status indicates that the iChain Proxy Services configuration might be incomplete or wrong.

Services Configuration: Reports the overall configuration status of all proxy services.

System Time Configuration: Reports the current NTP status.

IP Address Configuration: Reports the status of IP address assignments to network interfaces. iChain Proxy Services requires at least one IP address assignment for proper operation.

Gateway Configuration: Reports the status of the next hop gateway configuration. Without proper gateway configuration, the appliance cannot connect to origin Web servers.

DNS Configuration: Reports the status of DNS server configuration and connectivity to configured DNS servers. Without access to a DNS server, proxy services cannot function properly.

Client Accelerator Configuration: Reports the status of the forward and transparent proxy configuration. If browser clients pointing to this appliance as their appliance have Web browsing problems, check the status here and in [“Services Tab” on page 349](#).

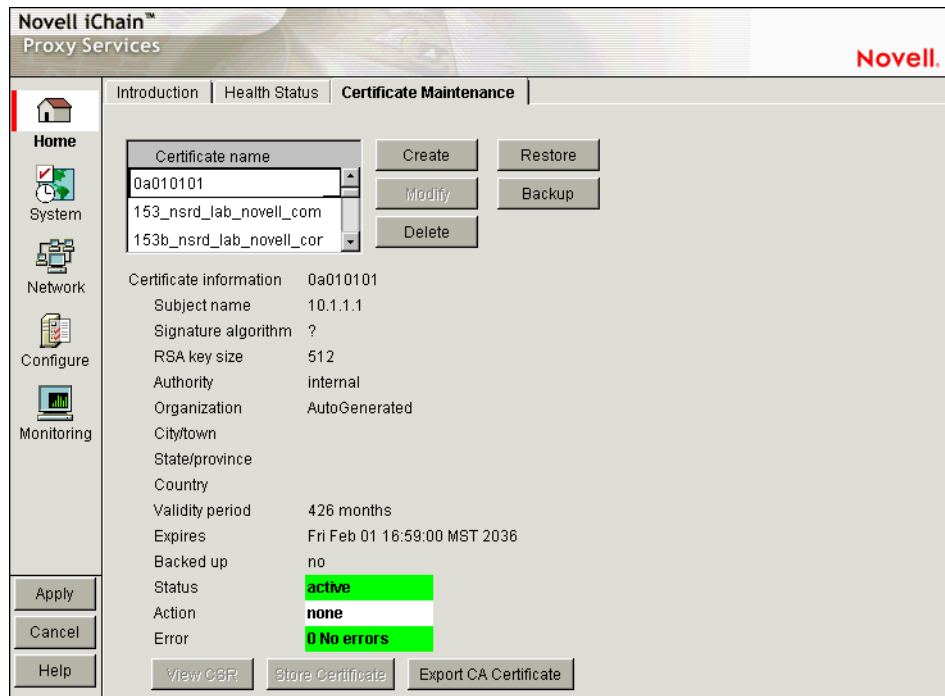
Web Accelerator Configuration: Reports the status of the Web server accelerator configurations. If browser clients have problems accessing a site being accelerated by this appliance, check the status of the Web server accelerator service in this section and in [“Services Tab” on page 349](#).

Filtering Configuration: Reports the status of filter service configurations. If filtering is enabled and waiting for a rating list to be downloaded, the status indicates that filtering is not active and the health status is yellow (warning).

Certificate Maintenance Tab

Path: Home > Certificate Maintenance

Figure 71 Certificate Maintenance Tab



The Certificate Maintenance tab lets you create, delete, back up, restore, and view authentication certificates stored on the appliance. This includes internal certificates generated by the appliance's certificate authority (CA) and external certificates generated by an external CA, such as VeriSign*. For more information, see [“Managing Appliance Certificates” on page 123](#).

Certificate Name: A list of certificates created on the appliance.

Certificate Information: Information for the certificate selected.

View CSR: Use this option to display the certificate signing request of a certificate you have created.

This is used to request a certificate from a certificate authority. For more information, see [“Obtaining a Certificate from an External CA” on page 126](#).

Store Certificate: Use this option to store certificate information received from a certificate authority. For more information, see [“Obtaining a Certificate from an External CA” on page 126](#).

Export CA Certificate: Use this option to display a certificate authority's certificate. For more information, see [“Viewing \(Exporting\) a Certificate's CA” on page 129](#).

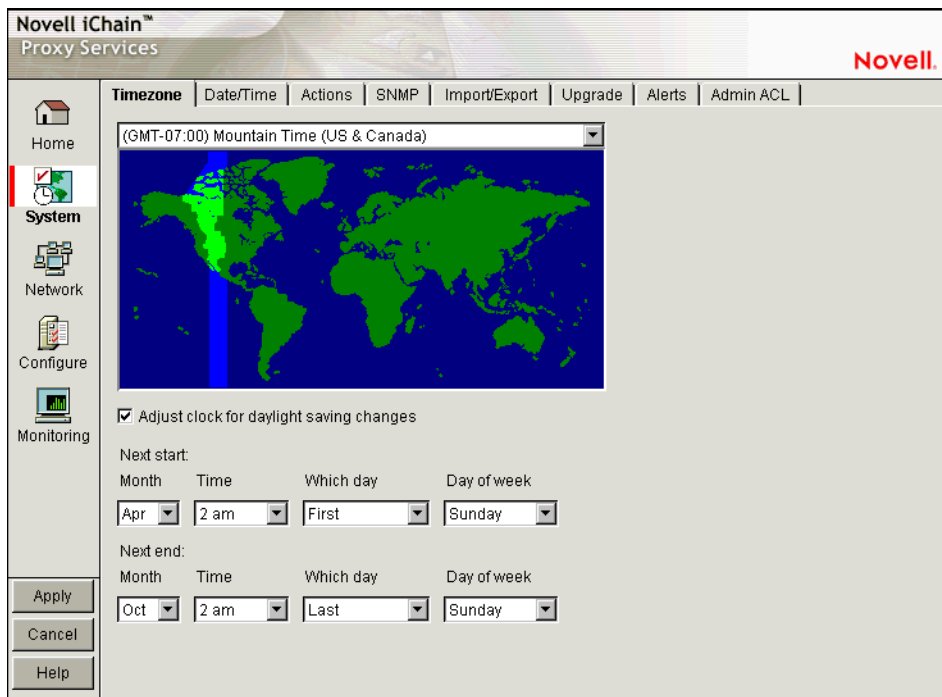
The System Panel

The System panel lets you perform actions that affect the appliance system in a general way. Use the tabs in this panel for changing and setting system time, changing the system password, restarting the appliance, upgrading the system, etc.

Timezone Tab

Path: System > Timezone

Figure 72 Timezone Tab



The Timezone tab lets you specify a time zone for the appliance. It also lets you specify exactly when daylight saving time begins and ends.

The Time Zone Map: Lets you select a time zone for the appliance by clicking the map. The granularity offered through this method is adequate for most appliance installations. Additional flexibility in setting time is available on this tab and from the command line. For more information on command line options, refer to the command line help for the set command and the time zone argument. See [“Command Line Reference” on page 355](#) for more information.

Adjust Clock for Daylight Saving Changes: If you check this option, the appliance clock begins daylight saving time and resumes standard time on the dates and times defined in the fields below Next Start and Next End. For example, most U.S. time zones begin daylight saving on the first Sunday of April at 2:00 a.m. and resume standard time on the last Sunday of October at 2:00 a.m.

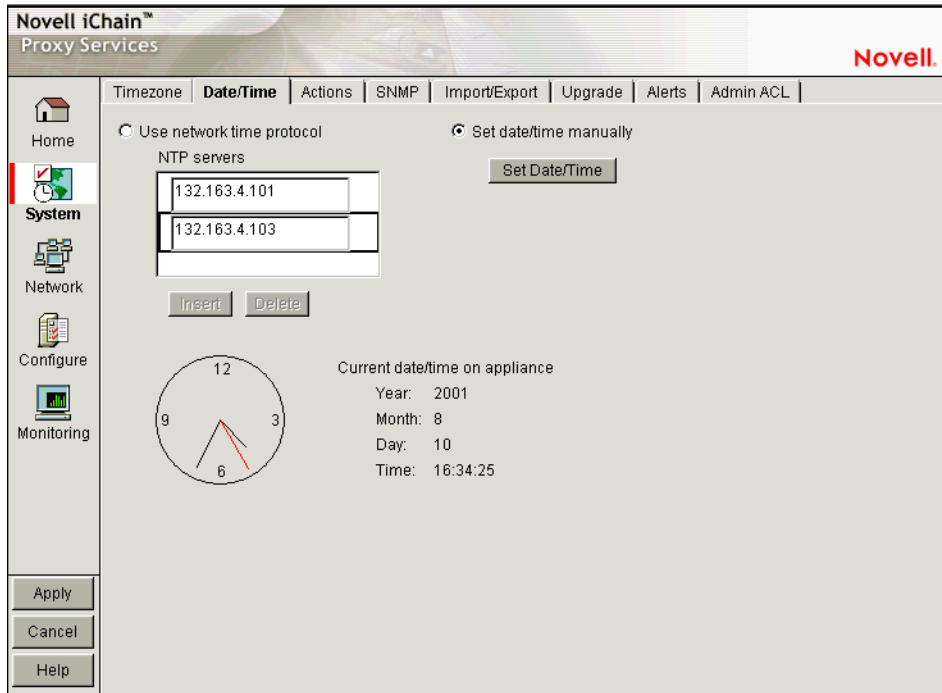
To set nonstandard daylight saving parameters in this tab, select the start and end field values for Month, Time, Which Day, and Day of Week in their respective drop-down lists.

To set nonstandard parameters from the command line, refer to command line help for the set command and the dsstart, dsend, and dstime arguments. See the instructions for using command line online help in [“Command Line Reference” on page 355](#) for more information.

Date/Time Tab

Path: System > Date/Time

Figure 73 Date/Time Tab



The Date/Time tab lets you set the appliance system time so that the time stamps in cache logs are accurate and valid. An ISP, for example, might bill customers based on their access to the appliance. Accurate log time stamps are essential to issuing credible billing statements.

NOTE: iChain Proxy Services stamps log entries with Greenwich Mean Time (GMT). If the appliance is using an NTP server, the GMT stamp comes from that server. If the appliance is using a manually set time, iChain Proxy Services assumes the time is accurate and calculates the GMT value based on the appliance's time zone and daylight saving settings.

Use Network Time Protocol: Checking this option turns the network time protocol on or off. This enables the appliance to synchronize its system time with an NTP server. Using an NTP server makes appliance cache log time stamps as reliable as possible. This can be especially important if you use the logs for customer billing. The appliance comes with two sample NTP servers: 132.163.4.101 and 132.163.4.103. You can remove these or add additional NTP servers.

IMPORTANT: When you specify an NTP server, synchronization between the NTP server clock and the appliance clock might not be immediate.

If the NTP server clock has an earlier time than the appliance clock, iChain Proxy Services will slow the appliance clock down until the two are synchronized. This provides for proper incrementation of log files and other time-sensitive information during the synchronization process.


If the NTP server clock is later than the appliance clock, synchronization between the two will generally be immediate. However, in certain situations you might observe the appliance clock incrementing by 600-minute intervals. This is normal system behavior.

The fact that the Apply button changes from Wait back to Apply indicates only that the NTP configuration change has been made, not that the appliance clock is fully synchronized with the NTP server.

If the above features are problematic in your situation, you can set appliance time manually to the target time and then re-enable the NTP feature.

Set Time Manually: The dialog box in [Figure 74](#) appears when you select this radio button and click Set Time. Set the date and time using the drop-down lists. Clicking OK immediately resets the system clock.

Figure 74 Set Time Dialog Box

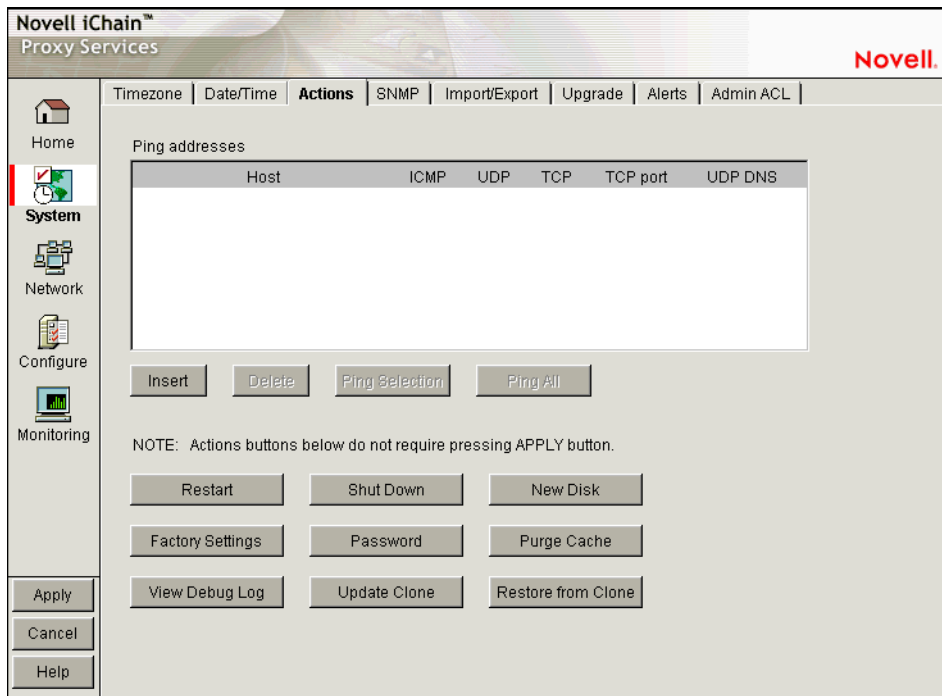


Use this option if NTP is not available to your appliance or you need to set a specific time for some reason.

Actions Tab

Path: System > Actions

Figure 75 Actions Tab



The Actions tab lets you perform tasks related to the appliance hardware and software.

NOTE: Most changes made in the browser-based management tool are not effective until you click Apply. However, changes made in the Actions tab are immediately effective.

Ping Addresses: You can check network connections using appliance ping functions by adding target hosts and port numbers to this list and then clicking Insert. Follow the address with a colon and a port number (an integer value from 0 to 65535) you want to ping. Using a port number lets you check whether a host has HTTP support (port 80), HTTP forward proxy support (port 8080), DNS support (port 53), ICP peer/parent support (port 3130), etc.

Restart: Shuts down the caching system and then restarts it. Configuration settings are retained but cached objects are removed.

Shut Down: Shuts down the caching system. The hardware remains turned on until manually powered off.

When the appliance has successfully shut down, a series of three beeps is repeated until the box is powered off.

New Disk: Scans for new disks that the system has not auto-detected.

Factory Settings: Resets the appliance to its original factory configuration as explained in [“Restoring Factory Settings” on page 167](#). Passwords are retained. If you want to preserve other settings for later use on this or another appliance, see [“Import/Export Tab” on page 291](#).

Password: See [“Password Dialog Box” on page 287](#).

Purge Cache: See [“Purge Cache Dialog Box” on page 288](#).

View Debug Log: When an appliance experiences an abnormal shutdown due to a configuration error or other problem, iChain Proxy Services logs critical history information associated with the shutdown. Clicking this button displays the log in a separate browser window. You can then save the log file locally, print it, and e-mail it to Technical Support.

Update Clones: Each appliance stores a clone image that, initially, is the same as the factory image. If the appliance experiences an abnormal shutdown four times within a half hour period, or if it is restarted six times within a half hour period, iChain Proxy Services assumes the current configuration is faulty and automatically replaces it with the clone image.

You can overwrite the default clone image with an alternate configuration by selecting this option.

IMPORTANT: This process reboots the appliance, causing a temporary interruption of services.

Restore from Clones: Selecting this option restores the appliance to the configuration of the clone image (either the original factory clone image or an alternate clone image you have saved using the Update Clones option).

IMPORTANT: This process reboots the appliance, causing a temporary interruption of services. If the image being restored is the original factory clone image, you will also need to reconfigure proxy services on the appliance or use a .NAS file to restore these. See [“Restoring the Appliance to the Clone Image” on page 168](#) and [“AUTOLOAD.NAS” on page 163](#).

Password Dialog Box

Path: System > Actions > Password

Figure 76 Password Dialog Box



IMPORTANT: It is critical that you assign system passwords when initially configuring the appliance. Otherwise, access through Telnet, FTP, and the browser-based management tool is not restricted.

You can specify passwords for two users with different access privileges.

Users logging in using the View user password can view everything in the browser-based management tool and execute get commands from the command line. The Apply function and the set command are not available.

Users logging in using the Config user password have full access to the browser-based tool and the command line interface.

Change: Immediately changes the password for the user selected.

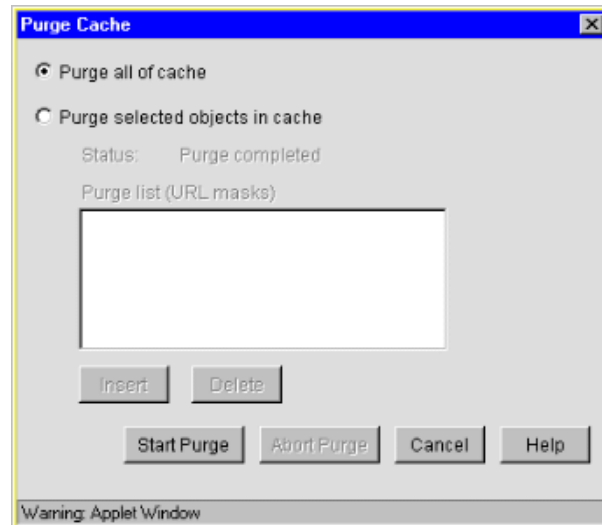
Remove: Removes (sets to null) the password for the user selected.

Appliance passwords are case-sensitive.

Purge Cache Dialog Box

Path: System > Actions > Purge Cache

Figure 77 Purge Cache Dialog Box



You can remove all cached objects from the appliance's cache, or you can perform a limited purging of cached objects based on URL masks. Purging cannot be undone.

Purge All of Cache: Starting the purge with this option selected will purge everything from the appliance's cache.

Purge Selected Objects in Cache: Selecting this option allows you to specify URL patterns or masks for the pages or sites whose objects you want to purge. When defining the masks, keep in mind that the appliance interprets everything in the URL mask between the asterisk wildcard (*) and the following delimiter as a wildcard. Delimiters include the forward slash (/), the period (.), and the colon (:) characters.

This option also allows purging of cache objects whose URL contains a specified query string or cookie. This mask is defined by placing a question mark (?) at the start of the mask followed by text strings and wildcards as necessary. String comparisons are not case sensitive. For example, ?*=SPORTS will purge all objects with the text "=SPORTS" or any other combination of upper and lower case letters for "=SPORTS" following the question mark in the URL.

SNMP Tab

Path: System > SNMP

Figure 78 SNMP Tab

The screenshot shows the Novell iChain Proxy Services web interface. The top navigation bar includes tabs for Timezone, Date/Time, Actions, **SNMP**, Import/Export, Upgrade, Alerts, and Admin ACL. The left sidebar contains icons for Home, System (selected), Network, Configure, and Monitoring. The main content area is titled 'SNMP' and contains the following configuration sections:

- Monitor state:**
 - No community may read
 - Specified community may read:
- Control state:**
 - No community may write
 - Specified community may write:
- Trap state:**
 - Do not send traps
 - Trap community name:
- IP addresses of management stations:**
- Node name for SNMP:**

At the bottom of the configuration area, there are three buttons: Hardware Description, Physical Location, and Human Contact. Below these buttons is a red note: "NOTE: The appliance will be restarted when APPLY button is pressed in order for the changes to take effect." At the very bottom of the interface are buttons for Apply, Cancel, and Help.

The SNMP tab lets you configure the appliance with basic SNMP information so the appliance can communicate with your SNMP management workstations.

The appliance's SNMP implementation follows the ISO SNMP version 1 standard outlined in *RFC 1067: A Simple Network Management Protocol*.

When SNMP-enabled appliance components start up, they register with the system. When the system receives a request for a specific SNMP parameter, it knows which component to contact to obtain the information.

Each appliance contains an iCHAIN.MIB file in the SYS:\ETC\PROXY\DATA directory. To see a list of standard SNMP parameters, retrieve this file using the FTP get command and compile it for use with your SNMP management software.

If you specify a trap community name and specify an SNMP management workstation in the SNMP tab, all alerts you check in the Alerts Tab (see “Alerts Tab” on page 295) are automatically sent as SNMP traps even if you have not configured syslog or e-mail alert notification on the Alerts tab.

Monitor State: Allows you to specify community Read access and the community name or password to be used. Community names must contain ASCII characters only and must not have spaces.

Control State: Allows you to specify community Write access and the community name or password to be used. Community names must contain ASCII characters only and must not have spaces.

IMPORTANT: The default name or password for the control community is No, meaning that control access is turned off. You can reset this value. However, this is not normally recommended, since the control community password is stored as clear text and could allow unauthorized write access to SNMP parameters on the appliance.

Trap State: Allows you to either specify that traps are not sent, or to specify a community (location, IP octets, or other identifier) from which traps are sent to the management stations you designate. Community names must contain ASCII characters only and must not have spaces.

IP Addresses of Management Stations: One or more management station IP addresses, separated by semicolons.

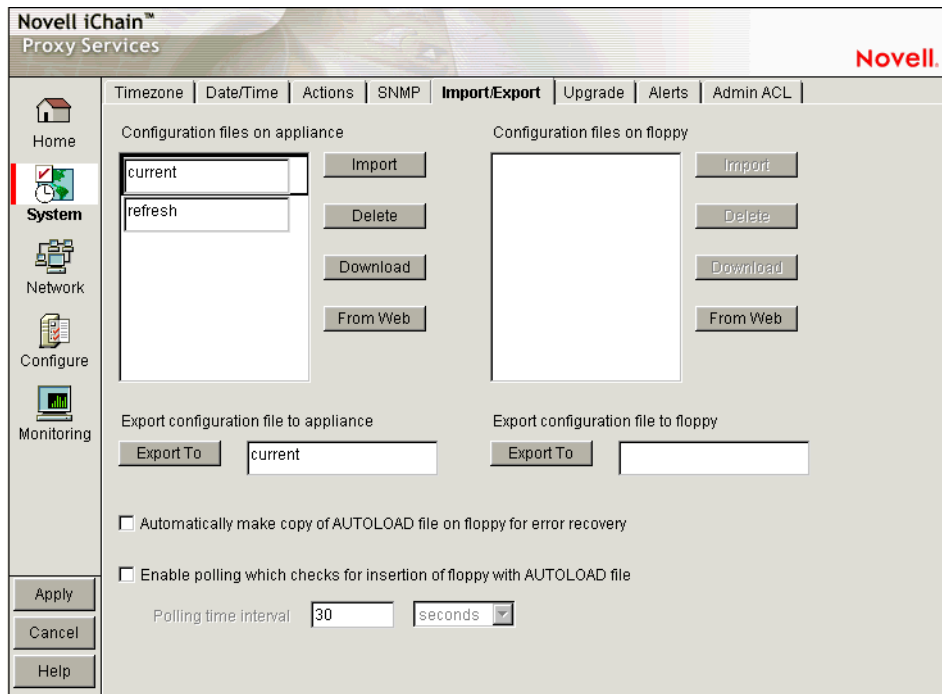
Node Name for SNMP: Lets you specify a node name for management of the appliance through SNMP.

The buttons below the node name field let you enter additional information regarding the hardware, the appliance’s physical location, and information regarding the person responsible for the appliance.

Import/Export Tab

Path: System > Import/Export

Figure 79 Import/Export Tab



The Import/Export tab lets you manage appliance configuration files on the appliance and on floppy disk.

IMPORTANT: You should have a backup configuration file named something other than AUTOLOAD.NAS. For further details, see [“Backing Up the Appliance Configuration” on page 166](#).

Configuration Files on Appliance: Displays a list of all of the configuration files stored on the appliance. These files are used to configure the appliance instantly, rather than using the GUI, command line, or Telnet to make individual changes. The appliance automatically updates the configuration file, CURRENT, each time you apply a change to iChain Proxy Services. The .NAS extension of these files is not shown in this list but is supplied by the server.

You can download, import, and delete any file in this list. You can also copy a configuration file from any URL to the appliance. The Download option opens the file in a separate browser window. The Import option changes the appliance configuration from its current settings to those contained in the selected configuration file. The Delete option removes the selected

configuration file from the appliance. The From Web option lets you specify the URL for the configuration file being copied to the appliance. If the file is in a secure area or is being downloaded using SSL (HTTPS:), you can also enter a username and password for authentication.

Configuration Files on Floppy: Displays a list of all the configuration files stored on the floppy disk located in the appliance's floppy drive. You can download, import, and delete any file in this list. You can also copy a configuration file from any URL to a floppy in the appliance's floppy drive. The previous section contains more detail regarding the Import, Delete, Download, and From Web options.

IMPORTANT: It is easy to confuse the diskette in the appliance's floppy drive with one located in your configuration workstation. Only the former is accessible through the browser-based management tool.

Export Configuration File to Appliance / Export Configuration File to Floppy: Clicking the button under one of these titles creates a configuration file on the appliance or on the diskette in the appliance's floppy drive.

Files saved using the Export feature contain the complete configuration of the appliance at the time of export. The default filename is CURRENT.NAS. You can specify any DOS-style eight-character name. Names are not case-sensitive. Each file has a .NAS extension that is not displayed in the list or specified when the file is created, but is automatically appended by the system.

Automatically Make Copy of AUTOLOAD File on Floppy for Error Recovery: Creates an AUTOLOAD file on a floppy diskette in the appliance when the configuration is changed. The appliance uses the AUTOLOAD file during error recovery to restore the system configuration after a successful restoration of the clone image. The AUTOLOAD file is also used when you invoke the import floppy command from a Telnet or command line session.

You can also export or import configuration files other than AUTOLOAD to a floppy diskette from the command line or Telnet interface.

Enable Polling which Checks for Insertion of Floppy with AUTOLOAD File: Allows the polling of the floppy diskette during normal operation of the appliance. If a floppy containing an AUTOLOAD.NAS is inserted when this option is on, iChain Proxy Services automatically applies the settings in the AUTOLOAD file.

Although this feature is useful, there is a short suspension of proxy activity while iChain Proxy Services polls the floppy drive.

Upgrade Tab

Path: System > Upgrade

Figure 80 Upgrade Tab

Novell iChain™ Proxy Services	
Timezone Date/Time Actions SNMP Import/Export Upgrade Alerts Admin ACL	
Home System Network Configure Monitoring	<input type="checkbox"/> Enable download Install from URL: <input type="text" value="http://137.65.82.73/data/nds_ichain/images/otwug/ichain21_2103"/> Download time: <input type="text" value="Immediately"/> <input type="checkbox"/> Enable install Install time: <input type="text" value="Immediately"/> Version being upgraded: 2001038 Description: iChain 2.1.038 Currently running version: 2001038 Last updated version: 2001037 Upgrade state: VersionUpgradeComplete
<input type="button" value="Apply"/> <input type="button" value="Cancel"/> <input type="button" value="Help"/>	<input type="button" value="Upgrade Log"/>

The Upgrade tab lets you set patch and upgrade parameters so you can download and install patches to the appliance. It also lets you uninstall the most recently applied patch.

Over-the-wire upgrades are secured through signing.

NOTE: We recommend you update the appliance's clone image after an upgrade. See ["Restoring the Appliance to the Clone Image" on page 168](#), ["Actions Tab" on page 285](#), and ["Upgrades" on page 364](#) for more information.

Enable Download: Lets you set the appliance to download updates automatically. If you check this box and enter the URL for the patch in the Install from URL field, it is downloaded as scheduled in the Download Time field. A valid entry for Install from URL is any valid URL or DNS name for a Web site.

Enable Install: Lets you set the appliance to install patches automatically. If you check this box, patches downloaded to the appliance are automatically installed as scheduled in the Install Time field.

Version Being Upgraded: Each update has a version number. The version of the current update appears in this field the moment the update process begins. You cannot upgrade the proxy server to a lower version than the one currently installed.

Description: A text name associated with the update file.

Currently Running Version: The update version number the appliance is currently running. Before installing the first update, this number is 0.

Last Updated Version: The update version number of the last update applied. For example, if you are currently running update version 3, this number might be 2.

Upgrade State: A state value indicating upgrade status. State values include Not Started, Download Pending, Version Download Complete, etc. The field is updated each time you click Upgrade.

Upgrade Log: Displays the text messages that have been generated by the upgrade process.

Alerts Tab

Path: System > Alerts

Figure 81 Alerts Tab

The screenshot shows the Novell iChain Alerts configuration interface. At the top, there's a navigation bar with tabs: Timezone, Date/Time, Actions, SNMP, Import/Export, Upgrade, Alerts (selected), and Admin ACL. The main content area is divided into several sections:

- Alert source name:** A text input field.
- Alert options:** Two checkboxes: Syslog and Email alert.
- Syslog servers:** A large empty text area with 'Insert' and 'Delete' buttons below it.
- Email recipients:** A large empty text area with 'Insert' and 'Delete' buttons below it.
- Email servers:** A large empty text area with 'Insert' and 'Delete' buttons below it.
- Syslog port:** A text input field containing the value '514'.
- Alert types:** A list of checkboxes for various system events:
 - Disk space shortage
 - Network receive buffers shortage
 - Oversized ping packets
 - SYN packet flooding
 - Oversized UDP packets
 - ICP parent down
 - Socks server down
 - System up
 - System down
 - Login failure
 - Configuration change

At the bottom left, there are three buttons: Apply, Cancel, and Help.

The Alerts tab lets you configure the appliance to send notification of generated system alerts to a network server hosting a Syslog service and to a list of e-mail recipients.

Alert Source Name: This identifies the appliance as the source of an alert. The system inserts this in the From field of an e-mail alert and in the Syslog alert message.

Syslog: Checking this box enables syslog alerts. Alert messages are then sent to one of the syslog servers.

E-mail Alert: Checking this box enables e-mail alerts. Alert messages are then sent to all of the e-mail recipients.

IMPORTANT: For this feature to work, e-mail servers must be able to relay e-mail from the appliance without authentication.

Because of increasing security risks, many e-mail servers have this feature disabled.

If you plan to have the appliance use e-mail alerts you must either ensure that the e-mail server can relay unauthenticated messages, or you must configure the server to accept mail from the appliance without authentication.

Syslog Servers: This is a list of syslog servers to which the appliance sends alerts. The appliance pings servers in the list, starting with the first server, until it receives an acknowledgement. It then sends a syslog alert using UDP to the responding server.

E-Mail Recipients: This is a list of e-mail recipients to whom the appliance sends alert e-mails. The appliance sends e-mails to all addresses in the list.

E-Mail Servers: This is a list of e-mail servers through which the appliance routes alert e-mails. E-mails are sent to the first e-mail server in the list. If the server doesn't respond, other servers are accessed in turn until the transmission is successful.

Syslog Port: This is the port the syslog server listens for syslog alerts on. The default port is 514, but this can be changed if required.

Alert Types: Appliance-generated alerts are sent for the following conditions. You enable or disable notification of generated alerts to the configured syslog server, and you e-mail recipients by checking or unchecking an alert type.

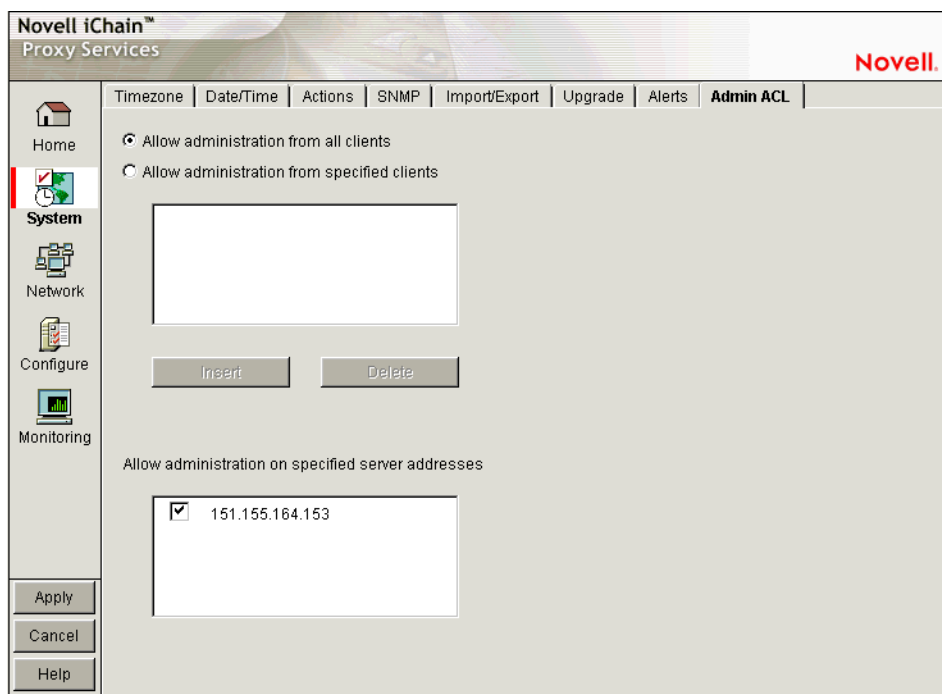
- ◆ *Disk Space Shortage:* The appliance generates this alert when disk space is low on the OS (SYS:) or Log (LOG:) volumes.
- ◆ *Network Receive Buffers Shortage:* The appliance generates this alert when the network receive buffers are low.
- ◆ *Oversized Ping Packets:* The appliance generates this alert when TCP/IP receives an oversized (greater than 10K) PING packet.
- ◆ *SYN Packet Flooding:* The appliance generates this alert when TCP/IP detects a SYN packet flooding attack (half-open connections).
- ◆ *Oversized UDP Packets:* The appliance generates this alert when TCP/IP receives an oversized (greater than 16K) UDP packet.
- ◆ *ICP Parent Down:* The appliance generates this alert when an ICP parent changes from an Up to a Down status.
- ◆ *SOCKS Server Down:* The appliance generates this alert each time it cannot communicate with a SOCKS server.
- ◆ *System Up:* The appliance generates this alert each time the appliance starts.

- ◆ *System Down*: The appliance generates this alert each time the appliance is shut down properly or restarted manually.
- ◆ *Login Failure*: The appliance generates this alert each time a login failure occurs from FTP or the browser-based management tool. The alert contains the IP address of the client making the unsuccessful attempt. Unsuccessful Telnet login failures are not detected.
- ◆ *Configuration Change*: The appliance sends this alert each time the appliance's configuration is changed and each time the appliance is initialized or re-initialized.

Admin ACL Tab

Path: System > Admin ACL

Figure 82 Admin ACL Tab



The Admin ACL tab lets you regulate access to appliance administrative functions in the browser-based management tool and the command line

interface. You can restrict administrative client access and limit the appliance IP addresses through which administrative access is allowed.

Allow Administration from All Clients: This option is selected by default and allows access to appliance administrative functions from any IP address.

Allow Administration from Specified Clients: When you select this option you must also insert at least one IP address from which IP administrative access is allowed. Otherwise, the system will deselect the option to prevent a global lockout.

NOTE: If you do not include the IP address from which you are specifying client access, and you click Apply, the address is not available for future administration sessions unless it is added later.

Allow Administration on Specified Server Addresses: This list contains all appliance IP addresses and indicates which are enabled for administrative access. The first addresses assigned to each network adapter are enabled for administration access by default. You change administrative access by checking and unchecking addresses in the list. The system doesn't allow unchecking all addresses. If this is attempted, the system reverts to the default setting by re-checking all first-assigned addresses.

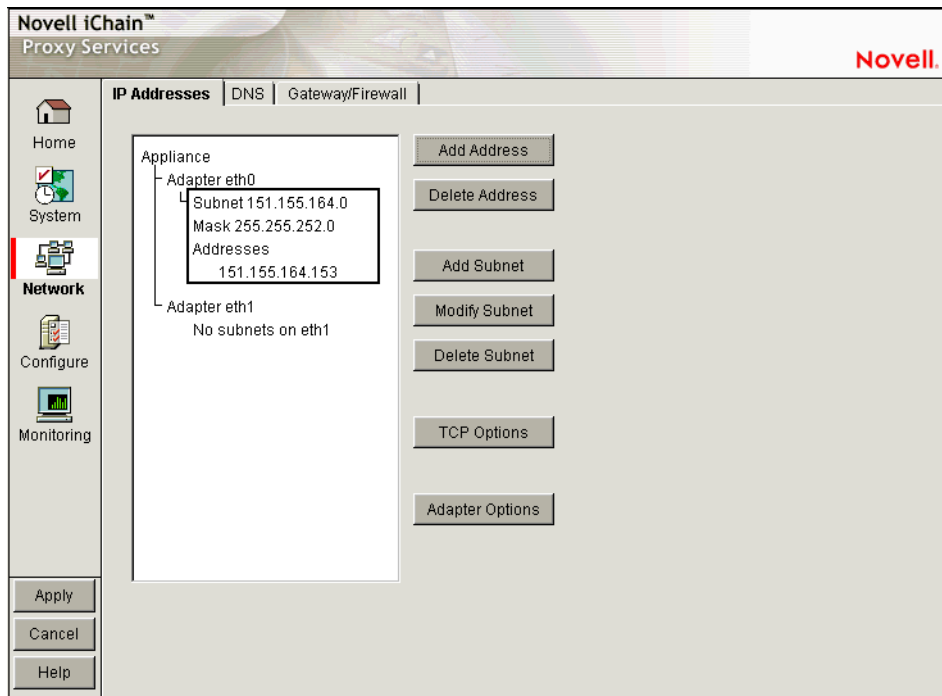
The Network Panel

The Network panel lets you configure the appliance to function on the network where it is installed.

IP Addresses Tab

Path: Network > IP addresses

Figure 83 IP Addresses Tab



The IP Addresses tab displays the network adapters, which are the physical connectors into the appliance, and the IP addresses associated with each adapter. The list reflects the current appliance hardware configuration.

Using the buttons to the right of the list, you can associate IP addresses with adapters and change IP address information. Each adapter can have multiple subnets associated with it, and each subnet will have one or more IP addresses associated with it. You can either define individual IP addresses and masks, or you can add a subnet address and mask and then add multiple IP addresses from that subnet range.

The IP address and the mask define a subnet. You cannot use the first or last address in any given subnet. You cannot create a subnet that collides with another subnet. You cannot create a subnet that spans multiple adapters.

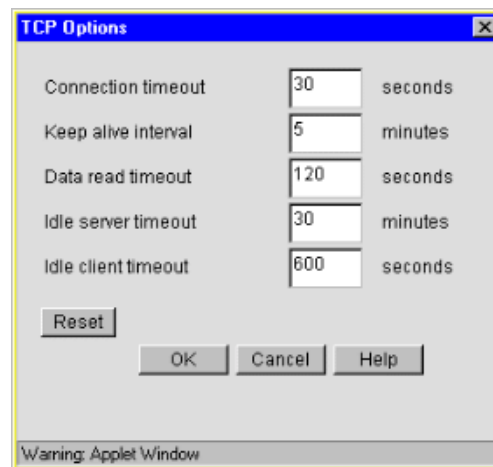
The following are valid appliance subnet masks (representing /1 through /31 in common router notation):

128.0.0.0	192.0.0.0	224.0.0.0	240.0.0.0	248.0.0.0
252.0.0.0	254.0.0.0	255.0.0.0	255.128.0.0	255.192.0.0
255.224.0.0	255.240.0.0	255.248.0.0	255.252.0.0	255.254.0.0
255.255.0.0	255.255.128.0	255.255.192.0	255.255.224.0	255.255.240.0
255.255.248.0	255.255.252.0	255.255.254.0	255.255.255.0	255.255.255.128
255.255.255.192	255.255.255.224	255.255.255.240	255.255.255.248	255.255.255.252
255.255.255.254				

TCP Options Dialog Box

Path: Network > IP Addresses > TCP Options

Figure 84 TCP Options Dialog Box



The parameters displayed in the TCP Options dialog box are standard TCP configuration settings. For more information on adjusting these parameters, see one of the TCP/IP references available at any bookstore carrying computer reference manuals.

Connection Timeout: The number of seconds the proxy server attempts to establish a connection before timing out because the other side has not

responded. You might want to increase this value if you notice that the remote server is reachable (the ping succeeds) but the load is heavy.

Keep Alive Interval: The number of minutes a connection is idle before the proxy server queries to check if the other server is still responding.

Data Read Timeout: The number of seconds the proxy server waits for expected data to begin arriving before it times out. You might want to increase this value if you notice that the browser receives incomplete data or the connection is disconnected in the middle of data transfer.

Idle Server Timeout: The number of minutes the proxy server keeps the TCP connection between the browser and the proxy server active, even if there is no data flow.

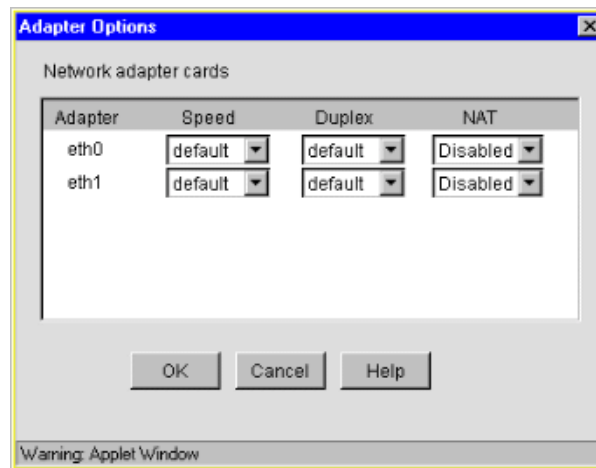
Idle Client Timeout: The number of seconds the proxy server keeps the connection to the origin Web server or another proxy server active, even if there is no data flow.

Reset: Resets the TCP configuration settings to the default values.

Adapter Options Dialog Box

Path: Network > IP Addresses > Adapter Options

Figure 85 Adapter Options Dialog Box



The Adapter Options dialog box lets you change settings for the network adapters on the appliance to ensure compatibility with an existing LAN.

Modify the default settings only if your LAN requires specialized adapter card changes.

Speed: Options include Default, 10 M, and 100 M.

Duplex: Options include Default, Half, and Full.

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

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

NAT: Options include Dynamic and Disabled.

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

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

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

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

DNS Tab

Path: Network > DNS

Figure 86 DNS Tab

The screenshot shows the Novell iChain Proxy Services configuration window, specifically the DNS tab. The window title is "Novell iChain™ Proxy Services" with the Novell logo in the top right. The interface is divided into a sidebar on the left and a main configuration area on the right. The sidebar contains icons for Home, System, Network (highlighted), Configure, and Monitoring, along with buttons for Apply, Cancel, and Help. The main area has tabs for IP Addresses, DNS (selected), and Gateway/Firewall. The DNS configuration includes:

- Domain:** A text field containing "nsrd.lab.novell.com".
- DNS server IP addresses:** Two text fields containing "137.65.1.4" and "137.65.1.3".
- Appliance domain name or alias:** A text field containing "153".
- Enable DNS proxy:** An unchecked checkbox.
- Advanced DNS Options:** A button.
- DHCP server IP addresses:** Three empty text fields.

 At the bottom of the main area are buttons for Apply, Cancel, and Help.

The DNS tab lets you configure the domain name service that the appliance will use, including setting a domain name for domain-relative address resolution.

DNS servers are searched in the order listed.

You must specify a domain name for the appliance to use relative domain names.

Domain: Specify the domain of your appliance. Valid ranges include all valid domain names.

DNS Server IP Addresses: Specify the IP addresses of the DNS servers you are using. You can enter up to three.

Appliance Domain Name or Alias: (Optional) Specify a unique domain name or alias for the appliance. This name is used in the Via headers that track packet routes across the network.

Enable DNS Proxy: Because of a potential security risk through the DNS port, the DNS proxy is disabled by default. You can enable the DNS proxy by checking this box.

Advanced DNS Options: See “Advanced DNS Options Dialog Box” on page 305.

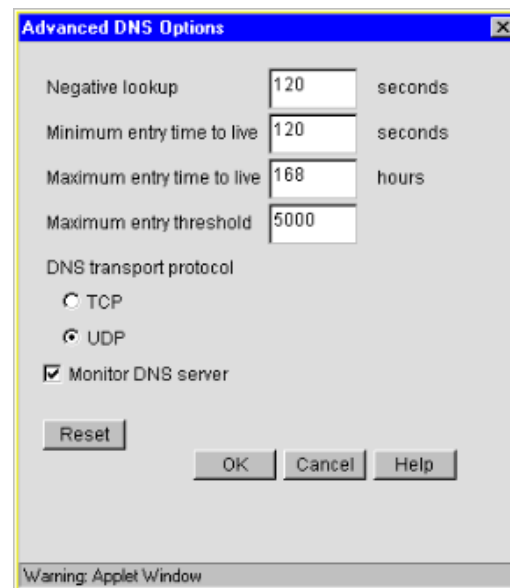
DHCP Server IP Addresses: Specify a list of DHCP servers to which the appliance will forward client DHCP requests.

This is critical if DHCP clients cannot directly access their designated DHCP servers. The appliance forwards the DHCP requests from the clients to the servers and forwards the replies back to clients. The appliance does not have to be enabled as a router to forward DHCP requests. However, the DHCP Server IP list must be filled in.

Advanced DNS Options Dialog Box

Path: Network > DNS > Advanced Options

Figure 87 Advanced DNS Options Dialog Box



The parameters displayed in the DNS Advanced Options dialog box are standard DNS configuration settings. For more information on adjusting these

parameters, see one of the TCP/IP references available at any bookstore carrying computer reference manuals.

Negative Lookup: How long a failed DNS lookup domain name remains in the proxy server cache. If the proxy server cannot resolve a domain name, it stores that information in its cache for the specified amount of time. If the proxy server receives requests for that domain name within this period, it sends a “Bad Gateway” error message to the browser and does not resolve the domain name again. Valid field values include 0 - 3600 seconds.

Minimum Entry Time to Live: The minimum amount of time that DNS entries remain in cache before they expire. This is the minimum value the appliance uses regardless of the value returned by the DNS name server. Valid field values include 0 - 3600 seconds.

Maximum Entry Time to Live: The maximum amount of time that DNS entries remain in cache before they expire. This is the maximum value the appliance uses regardless of the value returned by the DNS name server. Valid field values include 0 - 744 hours.

Maximum Entry Threshold: The maximum number of DNS cache entries. When this number is reached, the proxy server deletes old entries to make room for newer ones. The default is 5000. Valid field values include 2000 - 100000.

DNS Transport Protocol: The transport protocol DNS uses on the network where the appliance is installed.

Monitor DNS Server: The appliance normally monitors DNS server availability by pinging the configured servers every minute. This ensures timely handling of DNS requests. You should uncheck this item if the appliance accesses DNS through a connection that should not be kept continually open, such as a dial-up phone line or ISDN connection. Keep in mind, however, that unchecking the option will cause the DNS configuration on the **Health Status Tab** to fail.

Gateway/Firewall Tab

Path: Network > Gateway/Firewall

Figure 88 Gateway/Firewall Tab

The Gateway/Firewall tab lets you set up both default gateways as well as additional gateways for specific routing to hosts or networks. It also lets you specify RIP and SOCKS information for firewalls.

In order for the appliance to function, you must specify a default gateway (router) whether the appliance is originating packets that need to be routed (from proxy requests or scheduled downloads) or is serving as a router for packets that need to be routed externally.

Default Gateway IP Address: You must have at least one gateway defined for the appliance to function. This is the IP address of the gateway or router being used by the appliance.

Additional Gateways: The appliance uses these only if the Act As Router option is checked. See [“Additional Gateways Dialog Box” on page 309](#).

Enable RIP: Allows you to turn on Routing Information Protocol 1. Through this protocol, the appliance is able to learn routes.

The appliance can also work in a network that uses RIP 2, but you must manually add static routes using the [Routes Dialog Box](#).

Show Routes: See “[Routes Dialog Box](#)” on page 311.

Reset Learned Routes: Throws away all information acquired through RIP. RIP must be turned on for this to have any effect.

Act As Router: Check this box if the appliance will function as the default gateway for clients on the network. If you check this option, you can specify additional gateways.

Enable Gateway Monitoring: The appliance normally monitors gateway availability by pinging the configured gateways every minute. You should uncheck this item if the appliance accesses its gateways through a connection that should not be kept continually open, such as a dial-up phone line or ISDN connection. Keep in mind, however, that unchecking the option will cause the gateway configuration on the [Health Status Tab](#) to fail.

Enable SOCKS Client: SOCKS is a firewall communication protocol. If there is a firewall preventing the appliance from communicating directly, you can specify information for SOCKS4 or SOCKS5 servers.

Server IP Address: The address of the SOCKS server you want to use.

Server Port: The port number for SOCKS traffic on the network.

SOCKS V4: Enables the SOCKS4 protocol.

Username: Specify a username if the SOCKS4 server requires one for communication.

SOCKS V5: Enables the SOCKS5 protocol. The appliance currently supports only NULL and Username/Password authentications.

No Authentication: If you use SOCKS5 without verification, this box must be checked (where there is no username or password required).

Username/Password Authentication: Enables the entry of a SOCKS5 username and password if your SOCKS server requires authentication.

Username: Enter your SOCKS username.

Password: Enter your SOCKS password.

SOCKS Bypass Web Server List: If the SOCKS client is enabled, all HTTP and FTP server traffic is redirected to the SOCKS firewall. However, requests

to origin servers on an intranet within the firewall should not be routed through the SOCKS server. Requests to servers whose IP addresses are inserted into this list will not be sent to the SOCKS server.

Additional Gateways Dialog Box

Path: Network > Gateway/Firewall > Additional Gateways

Figure 89 Additional Gateways Dialog Box

Additional Gateways

Default gateway

Next hop address	Metric	Type
10 .1 .1 .1	1	Passive

Host gateways

Next hop address	Host address	Metric	Type
------------------	--------------	--------	------

Insert Delete

Network gateways

Next hop address	Subnet base address	Mask	Metric	Type
------------------	---------------------	------	--------	------

Insert Delete

OK Cancel Help

Warning: Applet Window

This dialog box lets you specify additional gateways. The appliance routes requests to specific destinations through these gateways. If a request could be routed through multiple gateways, the appliance chooses the gateway associated with the most restrictive mask (the smallest range of destination addresses). The default gateway is used only when no other routes apply.

IMPORTANT: The appliance uses additional gateways only when the Act As Router option is checked on the Gateway/Firewall tab.

Gateways fall within the following three basic groups:

- ◆ Host gateways for specific destination addresses
- ◆ Network gateways for destination addresses that fall within specific subnets
- ◆ The default gateway for destination addresses that aren't covered by host or network gateways

The syntax for this gateway is often expressed in router configuration tables as follows:

```
0.0.0.0 / 0.0.0.0 / iii.iii.iii.iii
```

The variable *i* represents the IP address of the default gateway.

IMPORTANT: If the appliance is acting as a router and you don't specify a default gateway, the appliance routes only those requests whose destination addresses are covered by a host or network gateway. Other requests are not routed.

The appliance uses Metric field values to alter the normal gateway use logic depending on a relative cost factor for using the gateway. The default field value is 1. A higher number indicates a higher cost associated with the gateway being referenced. This lets you configure the appliance in such a way that more expensive gateways are not used unless the default or less specific gateway is unavailable.

The appliance determines masking information when you enter the host or network information.

Default Gateway: The default gateway entered on the gateway panel. You can add a metric and specify whether the gateway is active or passive.

- ◆ *Next Hop Address:* The IP address of the gateway.
- ◆ *Metric:* A relative number indicating the bias you can add to the normal flow of gateway logic. Entering a number higher than 1 makes this resource more expensive and alters the gateway logic used. Valid numbers include 1 through 16.
- ◆ *Type:* Gateways can be active where they publish their presence, or passive where they do not.

Host Gateways: You can define one or more gateways to be used for packets being sent to specific hosts:

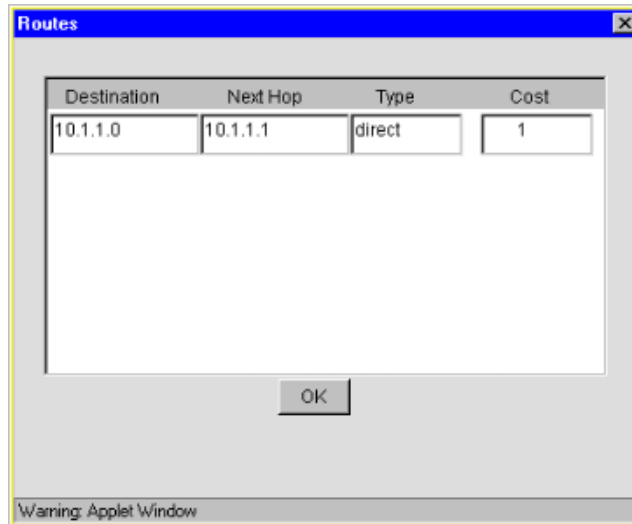
- ◆ *Next Hop Address*: The address of the host gateway that is to be used.
- ◆ *Host Address*: The IP address of the destination host. Valid addresses cannot be the first or last address of a class and must be unique.
- ◆ *Metric*: A value that alters the normal gateway use logic depending on a relative cost factor for using the gateways.
- ◆ *Type*: Gateways can be active where they publish their presence, or passive where they do not.

Network Gateways: You can define one or more gateways to be used for packets being sent to specific subnets.

- ◆ *Next Hop Address*: The address of the gateway that is to be used.
- ◆ *Subnet Base Address*: The subnet address for the destination IP address range. You can also enter a specific IP address on a given subnet and the appliance will calculate the subnet address using the mask.
- ◆ *Mask*: The subnet mask for the subnet or IP address above. A valid entry must be at least as large as a class mask where Class A Mask is 255.0.0.0, Class B Mask is 255.255.0.0, and Class C, D, E Masks are 255.255.255.0.
- ◆ *Metric*: A value that alters the normal gateway use logic depending on a relative cost factor for using the gateways.
- ◆ *Type*: Gateways can be active where they publish their presence, or passive where they do not.

Routes Dialog Box

Path: Network > Gateway/Firewall > Show Routes

Figure 90 Routes Dialog Box

This dialog box is useful for viewing and troubleshooting the routes the appliance is using. The list contains an entry for each defined gateway, each IP address assigned to an appliance network adapter, and routes discovered through RIP if the Enable RIP box is checked. Clicking Reset Learned Routes clears RIP entries from the list.

Destination: The default route is named and listed first. For other routes, the subnet address is shown.

Next Hop: This is the IP address of appliance network adapters, or the gateway address for all routes that are external to the appliance.

Type: Appliance network adapter routes are direct. All others are remote.

Cost: This is either the metric value you assigned to manually configured additional gateways (including the default gateway), or it is a relative cost factor assigned by the RIP function if the Enable RIP box is checked.

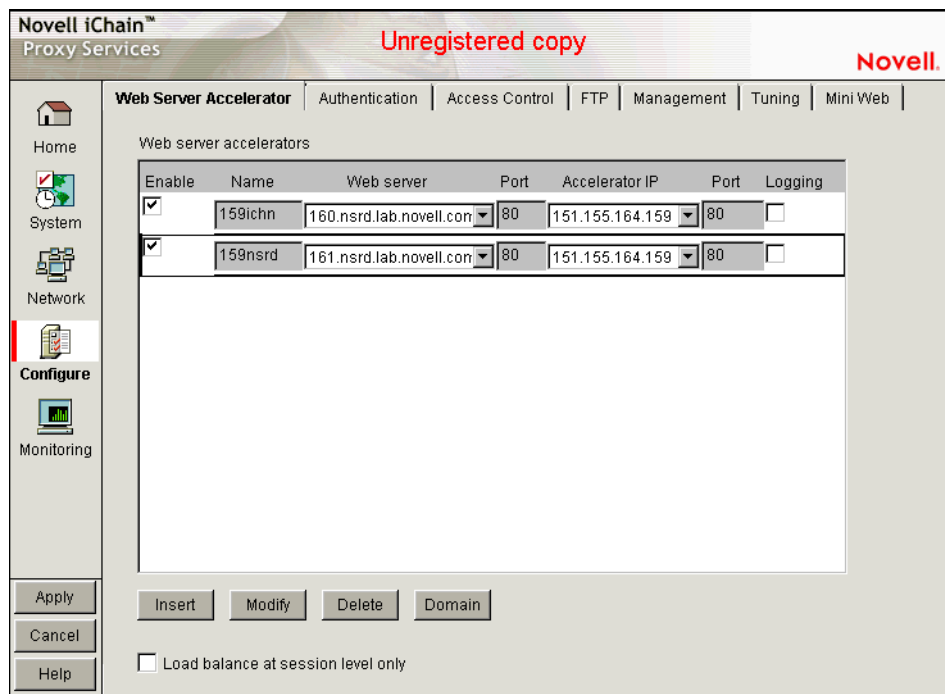
The Configure Panel

The Configure panel lets you set up Web acceleration. It also lets you configure authentication profiles and access control information, fine-tune caching services, and specify how error pages are vended to browsers.

Web Server Accelerator Tab

Path: Configure > Web Server Accelerator

Figure 91 Web Server Accelerator Tab



The Web Server Accelerator tab lets you add one or more Web server (reverse proxy) accelerators. The proxy server acts as the front end to Web servers on your Internet or intranet and off-loads frequent requests, thereby freeing up bandwidth. Using a Web server accelerator also increases security because the IP addresses of your Web servers are hidden from the Internet. For more information, see [“Overview of Web Server Acceleration” on page 60](#).

Field names shown in the Web Server Accelerators list are explained in [“Web Server Accelerator Dialog Box” on page 314](#).

Load Balance at Session Level Only: Checking this box causes the proxy server to use the same Web server for all fills during a session. This prevents eBusiness users from needing to log in multiple times. This setting affects all Web server accelerators configured on the proxy server.

Web Server Accelerator Dialog Box

Path: Configure > Web Server Accelerator > Insert

Figure 92 Web Server Accelerator Dialog Box

The Web Server Accelerator dialog box lets you create Web server accelerator services for handling requests to Web servers.

Enable This Accelerator: Specifies whether to enable the defined Web server accelerator after you have configured it. The default is Enabled.

Name: Each Web server accelerator service requires a name you create. For example, you can select a name that indicates which Web server is being serviced by the appliance, or alternately, a set of browsers configured to access the Web server accelerator as a proxy server. A valid name consists of a DOS-style, eight-character name with no special characters or spaces.

If logging is enabled, the appliance uses the Web server accelerator name as the directory name that the log files for the Web server accelerator are kept in.

DNS Name: The contents of this field depend on the type of accelerator you are using.

If you are accelerating multiple Web servers for multiple Web sites on the same IP address, you must create a Web server accelerator definition for each DNS name that is used in browser requests. This name must exactly match one of the names in the requests. If your infrastructure supports multi-homing, refer to [“Multi-homed Configurations” on page 115](#), [“Path-based Multi-homing” on page 118](#), and [“Path-Based Multi-homing Examples” on page 120](#) for further information.

If you are accelerating multiple Web servers for a single Web site and plan to use path-based multi-homing, you must use the same DNS name in every accelerator definition.

NOTE: If you are changing the DNS name on an accelerator which has authentication enabled, the existing cookie domain may not be valid with the new DNS name if the DNS name is not a subdomain of the cookie domain (resulting in the browser displaying a 403 Forbidden Error message). You can check the cookie domain at Authentication Options > Cookie Domain.

Cookie Domain: In Configure > Web Server > Modify > Authentication Options, there is a Cookie Domain field. Single authentication across two accelerators within the same iChain Proxy Server is not possible when the Cookie Domain is different on each accelerator. The Cookie Domain field can be used to allow single authentication across multiple accelerators where the cookie domain by default would be different for each accelerator. For example, if Accelerator One has a DNS name of `www.support.novell.com` and Accelerator Two has a DNS name of `www.developernet.novell.com`, the cookie domains would be `support.novell.com` for Accelerator One and `developernet.support` for Accelerator Two. The cookie domains are different for each accelerator, so users would be prompted for authentication when accessing Accelerator Two, even if they had already authenticated to Accelerator One. If the cookie domain is changed to `novell.com` on each accelerator, then users would not need to authenticate again when accessing Accelerator Two if they had already authenticated through Accelerator One. This would not be possible if one of the DNS names would have ended with `acme.com` and the other DNS name ended with `novell.com`. Both accelerators need to have a common subdomain for this to work.

Use Host Name Sent by Browser (Multi-homing Web Server): Checking this option preserves the host name in the HTTP header exactly as it came in the browser request.

Alternate Host Name: Checking this option causes the string specified to be substituted for the host name in the HTTP header before the request is forwarded to the Web server.

Return Error If Host Name Sent by Browser Does Not Match Accelerator DNS Host Name: Checking this option causes iChain Proxy Services to match the host name in the DNS header that came from the browser against the DNS name specified in this accelerator definition. If the names don't match, the request is not forwarded to the Web server. Instead, iChain Proxy Services returns an error to the requesting browser.

Act as a Tunnel: Normally an accelerator service processes HTTP requests in order to fill them. However, it is not unusual that some of the traffic coming through the appliance is not HTTP-based.

Web servers often handle SSL connections, and less frequently they might need to let Telnet, FTP, chat, or other kinds of traffic through without attempting to process it.

The Act as a Tunnel option lets you create one or more accelerator services for the specific purpose of tunneling non-HTTP traffic through the appliance to the origin Web server. When the option is checked, the accelerator sets up a tunnel for all incoming traffic.

When you check the Act as a Tunnel option, you have the additional option of having the accelerator service tunnel only SSL traffic.

Tunnel Only SSL Traffic: If you decide to have the accelerator act as a tunnel, you can elect to have it tunnel only SSL traffic. The service will then verify that the address and port being accessed are actually an SSL Web site. If verification fails, the service will tear down the connection.

NOTE: The SSL port number for the SSL tunnel is specified via the Accelerator proxy port and not the SSL listening port.

Forward Browser IP Address in Request Header [Enable X-Forwarded-For]: X-Forwarded-For headers are used to pass browser ID information along with browser request packets. If the headers are included, Web servers can determine the origin of browser requests they receive. If the headers are not included, browser requests have anonymity.

Checking the X-Forwarded-For option causes the appliance to either add information to an existing X-Forwarded-For or Forwarded-For header, or to create a header if one doesn't already exist.

Leaving the option unchecked causes the appliance to remove X-Forwarded-For headers from any Web accelerator requests passing through the appliance.

Deciding whether to check the option requires that you weigh the desires of browser users to remain anonymous against the desires of Web server owners (e-commerce sites, for example) to collect data about who is accessing their site.

Enable Authentication: Checking this box causes the appliance to require authentication of users wanting to use its Web server accelerator services. Clicking Authentication Options displays the Add Authentication Profiles dialog box.

Enable Logging for This Accelerator: Causes log files to be kept for this Web server accelerator. Clicking Log Options displays the Log Options dialog box.

Enable Secure Exchange: Selecting this option allows SSL to be used for HTTP requests between the client and the iChain box, and optionally between the iChain box and Web servers. The default is Disabled. Clicking Secure Exchange Options displays the Secure Exchange Options dialog box.

SSL Listening Port: Specifies the port that will be used for Secure Exchange communication.

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

Certificate: This drop-down list displays certificates you have stored on your appliance. System-generated certificates do not appear in the list.

Use this field after you have stored a certificate you created specifically for the Web server accelerator you are creating. This will prevent browsers from receiving certificate confirmation messages each time they access the appliance. For more information, see [“Managing Appliance Certificates” on page 123](#).

Enable Path-Based Multi-homing: This option is enabled only when you have created another accelerator definition and have not created a standard multi-homing relationship between previously defined accelerators on the iChain Proxy Services appliance. In other words, you don't have multiple accelerators sharing the same accelerator IP address and port.

Path-based multi-homing lets you configure the system so that a multi-homing master accelerator fills general requests from a site's main Web server and routes specific requests to specialized child accelerators that fill from other specialized Web servers. This option lets you create child accelerators for path-based multi-homing configurations.

When you enable path-based multi-homing for the accelerator you are defining, you must also click the Path Rule Options button and specify a path rule that the multi-homing master can use to route traffic to the accelerator you are defining. For more information, see [“Multi-homing Options Dialog Box” on page 324](#).

You will also notice that if you have created multiple accelerators that can function as multi-homing masters, when you select a name in the Multi-homing Master drop-down list, the DNS Name, Accelerator Proxy Port, and Accelerator IP Addresses selections dynamically change to match the accelerator whose name you have selected. For more information regarding path-based multi-homing, see [“Path-based Multi-homing” on page 118](#).

Multi-home Master: This drop-down list contains the names of accelerators you have defined that can function as multi-homing masters, meaning they are not configured as child accelerators to other multi-homing masters.

Custom Login Page Location (blank to disable): The field on the Web Server Accelerator dialog for the custom login page is populated with the directory from SYS:ETC\PROXY\DATA, where the login pages can be found. If a user enters NIKE in this field, the user is specifying that the custom login/logout pages will be found in SYS:ETC\PROXY\DATA\NIKE.

Web Server Port: The port number that the origin Web server is listening on for incoming connections. The default for HTTP is 80 (1 - 65535).

Web Server Addresses: The IP address or local DNS name of each Web server from which the appliance fills the cache for this Web server accelerator. The appliance must be able to fill all requests through any of these names or addresses unless path-based multi-homing is being used.

Accelerator Proxy Port: The port number that the proxy server is listening on for incoming connections. The default for HTTP is 80 (1 - 65535).

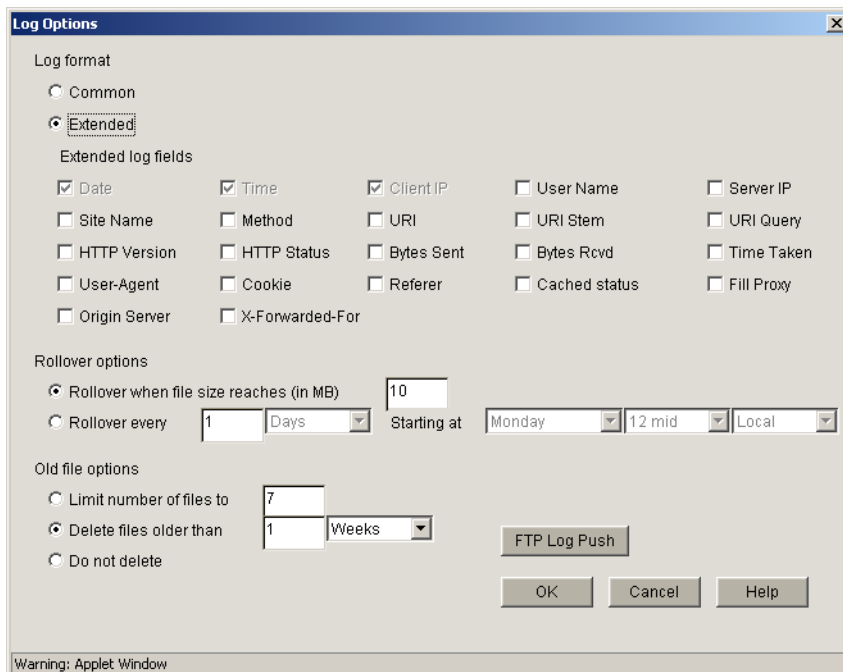
Accelerator IP Addresses: For normal accelerator situations, non-path-based multi-homing configurations, and accelerators configured as multi-homing masters, this is the appliance's IP addresses to which DNS resolves the Web server's (or Web site's) DNS name and on which the Web server accelerator listens for incoming connections from the Internet.

For child accelerators in path-based multi-homing configurations, this is the IP address or addresses to which the multi-homing master forwards browser requests that match the specified path rule.

Log Options Dialog Box

Path: Configure > Web Server Accelerator > Insert (Modify) > Log Options

Figure 93 Log Options Dialog Box



This dialog contains numerous settings for logging the operations and errors of the iChain Proxy Server.

Common: If this option is checked, only three parameters (Date, Time, and Client IP) will be enabled. No others will be selectable.

Extended: If this option is checked, all of the Extended log fields will be enabled. The Date, Time, and Client IP options are checked by default and may not be unchecked.

Rollover When File Size Reaches (in MB): If this option is selected, the log file will roll over (empty and start writing from the beginning) after the file reaches the size specified in the edit field.

Rollover Every: If this option is selected, the log file will roll over automatically over a specified time interval or starting at a certain day.

Limit Number of Files to: If this option is selected, the number of old log files kept on the server will be limited to the number entered in the edit field.

Delete Files Older Than: If this option is selected, old log files will be deleted on the time interval specified.

Do Not Delete: If this option is selected, none of the log files will be deleted.

FTP Log Push Button: Launches the FTP Log Push Configuration dialog box.

FTP Log Push Configuration Dialog Box

Path: Configure > Web Server Accelerator > Insert (Modify) > Log Options > FTP Log Push

Figure 94 Log Push Configuration Dialog Box

FTP Log Push Configuration

FTP log push enable

Host server

Login name

Password

Default directory

IP address

Delete log files from iChain proxy server after push

Log push result

Push logs when the logs rollover

Days to push the logs

1st of month Monday Tuesday

Wednesday Thursday Friday

Saturday Sunday

Time to push the logs

Log types

Accelerator Forward Streaming

Filter Cluster

OK Cancel Help

Warning: Applet Window

This dialog box enables you to set FTP log push options.

FTP Log Push Enable: If this option is checked, FTP log push is enabled.

Host Server: The name of the host server where the logs will be pushed.

Login Name: The user name on the host server to log in as. This user must have appropriate rights to be able to push logs to the server.

Password: The password for the user on the host server.

Default Directory: The directory on the host server where the logs will be pushed.

IP Address: The IP address of the host server.

Delete Log Files from iChain Proxy Server after Push: If this option is checked, the logs will be deleted from the iChain Proxy Server after they have been pushed to the specified host server.

Log Push Result: This field displays the result of the log push.

Push Logs When the Logs Rollover: If this option is checked, the logs will automatically be pushed to the host server when they roll over.

Days to Push the Logs: The Logs will be pushed to the host server every day that has been checked.

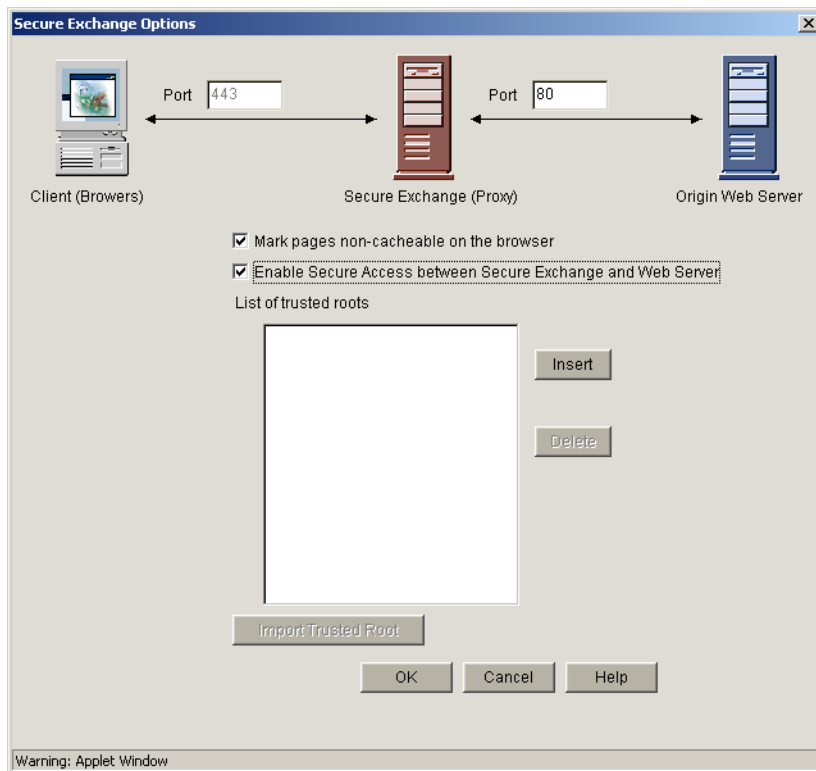
Time to Push the Logs: The logs will be pushed to the host server starting at the time specified by these two fields.

Log Types: Check each type of log that will be pushed to the host server.

Secure Exchange Options Dialog Box

Path: Configure > Web Server Accelerator > Insert (Modify) > Secure Exchange Options

Figure 95 Secure Exchange Options Dialog Box



This dialog box provides additional Secure Exchange options.

NOTE: When Secure Exchange is enabled, the Web server port must be configured under Secure Exchange Options.

Port (Client to Proxy): This field specifies which port will be used for communication between the client (browser) and the proxy server. The default is port 443.

Port (Proxy to Web Server): This field specifies which port will be used for communication between the proxy server and the origin Web server. The default is port 80. When Secure Exchange is enabled, this port must be configured here and not at the Web Server Accelerator dialog box.

Mark Pages Non-Cacheable on the Browser: If this option is checked, all pages will be marked as non-cacheable on the browser.

Enable Secure Access Between Secure Exchange and Web Server: If this option is checked, the connection between the proxy server and the origin Web server will be secured.

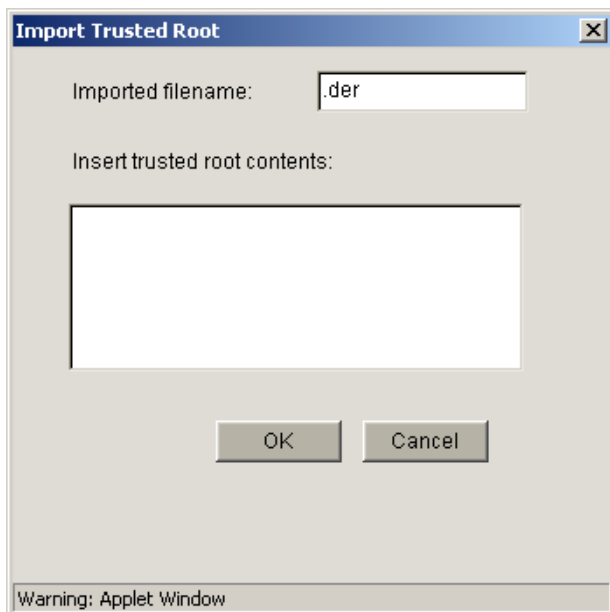
List of Trusted Roots: This list displays the current list of trusted roots in use by the proxy server.

Import Trusted Root: When you click Insert, the Import Trusted Root button is enabled. If you click this button, it will launch the Import Trusted Root dialog box.

Import Trusted Root Dialog Box

Path: Configure > Web Server Accelerator > Insert (Modify) > Secure Exchange Options > Import Trusted Root

Figure 96 Import Trusted Root Dialog Box



This dialog box allows the user to import trusted roots into the system for use during Secure Exchange.

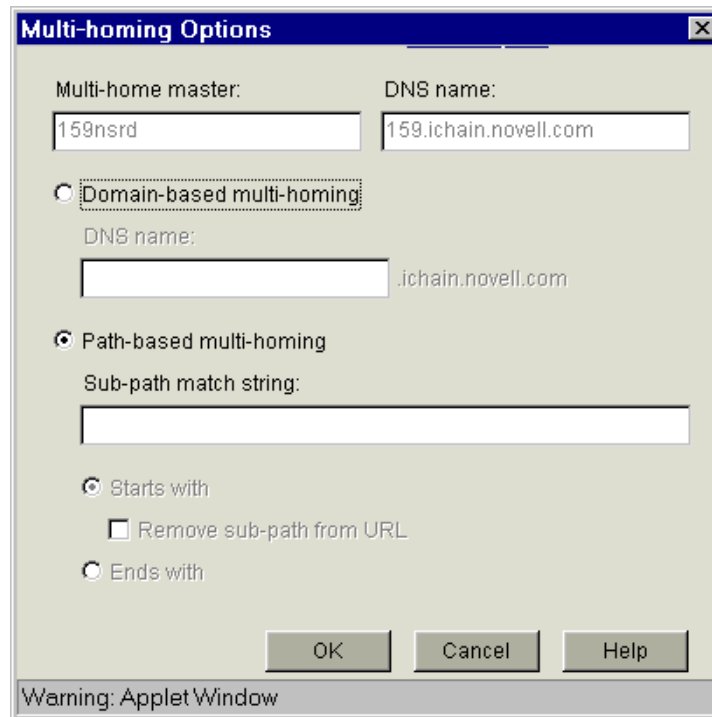
Imported Filename: The name of the trusted root file.

Inserted Trusted Root Contents: The contents of the trusted root file.

Multi-homing Options Dialog Box

Path: Configure > Web Server Accelerator > Insert > Enable Multi-homing > Multi-homing Options

Figure 97 Multi-homing Options Dialog Box



This dialog box lets you specify a string that, if present in the browser request, will cause the multi-homing master accelerator to route the request to the child accelerator being defined.

The string match can occur immediately following the DNS name (the Starts With option) or at the end of the URL (the Ends With option).

Sub-Path Match String: This is the string the multi-homing master will compare against the browser request. If the string is not found, the multi-homing master accelerator attempts to fill the request through the Web server addresses in its accelerator definition. If the string is found, the multi-homing master accelerator routes the request to the accelerator with the matching string.

Starts With: Checking this option indicates to the multi-homing master that the Sub-Path Match String field contains a path that might immediately follow the DNS name in the browser request. If the string matches the path in the request, the multi-homing master forwards the request to the child accelerator being defined. The initial forward slash (/) is required in the string to directly access the child page. If you don't include the initial forward slash, it will only find your child page (for example, <http://www.icweb.com/yahoo>) after you visit the main/master (<http://www.icweb.com>) page. Only when you have the initial forward slash can you access your child page directly. If you try to access the child page without the forward slash, you will get a Page Not Found error.

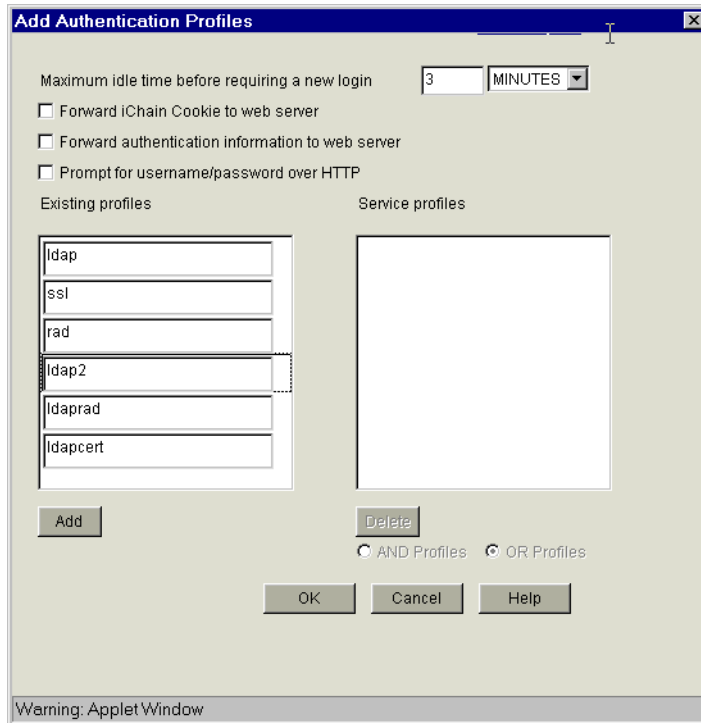
Remove Sub-Path from URL: You should check this option if the path string doesn't actually appear at the root of the Web server. If this option is checked, the string is stripped from the request before the request is sent to the Web server. This probably indicates that the object is at the root of the Web server. If this option is not checked, the matched string is retained in the request sent to the Web server.

Ends With: Checking this option indicates to the multi-homing master accelerator that the Sub-Path Match String field contains a file extension, such as gif, jpg, mpg, cgi, and so on. If a match is found at the end of the browser request, the multi-homing master will route the request to the child accelerator being defined.

Add Authentication Profiles Dialog Box

Path: Configure > Web Server Accelerator > Insert or Modify > Enable Authentication > Authentication Options

Figure 98 Add Authentication Profiles Dialog Box



The Add Authentication Profiles dialog box lets you select one authentication profile to authenticate the users of the proxy service from which you accessed this dialog box.

The Existing Profiles list shows all the authentication profiles you have created. See [“Authentication Dialog Box” on page 328](#). The Service Profiles list contains the profile that is active for the proxy service from which you accessed the Add Authentication Profiles dialog box.

Maximum Idle Time Before Requiring a New Login: The period of browser inactivity allowed before the proxy server requests a new login.

Forward Authentication Information to Web Server: If this option is checked, the username and password are sent to the Web server.

Prompt for Username/Password over HTTP: If this option is checked, authentication will be done over unencrypted HTTP instead of HTTPS.

Existing Profiles: A list of the authentication profiles you create in Cache > Authentication. For more information, see [“Authentication Tab” on page 327](#).

Service Profiles: The authentication profile that is active for the proxy service from which you accessed the Add Authentication Profiles dialog box. List content depends on whether And Profiles or Or Profiles is selected. You add a profile to the list by clicking a profile in the Existing Profiles list, then clicking Add. You can remove the profile from this list using the Delete button.

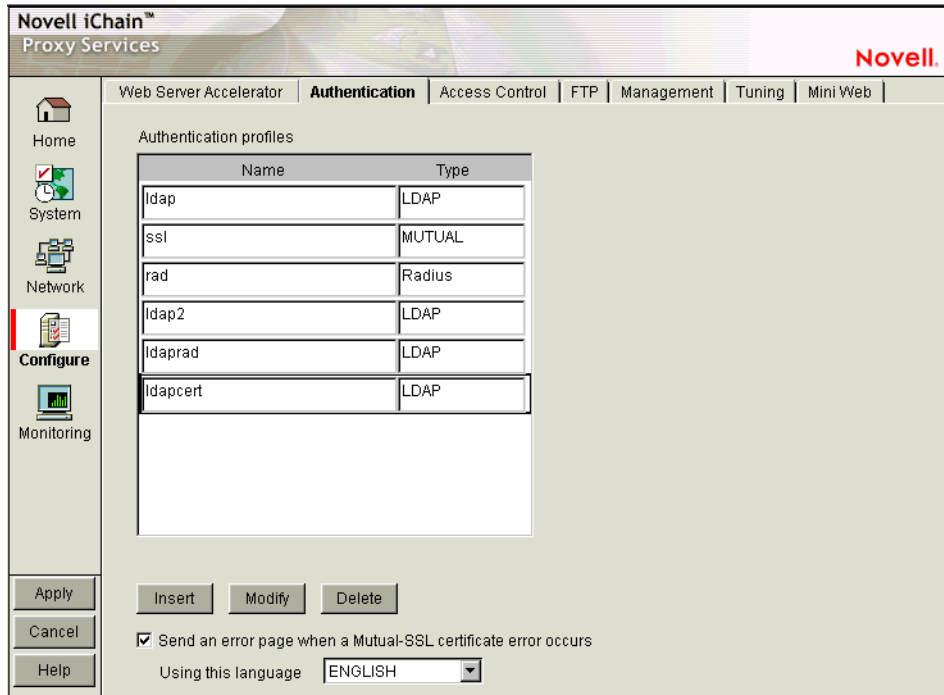
AND Profiles: If this button is selected, the profiles in the Service profiles list are used in conjunction with each other to authenticate to users accessing the Web servers specified by this accelerator.

OR Profiles: If this button is selected, either one of the profiles in the Service profiles list are used to authenticate the users accessing the Web servers specified by this accelerator.

Authentication Tab

Path: Configure > Authentication

Figure 99 Authentication Tab



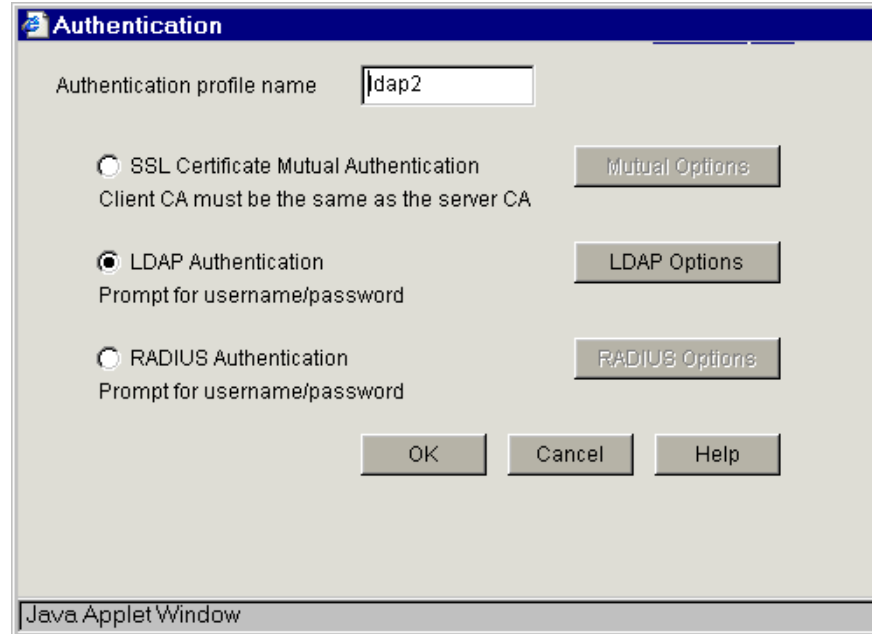
The Authentication tab lets you enable the appliance to authenticate users to an LDAP or RADIUS authentication source.

Authentication Profiles List: This lists the authentication profiles you have configured on the appliance using the Authentication dialog box.

Authentication Dialog Box

Path: Configure > Authentication > Insert under the Authentication Profiles list

Figure 100 Authentication Dialog Box



The Authentication dialog box lets you assign an authentication profile name and specify either LDAP or RADIUS as the authentication source.

IMPORTANT: iChain Proxy Services doesn't recognize case differences in profile names. MyProfile and myprofile are, effectively, the same profile name.

Also, iChain Proxy Services partially overwrites and concatenates previously created profiles without warning if a duplicate name is used. Therefore, if you create a profile named MyProfile and later create another profile named myprofile, iChain Proxy Services will remove the first name, concatenate parts the first profile with the second, and use the second name.

To avoid these problems, ensure that each profile has a unique name.

After selecting the authentication source, you must configure the source by clicking its respective Options button.

LDAP Options Dialog Box

Path: Configure > Authentication > Insert > LDAP Authentication > LDAP Options

Figure 101 LDAP Options Dialog Box

LDAP Options

LDAP server addresses

137 . 65 . 159 . 220
137 . 65 . 159 . 221
137 . 65 . 159 . 219

Insert
Delete

LDAP server listening port: 636

Enable secure access to LDAP server

LDAP server trusted root file: Import Trusted Root

Username/password for LDAP search - use blank username for anonymous bind

Username Password

LDAP login name format:

Use user's email
 Use distinguished name
 Use field name

LDAP contexts:

O=Novell

Insert
Delete

OK Cancel Help

Warning: Applet Window

Use the LDAP Options dialog box to configure the appliance so users can authenticate through an LDAP database.

LDAP Server Addresses: The IP addresses of the LDAP servers. Specifying multiple LDAP servers allows for failover LDAP servers. If the first LDAP server cannot be accessed, the next LDAP server on the list will be tried, an so-on down the line in the order specified.

LDAP Server Listening Port: The port number the LDAP server is listening on for requests from LDAP clients. The default is 389 for normal access. Use 636 for secure access.

Enable Secure Access to LDAP Server: Checking this box causes the data sent between the LDAP client and the LDAP server to be sent using SSL.

LDAP Server Trusted Root File: The path on the appliance to a trusted root file that contains the Certificate Authority (CA) used by the LDAP server in the profile you are creating.

The iChain Proxy Server fills this field with information for the trusted root file you create using the Import Trusted Root button. See the instructions found in [“Import Trusted Root Dialog Box” on page 332](#).

If the LDAP server uses a CA for which you have previously created a trusted root file, you can manually type the path and filename in this field. For example, you might be using the same LDAP server for multiple authentication profiles.

Import Trusted Root: Clicking this button opens the Import Trusted Root dialog box (see [“Import Trusted Root Dialog Box” on page 332](#)).

Username: Enter the username to use for accessing the LDAP server.

Password: Enter the password to use for accessing the LDAP server.

LDAP Login Name Format

The contents of this box change depending on the option selected.

Use User’s E-Mail: Select this option to have users log in using their e-mail name field in the LDAP database. You must provide one or more contexts in which the LDAP server will search for the e-mail name.

This option is somewhat redundant with Use Field Name because the e-mail name is simply an LDAP field name. E-mail is offered separately because it is used so often.

LDAP Search Base: Click Insert to enter the context of one or more LDAP containers from which the search for the e-mail name should begin.

Use Distinguished Name: Select this option to allow users to authenticate using their LDAP usernames. Users can use either their fully distinguished (full LDAP contexts) LDAP usernames, or you can provide a list of LDAP contexts so users only need to type their usernames.

IMPORTANT: Using this option with Netscape’s LDAP server requires a special setup procedure.

LDAP Contexts: This list contains specific contexts in which the LDAP server will look for usernames. This provides a shortcut to authentication of users by allowing them to type only their LDAP usernames.

The appliance searches each context until it either locates the name or exhausts the search. If duplicate names exist in different contexts, the appliance searches until the correct name/password match is found.

Use Field Name: Select this option to require that users enter a specific LDAP field name.

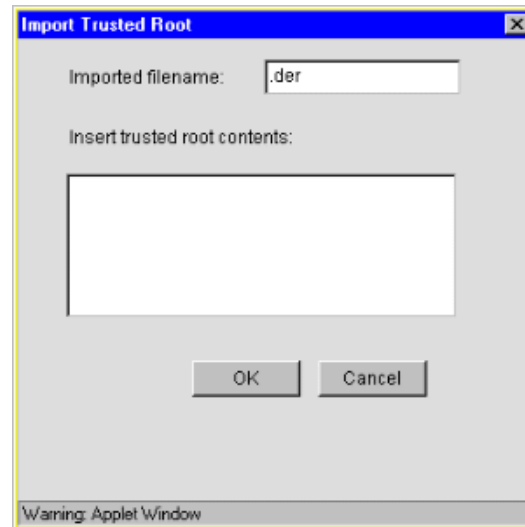
Field Name: The LDAP field name (such as User ID) through which users can authenticate.

LDAP Search Base: Click Insert to enter the context of one or more LDAP containers. The appliance will perform a subtree search in all containers in the list and in their subcontainers.

Import Trusted Root Dialog Box

Path: Configure > Authentication > Insert > LDAP Authentication > LDAP Options > Import Trusted Root

Figure 102 Import Trusted Root Dialog Box



The Import Trusted Root dialog box lets you create a trusted root file with a .DER extension that contains information identifying the Certificate Authority used by the LDAP server for the profile you are creating.

To create a trusted root file:

- 1 In the Imported Filename field, type a path and filename for the trusted root file.

The filename can contain up to eight alphanumeric characters. The appliance automatically appends the .DER extension if you don't include it.

IMPORTANT: Be sure you use a unique filename for each .DER file. The appliance overwrites files without warning if you use duplicate filenames.

Remember that iChain Proxy Services is not case-sensitive, so MyCert.DER and mycert.der are, effectively, the same filename.

The path must be a directory path that already exists. You cannot create directories on the appliance.

If you want to list your trusted root files later, use an FTP-accessible directory, such as SYS:\ETC\PROXY\DATA, as the path. Otherwise, you won't be able to list the files. For a list of FTP-accessible directories, see [“Limitations of the Appliance’s Mini FTP Server” on page 198.](#)

If you don't include a path with the filename, the proxy server creates the file at the root of the SYS: volume. You cannot see the root of the SYS: volume using FTP.

- 2 Using a text editor on your configuration workstation, open the .DER file for the Certificate Authority > select the file contents > paste the contents to the clipboard.

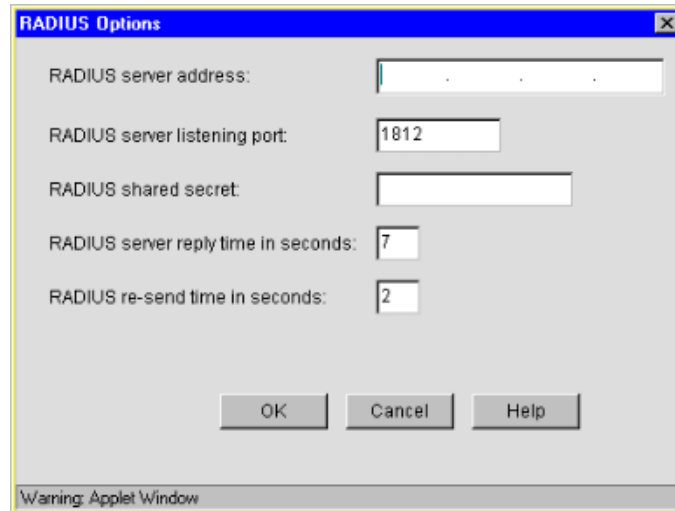
To obtain .DER files, contact a Certificate Authority vendor.

- 3 Return to the Import Trusted Root dialog box > paste the clipboard contents into the text box above the OK and Cancel buttons.
- 4 Click OK.

RADIUS Authentication

Path: Configure > Authentication > Insert > RADIUS Authentication > RADIUS Options

Figure 103 RADIUS Options Dialog Box



Use this dialog box to specify a RADIUS server the appliance can use for authentication.

NOTE: Port 1812 is the default RADIUS server port, however, the Novell RADIUS server defaults to port 1645. It is possible for a different default port to be specified when the Novell RADIUS server is loaded.

RADIUS Server Address: The IP address of the RADIUS server.

RADIUS Server Listening Port: The port number on which the RADIUS server listens for incoming authentication requests.

RADIUS Shared Secret: The string the RADIUS server uses to verify that the appliance can request authentication of users.

RADIUS Server Reply Time in Seconds: The total time the appliance will wait for a response from the RADIUS server before authentication fails. The default is 7 seconds.

RADIUS Resend Time in Seconds: The interval in seconds between appliance requests to the RADIUS server. The default is two seconds. This means that the appliance could send three requests before the 7-second default limit expires and the authentication request fails.

Third-Party RADIUS Server Support

iChain 2.1 supports RADIUS authentication via the RADIUS server supplied by Novell or a RADIUS server supplied by third-party vendors. The Novell RADIUS server uses NDS[®] eDirectory[™] 8.5 as the underlying database for storing user objects.

If the Novell RADIUS server is running on the iChain authorization server, it has the ability to indicate the fully distinguished name of the user object that was authenticated via RADIUS. iChain must use the fully distinguished name of the user in order to perform ACL rights checking when accessing a protected resource. A RADIUS server provided by a vendor other than Novell will store user objects in their own database.

iChain must be able to map this third-party name to a user object stored in NDS eDirectory 8.5 on the authorization server in order to determine if the user has rights to the protected resource he or she is attempting to access. This is done by performing an LDAP subtree search on the authorization server for the user name entered on the iChain login dialog. In order to perform this search, two fields must be changed on the ACLCheck authentication profile.

A search base must be specified that identifies an eDirectory container that defines where in the tree to begin the search and the anonymous bind feature must be set to No. This is done by entering the following commands on the iChain Proxy Services console screen:

```
add authentication aclcheck ldap searchbase = container name
set authentication aclcheck ldap bindanonymous = no
apply
```

If there is already a value specified for the searchbase, the set command must be used instead of the add command, as shown below:

```
set authentication aclcheck ldap searchbase = container name
```

The container name may specify an eDirectory organization or an organizational unit that appears on the authorization server. For example:

```
add authentication aclcheck ldap searchbase = o=novell
```

The implication here is that an eDirectory user object must be created on the authorization server whose CN is the same as the RADIUS user name and it must appear hierarchically in the tree somewhere below the container specified by the searchbase.

Also verify that iChain has defined an LDAP server, user, and password. This is needed to perform the searchbase lookup.

For example:

```
>get authentication aclcheck ldap
authentication aclcheck ldap address = 10.252.3.5
authentication aclcheck ldap bindusername = cn=admin;o=novell
authentication aclcheck ldap bindpassword = password
```

If the above values are empty, they can be input at the iChain browser GUI > Configure > ACL tab.

Access Control Tab

Path: Configure > Access Control

The Access Control tab lets you define the parameters necessary to set up access control. This tab includes the following fields:

ISO Object LDAP Name: Specifies the name of the iChain Service Object (ISO) containing parameter settings defining your iChain domain or infrastructure.

Password Management Servlet URL: Defines the URL of the password management servlet on the iChain Proxy Server. This servlet enables users within your iChain infrastructure to change their passwords. Use the full http URL. For example, `http://ichain.provo.novell.com/servlet/iChainPasswordMgr` where `/servlet/` is the servlet directory of your servlet engine.

LDAP Server Addresses: Specifies the IP addresses of the iChain LDAP access control servers. Specifying multiple LDAP servers allows for load-balancing of the LDAP servers. Requests to the LDAP servers for access control are done in a round-robin fashion.

Enable Secure Access to LDAP Server: Checking this box causes the data sent between the LDAP client and the LDAP server to be sent using SSL.

LDAP Server Trusted Root File: The path on the appliance to a trusted root file that contains the Certificate Authority (CA) used by the LDAP server in the profile you are creating. The system fills this field with information for the trusted root file you create using the Import Trusted Root button. If the LDAP server uses a CA for which you have previously created a trusted root file, you can manually type the path and filename in this field. For example, you might be using the same LDAP server for multiple authentication profiles.

Import Trusted Root: Clicking this button opens the Import Trusted Root dialog box.

LDAP Port: Specifies the port on which the LDAP server will listen for access control requests. The default is 389 for normal access. Uses 636 for secure access.

LDAP Proxy User: Specifies the LDAP username for the iChain Proxy Server to use when making requests for access control information from the iChain LDAP access control server.

Password: Specifies the LDAP password for the iChain Proxy Server to use when making requests for access control information from the iChain LDAP access control server.

Enable Unauthorized Access Rule Hit Logging: Specifies whether to enable logging for unauthorized access rule hits. The default is Disabled.

Enable Object Level Access Control: Specifies whether to enable object level access control.

IMPORTANT: When the access control profile is first generated, object level access control must be disabled. Apply the ACL profile, then mark Enable, then click Apply.

Enable Form Fill Authentication: Specifies whether to enable Single Sign-on with login HTML forms from origin servers.

Refresh ACLCheck: Refreshes the ACLCheck parameter settings.

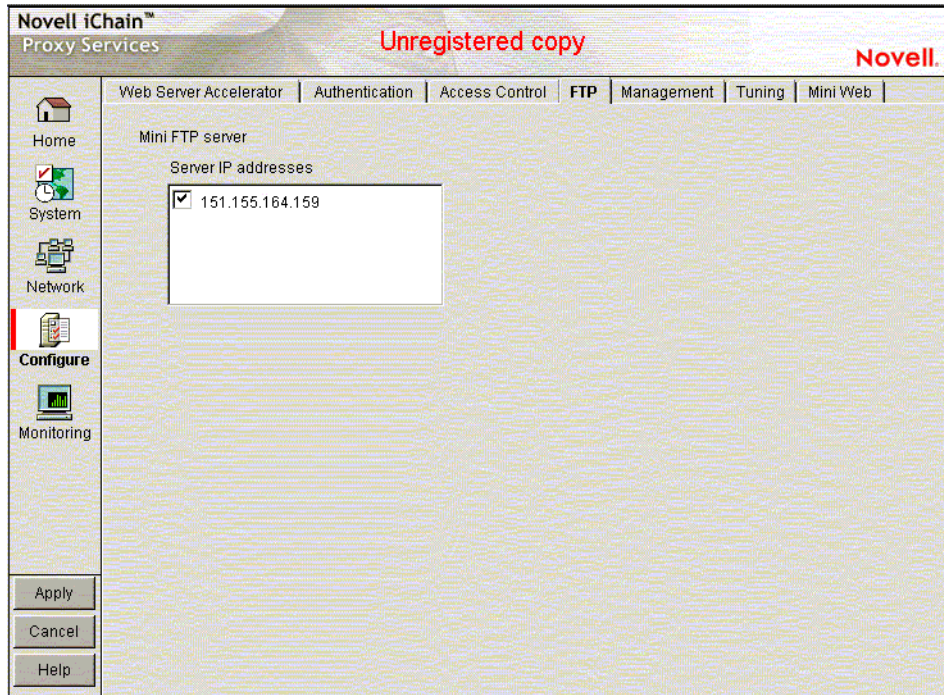
Refresh OLAC: If OLAC has been enabled, use Refresh to force an immediate reread of OLAC parameters.

Refresh Form Fill: Refreshes the connection to LDAP servers and Form Fill policy rules.

FTP Tab

Path: Configure > FTP

Figure 104 FTP Tab



The FTP tab lets you configure the appliance to provide an FTP listening address (Mini FTP Server) for appliance management, FTP forward proxy (client acceleration), and FTP reverse proxy (Web server acceleration) services.

Mini FTP Server: Check an address in the Server IP Addresses list to enable it for FTP listening. If an address is not checked, you cannot upload or download files using FTP.

Management Tab

Path: Configure > Management

Figure 105 Management Tab

Novell iChain™
Proxy Services

Novell.

Web Server Accelerator | Authentication | Access Control | FTP | **Management** | Tuning | Mini Web

Home
System
Network
Configure
Monitoring

Enable pin list

Default refresh frequency: Once Immediately | Default refresh time: 4 am :00

URL Mask	Pin Type	Pin Links	Pin Images	Refresh Frequency/Time
/servlet*	Bypass			

Insert | Modify | Delete

Enable dynamic bypass (Requires router to be enabled)

Dynamic bypass duration: 0 :05

Dynamic bypass error codes to enable: All | None

400 404 408 412 416 502
 401 405 409 413 417 503
 402 406 410 414 500 504
 403 407 411 415 501 505

Enable dynamic bypass logging
Log Options

Do not cache objects with a ? in the URL
Do not cache objects with /cgi in the path

Apply | Cancel | Help | Reset

The Management tab lets you identify objects that will be either pinned (explicitly downloaded and retained in cache as long as possible) or bypassed (explicitly not cached when requested by users). It also lets you specify how pinned objects are stored on the appliance (pin type).

Enable Pin List: Check this box to activate pinning on the appliance. For information regarding what pin lists are and how they function, see [“The Pin List” on page 202](#) and [“Pin List Examples” on page 209](#).

NOTE: Checking this option affects only the pinning of objects on the appliance. It has no effect on whether objects are cached unless the pin type is set to Bypass.

Default Refresh Frequency: Lets you specify when the appliance checks to see if items should be added to or removed from the list of objects being pinned. At refresh time the appliance re-evaluates the objects in cache for the URL list and downloads those objects that have changed. For absolute URLs, objects in links are also evaluated down to the link level specified. Choices range from one time only to any arbitrary time interval.

Default Refresh Time: Lets you specify the time of day or the time interval when pinned objects will be refreshed. To specify an interval, you must first select Timed Interval from the Default Refresh Frequency drop-down list.

The Pin List

The pin list lets you specify URL patterns for identifying objects on the Web. You configure each URL pattern in the list with specific handling instructions as explained below.

URL Mask: This can contain complete or partial URL patterns. A single URL mask might apply to a large set of URLs, or it might be so specific that only a single file on the Web matches it. For more information, see [“Pin List Examples” on page 209](#).

Pin Type: This specifies whether and how the appliance will cache objects that match the URL mask.

- ◆ *Normal:* iChain Proxy Services handles objects matching the mask in the same way it handles any other requested objects. In other words, objects are cached but not pinned.
- ◆ *Cache:* iChain Proxy Services keeps the pinned objects in cache as long as possible, although they might be written to the appliance’s hard disk.
- ◆ *Memory:* iChain Proxy Services keeps the pinned objects in memory as long as possible, writes them to disk when memory gets too full, and places them back in memory as soon as they are requested by a user of the cache.
- ◆ *Bypass:* iChain Proxy Services does not cache the objects. In other words, you can use this option to prevent objects from being cached.

Pin Links: This specifies how many link levels iChain Proxy Services will follow the pin type rule you’ve established.

Pin Images: This option is used to pin image files that reside on a different host than the page requested.

Refresh Frequency/Time: This lets you specify a refresh frequency and time for the URL that is different from the default values shown above the pin list. After inserting the URL, click Modify to see the dialog box that lets you change these fields. Functionality is the same as described for the default refresh frequency fields.

Enable Dynamic Bypass

The Dynamic Bypass feature lets you configure the appliance so that specific errors from Web sites are explicitly not cached and subsequent requests to the Web sites are simply passed through for a specific time period. For more information, see [“Dynamic Bypass” on page 173](#).

Enable Dynamic Bypass (Requires Router To Be Enabled): Checking this option enables the Dynamic Bypass feature.

Dynamic Bypass Duration: The period of time in hours and minutes that the URL remains in the Dynamic Bypass list, causing the appliance to transparently pass matching requests through to the origin Web server.

Dynamic Bypass Error Codes to Enable: This section lists all valid HTTP 1.1 errors. You can check and uncheck all the errors using the All and None buttons, or you can check only the errors you want enabled for dynamic bypass.

Enable Dynamic Bypass Logging: Checking this option enables logging of all dynamic bypass transactions. Log entries include the time of the log entry, the word Bypassed, and the URL added to the Dynamic Bypass list. Click Log Options to set log file rollover options and specify the handling of the oldest log files. For help with setting these options, see [“Log Rollover Options” on page 179](#). The information most relevant to the dynamic bypass feature is contained in [“Configuration Step 3: Specifying Rollover Options” on page 185](#), and [“Configuration Step 4: Specifying Handling of Older Files” on page 185](#).

Caching Based On URL Content

The two drop-down lists below the pin list let you specify object caching based on the following:

- ◆ Whether URLs contain a question mark
- ◆ Whether URLs have /cgi in the path

The Cache option lets you specify whether cacheable objects that meet the criteria are always cached.

The Cache option is sometimes misinterpreted to imply that objects meeting the criteria are always cached. That is not the case. The proxy server appliance will not cache objects that Web server administrators have marked non-cacheable.

By the same token, the Do Not Cache option is sometimes misinterpreted to imply that objects meeting the criteria are never cached. That is also not the case. Objects containing question marks and /cgi in the path might meet other criteria that cause them to be cached. This option only causes the proxy server to ignore question marks and /cgi in determining whether to cache objects.

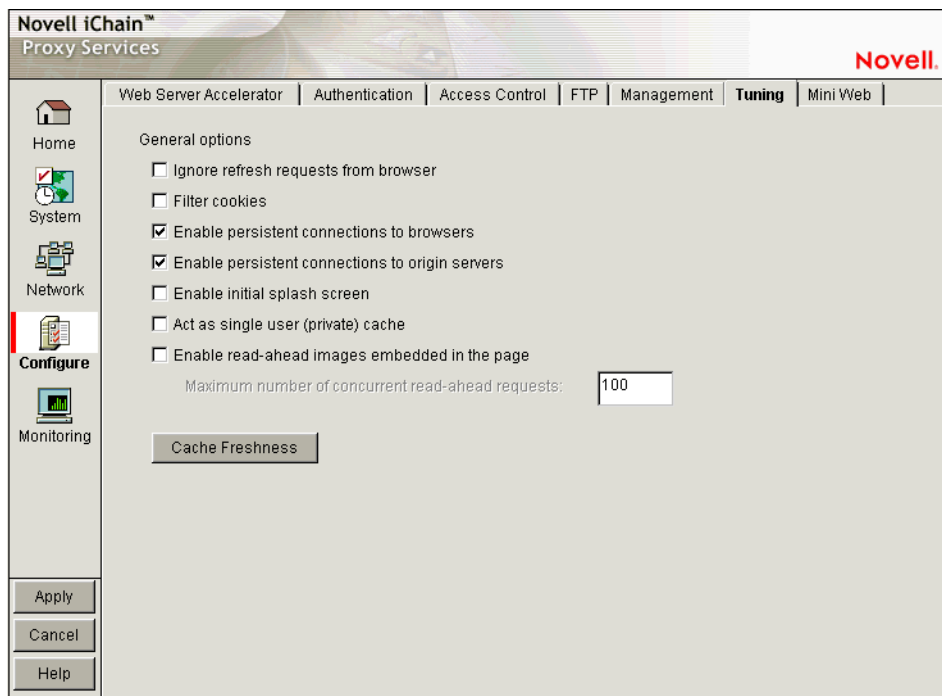
Resetting the Defaults

Clicking the Reset button resets the two drop-down lists back to their default (do not cache) settings.

Tuning Tab

Path: Configure > Tuning

Figure 106 Tuning Tab



The Tuning tab lets you restrict and enable functionality that affects all appliance operations. The implications for each option are explained below.

Ignore Refresh Requests from Browser: When a user clicks Refresh or Reload in the browser, the default system action is to send a new request to the Web server. Checking this option causes refresh or reload requests to be filled from the cache until cache freshness parameters are met. See [“Cache Freshness Dialog Box” on page 343](#).

Enable Persistent Connections to Browsers: If enabled, all connections from browsers to the appliance remain open. This makes the response time between the appliance and browsers faster.

Enable Persistent Connections to Origin Servers: If enabled, all connections from browsers to Web servers (through the appliance) remain open. This can cause some Web servers and their networks to crash, depending on the number of simultaneous connections they support.

Enable Initial Splash Screen: If enabled, browsers receive first-time and periodic notification that their requests are being processed by the appliance. The splash screen is customizable so that ISPs, for example, can advertise the fact that they are providing accelerated Web service. For information on customizing the splash screen, see [“Using FTP to Customize the Appliance Splash Screen” on page 201](#). The splash screen is disabled (turned off) by default.

Act as a Single User (Private) Cache: If enabled, iChain Proxy Services caches objects that have been flagged for private caches only.

Enable Read-Ahead Images Embedded in the Page: If checked, iChain Proxy Services will retrieve and cache objects that have been flagged Read-Ahead. You specify the maximum number of Read-Ahead objects the proxy server will retrieve in the Maximum Number of Concurrent Read-Ahead Requests field.

Cache Freshness Dialog Box

Path: Configure > Tuning > Cache Freshness

Figure 107 Cache Freshness Dialog Box

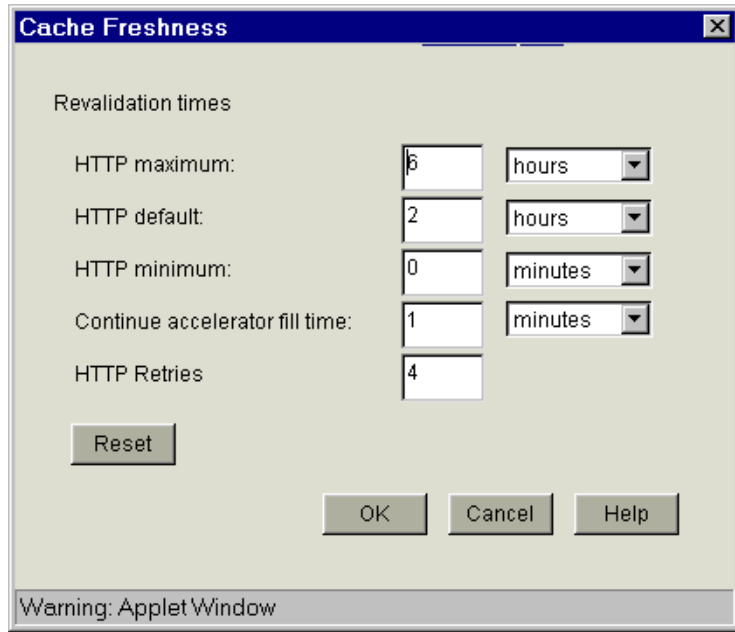
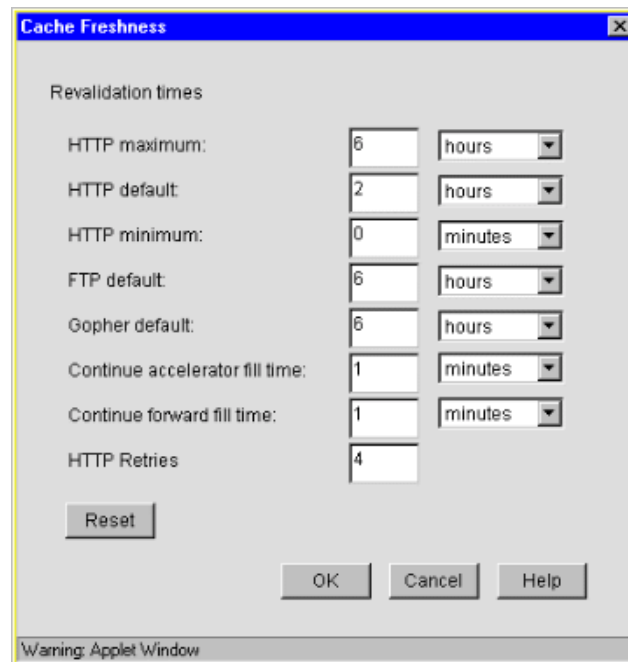


Figure 108 Cache Freshness Dialog Box



The Cache Freshness dialog box lets you set time values governing when iChain Proxy Services revalidates requested cached objects against those on their respective origin Web servers. If requested objects have changed, iChain Proxy Services re-caches them. Default field values are shown in [Figure 108](#).

iChain Proxy Services does not automatically recache objects when they expire. Expired objects are revalidated (and recached if they have changed) only when requested by browsers and in accordance with the time values set in the Cache Freshness dialog box. For more information on the appliance's cache-freshness features, see [“Cache Freshness” on page 155](#) and [“Managing Cache Freshness” on page 155](#).

HTTP Maximum: The maximum number of hours or days iChain Proxy Services will serve HTTP data from cache before revalidating it against content on the origin Web server. No object is served from cache after this value expires without being revalidated.

This overrides a freshness or Time to Expire header value specified by the Webmaster if he or she specified a longer time.

You use this value to reduce the maximum time iChain Proxy Services waits before checking whether requested objects need to be refreshed.

HTTP Default: The value iChain Proxy Services uses to determine when to revalidate requested objects for which Webmasters have not specified a freshness or Time to Expire header value.

HTTP Minimum: The minimum number of hours or minutes iChain Proxy Services will serve HTTP data from cache before revalidating it against content on the origin Web server. No requested object will be revalidated sooner than specified by this value.

This overrides the freshness or Time to Expire header value specified by the Webmaster if he or she specified a shorter time.

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

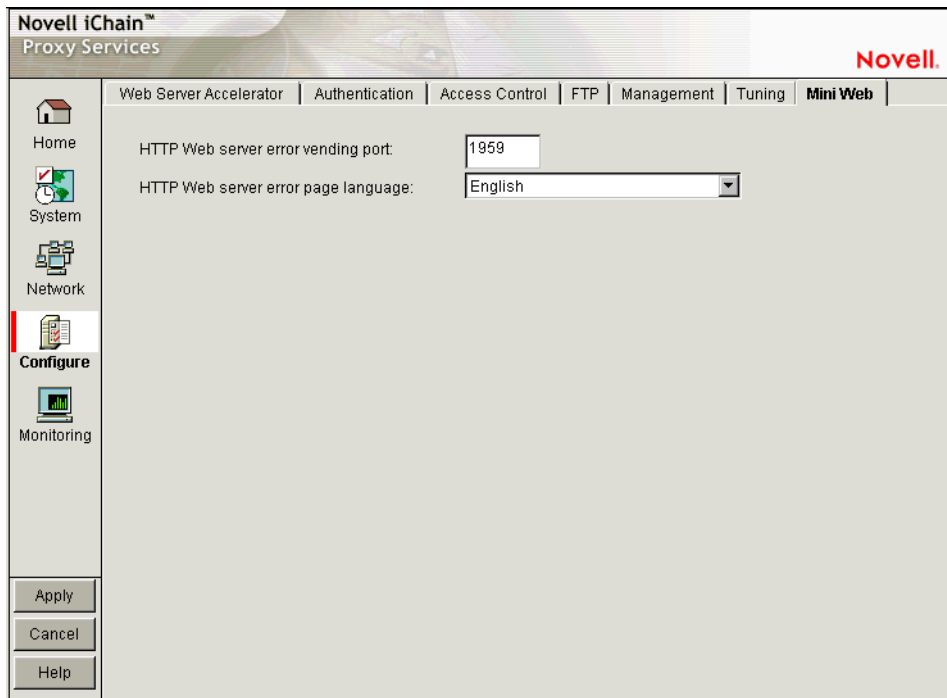
Continue Accelerator Fill Time: The number of minutes or hours that the appliance's Web acceleration services ignore browser request cancellations and continue downloading objects from the target Web server until the download is complete.

HTTP Retries: The number of retry requests that will be issued.

Mini Web Tab

Path: Configure > Mini Web (to see this tab, click the upper-right corner of the Tuning tab.)

Figure 109 Mini Web Tab



The Mini Web tab lets you configure how appliance-generated error pages are vended to browsers.

HTTP Web Server Error Vending Port: The port the browser will use when requesting objects that are part of the error pages. Changing this value does not affect the port for appliance administration, which is fixed at 1959.

HTTP Web Server Error Page Language: From the drop-down list, you can select a language for appliance-generated error messages to browsers.

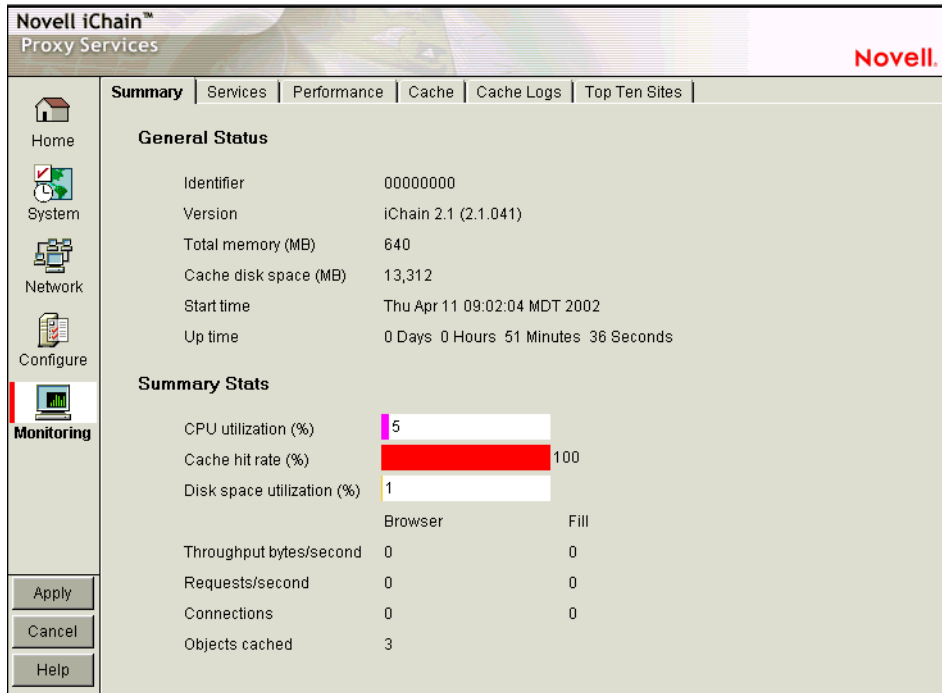
The Monitoring Panel

You can monitor various appliance activities and statistics in the browser-based tool.

Summary Tab

Path: Monitoring > Summary

Figure 110 Summary Tab



The Summary tab shows key appliance statistics at a glance. Statistics are refreshed every second.

Identifier: The make, model, and serial number of the appliance.

Version: The current system software version.

Total Memory (MB): Total available memory.

Cache Disk Space (MB): Total disk space available for caching. The amount shown is smaller than the total appliance disk space because it doesn't include the operating system and log partitions. Check this field to verify whether the proxy server has detected all disks installed on the appliance.

Start Time: The last time the appliance was started.

Up Time: Total time the appliance has been running since last started.

CPU Utilization (%): The current CPU utilization rate. Use this chart for capacity planning.

Cache Hit Rate (%): The current cache hit rate. A high cache hit rate indicates that the caching system is off-loading significant request processing from Web servers whose objects have been cached. Use this chart for capacity planning.

Disk Space Utilization (%): The percentage of caching disk space currently in use.

Throughput Bytes/Second: Current throughput.

Requests/Second: The rate at which browser clients are requesting Web objects.

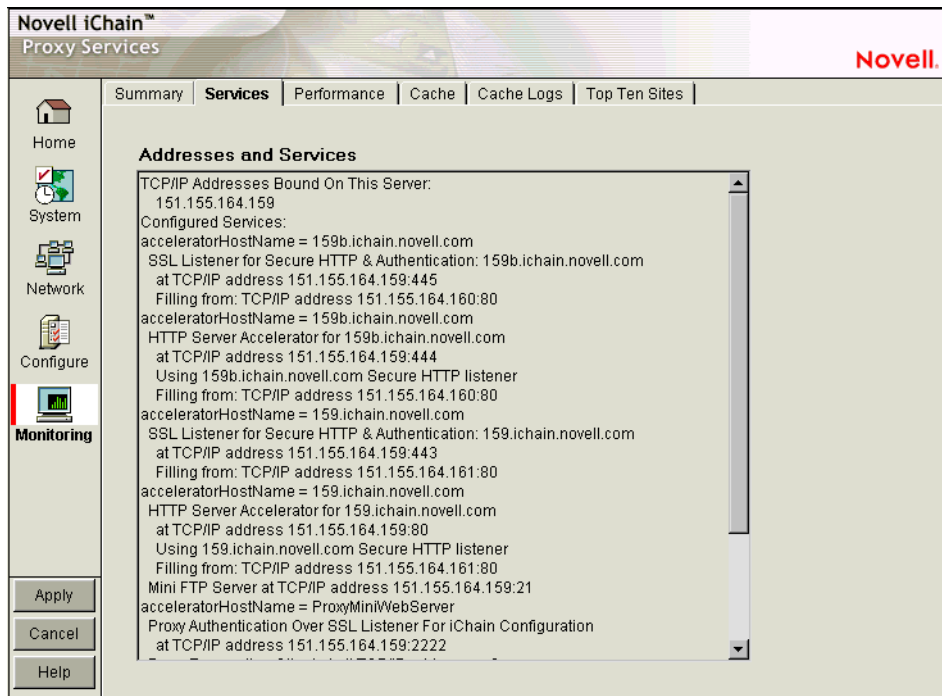
Connections: The total number of TCP connections that are active, idle, or closing.

Objects Cached: The total number of Web objects that have been cached.

Services Tab

Path: Monitoring > Services

Figure 111 Services Tab



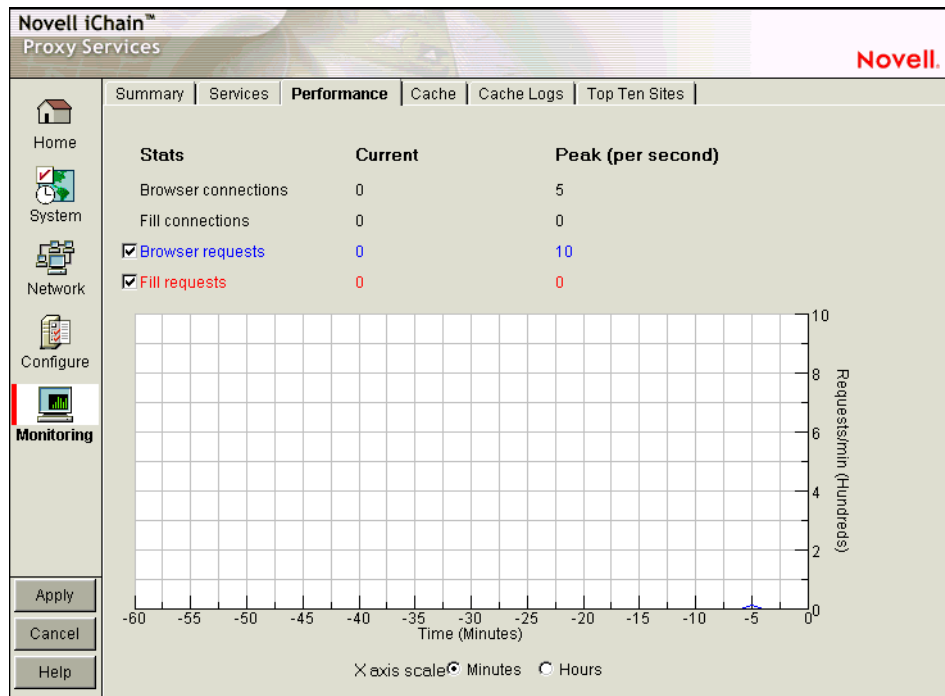
The Services tab shows you the IP addresses that are bound to appliance network cards and the services that are active. This information is refreshed every minute.

Addresses and Services: Use this list for troubleshooting problems with configured services. It shows active services along with the IP addresses and ports that they are running on. When the appliance detects errors, it displays appropriate error messages next to the services.

Performance Tab

Path: Monitoring > Performance

Figure 112 Performance Tab



The Performance tab shows current and peak levels of usage in terms of TCP connections and HTTP requests. The tab also displays a graph of HTTP requests from browsers to the appliance and from the appliance to origin Web servers.

Statistics are updated every ten seconds. The graph is updated once a minute.

Browser Connections: The current and peak numbers of browser connections to the appliance.

Fill Connections: The current and peak numbers of connections that the appliance has opened to origin Web servers.

Browser Requests: The current and peak numbers of browser HTTP requests per second made to the appliance. Check this box to enable graphing of browser requests.

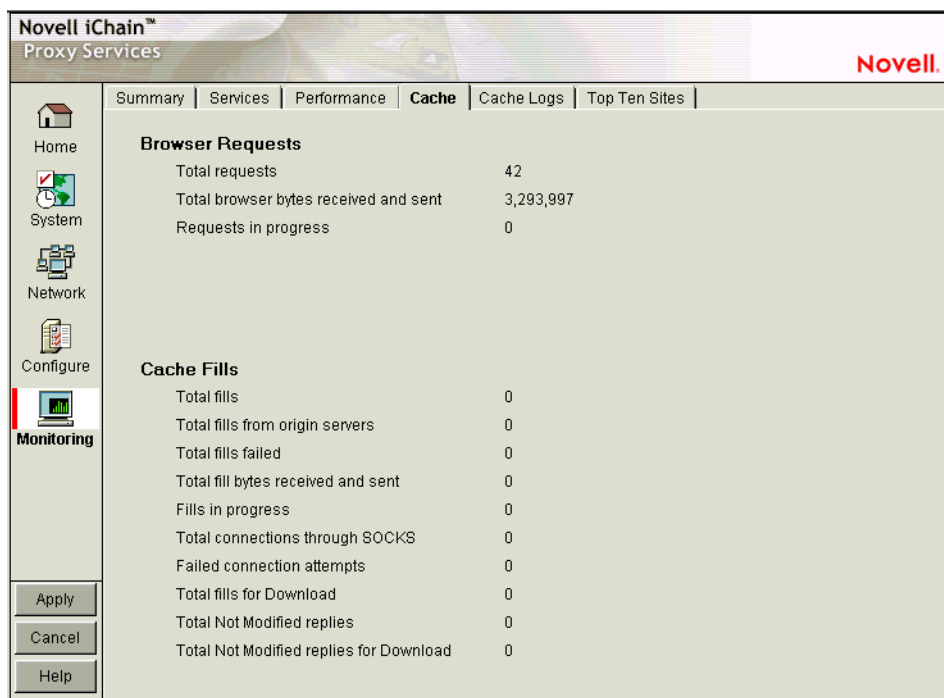
Fill Requests: The current and peak number of appliance requests per second to origin Web servers. Check this box to enable graphing of requests to origin Web servers.

Requests Graph: The HTTP browser requests to the appliance per minute (blue line) and HTTP fill requests to origin Web servers per minute (red line). Click Minutes or Hours to select the scale of the X axis. Minutes displays a one-hour history; Hours shows a 24-hour view.

Cache Tab

Path: Monitoring > Cache

Figure 113 Cache Tab



The Cache tab shows statistics for browser requests to the appliance and for appliance requests to origin Web servers. Statistics are refreshed every ten seconds.

Total Requests: The total number of requests that browser clients have made since the appliance was started.

Total Browser Bytes Received and Sent: The total bytes that browser clients have sent to and received from the appliance.

Requests in Progress: The number of active browser requests that are currently being processed by the appliance.

Total Fills: The total number of fill requests the appliance has made to origin Web servers and to neighbors in its cache hierarchy.

Total Fills from Origin Servers: The total number of fill requests the appliance has made to origin Web servers.

Total Fill Bytes Received and Sent: The total bytes the appliance has sent and received in order to fill its Web object cache.

Fills in Progress: The number of active fill requests that the appliance is waiting for.

Total Connections through SOCKS: The total number of connections the appliance has made through a firewall in order to fill its Web object cache.

Failed Connection Attempts: The total number of failed connection attempts the appliance has made while attempting to fill its Web object cache.

Total Fills for Download: The total number of fill requests sent to origin servers that originated from system read-ahead functionality and batch downloads.

Total Not Modified Replies: The total number of 304 Not Modified replies received for all fill requests to origin servers.

Total Not Modified Replies for Download: The total number of 304 Not Modified replies received for fill requests to origin servers that originated from system read-ahead functionality and batch downloads.

Cache Logs Tab

Path: Monitoring > Cache Logs

Figure 114 Cache Logs Tab

The screenshot shows the Novell iChain Proxy Services interface. The 'Cache' tab is selected, and the 'Cache Logs' sub-tab is active. The interface is divided into a left navigation pane and a main content area. The navigation pane includes links for Home, System, Network, Configure, and Monitoring (which is highlighted). Below the navigation pane are 'Apply', 'Cancel', and 'Help' buttons. The main content area displays two sections: 'Browser Requests' and 'Cache Fills', each with a table of statistics.

Browser Requests	
Total requests	42
Total browser bytes received and sent	3,293,997
Requests in progress	0

Cache Fills	
Total fills	0
Total fills from origin servers	0
Total fills failed	0
Total fill bytes received and sent	0
Fills in progress	0
Total connections through SOCKS	0
Failed connection attempts	0
Total fills for Download	0
Total Not Modified replies	0
Total Not Modified replies for Download	0

The Cache Logs tab provides access to logs by format and service.

Log Format: These options let you choose the format of the logs you want to download and view.

Select a Service: Clicking a service name displays the associated logs in the Log Files of the Chosen Service list.

Log Files of the Chosen Service: Contains a list of log files matching the format and service options you have selected.

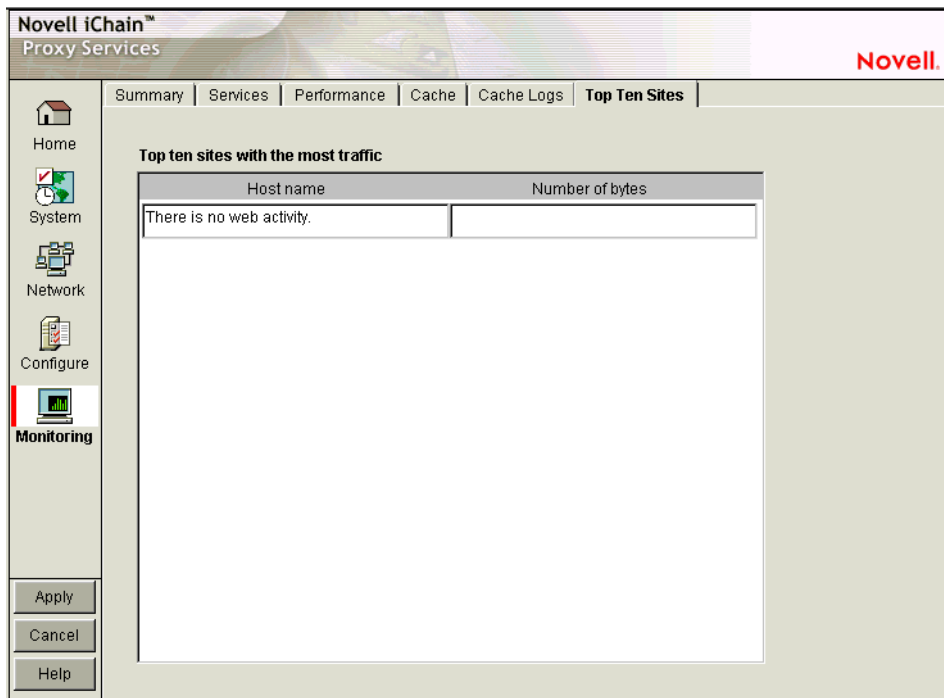
Download: Loads the log file into a separate browser window.

Delete: Removes the log file from the appliance.

Top Ten Sites Tab

Path: Monitoring > Top Ten Sites

Figure 115 Top Ten Sites Tab



The Top Ten Sites tab displays a list of origin Web servers with more than 0 bytes cached on the appliance. The ten sites with the most total bytes cached are sorted in descending order.

Command Line Reference

If you're working in a Telnet session or from the command line, you have 23 appliance commands available, several of which can be used with various parameters and values. These commands are listed in the table below.

iChain Proxy Services includes help for all commands and their associated arguments and parameters.

To see a list of commands, enter **help** at the command line.

To see a list of arguments for a command, enter **help command**.

To get help for a specific command/argument combination, enter **help command_argument**.

Command	Function	Requires Arguments?
<code>add</code>	Adds the new value to the current value.	Yes
<code>apply</code>	Applies the changes made at the command line. Some changes require a system restart, some merely suspend proxying and then restart, and others do not interrupt any process.	No
<code>cancel</code>	Discards all changes that are pending since the last <code>apply</code> command.	No
<code>clear</code>	Removes all items from a list or all settings from an argument.	Yes
<code>clearscreen</code>	Clears the current screen.	No
<code>createsessionbrokerkey</code>	Creates encryption key for Session Broker communication	No
<code>export</code>	Exports the named file.	Yes
<code>factorysettings</code>	Restores the appliance to original factory settings.	No
<code>get</code>	Displays current settings.	Yes
<code>health</code>	Displays the appliance health status.	No
<code>help</code>	Displays a list of available commands.	No
<code>identity</code>	Displays the appliance manufacturer, serial number, and hardware configuration.	No
<code>import</code>	Imports the named file.	Yes
<code>installsessionbrokerkey</code>	Installs Session Broker encryption key from floppy	No
<code>lock</code>	Re-locks the iChain Proxy Server console	No
<code>ping</code>	Sends a ping request to the addresses specified. Ports are optional.	Yes
<code>purgecache</code>	Purges the cache buffers.	No
<code>remove</code>	Deletes the specified value.	Yes

Command	Function	Requires Arguments?
<code>resetlearnedrout</code>	Resets the internal router table when the appliance is acting as a router.	No
<code>restart</code>	Restarts the appliance. All proxying ceases until the system restarts.	No
<code>restore</code>	Restores a configuration file to its directory	Yes
<code>restorefromclones</code>	Replaces the current appliance image with the clone image.	No
<code>scan</code>	Rescans disk drives on the appliance.	No
<code>set</code>	Sets an option by executing a clear command followed by an add command. (Existing settings are cleared when the set command is used.)	Yes
<code>shutdown</code>	Shuts down the appliance. All functionality ceases.	No
<code>updateclones</code>	Replaces the clone image with the current appliance image.	No
<code>version</code>	Displays the current version.	No

Troubleshooting the Command Line

Commands entered return an error

- ◆ Make sure you use the equal (=) sign when setting or adding, for example, `set forward enable=yes`
- ◆ Try the command again. Sometimes a command will fail the first time it's entered.

I made several changes that don't show when I use the GET command to display them

- ◆ You must conclude with the apply command to make the values take effect.

Connecting through Telnet

IMPORTANT: Telnet access is not secure unless a password is set. We strongly recommend you set system passwords as part of the appliance initialization process. For more information, see ["Password Dialog Box" on page 287](#).

You can manage the iChain Proxy Services by using commands from a workstation with a Telnet connection to the appliance.


Starting a Telnet Session

This section assumes you are using a Windows* 95/98 or Windows NT* 4 workstation. If you are using another type of workstation, use the following steps as general guidelines to configure your Telnet software to work with the appliance.

IMPORTANT: The appliance Telnet connection supports only the VT-100 terminal type.


- 1 Ensure that you have a network connection between the workstation that you are running Telnet on and the appliance.
- 2 Click Start > run TELNET.EXE.
- 3 From the Telnet screen, click Terminal > Preferences.

Figure 116 Telnet Preferences Dialog Box



- 4 Under Terminal Options, check VT 100 Arrows.
- 5 Under Emulation, check VT-100/ANSI.
- 6 Set the Buffer Size to 25.
- 7 (Optional) Set any of the other preferences that you desire.
- 8 Click OK.
- 9 Click Connect > Remote System.

Figure 117 Connect Dialog Box



- 10** Enter the appliance IP address in the Host Name field > select Telnet for the port and VT100 for the terminal.

The uppercase VT100 option usually works better with the appliance than the lowercase vt100 option.

- 11** Click Connect.
- 12** Type the Config user password.

All appliance passwords are case-sensitive.

Setting Up a Appliance Through Telnet

The Initial Installation Guide explains how to initialize an appliance and configure forward proxy services using the browser-based tool. The following procedures explain how to complete this task from the command line:

- 1** Start a Telnet session on the client machine.

For help starting a Telnet session, see [“Starting a Telnet Session” on page 358](#).

The starting address for Telnet should be 10.1.1.1, which is the address of eth0 on the appliance.

IMPORTANT: Versions later than 1.0 have no default password for Telnet access. Telnet is not secure unless a password is set for the Config user. (Telnet doesn't provide View user access.)

We strongly recommend you set system passwords as part of the initialization process. For more information, see [“Password Dialog Box” on page 287](#).

Telnet always prompts for a password. If you have not set a password for the Config user, enter a null password by pressing Enter.

After logging in to Telnet, you see the following prompt:

System Console

>

- 2** At the System Console prompt, enter the following:

```
set eth1 address=iii.iii.iii.iii, mask=mmm.mmm.mmm.mmm
set dns server=ddd.ddd.ddd.ddd
set dns domain=x
set gateway nexthop=ggg.ggg.ggg.ggg, metric=t
apply
```

The variables are i = the IP address, m = the subnet mask, d = the DNS server IP address, x = your domain name, g = the gateway IP address, and t = the number of hops to the next hop.

- 3** (Optional) Configure the appliance to provide forward proxy service.

At the System Console prompt, enter the following:

```
add forward address=iii.iii.iii.iii
apply
```

The variable i = the IP address you entered in [Step 2](#).

The appliance is now configured to begin providing forward proxy service. To configure client browsers to use the forward proxy service, see the Initial Installation Guide.

Additional Telnet Information

If the appliance has a monitor attached, you will notice that commands issued through a Telnet connection are echoed on the appliance monitor.

If you get a message asking whether you want the X-session displayed on a display other than the default, you have selected the wrong terminal type. Click Connect > click Disconnect > repeat the connection procedure starting with [Step 9 on page 358](#) and ensuring you have selected the VT100 terminal type in [Step 10 on page 359](#).

Establishing a Null-Modem Connection

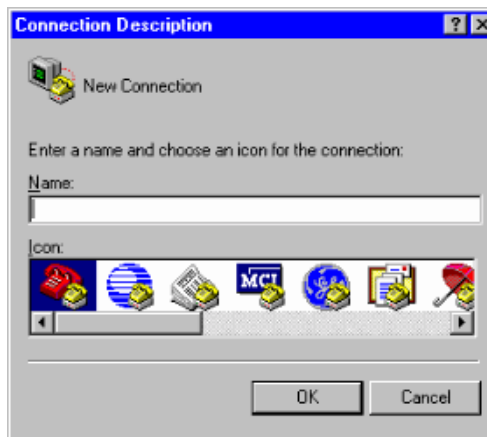
This section assumes you are using a Windows 95/98 or NT 4 workstation. If you are using another type of workstation, use the following steps as general guidelines to configure your terminal emulation software to work with the appliance.

IMPORTANT: The appliance null-modem connection supports only the ANSI terminal type.

To establish a null-modem connection with the appliance:

- 1 Connect a null-modem cable to the serial port on the workstation and the appliance.
- 2 Click Start > run hypertrm.exe.

Figure 118 Connection Description



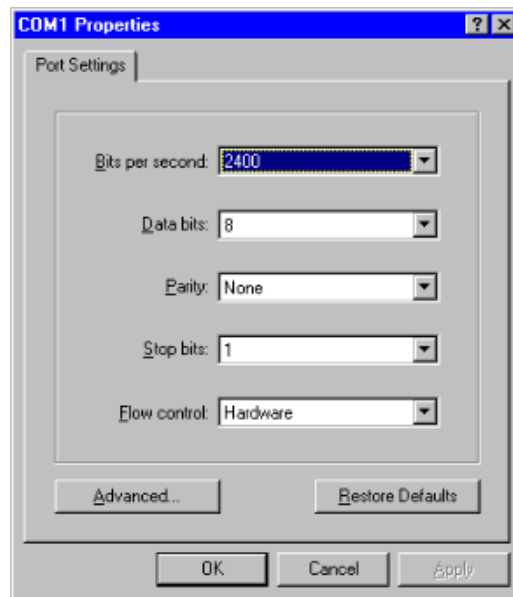
- 3 Enter a name for the connection > select an icon as instructed > click OK.

Figure 119 Direct to Com Option



- 4 Click the Connect Using drop-down list > select a Direct to Com option corresponding to the serial port connection on your workstation > click OK.

Figure 120 Com Properties Dialog Box



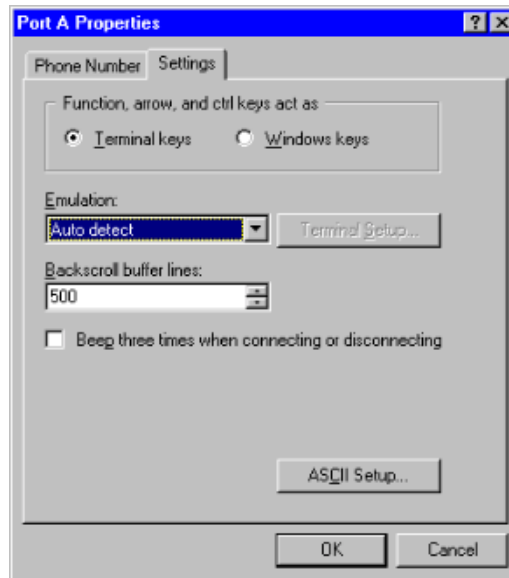
- 5 Set the properties according to the following table > click OK.

Property	Value
Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

- 6 Click File > Properties > Settings.

A dialog box similar to the following appears:

Figure 121 Properties Dialog Box



- 7** Click Terminal Keys.
- 8** Click the Emulation drop-down list > select ANSI.
- 9** (Optional) Specify cursor behavior by clicking Terminal Setup.
- 10** Click ASCII Setup > ensure that only Wrap Lines That Exceed Terminal Width is checked.
- 11** Click OK > OK.
- 12** Press Enter.

A command line prompt appears in Hyper Terminal.

You can now use all console commands to manage the appliance.

Additional Null-Modem Information

If the appliance has a monitor attached, you will notice that commands issued through a null-modem connection are not echoed on the appliance monitor.

The following commands are available when using a null-modem connection:

- ◆ `cls` clears the screen

- ◆ `_info` displays the Com port settings for the appliance

When accessing the appliance through a null-modem connection, arrow keys function in the following ways:

- ◆ The Back-arrow acts like a backspace or erase.
- ◆ The Forward-arrow acts like a space.
- ◆ The Up-arrow displays a history of previously executed commands (beginning with the most recent command).
- ◆ The Down-arrow scrolls forward through the command history and ends with a blank line.

Troubleshooting Telnet

Telnet never starts

To establish a Telnet connection, you must be able to ping the server.

Telnet starts after a long time

If you previously ran a session and turned on Transparent Proxy, Telnet might be very slow in starting.

Commands at the bottom of the screen look strange or don't make sense

The display might not update correctly. Enter `clearscreen` to get a new screen and start at the top.

Upgrades

You can apply over-the-wire upgrades to the appliance as they become available. These over-the-wire upgrades provide a quick and easy method to install patches and fixes from Novell. For more information, see [“Upgrade Tab” on page 294](#).

Critical Information

Making Sure You Update the Clone Image before and after Upgrading

Both before and after an upgrade or support pack installation, you must update the appliance's clone image to avoid having the upgraded system overwritten by the older clone image. For more information, see [Step 8 on page 366](#).

Preserving Configuration Settings

As mentioned previously, the system upgrade retains all appliance configuration settings. As a precaution, we recommend that you also export your appliance's configuration settings to a DOS-formatted floppy diskette prior to the upgrade.

For assistance, refer to [“Backing Up the Appliance Configuration” on page 166](#).

After the upgrade is completed, if you need to import the configuration file, refer to [“Apply a named configuration file from a floppy” on page 165](#).

Upgrading through a Firewall

In most cases upgrading through a firewall is not a problem. If your environment allows HTTP access to the Web, the appliance should be able to retrieve the upgrade files as easily as a browser downloads Web pages.

If normal HTTP access is restricted within your firewall, the appliance attempts to retrieve upgrade packages through firewalls in one of the following three ways:

1. First, the over-the-wire upgrade checks whether the appliance can use an ICP or CERN parent. If so, the appliance uses the parent to download the upgrade package.
2. If an ICP or CERN parent is not available, the over-the-wire upgrade checks whether the appliance is configured as a forward proxy with access through the firewall. If it is, the appliance tries the following two methods, in order:
 - a. If the firewall acts as a SOCKS server, you must configure the appliance as a SOCKS client. It can then retrieve the upgrade package from the origin server.

- b. If the firewall is not acting as a SOCKS server, you must create a hole through the firewall that allows the appliance to make HTTP connections to the origin server with the upgrade package.

Close the hole as soon as the upgrade is downloaded.

3. If neither of the previous two methods is available, the over-the-wire upgrade attempts to establish a direct connection with the origin server.

To enable this connection, you must create a hole through the firewall and close it as soon as the upgrade is downloaded.

Downloading and Installing the Upgrade

To upgrade your appliance:

- 1** In the browser-based management tool, click System > Actions > Update Clone > Update Clone.
- 2** When the update is complete, click the Upgrade tab.
- 3** Check Enable Download > type the URL for the upgrade.
The URL is available from your appliance vendor.
- 4** Click the Download Time drop-down list > select the time you want the download to occur.
- 5** Check Enable Install.
- 6** Click the Install Time drop-down list > select the time you want the upgrade to be installed.
- 7** Click Apply.
- 8** As soon as possible after the upgrade is installed, update the appliance's clone image by clicking System > Actions > Update Clones.

This is necessary to avoid the appliance automatically applying an earlier clone image which could make the proxy server unstable.

Uninstalling the Most Recent Upgrade

To uninstall the most recent upgrade:

- 1** In the browser-based management tool, click System > Upgrade.
- 2** Check Enable Uninstall.

3 Click Apply.

Only the most recently installed upgrade can be uninstalled.

368 Place Book Title Here

C

Rewriter Support

There are a wide variety of Web servers that may be hosted by Novell® iChain®. Some Web servers are powered by high-end applications running on Oracle* or iNotes*. Some Web servers are running portals like Novell Portal Services. These advanced Web servers usually have some conditions that can be solved by rewriter features.

This section examines what iChain can accomplish in rewriting data. It also examines a number of rewriting features that allow iChain to be a proxy for advanced backend servers.

Internal Rewriter

The Novell iChain internal rewriter is used to rewrite URL references to Web servers into the scheme (HTTPS or HTTP) and host names that are defined by an iChain accelerator.

For example, a company has a Web site with the host name of "internal.web.site.com". iChain provides access to this site through the host name of "mycompany.com". In many of the HTML pages there are HTML references (href=) to "http://internal.web.site.com/file1.html". If a user clicks on this reference and it has not been rewritten to "https://mycompany.com/file1.html", the page will not be found. This is because the administrator does not want "internal.web.site.com" to be in the company-wide DNS table.

The internal rewriter, when turned on, is used to rewrite URLs that contain backend Web server host names (internal.web.site.com) with the iChain accelerator host names (mycompany.com).

The internal rewrite is also used to rewrite the scheme for URL references. For example, a Web page may have a URL of "http://mycompany.com/file1.html". If the accelerator for mycompany.com is using SSL sessions between the

browser and iChain, then these URL references must be rewritten to "https://mycompany.com". Otherwise, each link that a user clicks on will bounce from HTTP to HTTP to re-establish a new SSL session.

The internal rewriter in iChain 2.0 can be used only on accelerators that had Secure Exchange enabled. This proved difficult for debugging problems because the data moving over the network was encrypted.

The internal rewriter in iChain 2.1 is enabled for all multi-homed children and multi-homed masters and on accelerators that have Secure Exchange enabled. This allows administrators to test the rewriting features over a clear (non-SSL) HTTP connection.

By default, the internal rewriter will not rewrite all URL references. The internal rewriter will try to rewrite all URL references that follow "href=" and "src=" within valid HTML syntax. The internal rewriter will also rewrite some URL references within the HTTP header (specifically, the "Set-Cookie" and "Location" tags).

Custom Rewriter

The Novell iChain custom rewriter is used to search and replace specific text within a text file. For example, a user may want to change all occurrences of "President Clinton" to "President Bush". The custom rewrite will not touch data within the HTTP header.

The custom rewriter must be loaded by running "REWRITE.NLM" from the system console. For more information, see [“Custom Rewriter — Rewrite Filter” on page 374](#).

The custom rewriter was developed because of the need to replace URL text pieces. For example, a document could have the following JavaScript variables:

```
var scheme = "http://"
var host = "internal.web.site.com"
var path = "/path/"
var file = "file1.html"
var URL = scheme + host + path + file;
```

These variables are used to build URL strings for the user. The internal rewriter can only rewrite true URL references such as http://internal.web.site.com/path/file1.html to the schema and host name of the

accelerator that accesses the internal Web site (<https://mycompany.com/path/file1.html>).

The custom rewriter has an overlapping region problem when used in conjunction with the internal rewriter. If both rewriters want to change the same characters of data within a file, it is undefined which rewriter will win out. In some cases, a couple of characters around the rewritten area will be lost or added. This problem happens when the custom rewrite is used to rewrite the same URLs that the internal rewriter also rewrites. The custom rewriter was not developed to search and replace full URL references; only URL pieces (as shown in the above example).

Internal Rewriter Features

At proxy startup time, the configuration file `sys://etc/proxy/rewriter.cfg` will be opened and read. The following are the sections of the configuration file. The configuration file does not support comments (`;` or `#`) and should contain two returns or lines between sections.

Rewrite All HTTP References

The most requested command in the rewriter config file is:

```
[http references]
rewrite=all
```

Currently the rewriter only rewrites the following references within an HTML document:

```
href="http://xxx.xxx.xxx"
src="http://xxx.xxx.xxx"
```

There are other `http://xxx.xxx.xxx` references within Java scripts, Applet parameters, or even text displays that are not rewritten. The internal rewriter, by default, will parse through all JavaScript without looking at any URL references. To rewrite these references, add the "rewrite=all" value to this section header.

This option will parse through the entire HTML file, looking for `http://` or `https://`. If the URL is a Web server host name reference, then the host name and scheme will be rewritten to the accelerator host name.

This option will not rewrite HTTP references within a query string because the data within the query string is considered part of the URL and only the first part of the URL (host name) is looked at.

HTML File Extensions

[Extension]

The internal rewriter will only rewrite the Mime Content-Type of text/HTML. The default file extensions are:

```
html
htm
shtml
jhtml
asp
jsp
NO EXTENSION
```

The format for adding extensions allows either a comma separating them on a line or each extension residing on a line of its own (see example below). Two carriage returns should separate the list from the next bracketed section. All of the extensions listed will be appended to the above list, so there is no need to list the seven extensions again. There is currently not a way to remove items from the above list when you add additional extensions.

```
[Extension]
home,myExtension
anotherLongExtension
a,b,c,d,e,f,g
```

Text/HTML Mime Type

```
[Mime Content-Type]
type=text/html
```

Some users fight the battle of a Web server having new extensions popping up every day. For example, a user might have a portal serving up pages or an Oracle server with named extensions that are named after the code that generated the page. This option will rewrite all of the pages that contain a Mime Content-Type of text/HTML regardless of the extension.

NOTE: The rewriter is built for parsing through HTML data that is formatted according to HTML syntax rules. Rewriting text or unformatted data may not achieve the desired

results. For this reason, the default rewrites the seven extensions discussed in ["HTML File Extensions" on page 372](#).

Alias Host Names

[Alias Host Names]

Sometimes a Web server references a host name that is not, nor can not, be defined as a Web server reference in an iChain accelerator. For example, a customer has the host name of home (<http://home/index.html>) but there may not be a DNS entry for "home", so the administrator cannot list this name in the Web server address box.

This section discusses how alias host names can be specified for an accelerator. On iChain 2.0 systems, the accelerator host name is used to reference the accelerator. On iChain 2.1 systems, the accelerator name will be enough to reference the accelerator.

The format may seem a little backwards from a normal search and replace format.

AcceleratorName=aliasName

```
[Alias Host Names]
www.cookie.com=home
www.cookie.com=alliance.state.com
www.cookie.com=intraigol.utahseg.com
www.cookie.com/child=TEST
```

The alias names are not case-sensitive because host names should not be case-sensitive.

Disabling the Rewriter

[Internal Rewriter]

AcceleratorName=disable

In iChain 2.1, the internal rewriter will always parse through HTML files, looking for any references that need to be rewritten. (This happens because the rewriting feature is important to the iChain product and customers traditionally want this feature on more than they want it disabled.)

In iChain 2.0, only HTML files that traveled over an SSL connection (using Secure Exchange) were rewritten. The problem with this approach is that it

was very difficult, if not impossible, for a Web master to debug the exact data that was being rewritten.

The internal rewriter can be an expensive operation in they way of looking for and replacing all full HTTP references. In some cases there no HTTP references that need to be rewritten because the name of the Web server is the accelerator DNS name.

Custom Rewriter — Rewrite Filter

The Novell iChain rewrite filter (REWRITE.NLM) provides a powerful tool which enables administrators to search and replace user-specified strings with new strings. The rewrite filter will replace all occurrences of user-specified string with a replacement string irrespective of the location original string in the data. This option is configured via a text file.

Rewrite Filter Configuration

Rewrite filter (rewrite.nlm) configuration is accomplished via text files. An example of a configuration file is given below:

```
[Name=oraclefilter]

[Extension]
html, htm

[URL]
sjf-siva.sjf.novell.com/oracle/*
sjf-siva.sjf.novell.com/database/OracleQuery.html
sjf-siva.sjf.novell.com/oraclereports/

[Replace]
var xsport      = "9003"<<>> var xsport      = "443"
var xsname      = "xyz01"<<>> var xsname      = "www"
var xconnectMode = "http"<<>> var xconnectMode = "https"
<PARAM name=serverHost value="xyz01.novell.com"><<>><PARAM
name=serverHost value="www.novell.com">
<PARAM name=serverPort value="9003"><<>><PARAM
name=serverPort value="443">
<PARAM name=connectMode value="https"><<>><PARAM
name=connectMode value="http"
```

[Name=<filtername>]: This is the name of the rewriter filter. It is used for filter identification in rwfilter commands.

[Extension]: The user can specify file extensions (in a single line) which need to be parsed. This section is optional, but if this section is not specified, all extension types will be parsed. Specifying this section will result in the parsing only specified extensions.

Conditions: All extensions should be specified in a single line (just after the Extensions section header). Individual extension elements should be separated by a comma.

For example:

```
[Extension]
html
```

Parses files with html extensions.

Parses `http://www.novell.com/index.html`.

Does not parse `http://www.novell.com/logo.gif`.

Does not parse `http://www.novell.com`.

```
[Extension]
html, htm, txt, ,
```

Parses files with extensions html, htm, txt, and any default files implied by a trailing slash (/).

Parses `http://www.novell.com/license.txt`.

Parses `http://www.novell.com/`.

Does not parse `http://www.novell.com/logo.gif`.

[URL]: This section lists all URLs to be parsed. The rules of URL specification are similar to the ACLCheck module as follows:

```
www.novell.com: Exact match
www.novell.com/contact/: All files in contact directory, but
not in subdirectories.
www.novell.com/contact/*: All URLs starting with
www.novell.com/contact, including subdirectories of contact.
```

[Replace]: The user can specify replacement pairs in this section. The string "<<>>" is used as a separator between the original string and the new string. The format is as follows:

```
[Replace]
<search string><<>><replacement string>
```

All the lines after this Replacement section header will be treated as replacement string pairs, so this should be the last section in the file.

Conditions: The Replace section should be the last one in the file. Each replacement pair (including the last one) must be terminated with the end of the line.

Usage

The name of the filter module is REWRITE.NLM. Configuration files could be supplied as command line parameters to filter.

Filename restrictions: Rewrite.nlm does not place any restrictions on the configuration filename format, but files which do not conform to DOS filename format (8.3) format) could cause errors while loading the file.

Load Time Command Line Options: The load time command line option is: `rewrite -s -f system/accell.rw -f etc/accel2.rw`.

Options:

- ◆ `-s`: Create separate screen for rewrite filter
- ◆ `-f`: Config file. This is the file from which configuration information is loaded. The user can specify as many config files as required. In the example above, `accell.rw` and `accel2.rw` are two configuration files in the specified format.

Runtime Console Commands

The following runtime console commands are also implemented:

- ◆ `rwfilter unload <filtername>`: Stops a filter (the filter name is as given in the configuration file).
- ◆ `rwfilter load <configFile>`: Loads a filter config file.
- ◆ `rwfilter list`: Lists all loaded rewrite filters by name.
- ◆ `rwfilter print <filtername>`: Displays rewrite configuration information (the filter name is as given in the configuration). Printed information should be similar to the configuration file.

Replacement Rules

The following are simple rules that the iChain rewrite filter uses:

1. String replacement is done as a single pass.
2. String replacement is not performed recursively. For example:


```
[Replace]
DOG<<>>CAT
A<<>>O
```

If the original string is DOG, the rewritten string will be CAT. That is, only one replacement will occur (CAT will not be further replaced with CAT).

3. Since string replacement is done in one pass, the string that matches first will take precedence. For example:

```
[Replace]
ABC<<>>XYZ
BCDEF<<>>PQRSTUVWXYZ
```

If the original string is ABCDEFGH, the replaced string is XYZDEFGH.

4. If two user-specified strings match the data portion, the original string of higher length will be used for replacement except in cases detailed in point 3 above. For example:

```
[Replace]
ABC<<>>XYZ
ABCDEF<<>>PQRSTUVWXYZ
```

If the original string is ABCDEFGH, the replaced string looks like PQRSTUVWXYZGH.

5. The rewrite filter module, REWRITE.NLM, applies only one matching filter (each configuration filter file constitutes a filter) for a request. If multiple filters match a request URL, the matching filter that was configured most recently will be used for data parsing.

378 Place Book Title Here

D

Upgrading Your iChain System

This appendix provides suggestions and steps on how to upgrade your current Novell® iChain® installation to iChain 2.1 Proxy Server and Authorization Server software.

The following topics are included in this appendix:

- ♦ “Internal Rewriter” on page 369
- ♦ “Upgrading from iChain 1.5” on page 383

Upgrading from iChain 2.0

This section discusses upgrading from the iChain 2.0 software version. The following steps should be considered:

- ♦ 1. Prepare the Current iChain Platform
- ♦ 2. Back Up the Existing iChain 2.0 Configuration
- ♦ 3. Upgrade eDirectory with the iChain 2.1 Schema Using the Install CD
- ♦ 4. Install ConsoleOne 1.3 and the iChain Snap-ins
- ♦ 5. Convert and Modify Existing ACL/ISO Definitions
- ♦ 6. Upgrade the Proxy Server to iChain 2.1
- ♦ 7. Test the System
- ♦ 8. Implement New Features

This section also addresses:

- ♦ Schema Differences Between 2.0 and 2.1

1. Prepare the Current iChain Platform

- 1 Prepare a test scenario with the customer (for each app, identify key profiles). Be aware of what each application requires as input (for example, simple authentication header, parameters passed using the command line).
- 2 Test the scenario on the running iChain 2.0 system and confirm that it is working.

2. Back Up the Existing iChain 2.0 Configuration

You will need to back up both the Authorization Server and the iChain Proxy Server.

To Back Up the Authorization Server

- 1 Back up eDirectory™.
- 2 Do an export to LDIF of the iChain objects (Access Control List, iChain Service Object, Communities).
- 3 Back up any custom tools or modules that may have been running on the Authorization server.
- 4 Rename the ConsoleOne® directory if ConsoleOne version 1.2x is installed (the iChain 2.1 Authorization Server CD ships with version 1.3). If this is not an option, rename the iChain snap-in and lib directories.

To Back Up the iChain Proxy Server

- 1 Do an export to a NAS file of the Proxy Server configuration and screen shot of all configuration screens.
- 2 Export or back up certificates that are being used by the proxy server.
- 3 Back up the following files:
 - ♦ /etc/hosts — contains host mappings to IP addresses
 - ♦ /etc/proxy/data — contains custom login pages (ca*.html)
 - ♦ /ichain/oac/oac.properties — contains advanced OLAC configuration settings
 - ♦ /etc/proxy/r_append.cfg — if any DNS search types changed
 - ♦ /system/appstart.ncf and /system/tune.ncf

- 4 Copy any tools or modules that you may have used from the server (for example, LSEARCH.NLM for LDAP testing, NETMON.NLM for taking traces).

NOTE: If possible, add an extra disk into the Proxy Server for the 2.1 install so that the 2.0 install can remain untouched while the 2.1 install is executed. Also, create a clone image of the 2.0 iChain build on the server before doing the upgrade.

Make sure you have everything needed to restore a valid iChain 2.0 Proxy Server image.

Please note that 2.1 schema is compatible with 2.0, meaning that if you leave your 2.0 iChain Service Object (ISO) untouched, you could have one proxy server running 2.0 while a second one is being upgraded. (This could help in doing a seamless migration and an easy roll back.)

3. Upgrade eDirectory with the iChain 2.1 Schema Using the Install CD

The install script will generate many BURP errors during this phase that can be ignored. These errors are generated because many of the modifications to the schema that the install script is trying to perform are already in place.

NOTE: If the tree you are upgrading also contains Novell® BorderManager® schema extensions, you will need to manually re-link the "brdsrvsOutgoingAcl" attribute with the object class named "brdsrvsACLRule". This is done easily in ConsoleOne schema manager, after applying the new schema and reloading ConsoleOne.

4. Install ConsoleOne 1.3 and the iChain Snap-ins

If it isn't already installed, install Console 1.3 and also install the iChain snap-ins from the Authorization Server CD. This is required for any RADIUS or token-based authentication setup.

5. Convert and Modify Existing ACL/ISO Definitions

Convert and modify existing Access Control List (ACL)/iChain Service Object (ISO) definitions to match new specifications in iChain 2.1.

The ConsoleOne snap-ins that ship with iChain 2.1 can detect iChain 2.0-formatted objects. After upgrading the Authorization Server from 2.0 to 2.1 and selecting properties of the original 2.0 ISO with the new 2.1 snap-ins, the ISO will be automatically extended with the new required attributes.

NOTE: If administrators are creating completely new objects, the following should be considered:

1. The ISO has many new attributes in 2.0. The most important of these involves ACLCHECK dynamic LDAP search attributes.
2. If you decide to recreate the ISO, the corresponding Rule Objects referencing the old ISO's protected resources must be recreated. If this is not done, ACLCHECK will report "old version" errors.

6. Upgrade the Proxy Server to iChain 2.1

- 1** Image the proxy server with iChain 2.1.
- 2** Unlock the Proxy Server system console by typing "unlock" at the prompt. You do not need to specify a password. Press Enter.
- 3** Import the NAS file by placing the floppy containing the CURRENT.NAS file into the proxy server. Type "import floppy".

Wait until the "completed execution of autoload" is displayed at the server console.
- 4** Import the server certificates that were backed up from the 2.0 server.

If problems exist accessing the proxy server GUI, do the following from the Internet Caching System console:
 - 4a** Run the `_kill` application to kill the java ServerApplication thread and all support modules.
 - 4b** Unload the CERT.NLM at the system console.
 - 4c** Reload CERT.NLM.
 - 4d** Execute APPSTART.NCF at the system console.
- 5** Restore the files backed up in 2-3 and 2-4.

NOTE: The OAC.PROPERTIES file will not be needed unless some non-default parameters were required for functionality in 1.5 (for example, increasing worker threads, synchronization interval).
- 6** Using the proxy server GUI, run the health check to make sure that all services are up and running.
- 7** Verify if the eDirectory server still has community objects (which shipped with 1.5, but not with 2.0 or 2.1) and rules based on community objects. If this is the case, modify the APPSTART.NCF to load ACLCHECK with the /M option.
- 8** Verify that you can access the iChain protected resource from the browser.

7. Test the System

- 1 Complete an offline test using your defined scenario.
- 2 Complete a production test.

8. Implement New Features

Only after you have confirmed that the old features are working should you enable any of the new iChain 2.1 features.

Schema Differences Between 2.0 and 2.1


The iChain 2.1 schema file is found on the Authorization Server CD in the \schema subdirectory. This file documents all iChain attributes and lists the new attributes added in version 2.1.

Upgrading from iChain 1.5

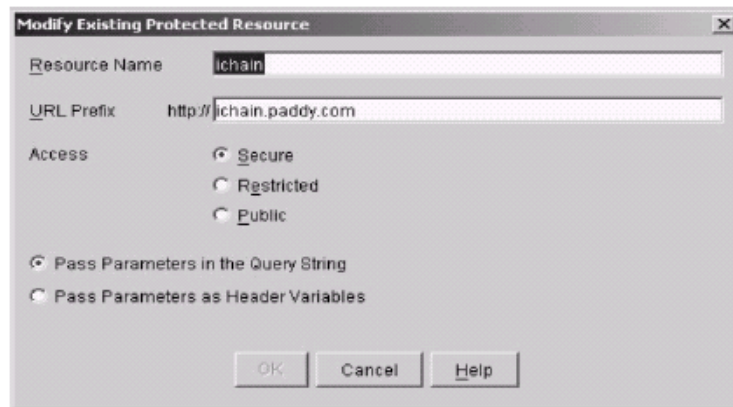
Upgrading from iChain 1.5 to iChain 2.1 will require the same steps as described in [“Internal Rewriter” on page 369](#), with a primary difference from what is described in [“5. Convert and Modify Existing ACL/ISO Definitions” on page 381](#). You should replace this section with the instructions below:

5. Converting and Modifying Existing ACL/ISO Definitions When Upgrading from iChain 1.5

The ConsoleOne snap-ins that ship with iChain 2.1 can detect iChain 2.0-formatted objects. After upgrading the Authorization Server from 1.5 to 2.1 and selecting properties of the original 1.5 ISO with the new 2.1 snap-ins, the admin will be presented with the following screen, asking whether the 1.5 ISOs should be upgraded:

Figure 122 iChain 1.5 Protected Resources Found

If the admin selects Yes, the ISO attributes will be converted to the 2.1 format where the ISO access mode will default to Secure (requiring users to authenticate before authorization to access the protected resource at the origin server is granted). If needed, modify the resources once the conversion is completed. The modify screen is shown here:

Figure 123 Modify Existing Protected Resource

If the admin selects No at the iChain 1.5 Protected Resources Found screen, it will leave the object as an iChain 1.5 protected resource and the resources on this ISO will continue to function as they did in iChain 1.5. For backward compatibility, you should make all changes to the protected resources with the iChain 1.5 snap-ins.

NOTE: If administrators are creating completely new objects, the following should be considered:

1. The ISO has many new attributes in 2.1. The most important of these involves the protected resource mode (public, secured, or protected). The defined protected resource needs a /* at the end for accessing resources in all subdirectories.

2. If you decide to recreate the ISO, the corresponding Rule Objects referencing the old ISO's protected resources must be recreated. If this is not done, ACLCHECK will report "old version" errors.

386 Place Book Title Here