


NetWare Enterprise Web Server

Administration Guide



NetWare[®] 5.1
NETWORKING SOFTWARE

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.

This product may require export authorization from the U.S. Department of Commerce prior to exporting from the U.S. or Canada.

Copyright © 1993-2000 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 Nos. 4,555,775; 5,157,663; 5,349,642; 5,455,932; 5,553,139; 5,553,143; 5,594,863; 5,608,903; 5,633,931; 5,652,854; 5,671,414; 5,677,851; 5,692,129; 5,758,069; 5,758,344; 5,761,499; 5,781,724; 5,781,733; 5,784,560; 5,787,439; 5,818,936; 5,828,882; 5,832,275; 5,832,483; 5,832,487; 5,859,978; 5,870,739; 5,873,079; 5,878,415; 5,884,304; 5,893,118; 5,903,650; 5,905,860; 5,913,025; 5,915,253; 5,925,108; 5,933,503; 5,933,826; 5,946,467; 5,956,718; 5,974,474. U.S. and Foreign Patents Pending.

Novell, Inc.
122 East 1700 South
Provo, UT 84606
U.S.A.

www.novell.com

NetWare Enterprise Web Server Administration Guide
January 2000
104-001219-001

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

Novell Trademarks

For a list of Novell trademarks, see the final appendix of this book.

Third-Party Trademarks

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

Contents

Preface	11
What's in This Documentation?	12
Novell Technical Support.	12
1 Managing Your Server	15
Using the General Administration Page	15
Accessing the General Administration Page	15
Remotely Accessing the General Administration Page.	16
Starting and Stopping a Web Server	16
Setting Up Multiple Web Servers	17
Using the General Administration Page	17
Using the Drop-Down List	18
Wildcards Used in the Drop-Down List	18
2 Managing Server Content	21
Setting the Primary Document Directory	21
Setting Additional Document Directories	22
Setting Virtual Document Directories	23
Establishing the Path to the Directory	23
Providing Public Access	23
Setting Server Access	24
Configuring User Document Directories	24
Creating the User with a User Home Directory	24
Creating a PUBLIC_HTML Directory	24
Adding Users' Contexts to the Search Contexts List	25
Restarting the NetWare Web Server	25
Activating User Document Directories	25
Providing Public Access	25
Configuring Document Preferences	26
Indexing Filenames	26
Directory Indexing	26
Server Home Page	27
Default MIME Type	27
Parsing the Accept Language Header	28
Setting Document Preferences	28
Forwarding URLs	28
Setting Up Hardware Virtual Servers	29
Securing a Hardware Virtual Server	30
Setting Up Software Virtual Servers	30

Assigning a Character Set	31
Specifying a Document Footer	33
Customizing Parsed HTML	34
Using Cache-Control Directives	35
Working with Configuration Styles.	36
Creating a Configuration Style	36
Editing a Configuration Style.	37
Applying a Configuration Style.	38
Removing a Configuration Style	38
Listing Configuration Style Assignments.	38

3 Configuring Server Preferences 41

Starting and Stopping the Server	41
Setting the Termination Timeout.	41
Restarting the NetWare Web Manager	42
Viewing Server Settings	42
Restoring Backup Configuration Files	43
Tuning Server Performance	43
Configuring Maximum Simultaneous Requests	44
Enabling Domain Name System Lookups	44
Configuring List Queue Size	45
Configuring the HTTP Persistent Connection Timeout	46
Configuring MIME Types	46
Configuring Network Settings	47
Changing the Server Name	47
Changing the Server Port Number.	48
Changing the Server Binding Address.	48
Changing the Server's MTA Host	48
Changing the Server's NNTP Host	48
Customizing Error Responses.	49
What Are the Errors?	49
Setting Up the Response	49
Restricting Access.	50
Setting Security	51

4 Controlling Access to Your Server. 53

Controlling Access with NDS	53
Controlling Access Using NetWare Web Access Controls	54
Using Access Control	54
User-Group Authentication	54
Host-IP Authentication	55
Access Control Files	56
How Does Access Control Work?	56
Restricting Access	58

Setting Access Control Actions	61
Specifying Users and Groups	61
Specifying Hostnames and IP Addresses	63
Setting Access Rights	64
Writing Customized Expressions	65
Changing Access Control.	65
Responding When Access Is Denied.	66
Examples.	66
Restricting Access to the Entire Server.	66
Restricting Access to a Directory	68
Restricting Access to a URI (Path)	69
Restricting Access to a File Type.	71
Restricting Access Based on Time of Day	72
5 Understanding ACL Files	75
ACL File Syntax	75
Authentication Statements	76
Authorization Statements.	77
Default ACL File	80
Referencing ACL Files in OBJ.CONF	81
6 Extending Your Server with Programs	83
Installing Server-Side Programs	84
Installing CGI Programs	84
Using the Query Handler	87
Installing Server-Side JavaScript Programs	88
Installing Client-Side Programs	94
Installing Client-Side JavaScript Programs	95
Installing Remote ODBC/JDBC Database Connectivity	95
Installing Drivers	95
Installing the Listener	96
Checking the Listener Installation	96
Adding a Data Source Name (DSN)	96
Using the ODBC Data Sources	96
Enabling the Novell Servlet Gateway	97
Registering a Remote UCS	98
7 Monitoring the Server	99
Working with Log Files	99
Viewing an Access Log File	99
Viewing the Error Log File	101
Setting Log Preferences	102
Archiving Log Files	104
Monitoring the Server Using HTTP.	105

Working with the Log Analyzer	106
Running the Log Analyzer from the Server Status Form	106
Running the Log Analyzer from the Command Line	108
Monitoring the Server Using SNMP	109
How Does SNMP Work?	110
The Enterprise Web Server MIB	111

8 Using Search 115

Configuring Text Search.	115
Controlling Search Access.	116
Mapping URLs.	116
Turning Search On or Off	117
Configuring the Search Parameters	118
Configuring Your Pattern Files.	120
Configuring Manually	121
Indexing Your Documents.	123
Collections	123
Collection Attributes	124
Creating a New Collection	126
Configuring an Existing Collection	127
Updating an Existing Collection	129
Maintaining an Existing Collection	130
Scheduling Regular Maintenance	131
Unscheduler Collection Maintenance.	131
Using the Enterprise Web Server to Simplify Your Search	132
Search Home Page	132
Standard Search Queries	132
Guided Search.	133
Advanced Search Query.	134
Search Results.	135
Displaying Collection Contents	136
Customizing the Search Interface	137
HTML Pattern Files	137
Search Function Syntax	139
Using Pattern Variables	142

9 Using Agents 149

Types of Agents	150
Document Agents	150
Directory Agents	150
Timer Agents.	150
Search Agents	151
Creating Authorized Users	151
Configuring Agent Services	151

Agent Information in the Configuration Files	153
Recovering Agent Files	154
How Agent Information Is Stored	154
Fixing Inconsistencies and File Corruption	154
Accessing Agent Services	157
Standard and Advanced Options.	157
Creating Agents	157
Creating Document Agents.	158
Creating Directory Agents	161
Creating Timer Agents	165
Creating Search Agents	168
Viewing and Managing Agents.	173
Modifying an Agent.	173
Deleting an Agent.	174
Disabling an Agent	174
Enabling a Disabled Agent	174
Managing Your Agents as a Group	175
10 Configuring Web Publishing	177
Setting Access Control for an Owner	178
Turning the Web Publishing On or Off	179
Turning WebDAV On or Off	179
Enabling or Disabling My Network and NDSDAV	180
Setting the Web Publishing Language	181
Maintaining Web Publishing Data	181
Changing the Link Management State.	183
Setting the Version Control Archive	184
Unlocking Files.	184
Adding Custom Properties	185
Managing Custom Properties	187
Indexing and Updating Properties	188
A Novell Trademarks.	191

Preface

The NetWare[®] Enterprise Web Server allows you to configure your NetWare server as a Web server. As a Web server, your NetWare server can serve Web pages using the HTTP protocol to respond to requests from Web browser clients. The NetWare Enterprise Web Server is managed using any browser interface such as Netscape* Navigator* or Communicator* or Internet Explorer. (To use the Web Folder feature, you must be running Internet Explorer 5.0 or higher.) Requests can come from an intranet (network behind a firewall) or from the World Wide Web.

The NetWare Enterprise Web Server allows administrators to

- ♦ Deploy departmental intranet servers using the existing NetWare network backbone
- ♦ Publish the existing data stored on the NetWare network for consumption by clients with a browser over the WAN
- ♦ Use the NetWare server as a WebDav/Web Folders server for use with the Microsoft* Office 2000 productivity applications without requiring client software
- ♦ Provide Web publishing services to users using Netscape Web Publishing
- ♦ Control how published documents are searched and updated
- ♦ Establish file access control using NDS and SSL-based security
- ♦ Provide users with the ability to quickly and easily create and manage their own home pages
- ♦ Establish security for users' systems
- ♦ Run server-side applications written using Java* servlets, CGI, Scripting, and Active Server Page technologies

- ♦ Deploy multiple Web server configurations on a single NetWare server using either hardware or software virtual servers

What's in This Documentation?

This documentation describes how to configure and use the NetWare Enterprise Web Server. Online help is available in the Enterprise Web Server interface. Documentation for all of NetWare appears on the *Documentation CD* included with your product, and on the Novell Documentation Web site (<http://www.novell.com/documentation>).

Novell Technical Support

The Novell[®] Support Connection[®] provides access to Novell's networking expertise through the Novell Support Connection Web site, the Novell Support Connection CD, and support programs for customers and partners.

By using the Novell Support Connection Web site or CD, you can connect to the same networking knowledge used by Novell technical support engineers. In addition, the Web site provides an open Internet-based forum for users and partners to share technical support information and solutions. The forums are staffed by volunteer System Operators (SysOps) who are invited and sponsored by Novell to answer questions posted in the forums. The Web site also offers information on Advanced Technical Training videos, CBTs and conferences.

For additional support, we encourage users to contact a Novell partner. Users can locate qualified partners using the Novell Support Connection Web site. Searches are based on geographic location, product expertise, or both.

Visit Novell Support Connection at:

- ♦ Novell Support for the Americas (<http://www.support.novell.com>)
- ♦ Novell Support for Europe, Middle East, and Africa (<http://www.support.novell.de>)
- ♦ Novell Support for Asia Pacific (<http://www.support.novell.com.au>)

or call:

- ♦ Americas (English) 1-800-858-4000/801-861-4000
- ♦ Europe, Middle East, Africa (English) (49) 211 5632 744

- ♦ French (49) 211 5632 733
- ♦ German (49) 211 5632 777
- ♦ Asia Pacific (English) (61) 2 9925 3133

See the Novell Support Connection Web site for a complete list of languages and support telephone numbers.

To order the Novell Support Connection CD, call 1-800-377-4136 or 1-303-297-2725 or visit the Novell Support Connection Web site.

1

Managing Your Server

This chapter describes how to configure and manage your server using the NetWare® Web Manager General Administration page.

During installation, you specified a port number for the NetWare Web Manager. The NetWare Web Manager helps you manage your NetWare® Enterprise Web Server® from the NetWare General Administration page. Click the Enterprise Web Server *servername* button to view the collection of forms used to change options and control your Web server.

Using the General Administration Page

You configure your NetWare Web Manager and access the configuration forms for other NetWare servers (including the NetWare Enterprise Web Server) with the NetWare General Administration page. This page contains links to the server managers for the NetWare servers you have installed.

You can perform the following Web server tasks from the General Administration page

- ♦ Choose a server to configure
- ♦ Start and stop a Web server

In addition, you can perform tasks for the NetWare Web Manager. For more information, see [Configuring NetWare Web Manager](#).

Accessing the General Administration Page

To launch the server and navigate to the General Administration page:

- 1 At the server console, type **nsweb**. The nsweb command executes an NCF file that runs the server.

Once the NetWare Enterprise Web Server is running, you can use any browser that supports frames and JavaScript* and has access to the NetWare Web Server Manager to configure your servers.

- 2 Open a browser and highlight the URL

http://servername:port number/

- 3 Substitute *servername* with the name you gave your server during installation. Substitute *port number* with the number assigned during installation.
- 4 When prompted, enter the username and password you chose during installation.
- 5 Click OK.

NOTE: The default installation modifies the AUTOEXEC.NCF to load the Web server whenever NetWare is restarted.

To disable autoloading, remove NSWEB from AUTOEXEC.NCF. To load and unload the Web server, type NSWEB and NSWEBDN respectively, at the system console.

Remotely Accessing the General Administration Page

To remotely access the General Administration page:

- 1 Using a browser that supports frames and JavaScript, such as Netscape* Communicator*, type the URL for the NetWare Web Manager

http://servername.your_domain.domain:port_number/

Use the port number for the NetWare Web Manager that you specified during installation. This is not the port number for the Web server.

- 2 Type the NetWare Web Manager username and password you specified during the installation > click OK.
- 3 Select the name of the server you want to configure.

Starting and Stopping a Web Server

You can start and stop the servers listed in the General Administration page by clicking the On/Off icon located to the left of the server's name.

Setting Up Multiple Web Servers

There are two ways you can have multiple Web servers on your NetWare server

- ♦ Hardware virtual servers
- ♦ Software virtual servers

Each approach has its strengths and weaknesses; you should choose the one that's right for your situation. The following are advantages and disadvantages of hardware and software virtual servers:

Hardware virtual servers allow you to map multiple IP addresses to multiple document roots. For example, if you have two IP addresses, you could map the first IP address to one document root and the second IP address to a second document root. While hardware virtual servers take fewer system resources than multiple instances of the server, they must share the same configuration information. For example, if one hardware virtual server has enabled security features or Web Publishing, they all must have it enabled.

Software virtual servers give you the ability to map a single IP address to multiple server names. Each software virtual server can have its own home page, which allows you to host multiple Web sites from one IP address. However, in order for software virtual servers to work correctly, the users accessing the server must use client software that supports the HTTP host header. Like hardware virtual servers, software virtual servers all must have the same configuration.

For more information, see “Setting Up Hardware Virtual Servers” on page 29 and “Setting Up Software Virtual Servers” on page 30.

Using the General Administration Page

The General Administration page provides access to the collection of forms you use to change options and control your server. From the General Administration page, which lists all the servers installed on your system according to identifier, access the NetWare Enterprise Web Server by clicking the Enterprise Web Server *servername* button (located next to the On/Off icon).

When you change server information, you must save and apply your changes in order for your changes to take place. After you submit a form, click Save and Apply.

You can return to the General Administration page by clicking the Admin button in the upper-right corner of the NetWare Enterprise Web Server forms.

Use the server configuration buttons in the top frame to configure the server. After clicking a button, you'll see a list of items on the left; click one of these links. The corresponding form is displayed in the main frame. If you need more information about a form, click Help for context-sensitive help.

Using the Drop-Down List

Most of the forms configure the entire server. Some forms can configure either the entire server or files or directories that the server maintains. These forms have a drop-down list at the top. The drop-down list allows you to specify what resource to configure.

Select a resource from the drop-down list for configuration. Click Browse to browse your primary document directory. Click Options to choose other directories. Click Wildcard to configure files with a specific extension.

Wildcards Used in the Drop-Down List

In many parts of the server configuration, you specify wildcard patterns to represent one or more items to configure. Note that the wildcards for access control and text search may be different from those discussed in this section.

Wildcard patterns use special characters. If you want to use one of these characters without the special meaning, precede it with a backslash (\) character.

Table 1 Drop-Down Wildcard Patterns

Pattern	Use
*	Match zero or more characters.
?	Match exactly one occurrence of any character.
	An <i>or</i> expression. The substrings used with this operator can contain other special characters such as an asterisk (*) or a dollar sign (\$). The substrings must be enclosed in parentheses, for example, (a b c), but the parentheses cannot be nested.

Pattern	Use
\$	Match the end of the string. This is useful in <i>or</i> expressions.
[abc]	Match one occurrence of the characters a, b, or c. Within these expressions, the only character that needs to be treated as a special character is the right bracket (<code>]</code>); all others are not special.
[a-z]	Match one occurrence of a character between a and z.
[^az]	Match any character except a or z.
*~	This expression, followed by another expression, removes any pattern matching the second expression.

Table 2 Drop-Down List Wildcard Examples

Pattern	Result
*.netscape.com	Matches any string ending with the characters .netscape.com
(quark energy).netscape.com	Matches either quark.netscape.com or energy.netscape.com
198.93.9[23].???	Matches a numeric string starting with either 198.93.92 or 198.93.93 and ending with any 3 characters.
.	Matches any string with a period in it.
~netscape-	Matches any string except those starting with netscape-
*.netscape.com~quark.netscape.com	Matches any host from domain netscape.com except for the single host quark.netscape.com

Pattern	Result
*.netscape.com~(quark energy neutrino).netscape.com	Matches any host from domain netscape.com except for hosts quark.netscape.com, energy.netscape.com, and neutrino.netscape.com
.com~.netscape.com	Matches any host from domain com except for hosts from the subdomain netscape.com

2

Managing Server Content

You can use the NetWare® Enterprise Web Server forms to help manage your server's content. You can create HTML pages and other files such as graphics, text, sound, or video and then store those files on your server. When clients connect to your server, they can view your files provided they have access to them. This chapter describes how you can configure and manage your server's content.

Setting the Primary Document Directory

You probably don't want to make all the files on your file system available to remote clients. An easy way to restrict access is to keep all of your server's documents in a central location, known as the document root or primary document directory.

Another benefit of the document directory is that you can move your documents to a new directory (perhaps on a different disk) without changing any of your URLs, because the paths specified in the URLs are relative to the primary document directory.

For example, if your document directory is `SYS:NOVONYX\SUITESPOT\DOCS`, a request such as `http://www.mozilla.com/products/info.html` tells the server to look for the file in `SYS:NOVONYX\SUITESPOT\DOCS\PRODUCTS\INFO.HTML`.

If you change the document root (by moving all the files and subdirectories), you only have to change the document root that the server uses, instead of mapping all URLs to the new directory or telling clients to look in the new directory.

To set your server's primary document directory:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Content Management > Primary Document Directory.
- 2 In the Primary directory field, type the full pathname of the directory that you want to make the primary document directory.
- 3 Click OK.
- 4 Click Save and Apply.

Setting Additional Document Directories

Most of the time you keep all of your documents in the primary document directory. Sometime you may want to serve documents from a directory outside of your document root. You can do this by setting additional document directories. By serving from a directory outside of your document root, you can let someone manage a group of documents without giving them access to your primary document root.

To add additional document directories:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Content Management > Additional Document Directories.
- 2 In the URL Prefix field, type the URL prefix or keyword you want to use to represent the path.

For example, the URL prefix could be *docs*.

- 3 In the Map to Directory field, type the absolute path of the directory you want the URL prefix to map to.

The command syntax is *vol:\directory\subdirectory*.

For example, the path could be

`\NOVONYX\MARKETING\PUBDOCS\INDEX.HTML`

- 4 You can select a configuration style to apply to this directory's configuration.
- 5 Click OK.

IMPORTANT: When you update information, but don't save and apply changes, your information is retained so that you can view and edit it, even though the changes have not taken effect.

Setting Virtual Document Directories

A virtual document directory allows you to serve documents from directories that do not reside on the file server where your Enterprise Web Server is running but exist in the same tree.

IMPORTANT: In order to use virtual directories on 4.x servers you must install IPX on your NetWare 5.1 server. To install IPX™, use the INETCFG utility at the system console or check IPX during the installation of the Enterprise Web Server.

Establishing the Path to the Directory

To establish the path to the directory:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Content Management > Additional Document Directories.
- 2** In the URL Prefix field, enter a key word, for example *text*, to represent the path to the virtual directory.
- 3** In the Map to Directory field, specify the path to your documents in the following format:

servername/vol:/directory/subdirectory

Providing Public Access

To provide public access to the virtual directory in NDS mode and restart the server:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Restrict Access.
- 2** Click Insert File > enter the path (or any portion of the path you want to be public) in the format
servername/vol:/directory/subdirectory
- 3** Click OK > Save Changes.
- 4** Click On/Off under Server Preferences to restart the server.

Setting Server Access

To give the file server running the Enterprise Web Server access to the directory structure of the server where the index file resides, the Enterprise Web Server file server must be configured as a trustee of that directory. Use NetWare Administrator or ConsoleOne™ to set the rights.

Configuring User Document Directories

User Document Directories allow you to set up document directories or home directories for every configured user.

IMPORTANT: In order to use User Document Directories on 4.x servers you must install IPX on your NetWare 5.1 server. To install IPX, use the INETCFG utility at the system console or check IPX during the installation of the Enterprise Web Server.

For every user that you want to provide a home page for, complete the following tasks:

- ♦ Create the user with a User Home Directory
- ♦ Create a PUBLIC_HTML directory in the user's home directory and copy an INDEX.HTML file to it
- ♦ Add the user's context to the Search Contexts List
- ♦ Restart the NetWare Web Manager
- ♦ Activate User Document Directories in the Enterprise Web Server
- ♦ Make the PUBLIC_HTML directory public

See the following sections for details on completing the above tasks.

Creating the User with a User Home Directory

Using NetWare Administrator, create new users in their appropriate contexts. Click Create Home Directory in the lower portion of the form to create their user document directories.

Creating a PUBLIC_HTML Directory

Create PUBLIC_HTML directories in the users' home directories and copy INDEX.HTML files to them.

NOTE: You can change the name of the PUBLIC_HTML directory. Should you choose to change it, make sure all references to this directory name are consistent.

Adding Users' Contexts to the Search Contexts List

- 1 From the General Administration page, click Global Settings.
The Configure Directory Services form appears.
- 2 Click Insert Context and enter the information for each new context in the New NDS[®] context window. Use the following format:

`ou=yourdepartment.o=yourcompany`

This information is added to the Search Contexts List.

If this context is already set in your AUTOEXEC.NCF file (set Bindery Context=) you don't need to add it here.
- 3 Click Save Changes.

Restarting the NetWare Web Server

Restart the server at the system console. Use the command NSWEBDN to down the server and NSWEB to restart it.

Activating User Document Directories

This step activates your users' home directories so that when the URL is entered all that is required is a slash (/) followed by *~username* in order to reach a particular user's home page.

- 1 From the General Administration page click Enterprise Web Server *servername* > Content Management > User Document Directories.

Under the form banner a message indicates that the service isn't active.
- 2 To activate the service, click OK.

Providing Public Access

To provide public access to the home directors and restart the server:

- 1 From the General Administration page, click Enterprise Web Server *servername* > Server Preferences > Restrict Access.
- 2 Click Insert File > enter the path (or any portion of the path you want to be public) in the format,

`servername/vol:/directory/subdirectory`
- 3 Click OK > Save Changes.
- 4 Click On/Off under Server Preferences to restart the server.

Configuring Document Preferences

You can configure the following document preferences from the General Administration page by clicking Enterprise Web Server *servername* > Content Management > Document Preferences.

Indexing Filenames

If a document name is not specified in the URL, and the server finds a file with this name in a document directory, it assumes that file is the index file. The server automatically displays this file when no specific file is requested. The defaults are INDEX.HTML and HOME.HTML. If more than one name is specified, the server searches in the order in which the names appear in this field until one is found. For example, if your index filenames are INDEX.HTML and HOME.HTML, the server first searches for INDEX.HTML, and if it doesn't find it, the server then searches for HOME.HTML.

Directory Indexing

In your document directory, you'll probably have several subdirectories. For example, you might create a directory called PRODUCTS, another called CLIENTS, and so on. It's often helpful to let clients access an overview (or index) of these directories.

The server indexes directories using the following process:

- ◆ The server first searches the directory for an index file called INDEX.HTML or HOME.HTML, which is a file you create and maintain as an overview of the directory's contents. (Note that these defaults are configurable for the whole server, so your server's files may vary. For more information, see "Indexing Filenames" on page 26.) You can specify any file as an index file for a directory by naming it one of these default names, which means you can also use a CGI program as an index if CGI is activated.
- ◆ If an index file isn't found, the server generates an index file that lists all the files in the document root. The generated index has one of the following formats:
 - ◆ Fancy directory indexing is fairly detailed. It includes a graphic that represents the type of file, the date the file was last modified, and the file size.

- ♦ Simple directory indexing is less detailed, but takes less time to generate.
- ♦ You can also specify that no dynamic directory listing be generated if the server looks for index files and cannot find any. If the server does not find any index files, it will not create a directory listing to show the user and will return an error message.

Server Home Page

When users first access your server, they usually use an URL such as `http://www.mozilla.com/`. When the server receives a request for this document, it returns a document called a home page. Usually this file has general information about your server and links to other documents.

By default the server finds the index file specified in the Index Filenames field and uses that for the home page. However, you can also specify a file to use as the homepage by selecting the Homepage icon (by the location field) and entering the filename for the home page in the field.

Default MIME Type

When a document is sent to a client, the server includes a section that identifies the document's type, so the client can present the document in the correct way. However, sometimes the server can't determine the proper type for the document because the document's extension is not defined for the server. In those cases, a default value is sent. For information about maintaining your server's Multi-Purpose Internet Mail Extension (MIME) types, see "Configuring MIME Types" on page 46.

The default is usually Text/Plain, but you should set it to the type of file most commonly stored on your server. Some common MIME types include

text/plain	text/html
text/richtext	image/tiff
image/jpeg	image/gif
application/x-tar	application/postscript
application/x-gzip	audio/basic

Parsing the Accept Language Header

When clients contact a server using HTTP 1.1, they can send header information describing the languages they accept. You can configure your server to parse this language information.

For example, if you store documents in Japanese and English, you could choose to parse the Accept Language header. When clients that have Japanese as the Accept Language header contact the server, they receive the Japanese version of the page. When clients that have English as the Accept Language header contact the server, they receive the English version.

If you do not support multiple languages, you should not parse the Accept Language header.

Setting Document Preferences

To set document preferences:

- 1 Click Enterprise Web Server *servername* > Content Management > Document Preferences.
- 2 In the Index Filenames field, type a new index filename or add a file.
- 3 Select the kind of directory indexing you want.
- 4 Select whether you want users to see a specified home page or an index file when they access your server. If you choose the homepage option, in the Index File field, type the filename of the home page you want.
- 5 In the Default MIME Type field, type the default MIME type you want the server to return if a client accesses a file with an extension that has not been set up as a MIME type on your server.
- 6 Select whether to parse the accept language header or not.
- 7 Click OK > Save and Apply.

Forwarding URLs

Redirection is a method for the server to tell a user that a URL has changed, for example, if you have moved files to another directory or server. You can also use redirection to seamlessly send a person who requests a document on one server to a document on another server.

To map a URL to another server, you must first specify the prefix of the URL you want the server to redirect. Then, you need to choose which URL to

redirect to. You can redirect to a URL prefix if the directory on the new server is the same as in the mapped URL; you can also redirect to a fixed URL (hostname, directory, and filename).

To forward URLs:

- 1** Click Enterprise Web Server *servername* > Content Management > URL Forwarding.
- 2** In the URL Prefix field, type the URL prefix you want to redirect. For example, if the URL you want to map is `http://www.netscape.com/info/movies`, you would type `/info/movies` in the field.
- 3** Select whether you want to forward requests to a URL prefix or to a fixed URL.

If you forward to a URL prefix, the forwarding keeps the full pathname and substitutes one prefix for another. For example, if you forward `http://www.netscape.com/info/movies` to a prefix `mozilla.com`, the URL `http://www.netscape.com/info/movies` redirects to `http://mozilla.com/info/movies`.

However, if the directory structure on the new server is not the same as in the mapped URL, you could forward the URL to a fixed URL. For example, you could forward `http://www.netscape.com/info/movies` to `http://mozilla.com/new-files/info/movies`.

Sometimes you may want to redirect requests for all the documents in one subdirectory to a specific URL. For example, if you had to remove a directory because it was causing too much traffic or because the documents were no longer to be served for any reason, you could direct a request for any one of the documents to a page explaining why the documents were no longer available. For example, a prefix on `/info/movies` could be redirected to `http://www.netscape.com/explain.html`.

- 4** Click OK > Save and Apply.

Setting Up Hardware Virtual Servers

A hardware virtual server lets your server respond to multiple IP addresses without your having to install multiple servers. With hardware virtual servers, you map multiple IP addresses to multiple document roots. For example, if you have two IP addresses, you could map the first IP address to one document root and the second IP address to another document root.

To set up hardware virtual servers:

- 1** Load and bundle all IP addresses.
- 2** Execute the following command on the server console:
add secondary IPaddress <IP address>
- 3** Add the above command to the AUTOEXEC.NCF file after the load and bind statements or after INITSYS.NCF if INETCFG is being used to configure the server.
- 4** Click Enterprise Web Server *servername* > Content Management > Hardware Virtual Servers.
- 5** In the IP Address field, type the secondary IP address.
- 6** In the Document Root field, type the document root, for example, SYS:NOVONYX\SUITESPOT\DOCS.
- 7** Click OK > Save and Apply.
- 8** Repeat the previous steps for each hardware virtual server.

Securing a Hardware Virtual Server

To secure a hardware virtual server check the box marked Encryption. While the Enterprise Web Server doesn't have to be secured for a hardware virtual server to be secured, you do have to specify a Key Material Object (KMO) during installation to use Encryption. Once the KMO is created, use Server Preferences > Encryption On/Off to select a KMO.

For more information on security, refer to "Setting Security" on page 51

IMPORTANT: Once you have turned Encryption on you must use https:// to contact this server rather than http://.

Setting Up Software Virtual Servers

A software virtual server is a way to host several Web sites on one computer without needing to have more than one IP address on the computer. For example, you can set up your system so that both www.mozilla.com and www.netscape.com resolve to 192.3.4.5, then set up software virtual servers to handle both server names, for example, <http://www.mozilla.com/> and <http://www.netscape.com>. The server can respond differently to requests depending upon the URL, even though the server only has one IP address.

For example, an Internet service provider (ISP) installs a Web server and then wants to set up a software virtual server for each of its customers (for example, customers aaa, bbb, and ccc) so that each customer can have an individual domain name.

The ISP first configures the Domain Name System (DNS) to recognize that a customer's URL, `www.aaa.com`, resolves to the ISP's IP address. The ISP then creates a subdirectory for each company (aaa, bbb, and ccc) in the document root. These subdirectories contain the files for that company, including the home page, `aaa/HOME.HTML`. Next, the ISP sets up software virtual servers. The URL host would be `www.aaa.com` and the homepage would be `aaa/HOME.HTML`. The ISP would do this for all the companies.

Because software virtual servers use the HTTP host header to direct the user to the correct page, not all client software works with software virtual servers.

NOTE: Only client software (such as Netscape* Communicator*), which supports the HTTP host header, works. In the previous example, the ISP would set up the `INDEX.HTML` file in the document root to be an index page that links to all the virtual servers hosted by the system, so all users could access the home pages.

To set up a software virtual server:

- 1** Click Enterprise Web Server *servername* > Content Management > Software Virtual Servers.
- 2** Create a directory under the docs directory, for example, `SYS:NOVONYX\SUITESPOT\DOCS\TEST`.
- 3** In the URL Host field, type the URL host whose custom home page you want to set up, for example *test/*.
- 4** In the Homepage field, type the path to the home page you want to use for this virtual server, for example `INDEX.HTML`. If you type a full path, the server uses that specific document. If you type a partial path, the server interprets it as relative to your primary document directory.
- 5** Click OK > Save and Apply.
- 6** If you want to modify preferences on the default homepage, click Edit the Default Home Page at the top of the form.

Assigning a Character Set

The character set of a document is determined in part by the language it is written in. You can override the Netscape Navigator's default character set

setting for a document, a set of documents, or a directory by selecting a resource and entering a character set for that resource.

Netscape Communicator can use the MIME type charset parameter in HTTP to change its character set. If the server includes this parameter in its response, Netscape Communicator changes its character set accordingly. The following are some character set examples:

Content-Type: text/html; charset=iso-8859-1

Content-Type: text/html; charset=iso-2022-jp

The charset names recognized by Netscape Communicator are specified in RFC 1700 (except for the names that begin with x-). These charset names include the following:

us-ascii	iso-8859-1
iso-2022-jp	x-sjis
x-euc-jp	x-mac-roman

Additionally, the following aliases are recognized for us-ascii:

ansi_x3.4-1968	iso-ir-6
ansi_x3.4-1986	iso_646.irv:1991
ascii	iso646-us
us	ibm367
cp367	

The following aliases are recognized for iso_8859-1:

latin1	iso_8859-1
iso_8859-1:1987	iso-ir-100
ibm819	cp819

To change the character set:

- 1** Click Enterprise Web Server *servername* > Content Management > International Characters.
- 2** Click the Editing drop-down list > select the server resource for which you want to change the character set.
- 3** Click Browse to view the different server resources.
- 4** Click Wildcard to type the pattern you want to edit.
- 5** In the Character Set field, type one of the character sets mentioned in the previous paragraphs.
- 6** Click OK Save and Apply.

Specifying a Document Footer

You can specify a document footer, which can include the last-modified time, for all the documents in a certain section of your server without using server-parsed HTML. This footer works for all files except the output of CGI scripts or parsed HTML (.SHTML) files. If you need your document footer to appear on CGI-script output or parsed HTML files, enter your footer text into a separate file and add a line of code or include another server-side to append that file to the page's output.

To specify a document footer

- 1** Click Enterprise Web Server *servername* > Content Management > Document Footer.
- 2** Click the Editing drop-down list > select the resource to which you want to apply the document footer.
- 3** Click Browse to view the different server resources.
- 4** Click Wildcard to enter the pattern you want to edit.
- 5** In the For Files of Type field, type the kind of files that you want to include in the footer. The default is text/html.
- 6** Select the time format from the drop-down list.
or
Type a date in the Custom Date Format field.
- 7** In the Footer Text field, type the footer text.

The maximum number of characters for a document footer is 765. Type the string :LASTMOD: if you want to include the date the document was last modified.

8 Click OK > Save and Apply.

9 To change the footer text, click Deactivate Custom Trailer.

When you change the document footer for an HTML page, the last-modified date doesn't change.

Customizing Parsed HTML

HTML is normally sent to the client exactly as it exists on disk without any server intervention. However, the server can search HTML files for special commands (that is, parse the HTML) before sending documents. If you want the server to parse these files and insert request-specific information or files into documents, you must first enable HTML parsing.

To customize parsed HTML:

- 1** Click Enterprise Web Server *servername* > Content Management > Parse HTML.
- 2** Click the Editing drop-down list > select the server resource to edit.
- 3** Click Browse to view the different server resources.
- 4** Click Wildcard to type the pattern you want to edit.
- 5** Select whether or not you want to activate parsed HTML.

If you activate it, you need to choose whether to activate it with or without the exec tag. The Exec tag allows an HTML file to execute an arbitrary program on the server. You might not want to allow the Exec tag for security or performance reasons.

- 6** Select which files to parse.

The default choice is to parse only files with the extension .SHTML. In this case, all files you want to parse must have the .SHTML extension. You can have the server parse all of its HTML files. Choosing this option can slow your server's performance.

- 7** Click OK > Save and Apply.

Using Cache-Control Directives

Cache-control directives are a way for the NetWare Enterprise Web Server to control what information is cached by a proxy server. By using cache-control directives, you override the default caching of the proxy to protect sensitive information from being cached and perhaps retrieved later. For these directives to work, the proxy server must comply with HTTP 1.1.

For specific directories in your server, you can set the cache-control directives to one of the following levels:

- ♦ **Public:** The response is cacheable by any cache.
- ♦ **Private:** The response is only cacheable by a private (non-shared) cache.
- ♦ **No Cache:** The response must not be cached anywhere.
- ♦ **No Store:** The cache must not store the request or response anywhere in nonvolatile storage.
- ♦ **Must Revalidate:** The cache entry must be revalidated from the originating server.
- ♦ **Maximum Age (in seconds):** The client does not accept a response that has a greater age than the maximum age.

To set the cache-control directives:

- 1** Click Enterprise Web Server *servername* > Content Management > Cache Control Directives.
- 2** Click the Editing drop-down list > select the directory or directories for which you want to set cache-control directives.
- 3** Click Browse to view the different server resources.
- 4** Click Wildcard to type the pattern you want to edit.
- 5** Select the level of control you want to set.

The default is public.

- 6** Click OK.

For more information on HTTP 1.1, see the Hypertext Transfer Protocol (HTTP/1.1 specification [RFC 2068]) at (<http://www.ietf.org/html.charters/http-charter.html>).

Working with Configuration Styles

Configuration styles are an easy way to apply a set of options to specific files or directories that your server maintains. For example, you can create a configuration style that sets up access logging. When you apply that configuration style to the files and directories that you want to log, you don't have to individually configure access logging for all the files and directories.

Creating a Configuration Style

To create a configuration style:

- 1** Click Enterprise Web Server *servername* > Configuration Styles > New Style.
- 2** In the Style Name field, type the name you want to give the configuration style.
- 3** Click OK.
- 4** Click the Style drop-down list > select a configuration style to edit > click Edit This Style.
- 5** From the list of links available, click the category you want to configure for your style. You can configure the following information:
 - ♦ CGI File Type: Allows you to activate CGI as a file type. For more information about CGIs, see “Installing CGI Programs” on page 84.
 - ♦ Character Set: Allows you to change the character set for a resource. For more information about character sets, see “Assigning a Character Set” on page 31.
 - ♦ Default Query Handler: Allows you to set a default query handler for a server resource. For more information about query handling, see “Using the Query Handler” on page 87.
 - ♦ Document Footer: Allows you to add a document footer to a server resource. For more information about document footers, see “Specifying a Document Footer” on page 33.
 - ♦ Error Responses: Allows you to customize the error responses that clients see when they encounter an error from your server. For more information about error responses, see “Customizing Error Responses” on page 49.

- ♦ Log preferences: Allows you to set preferences for access logs. For more information about log preferences, see “Setting Log Preferences” on page 102.
 - ♦ Restrict Access: Allows you to restrict access to the entire server or parts of it. For more information about access control, see “Restricting Access” on page 58.
 - ♦ Server Parsed HTML: Allows you to specify whether the server parses files before they are sent to the client. For more information about using parsed HTML, see “Customizing Parsed HTML” on page 34.
- 6** Fill out the form that appears > click OK.
 - 7** Repeat Steps 5 and 6 to make any other changes to the configuration style.
 - 8** Click OK on the form you modified.
 - 9** Click OK on the Edit a Style form.
 - 10** Click Save and Apply.

Editing a Configuration Style

To edit a configuration style:

- 1** Click Enterprise Web Server *servername* > Configuration Styles > Edit Style.
- 2** Click the Style drop-down list > select a configuration style to edit.
- 3** Click Edit This Style.
- 4** From the list of links available, click the category you want to configure for your style. For more information on these categories, see “Creating a Configuration Style” on page 36.
- 5** Fill out the form that appears > click OK.
- 6** Repeat Steps 4 and 5 to make any other changes to the configuration style.
- 7** Click OK on the form you modified.
- 8** Click OK on the Edit a Style form.
- 9** Click Save and Apply.

Applying a Configuration Style

Once you've created a configuration style, you can apply it to files or directories in your server. You can specify either individual files and directories or wildcard patterns, such as *.GIF.

To apply a configuration style:

- 1 Click Enterprise Web Server *servername* > Configuration Styles > Assign Style.
- 2 In the URL Prefix Wildcard field, type the prefix of the URL to which you are applying this configuration style.

If you select a directory inside the document root, only type the path after the document root. If you type /* after the directory, you apply the configuration style to all of the directory's contents.
- 3 Click the Style drop-down list > select the configuration style you want to apply.
- 4 Click OK > Save and Apply.

Removing a Configuration Style

Before removing a configuration style, apply the None configuration style to any files or directories that had the configuration style applied to them. If you do not apply None before removing the configuration style, you must manually edit your OBJ.CONF file, search for the configuration style in the file, and replace it with None. If you don't do this search and replace, anyone who accesses the files or directories to which the deleted configuration style was applied will get a server configuration error message.

To remove a configuration style:

- 1 Click Enterprise Web Server *servername* > Configuration Styles > Remove Style.
- 2 Click the Remove drop-down list > select the configuration style you want to remove.
- 3 Click OK > Save and Apply.

Listing Configuration Style Assignments

After you have created configuration styles and applied them to files or directories, you can get a list of the configuration styles and where you applied them.

To list configuration style assignments

- 1** Click Enterprise Web Server *servername* > Configuration Styles > List Assignments.
- 2** To edit a configuration style assignment, click Edit Style Assignment next to the configuration style name.

3

Configuring Server Preferences

This chapter describes how to configure server preferences for your NetWare® Enterprise Web Server by using the Server Preferences configuration forms.

Starting and Stopping the Server

Once installed, the server runs constantly, listening for and accepting requests. If your server is running, you'll see the On icon and its green light (to the left of the Enterprise Web Server *servername* button) in the General Administration page. You can start and stop the server by clicking the icon. You can also start, restart, and stop the server from the Server Preference form.

To start or stop the server:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > On/Off.
- 2 Click Server On or Server Off.

If your server is on and you click Server On, the server will restart.

After you shut down the server, it may take a few seconds for the server to complete its shut-down process and for the status to change to Off.

If your machine crashes or is taken offline, the server stops and any requests it was servicing may be lost.

Setting the Termination Timeout

When you stop your server, the server stops accepting new connections. Then it waits for all outstanding connections to complete. The time the server waits before timing out is configurable in the MAGNUS.CONF file. By default it is

set to 3 seconds. You probably do not need to change this value. If you do need to change the value, add in `MAGNUS.CONF` `TerminateTimeout seconds`, where *seconds* represents the number of seconds you want the server to wait before timing out.

The advantage to configuring this value is that you can wait longer for connections to complete. However, because most servers have connections open from non-responsive clients, if you increase the time the server waits, you will almost always have to wait the full time before your server shuts down.

Restarting the NetWare Web Manager

To restart the server:

- 1 At a console command, type **`nvxwebdn`** to unload the NLM.
- 2 To restart the server, type **`nvxwebup`** to execute an NCF file that runs the server.

Viewing Server Settings

From Server Preferences, you can view your server's technical and content settings and see if your server is running. The technical settings come from `MAGNUS.CONF`, and the content settings come from `OBJ.CONF`. These files are located in the server root, in the directory `HTTP-servername CONFIG`. For more information about the `MAGNUS.CONF` and `OBJ.CONF` files, see the Novell Developer Kit Web site (<http://www.developer.novell.com/ndk/nscomp.htm>).

The following list explains the server's technical settings:

- ♦ **Server Root:** The directory where the server binaries are kept. You first specified this directory during installation.
- ♦ **Hostname:** The URL clients use as a hostname to access your server.
- ♦ **Port:** The port on your system that the server monitors for HTTP requests.
- ♦ **Error Log:** The name and path of the server's error log file.
- ♦ **MTA Host:** The name of the mail server (used by agents).
- ♦ **NNTP Host:** The name of the news server (used by agents).
- ♦ **DNS:** Indicates whether DNS is enabled or disabled.
- ♦ **Security:** Indicates whether SSL is enabled or disabled.

- ♦ **Asynch DNS:** Indicates whether asynchronous DNS is enabled or disabled.

The server's content settings depend on its configuration. Common server content settings include the server's document directory, its index filenames, name and location of its access log, and default MIME type.

Restoring Backup Configuration Files

You can view or restore a backup copy of your configuration files (HTTPS-SERVER_ID.ACL, MAGNUS.CONF, OBJ.CONF, WEBPUB.CONF, AGENT.CONF, MIME.TYPES, .ACL files, RDM.CONF, CSID.CONF, PROCESS.CONF, ROBOT.CONF, and FILTER.CONF).

To view or restore a backup copy of your configuration files:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Restore Configuration.
- 2** In the Set Number of Sets of Backups field, type the number of backups displayed on the form and click Change.
- 3** Click Restore if you want to restore a backup version.
To restore all files to their states at a particular time, click Restore to Date, which lists the specific time to which you want to restore.
- 4** Click OK > Save and Apply.
- 5** Click View next to the backup version you want to view.

Tuning Server Performance

You can configure the server's technical options, including the number of maximum simultaneous requests, listen-queue size, and DNS usage.

To get the number of simultaneous requests, the server counts the number of active requests, adding 1 to the number when a new request arrives and subtracting 1 when a request is finished. When a new request arrives, the server checks to see if it is already processing the maximum number of requests. If it has reached the limit, it defers processing new requests until the number of active requests drops below the maximum amount.

Configuring Maximum Simultaneous Requests

You can set the number of maximum simultaneous requests, which is the number of active requests allowed for the server at one time. If your site is processing many requests that take many seconds, you may need to increase the number of maximum simultaneous requests. However, for general Internet or intranet use, you probably will not need to change the default value (128 requests).

If you need to change the number of maximum simultaneous requests, set the number before starting the server.

To reset the number:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Performance Tuning.
- 2 In the Maximum Simultaneous Requests field, type the number of requests.
- 3 Click OK > Save and Apply.

Enabling Domain Name System Lookups

You can configure the server to use Domain Name System (DNS) lookups during normal operation. By default, DNS is not enabled; if you enable DNS, the server looks up the hostname for a system's IP address. Although DNS lookups can be useful for server administrators when looking at logs, they can affect performance. When the server receives a request from a client, the client's IP address is included in the request. If DNS is enabled the server must look up the hostname for the IP address of each client that makes a request.

IMPORTANT: If you turn off DNS lookups on your server, hostname restrictions won't work, and hostnames won't appear in your log files. Instead, you'll see the IP addresses.

You can also specify whether to cache the DNS entries. If you enable the DNS cache, the server can store hostname information after receiving it. In the future, if the server needs information about the client, the information is cached and available without further queries. You can specify the size of the DNS cache and an expiration time for DNS cache entries. The DNS cache can contain from 32 to 32768 entries; the default value is 1024 entries. Values for the time it takes for a cache entry to expire can range from 1 second to 1 year (specified in seconds); the default value is 1200 seconds (20 minutes).

To modify DNS settings:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Performance Tuning.
- 2 Select No or Yes to enable DNS.
- 3 Select No or Yes to enable Async DNS.
- 4 Select No or Yes to cache DNS entries.
- 5 If you cache DNS entries, type the number of entries that you want cached in the Size of DNS Cache field. In the Expire Entries field, type the number of seconds at which a cache entry will be deleted.
- 6 Click OK > Save and Apply.

Configuring List Queue Size

The listen queue size is a socket-level parameter that specifies the number of incoming connections the system will accept for that socket. The default setting is Incoming Connections.

Normally, you should not change the listen queue size. The default setting is sufficient in most cases.

If you manage a heavily used Web site, you should make sure your system's listen queue size is large enough to accommodate the listen queue size setting from the Server Preferences form. If you do change the listen queue size, make sure that your system supports the new size. The listen queue size set from the Server Preferences form changes the listen queue size requested by the server. If the server requests a listen queue size larger than the system's maximum listen queue size, the size defaults to the system's maximum.

IMPORTANT: Setting the listen queue size too high can degrade server performance. The listen queue size was designed to prevent the server from becoming overloaded with connections it cannot handle. If your server is overloaded and you increase the listen queue size, the server will only fall further behind.

To modify the listen queue size:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Performance Tuning.
- 2 In the Listen Queue Size field, type the listen queue size you want.
- 3 Click OK > Save and Apply.

Configuring the HTTP Persistent Connection Timeout

With HTTP 1.1, a connection can be set to be persistent (similar to Keep Alive in HTTP 1.0). However, even if a connection is persistent, it still needs to have a timeout setting, or it may consume system resources.

Normally, you should not change the persistent connection timeout. The default setting is sufficient in most cases.

To change the setting:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Performance Tuning.
- 2 In the HTTP Persistent Connection Timeout field, type a number in seconds.
- 3 Click OK > Save and Apply.

Configuring MIME Types

Multi-Purpose Internet Mail Extension (MIME) types control what types of multimedia files your e-mail system supports. You can also use MIME types to specify what file extensions belong to certain server file types, for example, to designate what files are CGI programs. For more information on using file extensions with programs, see “Installing CGI Programs” on page 84.

To add a new MIME type:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > MIME Types.

The Global MIME Types form shows all the MIME types listed in the FastTrack Server’s MIME.TYPES file.

- 2 Click the Category drop-down list > select the category.

Type is in the file or application type, Enc is the encoding used for compression, and Lang is the language encoding.

- 3 In the Content-Type field, type the context type that will appear in the HTTP header.

The receiving client uses the header string to determine how to handle the file. The standard strings are listed in RFC 1521.

- 4 In the File Suffix field, type the file suffix.

This is the file extension that maps to the MIME type. To specify more than one extension, separate the entries with a comma and do not include any spaces. Do not map one file extension to two MIME types.

5 Click New Type.

To edit a MIME type:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > MIME Types.
- 2** Click Edit next to the category you want to edit.
- 3** In the Content-Type field, type the context type.
- 4** In the File Suffix field, type the file suffix.
- 5** Click Change MIME Type.
- 6** Click Save and Apply.

IMPORTANT: Do not type spaces between the file suffixes when you add or edit a MIME type. If you put a space between them, you may receive an error or your server may not restart. If this happens, edit your MIME.TYPES file to delete the space. The MIME.TYPES file is in your server root in the HTTPS-*servername*/CONFIG directory. After you have edited the file, from Server Preferences, click Apply.

To remove a MIME type:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > MIME Types.
- 2** Click Remove next to the category you want to remove.
- 3** Click Save and Apply.

Configuring Network Settings

You can change your server's network settings by using the Server Preferences. Following is a brief introduction to each setting and instructions to modify your network settings.

Changing the Server Name

The server name is the full hostname for your server. When clients access your server, they use this name. The format for the server name is *servername.yourdomain.domain*.

For example, if your full domain name is novell.com, you could install a server with the name www.novell.com.

If your system administrator has set up a DNS alias for your server, use that alias.

Changing the Server Port Number

The server port number specifies the TCP port to which the server listens. The port number you choose can affect your users. If you use a nonstandard port, then anyone accessing your server must specify a server name and port number in the URL. For example, if you use port 8090, the users would specify the following:

```
http://www.novell.com:8090
```

The standard unsecure Web server port number is 80; the standard secure Web server port number is 443. Technically, the port number can be any port from 80 to 65,535.

Changing the Server Binding Address

At times you'll want the server to answer to two URLs. Your system must already be set up to listen to multiple IP addresses. For information on configuring multiple IP addresses, refer to "Setting Up Hardware Virtual Servers" on page 29.

Changing the Server's MTA Host

The server's Message Transfer Agent (MTA) host is the name of the Simple Mail Transfer Protocol (SMTP) mail server.

Changing the Server's NNTP Host

The Network News Transfer Protocol (NNRP) host is the name of the news server.

To modify your network settings:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Network Settings.
- 2** In the Server Name field, type the full hostname or DNS alias of your server.
- 3** In the Server Port field, type the port number of your server.

- 4** In the Bind to Address field, type the IP address that is associated with the specified hostname.
- 5** In the MTA Host field, type the name of your SMTP mail server.
You must enter a valid MTA if you want to use the agent e-mail function.
- 6** In the NNTP Host field, type the name of your server.
You must enter a valid NNTP host if you want to use agents with the capability to post to news.
- 7** Click OK > Save and Apply.

Customizing Error Responses

You can specify a custom error response that sends a detailed message to clients when they encounter errors from your server. You can specify a file to send or a CGI program to run. Instead of sending back the default file, you might want to send a custom error response. For example, if a client repeatedly tries to connect to a part of your server protected by access control, you might return an error file with information on obtaining an account.

What Are the Errors?

You can customize the response to the following kinds of errors:

- ♦ **Unauthorized:** Occurs when users without access permission try to access a document on the server that is protected by access control.
- ♦ **Forbidden:** Occurs when the server doesn't have file system permissions to read something, or if the server is not permitted to follow symbolic links.
- ♦ **Not Found:** Occurs when the server can't find a document or when it has been instructed to deny the existence of a document.
- ♦ **Server Error:** Occurs when the server is not configured properly or when a catastrophic error occurs, such as the system running out of memory or producing a core dump.

Setting Up the Response

Before you can set up the response, you need to write the HTML file to send or create the CGI program to run.

To set the response:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Error Responses.
- 2** Click the Editing drop-down list > select the server resource you want to configure.
- 3** Click Browse to choose a specific part of your server.
- 4** Click Options to browse files and directories on your server.
- 5** Click Back to return to the Custom Error Responses form.
- 6** Click Wildcard to type the wildcard pattern to edit.
- 7** Select the error response you want to customize.
- 8** In the appropriate field, type the absolute path name to the file or CGI script that you want to return for that error code.
- 9** Check the CGI box if the file is a CGI program that you want to run.
- 10** Repeat this process for each of the error responses you want to customize.
- 11** Click OK.

To remove a customization, return to the form and delete the filename from the field next to the error code.

Restricting Access

Use the Restrict Access form to configure several features.

When you use Public Directory Designation you're actually specifying what files and directories you want to allow public access to. The Public Directory Designation box lists directories and files that are currently public with associated prefixes. Examples of prefixes are those selected when you configure a Map to Directory for Additional or Virtual Document Directories or User URL prefixes selected when you configure User Document Directories. For information on Additional or Virtual Document Directories, refer to "Setting Additional Document Directories" on page 22 and "Setting Virtual Document Directories" on page 23. For information on User Document Directories, refer to "Configuring User Document Directories" on page 24.

The Password Redirection File allows you to create and display a file that would alert users that their password has expired and that they are using grace

logins. When the user accesses this URL this redirected file appears rather than INDEX.HTML.

When you are in NDS mode file access is determined by NDS rights granted to users. Rights Checking Mode allows you to do rights checking at a more granular level. If you have the system check rights at the file level system performance will be effected.

To make a file or directory public:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Restrict Access.
- 2** Click Insert File to insert a file or directory.
- 3** Click Save to save your changes.

To display a file that indicates to users that their password has expired:

- 1** Enter the path to the location where you have saved your password expiration notification file in the Password Expiration Redirection File field. The default is
\\NOVONYX\\SUITEPOT\\DOCS\\NDSDIRECT.HTML
- 2** Click Save to save your changes.

If you are using NDS, you can use Rights Checking Mode to determine at what level you want rights checked. Checking File will effect performance. Click Save to save your changes.

Setting Security

After completing installation you must configure security or usernames and passwords will be sent across the wire in clear text.

When you install the Novell Certificate Server (during the NetWare installation), a Key Material Object (KMO) was created by default. A KMO, also called a server certificate object, includes a server certificate and key pair files.

For related information on securing the NetWare Web Manager, refer to Configuring NetWare Web Manager.

For more information on installing and configuring the Novell Certificate Server, refer to the Novell Documentation Web site (<http://www.novell.com/>)

documentation) for the NetWare 5.1 Installation Guide and Novell Certificate Server Administration Guide.

If you created a KMO the following options are available.

To enable security on the Enterprise Web Server:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Encryption On/Off.

This option is only available if you have created a KMO.

- 2** Under Encryption, select On.

The Port Number field displays 443.

- 3** Click the KMO drop-down list > select the KMO you want to use for encryption.

- 4** Click OK.

To restart the server, click the Enterprise Web Server *servername* button on the General Administration page to Off, and then On again.

NOTE: Once you have turned security on the NetWare Web Manager or Enterprise Web Server, you must use `https://` in the URL to access them.

For information on securing a hardware virtual server, refer to “Securing a Hardware Virtual Server” on page 30.

4

Controlling Access to Your Server

You can control who accesses the files on your Web site. This chapter discusses the various methods you can use to determine who has access to specific files or directories on your Web site. If you want to control who can configure the Web server itself and who can access the server configuration files, refer to *Configuring NetWare Web Manager*.

NetWare® Enterprise Web servers use NDS® by default. You can also use either local or remote Lightweight Directory Access Protocol (LDAP). For more information on access control see, Chapter 5, “Understanding ACL Files,” on page 75.

With NDS, you manage access control through the NetWare file system trustees. The following section explains how to manage access control (ACL) with NDS.

Controlling Access with NDS

Novell® file system trustee assignments allow you to restrict access to files, but do not allow you to restrict access based on IP address or other criteria. If access must be restricted based on IP address or other parameters, you must either change modes to leverage LDAP or find an alternative method to restrict access, such as a firewall.

You manage Novell file system trustee assignments with NetWare administration utilities (such as NetWare Administrator or ConsoleOne™) and allow an administrator to restrict access to files based on client identity and a series of rights. See *Configuring Users and Groups* for more information on NetWare file system trustee assignments.

Controlling Access Using NetWare Web Access Controls

Using Access Control

Access control lets you determine who can access the server. You can use two attributes for controlling access:

- ♦ **User-Group:** Requires users to enter a username and password before accessing the server. Or the server can use client authentication by checking an LDAP directory for a security certificate before giving access to a file or set of files on your Web site.
- ♦ **Host-IP:** Requires the user to view your Web site from a specific computer, where the server recognizes the computer by either its hostname or its IP address.

User-Group Authentication

You can require users to authenticate themselves before getting access to your Web site. Authentication means that users verify their identity either by entering a username and password or by using a client certificate installed in their network browser, such as Netscape* Navigator* or Netscape Communicator*. The first method of getting the username and password is the Basic method, which can be done with or without encryption. The second method of using client certificates is the SSL method, which must be done with encryption on.

For more information on installing and configuring the Novell Certificate Server, refer to the Novell Documentation Web site (<http://www.novell.com/documentation>) for the NetWare 5.1 Installation Guide and Novell Certificate Server Administration Guide.

Username and Password

If you require users to enter a username and password to get access to your Web site, you store the list of users and groups in an LDAP database, which can be either a file stored on the Web server computer or an LDAP server on a remote computer, for example, NDS via LDAP or by using NDS directly.

When users attempt to access a file or directory that has User-Group authentication, the Web browser displays a dialog box asking the user to enter a username and password. The server can access the encrypted information or not, depending on whether encryption is turned on for your server.

After entering the username and password, the user either sees the requested file or directory listing or a message denying them access. (You can customize the access denied message that they see.)

IMPORTANT: If your server doesn't use encryption, the username and password that the end user types are sent unencrypted across the network. Someone could intercept the network packets and read the username and password being sent to the Web server. For this reason, User-Group authentication is most effective when combined with encryption, Host-IP authentication, or both.

Client Certificate Authentication

You can confirm users' identities with security certificates before giving the users access to your Web site. You can do this in the following two ways:

- ♦ The server can use the information in the certificate as proof of identity.
- ♦ The server can verify the certificate, provided the certificates are published in an LDAP directory.

When a request comes in and you have client authentication on, the server performs these actions:

1. The browser sends the certificate.
 2. The server checks whether the certificate is from a trusted Certificate Authority (CA).
 3. If the certificate isn't from a trusted CA, the server ends the transaction.
- ♦ If the certificate maps correctly, then the Web server follows the ACL rule specified for that user. The rule can deny or allow the request.

Host-IP Authentication

You can limit access to files and directories on your Web site by making them available only to people using specific computers. You specify hostnames or IP addresses for the computers that you want to allow or deny. You can use wildcard patterns to specify multiple computers or entire networks. If you want to use this feature, you must have DNS running in your network and your computer must be configured to use it.

End user access to a file or directory using Host-IP authentication appears seamless. Users can access the files and directories immediately without entering a username or password. If the computer doesn't have access, the user will get a message denying access. You can also customize this message.

It's possible for more than one person to have access to a particular computer. For this reason, Host-IP authentication is most effective when combined with User-Group authentication. If both methods of authentication are used, the end user will have to enter a username and password before getting access.

Access Control Files

When you use access control on your Web server, the settings are stored in a file with the extension .ACL. Access control files are stored in the directory *server_root/server_type/ACL* where *server_type* is the name of the server. For example, the NetWare Web Manager uses the directory ADMINACL. The main ACL filename is GENERATED-HTTPS-SERVER-ID.ACL. The temporary working file is called GENWORK-HTTP-SERVER-ID.ACL. If you use the Server Preference form to restrict access, you'll have these two files. However, if you want more complex restrictions, you can create multiple files and reference them from the MAGNUS.CONF file. There are also a few features available only by editing the files. For example, you can restrict access to the server depending on the time of day or day of the week.

You also manually create and EDIT.ACL files if you want to customize access control. For example, you might want to use an Oracle* or Informix* database of users instead of an LDAP database. To do this type of customizing, you need to use the access control API to program a hook into the server's access control structure. This API is written in C. For more information on the API, see Netscape DevEdge Online* site (<http://developer.netscape.com>).

How Does Access Control Work?

You can control access to the entire server or to parts of the server (directories, files, file types). When the server evaluates an incoming request, it determines access based on a hierarchy of rules called access control entries (ACEs), and then it uses the matching entries to determine if the request is allowed or denied. Each ACE specifies whether or not the server should continue to the next ACE in the hierarchy. The collection of ACEs is called an access control list (ACL).

When a request comes in to the server, the server looks in OBJ.CONF for a reference to an ACL, which is then used to determine access. By default, the server has one ACL file that contains multiple ACLs.

For example, suppose someone requests the following URL:

http://www.novell.com/my_stuff/web/presentation.html

The server would first check access control for the entire server. If the ACL for the entire server was set to continue, the server checks to see if there is an ACL for the file type (*.HTML). Then it checks for an ACL for the directory MY_STUFF. If one exists, it checks the ACE and then moves on to the next directory. The server continues traversing the path either until it reaches an ACL that says not to continue or until it reaches the final ACL for the requested URL (in this case, the file PRESENTATION.HTML).

To set up access control for this example using the Server Preferences forms, you could create an ACL for the file only or for each resource leading to the file, for example, one for the entire server, one for the MY_STUFF directory, one for the MY_STUFF/WEB directory, and one for the file.

The following sample ACL file illustrates one way to control access to this resource.

```
# File automatically written
#
# You may edit this file by hand
#
version 3.0;

# This ACL denies all access to the my_stuff directory
acl "path=C:\Netscape\SuiteSpot\docs\my_stuff";
deny (all)
    user = "anyone";

# This ACL allows access to anyone in the user database
acl "path=C:\Netscape\SuiteSpot\docs\my_stuff\web";
allow (all)
    user = "anyone";

# This ACL allows everyone in the local database or LDAP
directoryacl "agents";authenticate (user,group) {prompt =
"EnterpriseFastTrack Server";};deny (all)    user =
"anyone";allow absolute (all)    user = "all";# This ACL
allows access to the file to anyone in the "my_group" group
acl
"path=C:\Netscape\SuiteSpot\docs\my_stuff\web\presentatio
n.html";
allow (all)
    user = "anyone";
    group = "my_group"

# This is the default ACL and denies access to anyone
acl "default";
deny (all)
```

Restricting Access

This section takes you through the process of restricting user access to documents on your Web site. The sections following this one describe, in detail, each option available when using access control. Keep in mind that most access control rules use only a subset of the available options.

There is also a section of examples on restricting different resources. You can review these examples in “Restricting Access to the Entire Server” on page 66.

To create an access control rule:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Restrict Access.

A form appears where you select and edit an existing access control rule or specify a new rule by either choosing the resource you want to apply to the rule (the file, directory, or wildcard pattern you want to control) or typing a name to assign to the ACL.

There are three sections to this main form:

- ◆ **Pick a Resource:** Allows you to specify a wildcard pattern for files or directories to restrict access to (such as *.HTML) or to specify a directory or a filename to restrict. You can also browse for a file or directory by clicking Browse.
 - ◆ **Pick an Existing ACL:** Allows you to select an ACL that you’ve created from the drop-down list.
 - ◆ **Type in the ACL Name:** Allows you to create named ACLs. Use this option only if you’re familiar with ACL files and the OBJ.CONF configuration file. You’ll need to manually edit OBJ.CONF if you want to apply named ACLs to resources.
- 2** In the section you want to modify, from the editing box select the part of your Web site (the resource) that you want to control. For example, you can select Entire Server to set up access control for your entire server.

The following items are some common examples of resources you might use for access control

Resource wildcard	What It Means
default	A named ACL created during installation that restricts write access, so only users in the local database or LDAP directory can publish documents, for example, by using the Web Publisher.
Entire Server	One set of rules determines the access to your entire Web site, including any virtual servers you have running. To restrict access to a virtual server, specify the path of its document root.
*.html	Controls access to all files with the .HTML extension.
*.cgi	Controls access to all files with the .CGI extension.
usr/ns-home/cgi-bin/*	Controls access to all files and directories in the CGI-BIN directory. Note that the path is absolute. On NT, the path must include the drive letter.
agents	A named ACL that restricts access to all agents. The Web server contains this ACL by default.
uri="/sales"	Controls access to the sales directory in the document root. To specify URIs, create a named ACL.

3 Click Edit ACL Access Control.

The screen divides into two frames that you use to set the access control rules. If the resource you chose already has access control, the rules will appear in the top frame.

The ACL form contains links that, when clicked, display another form in the bottom frame.

4 Click New Line.

This adds a default ACL rule to the bottom row of the table. You can use the up and down arrows in the left column to move the rule, if needed.

- 5** Click Deny to select the action you want to apply to the rule.

The bottom frame displays a form where you can select if you want to allow or deny access to the users, groups, or hosts you'll specify in the following steps. Select the option you want > click Update.

- 6** Click Anyone to specify user-group authentication listed under the Users/Groups column.

The bottom frame displays a form for configuring User-Group authentication. By default, there is no authentication, meaning anyone can access the resource.

Select the options you want > click Update. See "Specifying Users and Groups" on page 61 for detailed information on the options.

- 7** Click Anyplace to specify the computers you want to include in the rule.

The bottom frame displays the From Host form, where you can enter wildcard patterns of hostnames or IP addresses to allow or deny.

Select the options you want > click Update. See "Specifying Hostnames and IP Addresses" on page 63 for more detailed information on the options.

- 8** Click All to specify the access rights you want to include in the rule.

Check the access rights in the bottom frame > click Update.

- 9** Click X under the Extra column if you want to enter a customized ACL entry > click Update.

- 10** Check the appropriate box in the Continue column if you want the access control rule to continue in a chain.

This means the next line is evaluated before the server determines if the user is allowed access. When creating multiple lines in an access-control entry, it's best to work from the most general restrictions to the most specific ones.

- 11** Check to see if access control is on.

See "Changing Access Control" on page 65 for more detailed information.

- 12** Check Response When Denied if you want the user to be redirected to another URL if their request is denied.

13 Select Respond with the Following URL > type the URL in the field.

14 Click Update.

See “Responding When Access Is Denied” on page 66 for more information.

15 Repeat Steps 4 through 10 for each rule you need.

16 Click Submit to store the new access control rules in the ACL file.

If you click Revert, the server removes any changes you made to the rules from the time you first opened the two-frame window.

WARNING: Be cautious when using Revert because you can't restore your edits. In most cases, it's probably better to delete the rule lines individually.

17 Click Save and Apply.

The following sections describe the options that appear in the bottom frame of the access control window.

Setting Access Control Actions

You can specify the action the server takes when a request matches the access control rule.

- ♦ Allow: The users or computers can access the requested resource.
- ♦ Deny: The users or computers cannot access the requested resource.

The server goes through the list of ACEs to determine the access. For example, the first ACE is usually to deny everyone. If the first ACE is set to continue, the server checks the second ACE in the list. (If Continue is not checked, everyone would be denied access to the resource.) If the second entry matches, then the next ACE is used. The server continues down the list until it reaches either an ACE that doesn't match or that matches but is set to not continue. The last ACE that matches is used to determine if access is allowed or denied.

Specifying Users and Groups

You can restrict access to your Web site based on the user who requests a resource. With user and group authentication, users are prompted to enter a username and password before they can access the resource specified in the access control rule.

The Web server uses a list of users, who might be sorted into groups, to determine access rights for the user requesting a resource. The list of users

(and the groups they are included in) are stored either in a database on the Web server computer or in an LDAP server, such as Netscape Directory Server. You should make sure the database has users and groups in it before you set access control.

You can allow or deny access to everyone in the database, or you can allow or deny specific people by using wildcard patterns or lists of users or groups.

To configure access control with users and groups, follow the general directions for restricting access. When you click the Users/Groups column, a form appears in the bottom frame. The following list describes the options in the form.

- ◆ **Anyone (No Authentication):** Anyone can access the resource without having to enter a username or password. However, the user might be denied access based on other settings, such as hostname or IP address. This is the default setting.
- ◆ **Authenticated People Only:** All users requesting the resource will have to type a username and password before getting access. If the username they enter isn't in the database, the access control rule won't apply to them. However, if the rule says deny and then a group is listed, that group is denied, but everyone else in the database could be allowed depending on if there is another ACL that matches their request.
- ◆ **All in the Authentication Database:** Matches any user who has an entry in the database. To use this option, you must also select **Authenticated People Only**.
- ◆ **Only the Following People:** Allows you to specify certain users and groups to match. You can list the users and groups of users individually by separating the entries with commas. Or you can enter a wildcard pattern. To use this option, you must also select **Authenticated People Only**.
 - ◆ **Group:** Matches all users in the groups you specify.
 - ◆ **User:** Matches the individual users you specify.
- ◆ **Prompt for Authentication:** Allows you to specify message text that appears in the authentication window. You can use this text to describe what the user needs to enter. Netscape Navigator and Netscape Communicator cache the username and password and associate them with the prompt text. This means that if the user accesses areas (files and directories) of the server that have the same prompt, the user won't have to retype usernames and passwords. Conversely, if you want to force

users to reauthenticate for various areas, you simply need to change the prompt for the ACL on that resource.

- ♦ **Authentication Methods:** Specifies the method the server uses when getting authentication information from the client.
 - ♦ **Default:** Uses the default method you specify in the OBJ.CONF file, or Basic if there is no setting in OBJ.CONF. If you select Default in this form, the ACL rule doesn't specify a method in the ACL file. Default is the best choice because you can easily change the methods for all ACLs by editing one line in the OBJ.CONF file.
 - ♦ **Basic:** Uses the HTTP method to get authentication information from the client. The username and password are encrypted only if encryption is turned on for the server.
 - ♦ **SSL:** Uses the client certificate to authenticate the user. If you use this method, SSL must be turned on for the server. If you have encryption on, you can combine Basic and SSL methods.
 - ♦ **Other:** Uses a custom method you create using the access control API.
- ♦ **Authentication Database:** Allows you to select a database that the server uses to authenticate users. The default setting means the server looks for users and groups in either the local database or an LDAP directory, depending on the setting specified in the NetWare Web Manager. However, you can configure individual ACLs to use different databases. You can specify different databases and LDAP directories in the file *server_root/USERDB/DBSWITCH.CONF*, then you can choose the database you want to use in the ACL by selecting it in the drop-down list. If you use the access control API to use a custom database (for example, to use an Oracle or Informix database), you can type the name of the database in the Other field in the Users & Groups form.

Specifying Hostnames and IP Addresses

You can restrict access to your Web site based on which computer the request comes from. You specify this restriction by using wildcard patterns that match the computers' hostnames or IP addresses. For example, to allow or deny all computers in a specific domain, you would enter a wildcard pattern that matches all hosts from that domain, such as *.novell.com.

To specify users from hostnames or IP addresses, follow the general directions for restricting access. When you click anywhere, the From Host field appears in the bottom frame. Select the Only From option, then type either a wildcard

pattern or a comma-separated list of hostnames and IP addresses. Restricting by hostname is more flexible than by IP address. If a user's IP address changes, you won't have to update this list. Restricting by IP address, however, is more reliable-if a DNS lookup fails for a connected client, hostname restriction cannot be used.

The hostname and IP addresses should be specified with a wildcard pattern or a comma-separated list. The wildcard notations you can use are specialized; you can only use an asterisk (*). Also for the IP address, an asterisk must replace an entire byte in the address. For example, 198.95.251.* is acceptable, but 198.95.251.3* is not. When an asterisk appears within an IP address, it must be the farthest right character. For example, 198.* is acceptable, but 198.*.251.30 is not.

For hostnames an asterisk must also replace an entire component of the name. For example, *.netscape.com is acceptable, but *sers.netscape.com is not. When an asterisk appears in a hostname, it must be the farthest-left character. For example, *.netscape.com is acceptable, but users.*.com is not.

Setting Access Rights

You can set access rights to files and directories on your Web site. In addition to allowing or denying all access rights, you can specify a rule that allows or denies partial access rights. For example, you can give people read-only access rights to your files, so they can view the information, but not change the files. This is particularly useful when you use the Web Publisher feature to publish documents.

When you create an access control rule, the default access rights are set to all access rights. To change access rights, click the appropriate link in the Rights column in the top frame, then check or uncheck the access rights you want to set for a particular rule. The following list describes each access right you can check.

- ♦ Read Access: Lets users view a file. This access right includes the HTTP methods GET, HEAD, POST, and INDEX.
- ♦ Write Access: Lets users change or delete a file. This access right includes the HTTP methods PUT, DELETE, MKDIR, RMDIR, and MOVE.
- ♦ Execute Access: Applies to server-side applications, such as CGI programs, Java* applets, and agents.
- ♦ Delete Access: Lets users delete a file or directory.

- ♦ List Access: Lets users have directory information. That is, they can get a list of the files in that directory. This applies to Web Publisher and to directories that don't contain an INDEX.HTM file.
- ♦ Info Access: Lets users see headers (http_head method). This is mainly used by the Web Publisher.

Writing Customized Expressions

You can enter custom expressions for an ACL. There are a few features available only by editing the ACL file or creating custom expressions. For example, you can restrict access to your server depending on the time of day, day of the week, or both.

The following customized expression shows how you could restrict access by time of day and day of the week. This example assumes you have two groups in your LDAP directory: the regular group gets access Monday through Friday, 8:00 a.m. to 5:00 p.m. The critical group gets access all the time.

```
allow (read)
{
  (group=regular and dayofweek="mon,tue,wed,thu,fri");
  (group=regular and (timeofday>=0800 and timeofday<=1700));
  (group=critical)
}
```

For more information on valid syntax and ACL files, see the online help.

Changing Access Control

You can turn off access control for any part of the server that a user accesses. For example, you could create an ACL that restricts access to the resource *html, then you could have an ACL for the entire server that is turned off. In this case, the only time access control is used is when a user requests any file or directory in the .HTML extension.

When you uncheck the option, you'll get a prompt asking if you want to erase records in the ACL. When you click OK, the server deletes the ACL entry for that resource from the ACL file.

If you want to deactivate an ACL, you can comment out the ACL lines in the file GENERATED-HTTPS-SERVER-ID.ACL by putting pound (#) signs at the beginning of each line.

Responding When Access Is Denied

You can choose the response a user sees when denied access. You can vary the message for each access control object. By default, the user is sent a message that says the file wasn't found. The HTTP error code "404 Not Found" is also sent.

To change what message is sent for a particular ACL:

- 1** In the ACL form, click Response When Denied.
- 2** In the lower frame, select Respond with the Following URL.
- 3** In the text field, type a URL or URI to a text or HTML file in your server's document root that you want to send to users when they are denied access.

Make sure the file doesn't contain references to other files, such as style sheets or images, because they won't be sent.

- 4** Click Update.

Make sure any users who get the response file have access to that file. If you have access control on the response file and the user is denied access to both the original resource and the response file, the server will send the default Denied response.

- 5** Click Submit in the top frame to submit the access control rule.

Examples

This section describes some common examples for restricting access to a Web server and its contents. Some of these examples assume you set up the default ACL to deny anyone access to the server. You can also add a deny all line as the first rule to each of these examples, as done in the example for the entire server.

Restricting Access to the Entire Server

This example allows access to users in a group called Employees, who access the server from computers in a subdomain. There are no access control rules for other resources on the server. You might use this example if you have a server for a department and you only want users to access the server from computers in a specific subdomain of your network.

To restrict access to the entire server:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Restrict Access.
- 2** In the section called Pick a Resource, click the Editing drop-down list > select the entire server.

The resource you select must be selected.

- 3** Click Edit Access Control.
- 4** Click New Line.

The default rule appears, which denies all access to the server. Typically you should deny all access to your server, then allow specific access to users, groups, and computers. However, you might change this if you want to deny access only to a small group of users or groups.

- 5** Click New Line again to create a second rule.
- 6** Click Deny in the second rule.
- 7** In the bottom form that appears, select Allow > click Update.
- 8** Click Anyone in the second rule.
- 9** In the bottom form, type the group that you want to have access to the server.

For this example, type **Employees** in the Group field. Note that the two options called Authenticated People Only and Only the Following People are checked automatically.

- 10** Click Update.
- 11** Click Anyplace in the second rule.
- 12** In the bottom form, type a wildcard pattern for the hostnames of the computers you want to allow.
- 13** Click Update.
- 14** Uncheck the Continue box in the second rule of the top frame > click Submit.
- 15** Click Save and Apply.

Be sure to restart the server for the changes to take effect. The following text is the ACL file for this example:

```
# File automatically written
#
```

```
# You may edit this file by hand
#
version 3.0;
acl "default";
    deny (all)
        user= "anyone";
    allow absolute (all)
        user = "employee" and
        dns = "*.emp.mozilla.com";
```

Restricting Access to a Directory

This example lets users in a group called Executives have read access to a directory and its subdirectories and files on the server. The user called CEO has full permissions to the directory.

You might use this example if you have a directory on your server that one person owns (he or she publishes to this directory) and you want one group of users to read the files. For example, you might have a project owner who publishes status information for the project team to review.

To restrict access to a directory on the server:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Restrict Access.
- 2** In the section called Pick a Resource, click Browse.
- 3** In the form that appears, click the link for the directory you want to restrict.

The directories listed in this form are in the server's document root. Once you click a link, the Editing drop-down list displays the absolute path to the directory.

If you want to view all files in your server root, click Options on the Choose a Part of Your Server form > check the List Files As Well As Directories field > click OK.

- 4** Click Edit Access Control.
- 5** Click New Line twice to create two rules.

Don't edit the default values for the first rule. They deny all access to the directory. You'll edit the second rule to allow read access to the executives group.

- 6** Click Deny in the second rule.
- 7** In the bottom form that appears, select Allow, and then click Update.
- 8** Click Anyone in the second rule.
- 9** In the bottom form, type the group you want to have access to the server.
For this example, type **Executives** in the Group field.
- 10** Click Update.
- 11** Click All in the top frame.
- 12** Uncheck the Write and Delete access rights. This means the users in the executives group can't add or remove files, but they can view them and run any applications in the directories.
- 13** Click Update.
- 14** Click New Line to create a rule for the CEO user.
- 15** Select Allow for the third rule > click Anyone.
- 16** In the bottom form, type **CEO** in the User field > click Update.
- 17** Uncheck Continue for both the second and the third rules.
This means that the server ignores any ACLs for directories or files under the directory you specified in Step 2.
- 18** Click Submit > save and apply your changes.

The entry in the generated.https-serverid.acl file for this example looks like this:

```
acl "path=d:/netscape/suitespot/docs/senior-staff/";
  deny (all)
    user = "anyone";
  allow absolute (read,execute,list,info)
    group = "executives";
  allow absolute (all)
    user = "ceo";
```

Restricting Access to a URI (Path)

This example uses a URI to control access to a single user's content on the Web server. URIs are paths and files relative to the server's document root

directory. Using URIs is an easy way to manage your server's content if you frequently rename or move all or part of it, for example, for disk space. It's also a good way to handle access control if you have additional document roots.

This example gives anyone read access to files and directories in the path specified by the URI /MY_DIRECTORY. Only one user ("me," in this example) has full access to the directories and files.

You might use this example if you have several users who publish their content on your server. The users want to have write access to their content, and they want anyone to have read/execute access.

To restrict access to a URI:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Restrict Access.
- 2** In the section called Type in the ACL name, type the URI you want to control.

For example, type **uri=/my_directory**.
- 3** Click Edit Access Control.
- 4** Click New Line to create the first rule that allows all users read access.
- 5** Click Deny.
- 6** In the bottom form that appears, select Allow > click Update.
- 7** Click All.
- 8** Uncheck the Write and Delete access rights.

This means users can't add or remove files, but they can view them and run any applications in the directories.
- 9** Click Update.
- 10** Click New Line to create a rule for the owner of the directory.
- 11** Select Allow for the second rule.
- 12** Click Anyone.
- 13** In the bottom form, type **me** in the User field > click Update.
- 14** Uncheck Continue for both the first and second rules.

This means that the server ignores any ACLs for other URIs, directories, or files under the URI you specified in Step 2.

15 Click Submit > Save and Apply.

The entry in the generated.https-serverid.acl file for this example looks like this:

```
acl "uri=/my_directory";  
    allow absolute (read,execute,list,info)  
        user = "anyone";  
    allow absolute (all)  
        user = "me";
```

Restricting Access to a File Type

This example controls write and delete access to all files with the extension .CGI. You might use this example if you only want specific users to create programs that run on your server. In this example, anyone can run the programs, but only users in the Programmers group can create or delete them.

To restrict access to a file type:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Restrict Access.
- 2** In the section called Pick a Resource, click Wildcard.
- 3** In the prompt that appears, type ***.CGI** > click OK.

This wildcard pattern matches any request that contains a file or directory with the .CGI extension.

- 4** Click Edit Access Control.
- 5** Click New Line to create the first rule that will allow all users read access.
- 6** Click Deny.
- 7** In the bottom form that appears, select Allow > click Update.
- 8** Click All.
- 9** Uncheck the Write and Delete access rights.

This means users can't add or remove files or directories with the .CGI extension.

- 10** Click Update.
- 11** Click New Line to create a rule that allows write and delete access to the programmers group.
- 12** Select Allow for the second rule.

13 Click Anyone.

14 In the bottom form, type **Programmers** in the Group field > click Update.

15 Click Submit > Save and Apply.

In this example, both Continue boxes are checked. This means that if a file is requested, the server will first look at the ACL for the file type, then it will continue to look for another ACL that matches, for example, an ACL on the URI or the path. The server checks ACLs in the following order:

1. Pathcheck functions in OBJ.CONF: For example, these could be wildcard patterns for files or directories. The entry in the ACL file would appear as follows: `acl "*.CGI";`
2. URIs: For example, a path relative to the document root. The entry in the ACL file would appear as follows: `acl "uri=/my_directory";`
3. Pathnames: For example, an absolute path to a file or directory. The entry in the ACL file would appear as follows:
`acl "path=d:\netscape\suitespot\docroot1\sales/";`

The entry in the GENERATED.HTTP-SERVERID.ACL file for this example looks like this:

```
acl "*.cgi";
    allow (read,execute,list,info)
        user = "anyone";
    allow (all)
        group = "programmers";
```

Restricting Access Based on Time of Day

This example restricts write and delete access to the server during working hours. You might use this example if you don't want people publishing documents at times when people might be accessing the files. This example allows users to publish during the evening hours of the week (between 6:00 p.m. and 6:00 a.m., Monday through Friday) and all times during the weekend.

To restrict access based on time of the day and day of the week:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > Restrict Access.
- 2** In the section called Pick a Resource, select the entire server from the Editing drop-down list.

You can select any resource.

- 3** Click Edit Access Control.
- 4** Click New Line.
- 5** Click Deny.
- 6** In the bottom form that appears, select Allow > then click Update.
- 7** Click All.
- 8** Uncheck the Write and Delete access rights.

This means that if a user wants to add, update, or delete a file or directory, this rule won't apply and the server will search for another rule that matches.

- 9** Click Update.
- 10** Click New Line to create a rule that restricts the write and delete methods.
- 11** Select Allow for the second rule.
- 12** Click X to create a customized expression.

In the bottom field, edit the existing lines to include the following:

```
user = "anyone" and  
dayofweek = "sat,sun" or  
(timeofday >= 1800 and  
timeofday <= 600)
```

You might want to select the entire text element and copy it to memory. If there are errors, you'll have to reenter the text > click Update.

The top form will display unrecognized expressions in the Users/Groups and From Host columns, because you created a custom expression.

- 13** Click Submit.
If you made any errors in the custom expression, you'll get a JavaScript* alert. Correct any changes > click Submit again.
- 14** Click Save and Apply.
- 15** Restart your server for the changes to take effect.

5

Understanding ACL Files

This chapter describes the access control list (ACL) files and their syntax. ACL files are text files that contain lists that define who can access resources stored on your Web server. By default, the Web server uses one ACL file that contains all of the lists for access to your server. However, you can create multiple ACL files and reference them in the OBJ.CONF file.

You need to know the syntax and function of ACL files if you plan on customizing access control using the access control API. For example, you might use the access control API to interface with another database, such as an Oracle* or Informix* database. For more information on the API, see the Novell Developer Kit Web site (<http://www.developer.novell.com/ndk/nscomp.htm>).

With either local database or LDAP directories, you manage access control through the Netscape* access controls. With Novell® Directory Services (NDS), you manage access control through NetWare® file system trustees. For more information on managing access control with NDS, see “Controlling Access with NDS” on page 53.

ACL File Syntax

An ACL file is a text file containing one ACL or more. All ACL files must follow a specific format and syntax. All ACL files must begin with the version number they use. There can be only one version line and it can appear after any comment lines. For example:

```
version 3.0;
```

You can include comments in the file by beginning the comment line with the pound (#) sign.

Each ACL in the file begins with a statement that defines its type. ACLs can follow one of three types:

- ♦ Path ACLs: Specify an absolute path to the resource they affect.
- ♦ URI (Uniform Resource Indicator) ACLs: Specify a directory or file.
- ♦ Named ACLs: Specify a name that is referenced in resources in the OBJ.CONF file. The server comes with a default named resource that allows read access to anyone and write access to users in the local database or LDAP directory. Even though you can create a named ACL from the Server Preference forms, you must manually reference the named ACLs with resources in the OBJ.CONF file.

The type line begins with the letters `acl` and then includes the type information in double quotation marks followed by a semicolon. Each type information for all ACLs must be a unique name, even among different ACL files. The following lines are examples of several different types of ACLs:

```
acl "path=C:\Netscape\SuiteSpot\docs\mydocs\";  
acl "*.html";  
acl "default";  
acl "uri=/mydocs/";
```

After you define the type of ACL, you can have one or more statements that define the method used with the ACL (authentication statements) and the people and computers who are allowed or denied access (authorization statements). The following sections describe the syntax for these statements.

Authentication Statements

ACLs can optionally specify the authentication method the server must use when processing the ACL. There are two general methods:

- ♦ Basic: Requires users to enter a username and password before accessing a resource.
- ♦ SSL: Requires the user to have a client certificate. For this method to work, the Web server must have encryption turned on.

By default, the server uses the basic method for any ACL that doesn't specify a method. You can change the default setting by editing the following line in the MAGNUS.CONF file:

```
Init fn=acl-set-default-method method=SSL
```

Each authenticate line must specify what list (users, groups, or both) the server should use when authenticating users. The following authentication statement, which would appear after the ACL type line, specifies basic authentication with users matched to individual users in the database or directory:

```
authenticate (user) {  
    method = basic;  
};
```

The following example uses SSL as the authentication method for users and groups:

```
authenticate (user, group) {  
    method = ssl;  
};
```

Any allow or deny statements must match the lists you specify in the authenticate line. If the line says authenticate (user), the allow or deny line must also specify users. The following example allows any user whose username begins with the letters sales:

```
authenticate (user)  
    allow (all)  
        user = sales*
```

If the last line was changed to group = sales, then the ACL would fail because there are no groups in the user lists.

Authorization Statements

Each ACL entry can include one or more authorization statements. Authorization statements specify who is allowed or denied access to a server resource. Use the following syntax when writing authorization statements:

```
allow|deny [absolute] (right[,right...]) attribute qualifier  
expression;
```

Start each line with either allow or deny. It's usually a good idea to deny access to everyone in the first rule or command you enter then specifically allow access for users, groups, or computers in subsequent rules. This is because of the hierarchy of rules. If you allow anyone access to a directory called /MY_STUFF, then you have a subdirectory /MY_STUFF/PERSONAL that allows access to a few users, the access control on the subdirectory won't work because anyone allowed access to the /MY_STUFF directory will also be allowed access to the /MY_STUFF/PERSONAL directory. To prevent this, create a rule for the subdirectory that first denies access to anyone and then allows it for the few users who need access.

However, in some cases if you set the default ACL to deny access to everyone, then your other ACL rules don't need a Deny All rule.

The following line denies access to everyone:

```
deny (all)
    user = "anyone";
```

Hierarchy of Authorization Statements

ACLs have a hierarchy that depends on the resource. For example, if the server receives a request for the document (URI) /MY_STUFF/WEB/PRESENTATION.HTML, the server first looks for an ACL that matches the file type or any other wildcard pattern that matches the request, then it looks for one on the directory, and finally it looks for an ACL on the URI. If there is more than one ACL that matches, the server uses the last statement that matches. However, if you use an absolute statement, then the server stops looking for other matches and uses the ACL containing the absolute statement. If you have two absolute statements for the same resource, the server uses the first one in the file and stops looking for other resources that match.

For example, using the ACL hierarchy with the request for the document /MY_STUFF/WEB/PRESENTATION.HTML, you could have an absolute ACL that restricts access to the file type *.HTML, then the server would use that ACL instead of looking for one that matches the URI or the path.

```
version 3.0;
acl "default";
authenticate (user,group) {
    prompt="Enterprise Server";
};
allow (write,delete)
    user="all";
acl "*.html";
    deny absolute (all)
        user="anyone";
acl "uri=/my_stuff/web presentation.html";
    deny (all)
        user="anyone";
    allow (all)
        user="anyone";
```

Attribute Qualifier Expressions

Attribute qualifier expressions define who is allowed or denied access based on their username, group name, hostname, or IP address. The following lines are examples of allowing access to different people or computers:

- ♦ user = "anyone"
- ♦ user = "smith*"
- ♦ group = "sales"
- ♦ dns = "*.organization.com"
- ♦ dns = "*.organization.com" or "*.accounting_mail.com"
- ♦ ip = "198.*"

You can also restrict access to your server by time of day (based on the local time on the server) by using the `timeofday` attribute qualifier. For example, you can use the `timeofday` attribute qualifier to restrict access to certain users during specific hours.

Use a 24-hour clock to specify times (for example, use 0400 to specify 4 a.m. or 2230 for 10:30 p.m.).

The following example restricts access to a group of users called `guests` between 8 a.m. and 4:59 p.m.

```
allow (read)
    (group="guests") and
    (timeofday<800 or timeofday>=1700);
```

You can also restrict access by day of the week. Use the following three-letter abbreviations to specify days of the week: Sun, Mon, Tue, Wed Thu, Fri, and Sat.

The following statement allows access for users in the `premium` group any day and any time. Users in the `discount` group get access all day on weekends and on weekdays anytime except 8:00 a.m.-4:59 p.m.

```
allow (read) (group="discount" and dayofweek="Sat,Sun") or
    (group="discount" and (dayofweek="mon,tue,wed,thu,fri" and
    (timeofday<0800 or timeofday>=1700)))
or
    (group="premium");
```

Operators for Expressions

You can use various operators in attribute qualifier expressions. You can use parentheses to delineate the order of precedence of the operators. With user, group, dns, and ip qualifiers, you can use the following operators:

- ♦ and
- ♦ or
- ♦ not
- ♦ = (equals)
- ♦ != (not equal to)

With timeofday and dayofweek qualifiers, you can use the following additional operators:

- ♦ > (greater than)
- ♦ < (less than)
- ♦ >= (greater than or equal to)
- ♦ <= (less than or equal to)

Default ACL File

After installing the server, the server uses the default settings in the file `SERVER_ROOT/HTTPACL/GENERATED.HTTPS-SERVERID.ACL`.

There is also a file called `GENWORK.HTTPS-SERVERID.ACL` that is a working copy that the server uses until you save and apply your changes when working with the user interface. When editing the ACL file, you might want to work in the `GENWORK` file and then use the Server Preferences to save and apply the changes.

The following text is from the default file:

```
# File automatically written
#
# You may edit this file by hand
#

version 3.0;

acl "agents";
authenticate (user,group) {
    prompt = "Enterprise Server";
```



```

};
deny (all)
    user = "anyone"
allow absolute (all)
    user = "all";

acl "default";
allow (read,execute,list,info)
    user = "anyone";
allow (write,delete)
    user = "all";

```

The default ACL file is referenced in MAGNUS.CONF as follows:

```
ACLFile absolutepath/generated.https-serverid.acl
```

You can reference multiple ACL files in MAGNUS.CONF and then use their ACLs for resources in OBJ.CONF. However, the server uses only the first ACL file with the Web Publisher and with evaluation of access control for objects that don't have specific ACLs listed in OBJ.CONF. If you're using the Server Preference form to do some access control, the first ACL file in MAGNUS.CONF should point to the file GENERATED.HTTPS-SERVERID.ACL. See "Referencing ACL Files in OBJ.CONF" on page 81 for more information.

General Syntax Items

Input strings can contain the following characters:

- ♦ Letters a through z
- ♦ Numbers 0 through 9
- ♦ Period and underscore

If you use any other characters, you need to use double quotation marks around the characters.

A single statement can be placed on its own line and be terminated with a semicolon. Multiple statements are placed within braces. A list of items must be separated by commas and enclosed in double quotation marks.

Referencing ACL Files in OBJ.CONF

If you have named ACLs or separate ACL files, you can reference them in the OBJ.CONF file. You do this in the PathCheck directive using the check-acl function. The line has the following syntax:

```
PathCheck fn="check-acl" acl="aclname"
```

The aclname is a unique name of an ACL as it appears in any ACL file.

For example, you might add the following lines to your OBJ.CONF file if you want to restrict access to a directory using the ACL named testacl:

```
<Object ppath="/usr/ns-home/docs/test/*">  
  PathCheck fn="check-acl" acl="testacl"  
</Object>
```

In this example, the first line is the object that states which server resource you want to restrict access to. The second line is the PathCheck directive that uses the check-acl function to bind the name ACL (testacl) to the object in which the directive appears. The testacl ACL can appear in any ACL file referenced in MAGNUS.CONF.

6

Extending Your Server with Programs

In addition to serving HTML documents, your server can run programs that interact with clients. These applications that run on the server are called server-side applications. Client-side applications are downloaded to the client and run on the client machine.

Your server can run these types of server-side applications:

- ♦ Common Gateway Interface (CGI) programs
- ♦ JavaScript* applications
- ♦ Plug-in programs that use the server plug-in APIs, such as the Netscape* Server Plug-In (NSAPI)

This chapter describes how to install Java* applets, CGI programs, and JavaScript applications onto your server. Plug-ins extend or replace your server's features. For example, you can use plug-ins to provide a different way to control access or to login.

For information on writing and installing plug-ins, see the Novell Developer Kit Web site (<http://www.developer.novell.com/ndk/doc.htm>).

Additionally, your server can send server-side JavaScript programs to clients. This chapter deals mainly with the installation and configuration of server-side programs.

This chapter also describes the steps for specifying a default query handler CGI program. A query handler processes text sent to it via the ISINDEX tag in an HTML file.

Installing Server-Side Programs

JavaScript applications and CGI programs have different strengths and uses. CGI programs can be written in C, PERL, or other programming languages. All CGI programs have a standard way to pass information between clients and servers. JavaScript applications are written in JavaScript, an object-based scripting language that is easier to learn than an object-oriented programming language and lends itself to rapid application development.

Each type of program is installed onto the server differently. The following list summarizes the procedures:

- ♦ For CGI programs, configure your server to recognize certain files as CGI—all files with certain filename extensions or all files in specified directories.
- ♦ For JavaScript applications, check in each application individually through the Application Manager, which you can access from the Programs form or separately.

These installation instructions are described in the following sections.

Installing CGI Programs

Common Gateway Interface (CGI) programs can be created with any number of programming languages. On a Unix* machine, you're likely to find CGI programs written as Bourne shell or PERL scripts. On a Windows* computer, you might find CGI programs written in C++ or batch files. On NetWare®, you might find CGI programs written in NetBasic*, PERL, or LCGI NLM™ applications.

Regardless of the programming language, all CGI programs accept and return data in the same manner.

There are two ways to store CGI programs on your server:

1. Specify a directory that contains only CGI programs. All files are run as programs regardless of the file extensions.
2. Specify that CGI programs are all a certain file type. They will all use the file extensions .CGI, .EXE, .NLM, or .BAT. The programs can be located in any directory that the server can serve from.

There are benefits to either implementation. If you want only a specific set of users to be able to add CGI programs, keep the CGI programs in specified directories and restrict access to those directories. If you want to allow anyone

who can add HTML files to be able to add CGI programs, use the file type alternative. Users can keep their CGI files in the same directories as their HTML files.

If you choose the directory option, your server will attempt to interpret any file you place in that directory as a CGI program. Similarly, if you choose the file type option, your server will attempt to process any files with the file extensions .CGI, .EXE, .NLM, or .BAT as CGI programs. If a file has one of these extensions but is not a CGI program, an error occurs when a user attempts to access it.

Specifying a CGI Directory

To specify a CGI-only directory:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Programs > CGI Directory.
- 2** In the URL Prefix field, type the URL prefix you want to use for this directory.

The text you type appears as the directory for the CGI programs in URLs. For example, if you type **cgi-bin** as the URL prefix, then all URLs to these CGI programs have the following structure:

`http://yourserver.domain.com/cgi-bin/program-name`

The URL prefix you specify can be different from the real CGI directory you specify in Step 3 on page 85.

- 3** In the CGI Directory field, type the location of the directory as an absolute path.

This directory doesn't have to be under your document root. This is the reason that you need to specify a URL prefix in the previous step.

- 4** Click OK.
- 5** Click Save and Apply.

To edit an existing CGI directory:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Programs > CGI Directory.
- 2** Under Current CGI directories, click Edit next to the directory you want to edit.
- 3** In the URL prefix field, type the new prefix.

4 In the CGI directory field, type the new directory.

5 Click OK > Save and Apply.

To remove an existing CGI directory:

1 From the General Administration page, click the Enterprise Web Server *servername* button > Programs > CGI Directory.

2 Under Current CGI directories, click Remove next to the directory you want to remove.

3 Click OK > Save and Apply.

Copy your CGI programs into the directories you've specified.

Remember that any files in those directories will be processed as a CGI file, so you don't want to put HTML files in your CGI directory.

Specifying CGI as a File Type

To specify CGI programs as a file type:

1 From the General Administration page, click the Enterprise Web Server *servername* button > Programs > CGI File Type.

2 Click the Editing drop-down list > select the resource you want to apply this change to.

3 Click Browse to choose a part of your server.

4 Click Options to browse files and directories on your server.

5 Click Back to return to the CGI as a File Type form.

6 Click Wildcard to type the wildcard pattern to edit.

7 Select Yes to activate CGI as a file type.

8 Click OK > Save and Apply.

The CGI files must have the file extensions .BAT, .EXE, .NLM, or .CGI. Any non-CGI files with those extensions will be processed by your server as CGI files and will cause errors.

Downloading Executable Files

If you're using .EXE as a CGI file type, users will not be able to download .EXE files as executables.

One solution to this problem is to compress the executable files that you want users to be able to download, so that the extension is not .EXE. This solution has the added benefit of shortening the download time.

Another possible solution is to remove .EXE as a file extension from the MAGNUS-INTERNAL/CGI type and add it to the APPLICATION/OCTET-STREAM type (the MIME type for normal downloadable files). You can do this by clicking Enterprise Web Server *servername* > Server Preferences > MIME Types. However, the disadvantage to this method is that after making this change you cannot use .EXE files as CGI programs.

Another solution is to edit your server's OBJ.CONF file to set up a download directory, where any file in the directory is downloaded automatically. The rest of the server won't be affected. For directions on setting up this directory, see the technical note at

Netscape Technical Support (<http://help.netscape.com/kb/server/960513-130.html>)

CGI Scripting

Refer to the Novell Developer Kit Web site (<http://www.developer.novell.com/ndk/doc.htm>) for information on PERL, NetBasic, and LCGIs.

Using the Query Handler

You can specify a default query handler CGI program. A query handler processes text sent to it via the ISINDEX tag in an HTML file.

ISINDEX is similar to a form text field in that it creates a text field in the HTML page that can accept typed input. Unlike the information in a form text field, however, the information in the ISINDEX box is immediately submitted when the user presses Return. When you specify your default query handler, you tell your server the program to direct the input to. For an in-depth discussion of the ISINDEX tag, see an HTML reference manual.

To set a query handler:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Programs > Query Handler.
- 2** Click the Editing drop-down list > select the resource you want to set a default query handler to. If you choose a directory, the query handler you specify runs only when the server receives a URL for that directory or any file in that directory.

- 3** Click Browse to choose a part of your server.
- 4** Click Options to browse files and directories on your server.
- 5** Click Back to return to Query Handler form.
- 6** Click Wildcard to type the wildcard pattern to edit.
- 7** In the Default Query Handler field, type the full path for the CGI program you want to use as the default for the resource you chose.
- 8** Click OK > Save and Apply.

Installing Server-Side JavaScript Programs

To install server-side JavaScript programs, you need to activate server-side JavaScript for your server and use the Application Manager. This section includes information on accessing and using the Application Manager to install server-side JavaScript applications as well as to perform other functions.

For more information about writing JavaScript applications, see the Novell Developer Kit Web site (<http://www.developer.novell.com/ndk/doc.htm>).

You must activate server-side JavaScript before you can use the Application Manager. Also, put JSAC.EXE and LIBESNSPR20.DLL in your system directory so that they are in the search path. These files are found in the NOVONYX/SUITESPOT/BIN/HTTPS directory.

Activating Server-Side JavaScript

If you are using server-side JavaScript applications, you must first activate server-side JavaScript for your server.

To enable server-side JavaScript:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Programs > Server Side JavaScript.
- 2** Select Yes to require administration server password for Server Side JavaScript Application Manager.
- 3** Click OK > Save and Apply.
- 4** When the Activate Server Side JavaScript form appears click the link to use the Application Manager.

- 5 Enter the NetWare Web Manager username and password to use the Application Manager. For more information, see “Securing the Application Manager” on page 90.

For applications written in server-side JavaScript, you can perform many administrative tasks with the server-side JavaScript Application Manager. Using the Application Manager, you can do the following:

- ♦ Install a new JavaScript application; you must add an application before users can run it.
- ♦ Modify any of the attributes of an installed application, for example, its default home page, path to the .WEB file, and type of client-object maintenance.
- ♦ Stop, start, and restart an installed application.
- ♦ Run and debug an active application.
- ♦ Remove an installed application.

Running the Application Manager

To run the Application Manager, click the Enterprise Web Server *servername* button > Programs > Server Side JavaScript > click the link to the Application Manager. You can also run the Application Manager by loading the following URL in Netscape Navigator:

`http://server.domain/appmgr`

The Application Manager displays all applications currently installed on the server in a scrolling list in the left frame. Click an application in the scrolling list.

For the selected application, the right frame displays the following:

- ♦ Application name at the top of the frame.
- ♦ Path of the application .WEB file on the server.
- ♦ Default and initial pages for the application.
- ♦ Number of built-in maximum database connections allowed.
- ♦ External libraries used by the application (if any).
- ♦ Client object maintenance technique.
- ♦ Status of the application: Active or Stopped. Users can run only active applications.

To modify applications:

- 1** Click the drop-down list in the left column > select the application you want to modify.
- 2** Click the task button that indicates the action you want performed
 - ♦ **Start:** Activates the application. See “Starting, Stopping, and Restarting a Server-Side JavaScript Application” on page 93.
 - ♦ **Stop:** Stops the application. See “Starting, Stopping, and Restarting a Server-Side JavaScript Application” on page 93.
 - ♦ **Restart:** Restarts the application that was previously started and then stopped. See “Starting, Stopping, and Restarting a Server-Side JavaScript Application” on page 93.
 - ♦ **Run:** Retrieves application-Home form. See “Running a Server-Side JavaScript Application” on page 94.
 - ♦ **Debug:** Retrieves application-Home form.
 - ♦ **Modify:** Retrieves the specified application form. See “Modifying Installation Parameters” on page 93.
 - ♦ **Remove:** Removes the application. See “Removing a Server-Side JavaScript Application” on page 93.

The following explains the buttons on the green banner that runs across the top of the screen:

- ♦ **Configure:** Configures the default settings for Application Manager
- ♦ **Add Application:** Installs a new JavaScript application
- ♦ **Documentation:** Provides further documentation on server-side JavaScript
- ♦ **Help:** Provides instructions for using Application Manager

Securing the Application Manager

Your Application Manager runs on your Web server rather than on the Web Manager. The Application Manager is installed into the directory. You can access it without the Enterprise Web Server forms with this URL:

<http://yourserver.domain.com/appmgr>

Consequently, you may want to restrict access to the Application Manager URL and the application URI (Uniform Resource Identifier—an abbreviated

URL that is used for security) so that only you and trusted administrators can access them. If you don't restrict access to the Application Manager, anyone can add, remove, modify, start, and stop applications on your server.

If your server does not use the Secure Sockets Layer (SSL), the username and password for the Application Manager are transmitted unencrypted over the network. Any intruder who intercepts this data may be able to access the Application Manager. If you use the same password for your administration server, the intruder can also control your server. For security reasons, do not use the Application Manager from outside of your firewall (a network configuration that protects networked computers with an organization from outside access) unless you are using SSL.

Installing Server-Side JavaScript Applications

You must install (add) an application with the Application Manager before you can run it. You can install up to 120 JavaScript applications on one server.

To install a new application:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Programs > Server Side JavaScript.
- 2** Click the link to the Application Manager.
- 3** Click Add Application from the top of the page.
- 4** In the Name field, type the name of the JavaScript application. For specific information on application URLs, see "Application URLs" on page 92.
- 5** In the Web File Path field, type the absolute path to the .WEB file for the application.

This is a required field.

- 6** In the Default Page field, type the absolute path of the file to send to clients who do not indicate a specific page for the application.

This page is analogous to INDEX.HTML for a standard URL. This is a required field.

- 7** In the Initial Page field, type the absolute path of the page to run when the application is first started.

This page only runs once during the life of the application and is used to initialize values and establish database connections. This is an optional field.

8 In the Built-in Maximum Database Connections field, type the maximum number of database connections that this application can have at one time if you are using the built-in database object.

9 In the External Libraries field, type the absolute paths of any libraries to be used with the application.

This is an optional field. Libraries installed for one application can be used by all applications on the server.

10 In the Client Object Maintenance field, select the mode for maintaining the client object.

This can be Client-Cookie, Client-URL, Server-IP, Server-Cookie, or Server-URL.

11 Click OK > Save and Apply.

IMPORTANT: Do not give any JavaScript applications the same names as any subdirectories of your primary document directory. If you do, the server will no longer correctly process requests from the directory. For example, if you have a directory *server_root/DOCS/BUG* and a JavaScript application named Bug, all requests for any files in the BUG directory (or any of its subdirectories) will attempt to launch the JavaScript application Bug. The JavaScript application URI takes precedence.

Application URLs

When you install a server-side JavaScript application, you must enter a name for it. This name determines the application URL, the URL that clients use to access a JavaScript application. Application URLs are of the form *http://server.domain/appName/page.html*. *Server* is the name of the HTTP server, *domain* is the Internet domain (including the subdomains), *appName* is the application name you enter when you install it, and *page* is the name of the page in the application.

For example, if your server is named myserver, your domain name is mozilla.com, and the application is called Hello World, the application URL is *http://myserver.mozilla.com/world/hello.html*

This is a required field, and the name you type must be different from all other application names on the server. The name must include only alphanumeric characters and cannot include spaces.

IMPORTANT: Before you install an application, make sure the application name you choose does not usurp an existing URL on your server. All client requests for URLs that match the application URL are routed to the directory specified for the .WEB file, circumventing the server's normal document root.

Controlling Access to a Server-Side JavaScript Application

When you install an application, you may want to restrict its use to only certain users. You can do this by applying a configuration style to the application. For more information, see “Working with Configuration Styles” on page 36. For more information on restricting access to part of your server, see Chapter 4, “Controlling Access to Your Server,” on page 53.

Modifying Installation Parameters

You can change any of the parameters defined when you installed the application, except the application name. To change the name of an application, you must remove the application and then reinstall it.

If you modify the parameters of a stopped application, the Application Manager automatically starts it. When you modify parameters of an active application, Application Manager automatically stops and restarts it.

Removing a Server-Side JavaScript Application

Clicking Remove removes the application from the Application Manager, but does not delete files from the server. At this point, clients can no longer access the application.

If you delete an application, and you subsequently want to run it, you must install it again.

Starting, Stopping, and Restarting a Server-Side JavaScript Application

- ♦ Start: Starts an installed application that is stopped. If the application starts successfully, clients can run the application.
- ♦ Stop: Stops an active application. The application’s status changes to stopped, and clients can no longer run the application. You must stop an application if you want to move the .WEB file or update an application from a development server to a deployment server.
- ♦ Restart: Restarts a running application. For any changes you have made to take effect, you must restart an application after you compile it.

You can also start, stop, and restart an application by entering a special URL of the form:

```
http://server.domain/appmgr/  
control.html?name=appName&cmd=action
```

where *appName* is the application name and *action* is either stop, start, or restart.

Running a Server-Side JavaScript Application

There are two ways to run an installed application:

- ♦ Select the application name in the Application Manager > click Run. A new Navigator window accesses the application.
- ♦ Type the application URL in Navigator.

If you attempt to run a stopped application (one that is not active), then the Application Manager tries to start it first.

WARNING: The server should not be unloaded while a server-side JavaScript application is running because it can leave it in an unpredictable state.

Configuring Default Settings

When you install a new application, the default installation parameters are used for the initial settings.

You can specify the following default settings:

- ♦ Installation parameters of .WEB file path, default page, initial page, maximum number of built-in database connections, external libraries, and client object maintenance technique. You can specify a default directory path for your development area and native executables libraries.
- ♦ Prompts to confirm your action when you remove, start, stop, or restart an application.
- ♦ The application trace to appear, when debugging an application, in the same window as the application, but in another frame or in a window separate from the application.

Installing Client-Side Programs

Installing client-side programs in your server is relatively easy. There are two types of client-side programs: Java applets and JavaScript programs. Client-side Java applets are executable files identified in an HTML document, retrieved from the server, and executed on the client. The applets can reside anywhere under your server's primary document root. Client-side JavaScript programs are embedded in HTML files and executed on the client.

Installing Client-Side JavaScript Programs

Client-side JavaScript programs are created by lines of JavaScript code embedded in HTML files. The HTML files travel from the server to the client. Once the files reach the client, Navigator interprets the JavaScript code and performs the specified actions.

With LiveConnect you can connect server-side Java and JavaScript applications or client-side Java and JavaScript applications. For more information on LiveConnect, on embedding JavaScript in HTML, and on using client-side JavaScript with other programs, see the Novell Developer Kit Web site (<http://www.developer.novell.com/ndk/doc.htm>).

Installing Remote ODBC/JDBC Database Connectivity

The ODBC data sources store information about how to connect to an ODBC data provider. If you want to connect to a remote database running on Windows[®] NT[®], install the ODBC and JDBC drivers.

For more information on SQL Connector*, refer to the Novell Documentation Web site (<http://www.novell.com/documentation>).

In order to enable remote ODBC/JDBC database connectivity you must first activate the following services:

- ♦ Novell Servlet Gateway
- ♦ Server Side JavaScript

Once these services are enabled, install the ODBC and JDBC drivers to enable connections to a remote database running on Windows NT.

Installing Drivers

If you are running a database on Windows NT, install the ODBC and JDBC drivers.

To run the ODBC driver, run SETUP.EXE located in SYS:\PUBLIC\SQLC\CLIENT\ODBC\DISK1.

To run the JDBC driver, run SETUP.EXE located in SYS:\PUBLIC\SQLC\CLIENT\ODBC.

Installing the Listener

You must also install a Listener for remote ODBC database connectivity. To install the Listener, run SETUP.EXE, located in
SYS:\PUBLIC\SQLC\DRIVERS\NT\ODBC-OUT\DISK1.

Checking the Listener Installation

To check if the Listener was correctly installed go to the Windows NT Control Panel > Services. In the Services dialog check to make sure the SQLC-Server Listener Status is listed as Started.

Adding a Data Source Name (DSN)

The following instructions allow you to associate the ODBC/JDBC driver with the database file.

Go to the Windows NT Control Panel > ODBC Data Source. In the ODBC Data Source Administration click System DSN > Add DSN. Use the install wizard to complete the installation.

Using the ODBC Data Sources

For more information on data sources, refer to the SQL Connector, refer to the Novell Documentation Web site (<http://www.novell.com/documentation>).

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Programs > Novell Servlet Gateway.

- 2** Click Yes to Activate the Novell Servlet Gateway.

You must activate the Novell Servlet Gateway in order to use the Administer selection on the ODBC Data Sources form.

- 3** Click OK > Save and Apply.

- 4** Click Server Side JavaScript.

- 5** Click Yes to Activate the Server Side JavaScript Application Environment.

You must activate the Server Side JavaScript Application Environment in order to use the Test selection on the ODBC Data Sources form.

- 6** Click OK > Save and Apply.

- 7** Click ODBC Data Sources > Administer to launch the SQL Connector Manager.

The Administer Data Sources form appears.

- 8** To create a data source, click Select a Data Source for Local Oracle or Select a Data Source for Remote ODBC.

or

If you've already created a data source, click on it.

- 9** Use the buttons on this form to administer your data source.

Enabling the Novell Servlet Gateway

The Novell Servlet Gateway enables the NetWare Enterprise Web Server to execute Java servlets. A servlet can be thought of as a server-side applet without a user interface. The Novell Servlet Gateway provides Web application developers with additional functionality. For example, a servlet could be written and deployed to process data obtained from a client via an HTML form and the server side data processing could manipulate the data and store results in a database. Servlets provide an alternative to CGI.

Websphere* which contains its own servlet functionality, uses the `/servlet/` prefix by default. So, if you install both the NetWare Enterprise Web Server and Websphere you can use both the `/nwservlet/` and the `/servlet/` prefixes to execute servlets. In this case, Websphere will execute servlets using the `/servlet/` prefix and the Novell Servlet Gateway will execute those specified using `/nwservlet/`.

To enable the Novell Servlet Gateway:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Programs > Novell Servlet Gateway.
- 2** Click Yes to Activate the Novell Servlet Gateway.
- 3** Check Have the Servlet Gateway Handle `/servlet/` URLs box if:

Websphere is not installed and/or you want the Novell Servlet Gateway to handle requests via the `/servlet/` prefix as well as `/nwservlet/`.

Leave the box unchecked if:

Websphere is installed and you want it to handle servlets requested using `/servlet/`, or you do not want to use the `/servlet/` prefix.

- 4** Click OK > Save and Apply.

Directives are added to the OBJ.CONF file for handling servlet requests using the URL `/servlet/servletname`. These directives will always be added, even if Websphere is installed.

- 5** Click the link at the top of the form to the Novell Servlet Gateway Manager to configure the service.

Registering a Remote UCS

The Remote Universal Component System (UCS) form allows you to enter IP addresses or host names of Windows NT servers that can provide Active X components for your NetWare server. By entering this information entries are added to the UCS.INI file. These entries are used by the UCS remote NLM to connect your server with a Windows NT server that corresponds to an IP address or host name in the list.

Scripting languages that support UCS include JavaScript, Perl or Novell Script for NetWare.

To register a remote UCS:

- 1** From the Remote UCS form, enter an IP address or host name of a Window NT server that can provide Active X components in the Register a New Remote Slave field.

If you use a host name, it must be registered in DNS.

- 2** Click OK.

7

Monitoring the Server

You can monitor your server's activity using several different methods. You can view the server's status in real time—what is happening while you view it, compared to past performance—by using the Hypertext Transfer Protocol (HTTP) or the Simple Network Management Protocol (SNMP). You can also monitor your server by recording and viewing log files.

IMPORTANT: This is a departure from error reporting in previous versions of NetWare where the error reporting appeared on the system console.

Working with Log Files

Server log files record your server's activity. You can use these logs to monitor your server and to help you when troubleshooting. Both the error log file and the access log file are located in `/NOVONYX/SUITESPOT/HTTPS-SERVERNAME/LOGS`. The error log file lists all the errors the server has encountered, and the access log file records information about requests to the server and the responses from the server. You can use the Server Status form to specify what to include in the access log file. Use the log analyzer to generate server statistics. You can back up server error and access log files by archiving them.

Viewing an Access Log File

You can view the server's active and archived access log files from the Server Status form.

To view an access log:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Server Status > View Access Log.

- 2 Click the View This Log File drop-down list > select the access log file you want to see.

Active log files for resources and archived log files appear in the list.

- 3 To limit how much of the access log you see, type the number of lines you want to see in the Number of Entries field.

The order of the log entries on the screen is the order in which they were recorded in the log.

- 4 If want to filter the access log entries for a particular word, type the word in the Only Show Entries With field.

Case is important; make sure the case for your entry matches the case of the word you're searching for.

- 5 Click OK.

Here is an example of an access log in the Common Logfile Format:

```
wiley.a.com - - [16/Feb/1996:21:18:26 -0800] "GET / HTTP/1.0"
200 751
wiley.a.com - - [17/Feb/1996:1:04:38 -0800] "GET /docs/grafx/
icon.gif HTTP/1.0" 204 342
wiley.a.com - - [20/Feb/1996:4:36:53 -0800] "GET /help HTTP/
1.0" 401 571
arrow.a.com - john [29/Mar/1996:4:36:53 -0800] "GET /help
HTTP/1.0" 401 571
```

The following table describes the last line of the sample access log.

Table 3 Fields in the Last Line of the Sample Access Log File

Access Log Field	Example
Hostname or IP address of client	arrow.a.com (In this case, the hostname is shown because the Web server's setting for DNS lookups is enabled; if DNS lookups were disabled, the client's IP address would appear.)
RFC 931 information	- (RFC 931 identity not implemented)
Username	john (username entered by the client for authentication)
Date/time of request	29/Mar/1996:4:36:53 -0800

Access Log Field	Example
Request	GET/help
Protocol	HTTP/1.0
Status code	401
Bytes transferred	571

Here is an example of an access log using the flexible logging format:

```
wiley.a.com - - [25/Mar/1996:12:55:26 -0800] "GET /index.htm
HTTP/1.0" "GET" "/?-" "HTTP/ 1.0" 304 0 - Mozilla/2.0
(WinNT; I)
wiley.a.com - - [25/Mar/1996:12:55:26 -0800] "GET / HTTP/1.0"
"GET" "/?-" "HTTP/1.0" 304 0 - Mozilla/2.0 (WinNT; I)
wiley.a.com - - [25/Mar/1996:12:55:26 -0800] "GET / HTTP/1.0"
"GET" "/?-" "HTTP/1.0" 304 0 - Mozilla/2.0 (X11; I; IRIX
5.3 IP22)
```

The access log in the flexible logging format looks similar to the access log using the Common Logfile Format.

Viewing the Error Log File

The error log file contains errors the server has encountered after the log file was created; it also contains information about the server, such as when the server was started. Incorrect user authentication is also recorded in the error log. Use the error log to find broken URL paths or missing files.

To view the error log file:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Server Status > View Error Log.
- 2 If you want to see more or less than 25 lines of the error log, use the Number of Errors to View field to enter the number of lines you'd like to see.

The order of the log entries on the screen is the order in which they were recorded in the log.

- 3 If you'd like to filter the error messages for a particular word, type the word in the Only Show Entries With field.

Case is important; make sure the case for your entry matches the case of the word you're searching for.

4 Click OK.

Here is an example of an error log:

```
[13/Feb/1996:16:56:51] info: successful server startup  
[20/Mar/1996 19:08:52] warning: for host wiley.a.com trying  
to GET /report.html,      append-trailer reports: error  
opening  
[30/Mar/1996 15:05:43] security: for host arrow.a.com trying  
to GET /, basic-ncsa      reports: user jane password did not  
match database
```

In this example, the first line is an informational message—the server started successfully. The second log entry shows that the client wiley.a.com requested the file REPORT.HTML, but the file wasn't in the primary document directory on the server. The third log entry shows that the password entered for the user jane was incorrect.

Setting Log Preferences

You can customize access logging for any resource by specifying whether to log accesses, which format to use for logging, and whether the server should spend time looking up the domain names of clients when they access a resource.

Server access logs can be in Common Logfile Format, flexible log format, or your own customizable format. The Common Logfile Format is a commonly supported format that provides a fixed amount of information about the server. The flexible log format allows you to choose (from the Server Status form) what to log. A customizable format uses parameter blocks that you specify to control what gets logged. Once an access log for a resource has been created, you can't change its format unless you archive it or create a new access log file for the resource.

To set access logging preferences:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Server Status > Log Preferences.
- 2 Click the Editing drop-down list > select the resource to which you would like to apply custom logging.
- 3 Select whether to log client accesses.
- 4 Type the full path in the Log File field.

As a default, the log files are kept in the logs directory in the server root directory. If you specify a partial pathname, the server assumes the path is the logs directory in the server root.

- 5** In the Record field, select Domain Names or IP Addresses.
- 6** In the Format list, select Common Logfile Format, flexible log format (Only Log option), or Custom Format. If you select Only Log, you can check any or all of the following flexible log format items from the checklist:
 - ◆ Client Host Name: The hostname (or IP address if DNS is disabled) of the client requesting access.
 - ◆ Authenticate Username: The authenticated username is listed in the access log if authentication is necessary.
 - ◆ System Date: The date and time of the client request.
 - ◆ Full Request: The exact request the client made.
 - ◆ Status: The status code the server returned to the client.
 - ◆ Content Length: The content length, in bytes, of the document sent to the client.
 - ◆ HTTP Header, “Referer”: The referer specifies the page from which the client accessed the current page. For example, if a user was looking at the results from a text search query, the referer would be the page from which the user accessed the text search engine. Referers allow the server to create a list of backtracked links.
 - ◆ HTTP Header, “User-Agent”: The user-agent information—which includes the type of browser the client is using, its version, and the operating system it’s running on—comes from the User-Agent field in the HTTP header information the client sends to the server.
 - ◆ Method: The request method used.
 - ◆ URI: Universal Resource Identifier. The location of a resource on the server. For example, for `http://www.a.com:8080/special/docs`, the URI is `special/docs`.
 - ◆ Query String of the URI: Anything after the question mark in a URI. For example, for `http://www.a.com:8080/special/docs?find_this`, the query string of the URI is `find_this`.

- ♦ Protocol: The transport protocol and version used.
- 7** If you choose a custom format, type your custom format in the Custom format field. For more information about the parameters you should use, see Netscape DevEdge* Online at
(<http://developer.netscape.com/docs/manuals/doclist.html>)
- 8** If you don't want to log client access from certain hostnames or IP addresses, type the hostname or IP address in the Hostnames and IP Addresses fields. Type a wildcard pattern of hosts the server should ignore when recording accesses. For example, use *.netscape.com if you don't want to log accesses from people whose domain is netscape.com; you can type wildcard patterns for hostnames, IP addresses, or both.
- 9** Click OK.

Archiving Log Files

You can archive the access and error log files and have the server create new ones.

When you archive log files, the server renames the current log files and then creates new log files with the original names. You can back up or archive, or delete, the old log files, which are saved as the original filename followed by the date and time the file was rotated. For example, ACCESS might become ACCESS.24APR-04AM

You can archive log files immediately or have the server archive log files at a specific time on specific days. This information is stored in /NOVONYX/SUITESPOT/https-*servername*/LOGS

Before running the log analyzer, you should archive the server logs.

To archive log files:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Status > Archive Log.
The Archive Log Files form appears.
- 2** Click Archive if you want to rotate the log files immediately.
If you want archiving to occur at specific times on specific days, check the Rotate Log At button > select a time from the drop-down menu > check the days for archiving to occur.
- 3** Click OK.

Monitoring the Server Using HTTP

You can monitor your server's usage with the interactive server monitor. You can see how many requests your server is handling and how it is handling these requests. If the interactive server monitor reports that the server is handling a large number of requests, you might need to adjust the server configuration or the system's network kernel to accommodate the requests.

To monitor your server:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Status > Monitor Current Activity.
- 2** Click Monitor Server Activity on Port *port_number*.

The interactive server monitor reports the totals for the following server values on a new screen:

- ◆ Bytes Transferred: The number of bytes the server is transferring
- ◆ Total Requests: The number of requests the server is handling
- ◆ Bad Requests: The number of bad requests the server is handling
- ◆ 2xx: The number of status codes ranging from 200 to 299 that the server is handling
- ◆ 3xx: The number of status codes ranging from 300 to 399 that the server is handling
- ◆ 4xx: The number of status codes ranging from 400 to 499 that the server is handling
- ◆ 5xx: The number of status codes of 500 and higher that the server is handling
- ◆ xxx: The total number of 2xx, 3xx, 4xx, and 5xx status codes the server is handling minus timeouts and other errors that did not return an HTTP status code
- ◆ 200: The number of successful transactions the server is processing
- ◆ 302: The number of relocated URL status codes the server is processing
- ◆ 304: The number of requests for which the server tells the client to use a local copy of a URL instead of retrieving a newer version from the server
- ◆ 401: The number of unauthorized requests the server is handling
- ◆ 403: The number of forbidden URL status codes the server is handling.

Working with the Log Analyzer

Use the log analyzer to generate statistics about your server, such as a summary of activity, most commonly accessed URLs, times during the day when the server is accessed most frequently, and so on. You can run the log analyzer from the Server Status form or the command line.

Before running the log analyzer, you should archive the server logs. For more information about archiving server logs, see “Monitoring the Server Using SNMP” on page 109.

Running the Log Analyzer from the Server Status Form

To run the log analyzer:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Server Status > Generate Report.
- 2** Type the name of your server in the Server Name field.
This name appears in the generated report.
- 3** Select the output type—whether the report will appear in HTML or plain text format.
- 4** Select the Log File you want to analyze.
- 5** If you want to save the results in a file, type an output filename in the Output File field.

If you leave the field blank, the analyzer prints results on the screen. For large log files, you should save the results to a file because printing the output to the screen might take a long time.

- 6** Select whether to generate totals for certain server statistics. You can generate the following totals:
 - ♦ Total Hits: The total number of hits the server received after access logging was enabled.
 - ♦ 304 (Not Modified) Status Codes: The number of times the requesting client used a local copy of the requested document rather than retrieving it from the server.
 - ♦ 302 (Redirects) Status Codes: The number of times the server redirected to a new URL because the original URL moved.

- ♦ 404 (Not Found) Status Codes: The number of times the server couldn't find the requested document or the server didn't serve the document because the client was not an authorized user.
- ♦ 500 (Server Error) Status Codes: The number of times a server-related error occurred.
- ♦ Total Unique URLs: The number of unique URLs accessed after access logging was enabled.
- ♦ Total Unique Hosts: The number of unique client hosts who have accessed the server after access logging was enabled.
- ♦ Total Kilobytes Transferred: The number of kilobytes the server transferred after access logging was enabled.

7 Select whether to generate general statistics. You can generate the following general statistics:

- ♦ Top Number of One-Second Periods: You can generate the number of one-second periods during which requests were highest.
- ♦ Top Number of One-Minute Periods: You can generate the number of one-minute periods during which requests were highest.
- ♦ Top Number of One-Hour Periods: You can generate the number of one-hour periods during which requests were highest.
- ♦ Top Number of Users: You can generate the top number of users that accessed your server, provided that you included this as an item to log when you enabled access logging.
- ♦ Top Number of Referers: You can generate the number of referers that appear in your log analysis, provided that you included this as an item to log when you enabled access logging.
- ♦ Top Number of User Agents: You can generate the number of user agents that appear in your log analysis, provided that you included this as an item to log when you enabled access logging.
- ♦ Top Number of Miscellaneous Logged Items: You can generate the number of miscellaneous logged items that appear in your log analysis, provided that you included this as an item to log when you enabled access logging. These miscellaneous items include the request method, the URI, and the URI query.

To enable access logging, see “Setting Log Preferences” on page 102.

- 8 Select whether to generate a list of server access statistics. You can generate a list of the following:
 - ♦ Most Commonly Accessed URLs: You can have the log analyzer show the most commonly accessed URLs or URLs that were accessed more than a specified number of times.
 - ♦ Hosts Most Often Accessing Your Server: You can have the log analyzer show the hosts most often accessing your server or hosts that have accessed your server more than a specified number of times.
- 9 Type the order in which you want to see the results in the Output order field.
- 10 Click OK.

Running the Log Analyzer from the Command Line

To analyze access log files from the command line, run the flexanlg tool, which is in EXTRAS/FLEXANLG in your server root directory.

To run flexanlg, type the following command and options at the command prompt:

```
flexanlg [ -P ] [-n name] [-x] [-r] [-p order] [-i file]* [ -m metafile ]* [ o file] [ c opts] [-t opts] [-l opts]
```

The following describes the syntax. (You can get this information online by typing flexanlg -h at the command prompt.)

-P: proxy log format	Default: no
-n servername: The name of the server	
-x : Output in HTML	Default: no
-r : Resolve IP addresses to hostnames	Default: no
-p [c,t,l]: Output order (counts, time stats, lists)	
Default: ctl	
-i filename: Input log file(s)	Default: none
-o filename: Output log file	Default: stdout
-m filename: Meta file(s)	Default: none
-c [h,n,r,f,e,u,o,k,c,z]: Count these item(s) -	
Default: hnreuokc	
h: total hits	
n: 304 Not Modified status codes (Use Local Copy)	
r: 302 Found status codes (Redirects)	

```

f: 404 Not Found status codes (Document Not Found)
e: 500 Server Error status codes (Misconfiguration)
u: total unique URL's
o: total unique hosts
k: total kilobytes transferred
c: total kilobytes saved by caches
z: Do not count any items.
-t [sx,mx,hx, xx,z]: Find general stats -
Default:s5m5h24x10
s(number): Find top (number) seconds of log
m(number): Find top (number) minutes of log
h(number): Find top (number) hours of log
u(number): Find top (number) users of log
a(number): Find top (number) user agents of log
r(number): Find top (number) referers of log
x(number): Find top (number) for miscellaneous keywords
z: Do not find any general stats.
-l [cx,hx]: Make a list of - Default:
c+3h5
c(x,+x): Most commonly accessed URLs
(x: Only list x entries)
(+x: Only list if accessed more than x times)
h(x,+x): Hosts (or IP addresses) most often accessing
your server
(x: Only list x entries)
(+x: Only list if accessed more than x times)
z: Do not make any lists

```

Monitoring the Server Using SNMP

You can monitor your server in real time by using the Simple Network Management Protocol (SNMP). SNMP is a protocol used to exchange data about network activity. With SNMP, data travels between a managed device and a network management station (NMS) where users remotely manage the network.

A managed device is anything that runs SNMP (for example, hosts or routers). Your Novell[®] Enterprise Web Server is a managed device. An NMS is usually a powerful workstation with one or more network management applications installed. A network management application graphically shows information about managed devices (which device is up or down, which and how many error messages were received, and so on).

Every managed device contains an SNMP agent that gathers information regarding the network activity of the device. This agent is known as the

subagent. Each Netscape server (except the administration server) has a subagent.

Another SNMP agent exchanges information between the subagent and NMS. This agent is called the master agent. A master agent runs on the same host machine as the subagents to which it talks. You can have multiple subagents installed on a host machine. All of these subagents can communicate with the master agent.

Values for various variables that can be queried are kept on the managed device and reported to the NMS as necessary. Each variable is known as a managed object, which is anything the agent can access and send to the NMS. All managed objects are defined in a management information base (MIB), which is a database with a tree-like hierarchy.

How Does SNMP Work?

SNMP exchanges network information in the form of protocol data units (PDUs). PDUs contain information about various variables stored on the managed device. These variables, also known as managed objects, have values and titles that are reported to the NMS as necessary. Communication between an NMS and managed device can take place in one of two forms: NMS-initiated and managed-device-initiated.

NMS-Initiated Communication

NMS-initiated communication is the most common type of communication between an NMS and a managed device. In this type of communication, the NMS either requests information from the managed device or changes the value of a variable stored on the managed device.

The following steps make up an NMS-initiated SNMP session:

1. The NMS searches the server's MIB to determine which managed devices and objects need to be monitored.
2. The NMS sends a PDU to the managed device's subagent through the master agent. This PDU either requests information from the managed device or tells the subagent to change the values for variables stored on the managed device.
3. The subagent for the managed device receives the PDU from the master agent.
4. If the PDU from the NMS is a request for information about variables, the subagent gives information to the master agent and the master agent sends

it back to the NMS in the form of another PDU. The NMS then displays the information textually or graphically.

If the PDU from the NMS requests that the subagent set variable values, the subagent sets these values.

Managed-Device-Initiated Communication

This type of communication occurs when the managed device needs to inform the NMS of an event that has occurred. A managed device such as a terminal would initiate communication with an NMS to inform the NMS of a shutdown or startup. Communication initiated by a managed device is also known as a *trap*.

The following steps make up a managed-device-initiated SNMP session:

1. An event occurs on the managed device.
2. The subagent informs the master agent of the event.
3. The master agent sends a PDU to the NMS to inform the NMS of the event.
4. The NMS displays the information textually or graphically.

The Enterprise Web Server MIB

Each Enterprise Web Server has its own management information base (MIB). The Enterprise Web Server's MIB is a file called HTTP.MIB, which contains the definitions for various variables pertaining to network management for the Enterprise Web Server. These variables are known as managed objects. Using the Enterprise Web Server MIB and network management software, such as HP* OpenView*, you can monitor your Web server like all other devices on your network.

The Enterprise Web Server MIB has an object identifier of netscape 1 (`http OBJECT IDENTIFIER ::= { netscape 1 }`) and is located in the `server_root\PLUGINS\SNMP\MIBFILES\NETWARE` directory.

You can view administrative information about your Web server and monitor the server in real time using the Enterprise Web Server MIB. The following table lists and describes the managed objects stored in the HTTP.MIB.

Table 4 HTTP.MIB Managed Objects and Description

Managed Object	Description
httpEntityDescr	A description of the server (includes operating system information)
httpEntityId	The Enterprise subtree for vendors (for example, the MIB has an object identifier of 1.3.6.1.4.1.1450)
httpEntityProtocol	The HTTP version number
httpEntityVersion	The server software version number
httpEntityOrganization	The organization responsible for the server
httpEntityLocation	The full path for the server
httpEntityContact	The people responsible for the server and contact information
httpEntityAddress	The IP address of the machine the server is running on
httpEntityPort	The port number on which the server is listening
httpEntityName	The server's identifier name (for example, server2.a.com)
httpEntityType	The type of server
httpEntityMethods	The methods supported by the server (for example, GET, POST, PUT)
httpEntityMaxProcess	The maximum number of active processes on the server
httpEntityMinProcess	The minimum number of active processes on the server
httpEntityMaxThread	The maximum number of active threads on the server
httpEntityMinThread	The minimum number of active threads on the server
httpStatisticsPort	The portnumber on which this server is listening
httpStatisticsAddress	The IP address to which this server is bound
httpStatisticsStatus	The status of the server (up or down) The uptime of the server after it was started
httpStatisticsNum ProcessIdle	The number of idle threads

Managed Object	Description
httpStatisticsNumProcessProc	The number of threads that are processing requests
httpStatisticsNumProcessDns	The number of threads resolving hostnames
httpStatisticsRequests	The total number of requests received and generated
httpStatisticsRequestError	The total number of request errors detected
httpStatisticsInUnknowns	The total number of unknown messages received/generated
httpStatisticsInBytes	The total number of bytes received
httpStatisticsOutBytes	The total number of bytes sent by the server
httpStatisticsTimeOut	The total number of times the server timed out
httpStatisticsProcessNum	The number of running processes
httpStatisticsThreadNum	The number of threads running
httpStatisticsNumBytes	The total number of bytes sent by the server
httpStatisticsNum2xx	The number of 200-level status requests handled by the server
httpStatisticsNum3xx	The number of 300-level status requests handled by the server
httpStatisticsNum4xx	The number of 400-level status requests handled by the server
httpStatisticsNum5xx	The number of 500-level status requests handled by the server
httpStatisticsNum200	The number of 200 (Transfer OK) requests
httpStatisticsNum302	The number of 302 (Moved Temporarily) requests
httpStatisticsNum304	The number of 304 (Not Modified) requests
httpStatisticsNum401	The number of 401 (Unauthorized) requests
httpStatisticsNum403	The number of 403 (Forbidden) requests

8

Using Search

The NetWare[®] Enterprise Web Server search function provides you with the ability to search the contents and attributes of documents on your server. As the server administrator, you can create a customized text search interface that is tailored to your user community.

Server documents can be in a variety of formats, such as HTML, Microsoft* Word, Adobe* PDF, and WordPerfect*. The server converts many types of non-HTML documents into HTML as it indexes them so users can use your Web browser to view the documents that are found for their search.

Users can search through server documents for a specific word or attribute, obtaining a set of search results that list all documents that match the query.

As the server administrator, you can restrict which users and groups are authorized to use text search and which documents they can access, you can modify the configuration files that govern how text search operates, and you can customize the search query and results pages.

You need to identify the directory or directories of documents that you want prepared for searching and index the document information into a searchable database, called a collection. The next several sections discuss the details of configuring search and indexing collections.

Configuring Text Search

You can configure several aspects of the search function for your specific server, some of which are collection specific; others apply across all collections during a search. Collection-specific configuration affects how documents are indexed into a particular collection, so you must define the configurations before creating the collection. Other configuration actions can be defined at any time because they affect only the search.

Collection-specific configuration actions:

- ♦ Define URL mappings for the document directories that are to be indexed
- ♦ Define the pattern files to display for searches on a particular collection

Configurations that affect all collections:

- ♦ Establish access control for files and directories
- ♦ Define any words you want dropped from the search
- ♦ Define the search parameters
- ♦ Turn the search function off and on
- ♦ Restrict the amount of memory available for indexing operations

Controlling Search Access

The search function accesses the ACL database that is the default for your server. You can restrict access to the documents and directories on your server by defining explicit access control rules or you can rely on the default access control definitions. You can add users to your server's access control database through the Administration Server's Users & Groups function. See Chapter 4, "Controlling Access to Your Server," on page 53 for more information about setting access control.

You can set your server to check access permissions before displaying search results (through the Agents & Search > Search Configuration form discussed in "Configuring the Search Parameters" on page 118). When this is set, before returning the results of a search query, the server checks a user's access privileges and asks users to identify themselves before it displays any results.

Mapping URLs

When users search through a collection's files, the documents that are returned as search results use a partial URL, called a Uniform Resource Identifier (URI), to identify them. This is a security feature that prevents users from knowing the complete physical pathname for a file. A URI is set up by mapping a URL to an additional document directory.

For example, if the path for a file is *server_root/DOCS/MARKETING/BIZPLANS/PLANB.DOC*, you could set up a map that prevents users from seeing all but the last directory by defining a URL prefix of plans and mapping it to *server_root/DOCS/MARKETING/BIZPLANS*. From then on, users need only type *PLANS/PLANB.DOC* to locate the file. For more information, see Chapter 2, "Managing Server Content," on page 21.

The Enterprise Web Server provides four default mappings:

- ♦ /—the primary document directory (sometimes called the document root), which initially maps to *server_root*/DOCS.
- ♦ /SEARCH-UI—the directory for most of the search interface files.
- ♦ /WEBPUB-UI—the directory for most of the Web Publisher interface files.
- ♦ /PUBLISHER—the directory for most of the Web Publisher files.

When you create a collection, you must specify which document directory to index. You can only choose a directory that has a URL mapping or a subdirectory within a mapped directory. You can create your own mappings to define specific directories.

To create your own mapping:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Content Management > Additional Document Directories.
- 2** In the URL prefix field, type a keyword that maps the URL to the additional document directory you want to define.

For example, type **PLANS** as the name for your directory.
- 3** In the Map to Directory field, type the absolute physical path of the directory you want the URL to map to.

For example, SYS:\VOL1\DOCS\MARKETING\BIZPLANS
- 4** If you want to apply a style to the directory, select the style you want to apply in the Apply Style drop-down list. See Chapter 2, “Managing Server Content,” on page 21 for more information about styles.
- 5** Click OK.

Once you create a collection based on an additional document directory, you cannot change the URL mapping or the collection’s entries will target the URL mapping to the wrong physical file location.

Turning Search On or Off

You can turn search capabilities on and off for your server. Turning search off for a server where users do not use this function can improve server performance. You may also want to turn off the search function at certain

times when you know the server will experience heavy traffic, reserving this function for times when traffic is lighter.

If you turn it off, the search plug-in is not loaded when the HTTP server starts up. The default is for search to be turned on.

To turn off the search function:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Agents & Search > Search State.
- 2** Select Off.
- 3** Click OK > OK.

To turn search back on:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Agents and Search > Search State.
- 2** Select On.
- 3** Click OK > Save and Apply.

Configuring the Search Parameters

As server administrator, you can set the default parameters that govern what users see when they get search results.

To configure what users will receive after performing a search

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Agents and Search > Search Configuration.
- 2** In the Default Result Set Size field, type the default maximum number of search result items displayed to users at a time.

This cannot be larger than the value for the largest possible result set size, as defined in Step 3 on page 118. The default is 20.

- 3** In the Largest Possible Result Set Size field, type the maximum number of items in a result set.

The default is 5000.

For example, if you type 250 as the value, and there were 1000 documents that matched the search criteria, users would only be able to see the first 250 or the 250 top-ranked documents (for searches that rank results).

- 4** In the Date/Time String field, type the date and time in Posix format.

This will affect how the search results are displayed to users in the search results page. For example, the format %b-%d-%y %H:%M produces Oct-1-99 14:24. You can use the symbols listed in Table Table 5 on page 119.

Table 5 Common Posix Date and Time Formats

Format	Displayed Result (Example)
%a	Abbreviated week day (for example, Wed)
%A	Full week day (for example, Wednesday)
%b	Abbreviated month (for example, Oct)
%B	Full month (for example, October)
%c	Date and time formatted for current locale
%d	Day of the month as a decimal number (for example, 01-31)
%H	Hour as a decimal number, 24-hour time (for example, 00-23)
%m	Month as a decimal number (for example, 01-12)
%M	Minute as a decimal number (for example, 00-59)
%x	Date
%X	Time
%y	Year without century (for example, 00-99)
%Y	Year with century (for example, 1999)

5 In the Default HTML title field, type a default title for the document that is to be used if the document’s author has not included a title as part of the document, tagged with the HTML Title tag.

The typical default is Untitled, which appears in the search results page for HTML files.

- 6** If you want the user's access permission to be checked before displaying the search results, select Yes under the Check Access Permissions before Displaying Search Results field.
- 7** Click OK > Save and Apply.

Configuring Your Pattern Files

Pattern files are HTML files that define the layout of the text search interface. You can associate a pattern file with a search function and a set of pattern variables to create a specific portion of the interface. In the pattern file, you define the look and function of the text search interface. Pattern files use pattern variables that you can use to customize such things as background color, help text, and banners. In some cases, the values are pathnames to the files that contain the actual text and graphics that these variables represent; in other cases, the values represent text and HTML.

You can use the default pattern files, or you can create your own customized set of files and their respective links. See "Customizing the Search Interface" on page 137 for more information about how to change the user interface.

To define where the search function is to look for default pattern files associated with a particular search request:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Agents and Search > Search Pattern Files.
- 2** In the Pattern File Directory field, type the absolute path for the directory where you store your pattern files.

The default start (header), end (footer), and query page pattern files are located in this directory.

- 3** In the Default Start Pattern File field, type the relative pathname you want to use for the top of the search results page when a collection has no defined header file or when more than one collection is being searched.

Specify the path relative to the pattern file directory, as defined in Step 2 on page 120.

- 4** In the Default End Pattern File field, type the relative pathname you want to use for the footer of the search results page when a collection has no defined footer file or when more than one collection is being searched.

Specify the path relative to the pattern file directory, as defined in Step 2 on page 120.

- 5** In the Pattern File for Query Page, type the relative pathname for the pattern file you want to use for the search query page that appears when you start up the search function.

Specify the path relative to the pattern file directory, as defined in Step 2 on page 120.

- 6** Click OK.

Configuring Manually

The search function examines several configuration files to determine how search is configured on your server. These files define system settings, user-defined variables, and information about your search collections. You normally change this information through the Agents and Search form, but you can also modify the files manually with your own text editor. Some of the implications of changing the configuration files in order to customize the user interface are discussed in “Customizing the Search Interface” on page 137.

It is not recommended that you make any manual modifications to your configuration files, but if you do, you must restart the server for your modifications to take effect.

The configuration files that govern searching are WEBPUB.CONF, USERDEFS.INI, and DBLIST.INI.

- ♦ **WEBPUB.CONF**—This system configuration file contains system settings and file paths. In your server’s OBJ.CONF file, the search system initialization is mapped to the WEBPUB.CONF file. When you use the Search Configuration and Search Pattern Files forms, the data you input is reflected in the WEBPUB.CONF file. You can customize your server’s search configuration by changing some of the settings in the WEBPUB.CONF file, but in general, you can make the changes you need through the Agents and Search forms.
- ♦ **USERDEFS.INI**—This user definitions file defines the user-defined pattern variables. In the WEBPUB.CONF file, this is mapped to the USERDEFS.INI file for your language (English, German, Japanese, and so on). You can customize a search interface by creating and defining your own pattern variables in the USERDEFS.INI file that can be used throughout your pattern files (See “User-Defined Pattern Variables” on page 142 for details.)

- ♦ DBLIST.INI—This collection contents file describes collection-specific information. The WEBPUB.CONF file is mapped to the DBLIST.INI file. When you create and maintain collections, the DBLIST.INI file is updated for you with information about your collections.

Adjusting the Maximum Number of Attributes

Collections have different sets of default attributes that depend on their file format. For example, HTML files have Title and SourceType. You can also define META-tagged HTML attributes in your HTML files. Some file formats, such as PDF, have many default attributes. See “Collection Attributes” on page 124 and Table 6 on page 124 for more information about the attributes for each format.

You can use the Add Custom Property form to add additional properties. See Chapter 10, “Configuring Web Publishing,” on page 177 for directions on adding custom properties. These are the default maximum settings:

- ♦ Text: a maximum of 30, including all META-tagged attributes
- ♦ Numeric: a maximum of 5
- ♦ Date: a maximum of 5.

You can change the maximum settings for these in the WEBPUB.CONF file, although larger sets of attributes impact the performance of your server. You cannot set the maximums beyond 100 for text and 50 for dates and numbers.

To use the maximum settings, you need to manually edit the [NS-loader] section of the WEBPUB.CONF file to define the maximum numbers of attributes. For example, to change all three values, you could type these lines:

```
NS-max-text-attr = 50
```

```
NS-max-numeric-attr = 10
```

```
NS-max-date-attr = 10
```

You cannot use the additional attributes in existing collections, only in subsequently created collections. To use them in a search collection, you must use the Agents and Search > Maintain Collection form to remove the collection, then use the Agents and Search > New Collection form to create a new collection. If you want to use the new attributes in the Web Publishing collection, you must use your file system to remove both the WEB_HTM and LINK_MGR collection files from the search collections directory, then restart your server.

Restricting Memory for Indexing

You can set a limit on the amount of RAM available for indexing operations. To set the limit, you need to manually edit the [NS-loader] section of the WEBPUB.CONF file to add a line to define a maximum memory amount. For example,

```
NS-max-memory = 32000000
```

The default is for the server to use all of the available memory that the system can offer. Most typically, you need to limit the RAM used for indexing in these two cases:

- ♦ The Enterprise Web Server is installed on a machine that has less than the suggested minimum RAM requirement, 32 MB.
- ♦ For server administrators on Windows NT* servers that require a great deal of indexing, but who want to set aside some memory for other server operations.

Index File Size

Typically, an indexing operation requires approximately 1.5 MB per file, and because there are two files, one of which is temporary, you may need as much as 3 MB of disk space for indexing. Setting the file size to 1.5 MB per file puts a cap on how large each file can become.

Indexing Your Documents

Before users can execute searches, they need a database of searchable data against which they can target their searches. You need to create a database or collection that indexes and stores information about the documents such as their content and file properties. You can add or delete documents from a collection by optimizing, updating, and managing your collections as needed.

Collections

When your server administrator indexes all or some of a server's documents, information about the documents is stored in a collection. Collections contain such information as the format of the documents, the language they are in, their searchable attributes, the number of documents, the collection's status, and a brief description of the collection. For more details, see "Displaying Collection Contents" on page 136.

When you create a collection, you indicate the type of files that it contains: HTML, ASCII, news, e-mail, PDF, or multiple formats. The file format determines what happens during indexing: which attributes are indexed and what, if any, file conversion has to be done. You can index all the files in a directory or only those with a specific extension—for example, all the HTML, PDF, or *.DOC documents.

A collection has records with information about each document that has been indexed. If the document is deleted from the collection, only the collection's entry for that document is removed. The original document is not deleted.

Collection Attributes

Server documents can be in a variety of formats, such as HTML, Microsoft Excel, Adobe PDF, and WordPerfect. The following is a list of the types of files in which a search in Enterprise can perform:

- ♦ HTML
- ♦ ASCII text files
- ♦ Corel* WordPerfect Macintosh 3.x
- ♦ Corel WordPerfect 5.x, 5.x (Japanese), 6.x, 7.x, 8.x
- ♦ Microsoft Word 1.x, 2.x, 6.x, 97
- ♦ Corel Presentations* WPG 1.x, 2.x
- ♦ Corel Presentations Show 2.x, 3.x, 7.x
- ♦ Corel Quattro Pro* 5.x, 6.x, 7.x
- ♦ Lotus* Ami Pro*
- ♦ Rich Text Format (.RTF)
- ♦ Adobe PDF

Certain file formats have a default set of attributes that are indexed for files of that type, as shown in Table Table 6 on page 124.

Table 6 The Default Attributes Indexed for Each File Format

File Format	Attribute	Type	Description
ASCII	(none)	-	-
HTML	Title	text	The user-defined title of the file

File Format	Attribute	Type	Description
NEWS	SourceType	text	The original format of the document
	From	text	The source user ID of the news item
	Subject	text	The text from the subject field of the news item
	Keywords	text	Any keywords defined for the news item
E-MAIL	Date	date	The date the news item was created
	From	text	The source user ID of the e-mail
	To	text	The destination user ID of the e-mail
	Subject	text	The text from the e-mail's subject field
	Date	date	The date the e-mail was created

By default, HTML collections have Title and SourceType attributes, but they can be indexed to permit searching and sorting by up to 30 file attributes tagged with the HTML <META> tag. You can change the maximum settings for file attributes in WEBPUB.CONF, as discussed in “Adjusting the Maximum Number of Attributes” on page 122.

For example, a document could have the following lines of HTML code:

```
<META NAME="Writer" CONTENT="J. S. Smith">
<META NAME="PubDate" CONTENT="07-24-97">
<META NAME="Product" CONTENT="Communicator">
```

If this document was indexed with its META tags extracted, you could search it for specific values in the Writer, Publication Date, or Product fields.

Any attribute values in META-tagged fields are text strings only, which means that dates and numbers are sorted as text, not as dates or numbers. Also

illegal HTML characters in a META-tagged attribute are replaced with a hyphen.

Creating a New Collection

To create a new collection:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Agents and Search > New Collection.
- 2** Select any of the items in the drop-down list as a starting point for finding the directory you want to index.

If you want to index a different subdirectory, click View to see a list of resources. You can index any directory that is listed or you can view the subdirectories in a listed directory and index one of those instead. See “Mapping URLs” on page 116 for more information about additional document directories.

- 3** Return to New Collection and the Create Collection form (if you are currently in View).

The directory name appears in the Directory to Index field.

- 4** Specify the indexing criteria in the Documents Matching field.

You can index all files in the chosen directory by leaving the default *.HTML pattern in the Documents Matching field, or you can define your own wildcard expression to restrict indexing to documents that match that pattern.

You can define multiple wildcards in an expression. See Chapter 1, “Managing Your Server,” on page 15 for details of the syntax of wildcard patterns.

- 5** To also index the subdirectories within the specified directory, select Yes under Include Subdirectories.
- 6** In the Collection Name field, type a name for your collection.

The collection name is used for collection maintenance. This is the file’s physical name, so follow the standard directory naming conventions for your operating system. You can use any characters up to a maximum of 128, including spaces. Spaces are converted to underscores.

- 7** In the optional Collection Label field, type a user-defined name for your collection.

This is what users see when they use the text search interface. You can use any characters except single or double quotation marks, which prevent agent services from operating. You can use up to a maximum of 128 characters.

- 8** In the optional Description field, type a description for your collection.

You can use up to a maximum of 1024 characters, including spaces.

- 9** From the available options under Collection Contains, select the type of files the collection is to contain: HTML, ASCII, News, E-Mail, PDF, or Multiple Document Formats.

The kind of file format you choose indicates which default attributes are used in the collection and which, if any, automatic HTML conversion of the content is done as part of indexing. See “Collection Attributes” on page 124 and Table 6 on page 124 for information about the attributes for each format.

If you choose HTML as the file type and also try to index non-HTML files, the server creates the collection with the HTML set of default attributes and does not attempt to convert any non-HTML file it indexes. Regardless of the file type chosen, the content of the file is always indexed.

- 10** From the Extract Metatags field, select whether or not to extract META-tagged attributes from HTML files during indexing.

If you extract these attributes, you can search on their values. You can index on a maximum of 30 different user-defined META tags in a document. You cannot use this option for multiple-format collections.

- 11** Click the Documents Are In drop-down list > select the collection’s language.

The default is English, labeled “English (ISO-8859-1).” For more information on character sets, see Chapter 2, “Managing Server Content,” on page 21.

- 12** Click OK.

Configuring an Existing Collection

After you have created a collection, you can modify some of the initial settings. This data resides in the collection information file DBLIST.INI. When you reconfigure a collection, the DBLIST.INI file is updated to reflect your changes. See “Configuring Manually” on page 121 for more information about the configuration files.

The Configure Collection form allows you to modify some of the settings for the Web Publishing default collection, `web_html`, because you are not changing actual collection data. Avoid making unnecessary changes to this collection's settings.

To configure a collection:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Agents and Search > Configure Collection.
- 2** In the Choose Collection drop-down list, select the collection you want.

The Collection, Document Root, File Format, and Language should already be completed and correct. If not, see “Creating a New Collection” on page 126 for more information about creating new collections.

- 3** In the optional Description field, type a description for your collection up to a maximum of 1024 characters.
- 4** In the optional Label field, you can type a user-defined name for your collection.

This is what users see when they use the text search interface. You can use any characters except single or double quotation marks, which prevent agent services from operating. You can use up to a maximum of 128 characters.

- 5** In the URL for Documents field, type the new URL mapping for the collection's documents, if the mapping has changed.

See “Mapping URLs” on page 116 for more information about additional document directories.

- 6** In the Highlight Begin and Highlight End fields, type in the HTML tagging you want the server to use when selecting a search query word or phrase in a displayed document.

The default is to use bold, with the `<b>` and `</b>` tags.

- 7** In the Input Date Format field select the option that will input dates to be interpreted when using this collection.

You can define different default pattern files for displaying the search results: how the search result's header, footer, and list entry line are formatted, respectively. Initially, the pattern files are in `server_root/PLUGINS/SEARCH/UI/TEXTSERVER_ROOT\PLUGINS\SEARCH\UI\TEXT`

- 8** In the Record Pattern File, type the name of the file record you want recorded.
- 9** In the Result Pattern Files field, type the name of the pattern file you want to use when displaying a single selected document from the list of search results.
- 10** Click OK.

Updating an Existing Collection

After you have initially created a collection, you may want to add or remove files. If you are adding documents, the file contents are indexed (and converted if necessary), when their entries are added to the collection. If you are removing documents, the entries for the files are removed from the collection along with their metadata. This function does not affect the original documents, only their entries in the collection.

If you selected the Extract Metatags option when you created this collection, then the META-tagged HTML attributes are indexed whenever you add new documents to this collection.

To update a collection:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Agents and Search > Update Collection.
- 2** Click the Choose Collection drop-down list > select the collection you want to update.

The Selection Collection list shows which documents have index entries in the currently selected collection.

- 3** In the Documents Matching field, type a single filename or use wildcards to specify the type of files you want added to or removed from the collection.

If you type a wildcard such as *.HTML, only files with this extension are affected. You can indicate files within a subdirectory by typing the pathname as it appears in the list of files.

NOTE: Be careful how you construct wildcard expressions. For example, if you type INDEX.HTML you can add or remove the index file from the current collection. If you type `*/index.html` you can add or remove all INDEX.HTML files in the collection.

- 4** Under Include Subdirectories, select whether to index and add all matching documents from the subdirectories of the document directory that was originally defined for the collection.

If the collection originally indexed the /PUBLISHER directory, this option looks for documents matching the new pattern within all the subdirectories within /PUBLISHER. This does not apply for removing documents.

- 5** Click Add Docs to add the indicated files and subdirectories.
- 6** Click Remove Docs to remove the indicated files.

Maintaining an Existing Collection

Periodically, you may want to maintain your collections, especially if you do a great deal of indexing and updating collections. You can perform the following collection management tasks:

- ♦ **Optimize Collections:** You can optimize a collection to improve performance if you frequently add, delete, or update documents or directories in your collections. Optimizing is done automatically when you reindex or update a collection, so you would use this option infrequently, for examples, just before publishing it to another site or before putting it onto a read-only CD-ROM.
- ♦ **Reindex:** You can reindex a collection, which locates each file that already has an entry in the collection and reindexes its attributes and contents, extracting the META-tagged attributes if that option was selected when the files were originally indexed into the collection. This does not return to the original criteria for creating the collection, for example, *.HTML, and add any new documents that fit the original criteria. This option also removes collection entries when the source documents have been deleted and can no longer be found.
- ♦ **Remove:** You can remove a collection. This removes only the collection, not the original source documents.

To perform any of the collection management tasks:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Agents and Search > Maintain Collection.
- 2** Click the Choose Collection drop-down list > select the collection you want to manage.
- 3** Click Optimize to optimize the collection, Reindex to reindex the collection, or Remove to remove a collection.

Scheduling Regular Maintenance

You can schedule collection maintenance at regular intervals if you frequently add, delete, or update documents or directories in your collections.

Optimizing is done automatically when you reindex or update a collection, so you should not need to do additional optimizing. However some very active Web sites may require frequent reindexing if new documents are added on a daily basis.

To schedule the maintenance of your collection:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Agents and Search > Schedule Collection Maintenance.
- 2** Click the Choose Collection drop-down list > select the collection item you want.
- 3** Click the Choose Action drop-down list > select Reindex or Optimize.
You can set up separate schedules for reindexing and optimizing the same collection.
- 4** In the Schedule Time field, type in the time of day when you want the scheduled maintenance to take place.
Use 24-hour format (HH:MM). HH must be less than 24 and MM must be less than 60. You must type a time.
- 5** In the Schedule Day(s) of the Week checklist, check one or more of the day check boxes.
- 6** Click OK.

Unscheduling Collection Maintenance

If you have scheduled regular reindexing or optimizing of a collection, you can remove the scheduled maintenance when you no longer want the collection to be maintained at regular intervals.

To unschedule collection maintenance:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Agents and Search > Remove Scheduled Collection Maintenance.
- 2** Click the Choose Collection drop-down list > select a collection. This lists all your collections you have set up regular maintenance for.

- 3 Click the Choose Action drop-down list > select Reindex or Optimize.

The time and days of the week is listed on the screen for the maintenance you have scheduled.

- 4 Click OK.

Using the Enterprise Web Server to Simplify Your Search

Users are primarily concerned with entering search queries when searching document collections and getting a list of documents in return. When you install the Enterprise Web Server, a default set of search query and result forms are included.

The default installation of NetWare Enterprise Web Server includes three search query pages: standard and advanced HTML queries and a Java-based guided query.

Search Home Page

The search home page, at <http://yourserver/search-ui/examples> (*yourserver* is either your IP address or the name you have designated for it), provides individual links to each of the three search query interfaces as well as an online QuickStart tutorial on customizing the interface. The tutorial discusses the various pattern files and gives examples of how they can be changed to produce different results.

In order to enable search, the Web Publisher must be on.

Standard Search Queries

With the standard search query, you select a collection to search against and type a word or phrase to search for using the query language operators.

To conduct a standard search:

- 1 Type the following URL in the location field in your Web browser:

<http://yourserver/search>

Yourserver is either your IP address or the name you designated for it when Enterprise was installed.

- 2 In the search query page that appears, click the Search drop-down list > select the collection you want to search through.

- 3** In the For field, type the word or phrase for your search query.
- 4** Click Search to execute your query.

Guided Search

You can choose to use the Java-based guided search interface, which helps you construct the query. This is especially useful if you want to build a query that has several parts, such as searching for a word in the document contents as well as a specific attribute value. To enable Java, use the Languages option Preferences menu command (for Netscape browsers it is through Edit > Preferences > Advanced). If Java is not enabled, a modified form of the simple search will appear.

Depending on your server and its capabilities, a Java-enabled search may not be supported.

There are two ways to obtain the guided search page: through the search home page or through the standard search query page.

To access guided search through the search home page:

- 1** Type the following URL in the location field in your Web browser:

`http://yourserver/search-ui/examples`

Yourserver is either your IP address or the name you designated for it when Enterprise was installed.

- 2** Click Standard Search > Guided Search.

To access guided search through the standard search query page:

- 1** Go to the following URL:

`http://yourserver/search`

- 2** Click Guided Search on the standard search page.
- 3** Click the Search In drop-down list > select the collection you want to search through.
- 4** Click the For drop-down list > select the type of element you want to search for.
- 5** In the blank text field, type in the word you want to search for.
- 6** Click Add Line to add the first part of the query.

The word appears in the large text display box at the bottom of the form.

- 7** To add to your query, select another element from the For drop-down list.
A new drop-down list appears on the right side of the form, listing all attributes that are available for the chosen collection.
- 8** Select the attribute you want to search against.
- 9** Click the That drop-down list > select a query operator (Contains, Starts, Ends, Matches, Has a Substring) or logical operator (=, <, >, <=, >=).
- 10** In the blank text field, type the attribute value you want to search for.
- 11** Continue adding lines and modifying your query until it is complete.
Click Add Line to add another line for your query, Undo Line to remove the last line you added, or Clear to remove the entire query.
- 12** Click Search to execute the search.

Advanced Search Query

You can use the advanced HTML search interface, which helps you construct the query. This is especially useful if you want to create a query that searches through more than one collection or that produces results sorted by a specific attribute value. See “Guided Search” on page 133 for information about enabling Java.

There are two ways to obtain the advanced HTML search page: through the search homepage or through the standard search query page.

To access advanced HTML through the search home page:

- 1** Go to the following URL:
`http://yourserver/search-ui/examples`
Yourserver is either your IP address or the name you designated for it when Enterprise was installed.
- 2** Click the Advanced HTML Search link on the home page.

To access advanced HTML search through the standard search query page:

- 1** Go to the standard search query page by typing the following URL in the location field in your Web browser:
`http://yourserver/search`
- 2** Disable Java for your browser. To do this, use the Languages option Preferences menu command (for Netscape browsers click Edit > Preferences > Advanced).

- 3** Click Guided Search on the standard search page.
- 4** In the For field, type the word or phrase you want to search for.
- 5** Type one or more attributes to sort the results by.

The default is an ascending sort order, but you can indicate a descending sort order with a minus (-). See “Sorting the Results” on page 136 for more information about sorting.
- 6** Expand or limit the number of matching documents you want the search to return at a time.

Prev and Next allow you access to additional pages of documents if there are too many to fit on a page at once.
- 7** Click the Search drop-down list > select the collection you want to search.

You can select more than one collection by Ctrl + clicking another collection. All collections in a query must be in the same language.
- 8** Click Search to execute your query.

Search Results

There are two standard types of search results: a list of all documents that match the search criteria and the text of a single document that you selected from the list of matching documents.

Listing Matched Documents

When you execute a search from either the simple or advanced search query pages, you obtain a list of the documents that match your search criteria. The list gives some standard information about each file, depending on the collection’s format. For example, the default results page for e-mail collections give Subject, To, From, and Date for each entry and news collections give Subject, From, and Date for each entry.

The kind of file format in the collection indicates which default attributes are available for searching. See “Collection Attributes” on page 124 and Table 6 on page 124 for information about the attributes for each format.

For entries resulting from a search that checks for comparative proximity of words to each other or for the exactness of the match, the file’s ranking can be provided by showing a score, which is a percentage of how close the result was (ranked from 1-100%).

Sorting the Results

By default, or if you don't enter anything in the Sort By field on the advanced HTML query page, all documents matching the search are listed according to their relevance ranking (for queries that consider this) or their position in the server file database (for other queries).

If you enter an attribute name in the Sort By field, the documents are displayed in an ascending sort sequence. You can list the documents in a descending sort sequence by adding a minus sign (-) prefix to the attribute, as in -keywords or -title. You can do a multiple sort, by typing more than one field, as in Author,-PubDate.

Attribute values in META-tagged fields are text strings, which means that dates and numbers are sorted as text, not as dates or numbers. To convert the value into a date or number, you can create a new property in Enterprise Web Server *servername* > Web Publishing > Add Custom Property form and check the box that marks this property as a META-tagged attribute.

Displaying a Selected Document

In the default installation of NetWare Enterprise Web Server, when you obtain a list of the documents that match your search criteria, you can select a single document to view in your Web browser. Depending on how the pattern files are set up, the word you entered as your original search query can be selected in the displayed document with color, boldface text, or blinking text.

In the case of documents that have been converted into HTML, the URL points you to the original document. To get to the converted HTML document, click the document's title.

Displaying Collection Contents

You can display the contents of your collection database to see which attributes are set for each collection. The default installation of NetWare[®] Enterprise Web Server uses the HTML-DESCRIPTION.PAT file to display information about each of your collections that have been defined as displayable (NS-display-select = YES) in the DBLIST.INI file. The collection contents typically include

- ♦ Collection name, label, and description
- ♦ Collection format
- ♦ Number of attributes in the collection and a list of their names

- ♦ Number of documents in the collection
- ♦ Collection size and status
- ♦ Language and character set
- ♦ Input and output date formats

To display your collection database contents, go to the following URL:

<http://yourserver/search?NS-search-page=c>

Customizing the Search Interface

As server administrator, you can customize the search interface to meet specific user requirements. All of the HTML-based forms that the user sees are defined through a set of pattern files that set up display formats for the search results page header and footer as well as each search result record listed in response to a query. There are a set of pattern variables that you can use to construct the forms used for search input and output. Many of the variables are defined in the system and user configuration files (USERDEFS.INI, WEBPUB.CONF, and DBLIST.INI, which are discussed in “Configuring Manually” on page 121).

The Search home page at <http://yourserver/search-ui/examples> also provides an introduction to the search interface as well as an online QuickStart tutorial on customizing the interface. The tutorial discusses the various pattern files and gives examples of how they can be changed to produce different results.

HTML Pattern Files

A good place to begin customizing the interface is by modifying the existing pattern files. After you see how they work and you understand pattern variables, you can create your own pattern files and change the configuration files and other pattern files to point to them. In the default installation of NetWare Enterprise Web Server, the pattern files are in this directory: *server_root\PLUGINS\SEARCH\UI\TEXT*. (If you edit them, make copies of your original pattern files so you can restore them afterwards.)

There are pattern files for different kinds of collections: e-mail, news, ASCII, PDF, and HTML, as well as one for the Web Publishing collection. (The Web Publishing pattern file is a special case, using a several collection-specific attributes as variables in the *dblist.ini* file.) There are several general types of pattern files, each of which has a particular use:

- ♦ QUERY.PAT displays the Standard and Advanced Query pages.
- ♦ TOCSTART.PAT displays the header across the top of the Search Results page.
- ♦ TOCREC.PAT displays each document listed on the Search Results page.
- ♦ TOCEND.PAT displays the footer across the bottom of the Search Results page.
- ♦ RECORD.PAT displays a single selected document from the Search Results page (See “Displaying a Selected Document” on page 136 for more information.)
- ♦ DESCRIPTIONS.PAT displays the collection contents.

The pattern files contain HTML formatting instructions, which define how elements look, and HTML search arguments and variables, which define the text label or value that is displayed.

There are three kinds of pattern variables (discussed further in “Using Pattern Variables” on page 142):

- ♦ User-defined in the USERDEFS.INI file, with a \$\$ prefix. See “User-Defined Pattern Variables” on page 142.
- ♦ Defined in the configuration files WEBPUB.CONF and DBLIST.INI, with a \$\$NS- prefix. See “Configuration File Variables” on page 144, Table 9 on page 144, and Table 10 on page 145.
- ♦ Search macros and variables generated by a pattern file, with a \$\$NS- prefix. See “Macros and Generated Pattern Variables” on page 146 and Table 11 on page 146.

To see how these work together, the following are lines from the standard query pattern file, NS-QUERY.PAT:

```
<input type="hidden" name="NS-max-records" value="$$NS-max-records">

<td align=left colspan=2>$$logo</td>
<td align=right><h3>$$sitename</h3></td>

<td align=right><b>$$queryLabel</b></td>
<td align=left>&nbsp;<input name="NS-query" size=40
value="$$NS-display-query"></td>
```

Each line contains standard HTML tags and one or more variables with the \$\$ or \$\$NS- prefix. To examine each line more closely, see the configuration files mentioned in “Configuring Manually” on page 121.

- ◆ **NS-max-records:** Defined in the WEBPUB.CONF file. Because this field is hidden, users cannot change this value, which defines how many matching documents to return at a time. In the advanced HTML query pattern file, NS-ADVQUERY.PAT, this is a user-modifiable input field.
- ◆ **\$\$NS-max-records:** The search generates a variable from this field that can be used in subsequent searches to calculate how many result records to display at a time. Because this field is not modifiable here, the value is set to the one contained in the WEBPUB.CONF file. In the advanced query, this value could vary for each query.
- ◆ **\$\$logo:** Defined in the USERDEFS.INI file. This could be any image or text the user wanted to display on the form.
- ◆ **\$\$sitename:** Defined in the USERDEFS.INI file as the server’s hostname that is provided by the \$\$NS-host search macro.
- ◆ **\$\$queryLabel:** Defined in the USERDEFS.INI file as a text label for the query input field. In this case, the label on the form is the word (and punctuation) “For:”
- ◆ **NS-query:** Defined in this pattern file as the name of the input field.
- ◆ **\$\$NS-display-query:** Defined in the USERDEFS.INI file. The search generates a variable from this field that can be used in subsequent searches to determine which word or phrase to selected when an entire matching document is displayed.

Search Function Syntax

The search function uses standard URL syntax with a series of name-value pairs for the search arguments. The following is the basic syntax:

```
http://yourServer/  
search?name=value [&name=value] [&name=value]
```

As you use the HTML search query and results pages, you can see search functions and arguments displayed in the URL field of your browser. When entered directly into the URL field, these are sometimes called decorated URLs. You can also embed them in your pattern files with the HREF tag.

You can create a complete search function as an HREF element within a pattern file. The following example is from the HTML-DESCRIPTIONS.PAT file, which defines how collection information is

displayed. The following lines produce a heading for each collection for the label (“Collection:”) and provides a link to the actual collection file through the collection’s label (NS-collection-alias) that was defined in the DBLIST.INI file.

```
<td colspan=6><font size=+2><b>$$collectionLabel</b>
<a href=$$NS-server-url/search?NS-collection=$$NS-
collection>$$NS-collection-alias</a>
</font></td>
```

The HREF contains a complete search function by using the following elements:

- ♦ \$\$NS-server-url: A search macro that determines the user’s server URL.
- ♦ /search: The search command itself.
- ♦ ?: The query string indicator. Everything after the question mark is information used by the search function.
- ♦ NS-collection=\$\$NS-collection: This uses the search macro \$\$NS-collection to define the collection’s filename.

You can set up a search to use a variable conditionally so that if there is no value associated with the variable, nothing is displayed. The syntax is as follows:

```
variableName[conditionalized output]
```

For example, you could request that the document’s title be output if it exists. If there is no title for this document, not even the label “Title:” will be displayed. To display the title, type the following:

```
$$Title[<P>Title: <B>$$Title</B>]
```

URL Encodings

When you construct HTML instructions, whether in decorated URLs or within a pattern file, you need to follow the rules for URL encoding. Any character that might be misunderstood as part of a URL should be encoded with a code in the format of %*nn*, when *nn* is a hexadecimal code. Blanks are converted to the + symbol (plus sign) in queries or to %20 in output. Table 7 on page 141 shows the most commonly used URL codes.

Table 7 Common URL Encodings

Character	Description	Code
	Space	%20
;	Semicolon	%3B
/	Slash	%2F
?	Question mark	%3F
:	Colon	%3A
@	At sign	%40
=	Equal sign	%3D
&	Ampersand	%26

Required Search Arguments

Although you can customize almost every aspect of query and result pages, there are some arguments required for search functions to display the different types of search pages. These arguments are required whether the search function is in a decorated URL or embedded as an HREF in a pattern file.

Search functions that display the search query page require the following arguments:

- ♦ Search query (the word, phrase, or attribute you want to search on)
- ♦ Collection (can be specified more than once for multiple-collection searches)

Search functions that display the search results page require the following arguments:

- ♦ NS-search-page=results (or r, in upper or lowercase)
- ♦ Collection (can be specified more than once for multiple-collection searches)
- ♦ Search Query
- ♦ NS-search-page=document (or d, in upper or lowercase)
- ♦ Document path
- ♦ Collection (can be specified only once)

- ♦ Search query (necessary if you want to select the query data)

Search functions that display a highlighted document require these arguments:

NS-search-page=contents (or c, in upper or lowercase)

Using Pattern Variables

By using pattern variables, you can customize the search text interface and eliminate the need to update the actual HTML pages as user requirements change. For example, if the interface has graphics or text elements that change periodically, you can define a pattern variable that points to a pathname where that graphic or text is maintained and stored.

There are three categories of pattern variables:

- ♦ Variables defined in the USERDEFS.INI file, to which are added a \$\$prefix in decorated URLs and pattern files. For example, uidir, logo, and title become \$\$uidir, \$\$logo, and \$\$title.
- ♦ Variables defined in the configuration files WEBPUB.CONF and DBLIST.INI, which have an NS- prefix when they are defined in the configuration file and which have an \$\$NS- prefix when they are used in decorated URLs and pattern files. For example, NS-max-records, NS-doc-root, and NS-date-time become \$\$NS-max-records, \$\$NS-doc-root, and \$\$NS-date-time.
- ♦ Search macros and variables generated by a pattern file, which always have a \$\$NS- prefix. For example, \$\$NS-host, \$\$NS-get-next, and \$\$NS-sort-by.

User-Defined Pattern Variables

You can create any number of your own user-defined pattern variables in the user definitions file USERDEFS.INI, or you can modify existing definitions. When one of these variables is used in a pattern file, the \$\$ prefix is added to it. Variable names can have up to 32 characters or digits or combinations of both. Characters can be letters A-Z in upper or lowercase, hyphens (-), and underscores (_). Names are case-sensitive.

The default USERDEFS.INI file included with NetWare Enterprise Web Server contains variables that are used to define the search query page (labeled [query]) in the file, the results listing (labeled [toc]), the document display page, (labeled [record]), and the collection contents page (labeled [contents]). Each line begins with a variable name and is followed by a definition for that variable. Many are labels for screen elements, some are paths to other files,

and some have more complex contents. For example, the following lines are from the query section of the default USERDEFS.INI file.

```
[query]
help=/publisher/help/srchhelp.html
title=ES3.0 Sample Search Interface
queryLabel=Search&nbsp;for:
collectionLabel=Collection:
booleanLabel=Boolean:
sortByLabel= Sort&nbsp;for:
copyright = Copyright &#169; 1997 Netscape Communications
Corporation. All Rights Reserved.
```

The file also includes references to search macros, such as \$\$NS-server-url, and can also refer to other user-defined variables, as in the following lines:

```
uidir = $$NS-server-url/search-ui
icondir = $$uidir/icons
```

Search macros are described further in the section “Macros and Generated Pattern Variables” on page 146.

You can use any supported HTML character entity in your variable definitions. You can use entity names that are defined in the `&name;` format as well as those defined with the three-digit code in the `&#nnn;` format. In the USERDEFS.INI code sample, the entity ` ` inserts a nonbreaking space and `©` inserts a copyright symbol. Some of the more commonly used entities are listed in Table 8 on page 143.

Table 8 Common HTML Character Entities

Numeric Code	Entity Name	Description
 		Space
"	"	Quotation mark
$	\$	Dollar sign
:	-	Colon
<	<	Less than
>	>	Greater than
™	-	Trademark symbol
 	 	Nonbreaking space

Numeric Code	Entity Name	Description
©	©	Copyright symbol
®	®	Registered trademark

Configuration File Variables

Some variables are defined in the system configuration and the collection configuration files. These use a prefix of NS- in the configuration file to differentiate them from other markup tags in an HTML page. To use these variables as arguments to the search function, add another prefix \$\$ to the variable, as in \$\$NS-date-time and \$\$NS-max-records.

Variables that define defaults for all searches on a server are defined in the system configuration file WEBPUB.CONF. For example, the default installation of NetWare Enterprise Web Server includes the following variables in the WEBPUB.CONF file:

```
NS-max-records = 20
NS-query-pat = /text/NS-query.pat
NS-ms-tocstart = /text/HTML-tocstart.pat
NS-ms-tocend = /text/HTML-tocend.pat
NS-default-html-title = (Untitled)
NS-HTML-descriptions-pat = /text/HTML-descriptions.pat
NS-date-time = %b-%d-%y %H:%M
```

Although installations may vary depending on each server's configuration, the most commonly found variables from the WEBPUB.CONF file are listed in Table 9-9.

Table 9 Commonly Found Variables Defined in WEBPUB.CONF

Variable	Description
NS-default-html-title	The name given to HTML documents that do not contain a user-defined title. Typically set to "(Untitled)."
NS-date-time	The date and time format to use when displaying results.
NS-date-input-format	The format for inputting dates (the default is MMDDYY).
NS-HTML-descriptions-pat	The pattern file to use when displaying the contents of the collections.

Variable	Description
NS-largest-set	The maximum number of records that can be handled as matching the search criteria. The records are displayed in groups of NS-max-records.
NS-max-records	The maximum size of the result set displayed at one time.
NS-ms-tocend	The pattern file to use for the footer at the bottom of the Search Results page when searching multiple collections.
NS-ms-tocstart	The pattern file to use for the header at the top of the Search Results page when searching multiple collections.
NS-query-pat	The query pattern file used when creating a query page.
NS-search-type	The type of search to perform. Only Boolean is permitted.

Collection-specific variables are defined in the DBLIST.INI file. For example, the default installation of NetWare Enterprise Web Server includes variables for the Web Publishing collection. Among the variables defined there are

```
NS-collection-alias = Web Publishing
NS-doc-root = C:/Netscape/SuiteSpot/docs
NS-url-base = /
NS-display-select = YES
```

The variables in your DBLIST.INI file may differ according to the type of collections you are using. Table 10 on page 145 contains some of the more commonly found collection-specific variables.

Table 10 Commonly Found Variables in DBLIST.INI

Variable	Description
NS-collection-alias	The collection's label. Can be specified more than once to search multiple collections.
NS-doc-root	The root directory for the documents in the collection.
NS-display-select	This indicates whether the collection is displayed as part of the collection information listing, when NS-search-page=contents. The default is Yes.
NS-highlight-start	Begin highlighting at this point in the displayed document. Typically this selects the search query criteria.
NS-highlight-end	End highlighting at this point in the displayed document.

Variable	Description
NS-language	The language of the documents in the collection.
NS-record-pat	The pattern file to use when displaying a selected document page.
NS-tocend-pat	The footer pattern file associated with a collection to be used when formatting the search results.
NS-tocrec-pat	The record pattern file associated with a collection to be used when formatting the search results.
NS-tocstart-pat	The header pattern file associated with a collection to be used when formatting the search results.
NS-url-base	The base URL used when constructing the link used to locate the file.

Macros and Generated Pattern Variables

There are some search macros that you can use in your pattern files or decorated URLs, and the search function itself generates some pattern variables that you can use in subsequent search requests to define how the later output is to be displayed. These macros and variables have a prefix of `$$NS-` to indicate their use.

For example, after doing an initial search query that results in 24 documents on the Results page, you can reuse the search-generated `$$NS-docs-matched` and the `$$NS-doc-number` variables to help define a document page displaying one of the documents in detail. In this way, you can tell the user that this document is number 3 of 24 documents returned for the original search.

The search macros and the generated variables that you can use in a subsequent pattern file or decorated URL are listed in Table 11 on page 146.

Table 11 **Macros and Generated Pattern Variables**

Variable	Description
<code>\$\$NS-collection-list</code>	An HTML multiple select list of all the collections in DBLIST.INI, where NS-display-select is set to Yes.
<code>\$\$NS-collection-list-dropdown</code>	An HTML drop-down list version of NS-collection-list.

Variable	Description
\$\$NS-collections-searched	The number of collections searched for this request.
\$\$NS-display-query	The HTML-displayable version of the query that is generated for a results page.
\$\$NS-doc-href	The HTML HREF tag for the document. This provides a URL to the original source document. For e-mail, this is in the form mailbox:/boxname?id=messageID and for news, it is in the form news:messageID
\$\$NS-doc-name	The document's name.
\$\$NS-doc-number	The sequence number of the document in the results page list.
\$\$NS-doc-path	The absolute path to the document.
\$\$NS-doc-score	The ranked score of the document (ranges 0 to 100).
\$\$NS-doc-score-div10	The ranked score of the document (ranges 0 to 10).
\$\$NS-doc-score-div5	The ranked score of the document (ranges 0 to 5).
\$\$NS-doc-time	The creation time for a document in the results list. To obtain this value, you must set NS-use-system-stat = Yes in the WEBPUB.CONF file. By default it is set to No, because system statistics are expensive.
\$\$NS-doc-size	The size of the document rounded to the nearest KB. To obtain this value, you must set NS-use-system-stat = Yes in the WEBPUB.CONF file. By default it is set to No, because system statistics are expensive.
\$\$NS-docs-found	The actual number of documents that the search engine found for this request.
\$\$NS-docs-matched	The number of documents returned from the search (up to NS-max-records) for this request.
\$\$NS-docs-searched	The number of documents searched through for this request.
\$\$NS-get-highlighted-doc	This provides the URL for a highlighted document in order to be able to display the document as HTML text with highlights.
\$\$NS-get-next	Gets the next set of search results to be displayed. The set is equal to NS-max-records and is positioned by using NS-search-offset.
\$\$NS-get-prev	Gets the previous set of search results that has been displayed. The set is equal to NS-max-records and is positioned by using NS-search-offset.

Variable	Description
\$\$NS-host	The hostname.
\$\$NS-insert-doc	A placeholder used in the NS-record-pat pattern files for HTML to indicate where the source document is to be inserted.
\$\$NS-rel-doc-name	The relative name of the document to display creating a document page.
\$\$NS-search-offset	The offset into the set of records returned as search results. Used to determine which set of records are displayed when you use NS-get-next and NS-get-prev.
\$\$NS-server-url	The URL for the server.
\$\$NS-sort-by	The sort sequence for the items on the results page. You can select one or more of the available attributes for the collection. The default is an ascending sort.

9

Using Agents

NetWare[®] Enterprise Web Server allows you to use server-based agents to manage server files and folders. Agents act as watchdogs for you. They watch for a specific event or time then perform a task for you. For example, you can set up a document agent to notify you when a specific URL has been updated, or you can have a search agent execute each week at the same time to create a list of all Web publishing documents that have been updated during the week. The notification could be an e-mail message or a posting to a newsgroup.

An agent is stored on the server, so you must be connected to the server when you create the agent. The agent resides on the server until it is deleted or completes the assigned task. The server allows only users with the correct permissions, as recorded on the access control list (ACL), to submit an agent. An agent can perform only operations that you are authorized to perform, but it cannot access authenticated sites because agents don't send authorization data with their requests.

As server administrator, you can configure how your server manages agents. For example, you can define who has access to specific agent events and you can restrict who can create or disable agents.

Most agents are triggered by Web Publishing actions; therefore, when Web Publishing is disabled, agent services are also disabled. When you turn Web Publishing back on, agents that were turned off when you turned off Web Publishing are also turned back on. Agents that were disabled for other reasons are still disabled.

Because agents are stored on the server, you must have sufficient disk space available for all agents created on your server. A general rule is to allow 512 bytes per agent, which is approximately 70-100 MB of space for 100,000 agents.

Types of Agents

There are several types of agents, each of which has a particular use. When an event occurs that an agent is monitoring or when the specified time for activation occurs, the agent activates and begins to perform its assigned actions, until it is deleted or expires. Agents can be assigned such actions as the following:

- ♦ Sending an e-mail message
- ♦ Posting a news article to one or more newsgroups
- ♦ Performing an HTTP operation to post to a URL or to get a URL (advanced options only)

Document Agents

Document agents respond to document-related events that take place when something has occurred to a document on the server, such as the following:

- ♦ A document is changed, moved, copied, or deleted
- ♦ Someone views or modifies a document attribute
- ♦ A document is locked or unlocked

Directory Agents

Directory agents respond to directory-related events that take place when something has occurred to a directory on the server such as the following:

- ♦ A directory is changed, moved, copied, or deleted
- ♦ A directory is added
- ♦ Someone lists a directory
- ♦ Someone views or modifies a directory attribute

Document and directory agents are only triggered by HTTP server activities. Making a directory or deleting a document by any means other than Web Publishing will not trigger one of these agents.

Timer Agents

Timer agents respond to time-related events that occur as a result of the date or time. You can submit an agent to activate

- ♦ At a specified date and time

- ♦ At recurring times (for example, every Tuesday at 10 a.m.)
- ♦ At periodic intervals (for example, every five hours)

Search Agents

Search agents execute periodically, notifying the client of any documents that have been modified after the last time the search agent was executed. Search agents can also check the content of documents, pulling a list of all documents in the chosen collection that contain the specified search criteria or text string. You can limit the content search to recently modified documents or you can extend the search to include all documents. For example, in a specified collection you can

- ♦ Check the server at 5:00 a.m. every Monday morning for all documents in a collection that have been modified in the preceding week
- ♦ Check the server at 5:00 a.m. every Monday morning for all documents in a collection that contain the string “JavaScript” that have been modified in the preceding week
- ♦ Check the server at 5:00 a.m. on the first of each month for all new documents in a collection with the word *NetWare* in the title

Creating Authorized Users

Only users that you have added to an access control database are permitted to use agents. You must use the NetWare Web Manager to add users and groups.

Agents access the ACL database that is the default for your server. The local LDAP database is typically set as your default during server installation, but you can direct agent services to access a different ACL database. See “Controlling Access to Your Server” on page 53, for more information about setting access control.

Configuring Agent Services

As server administrator, you must begin by enabling and configuring agent services. You can enable or disable agent services as well as define many default characteristics of individual agents. Before you can use agent services on your server, you must define the MTA (mail) and NNTP (news) hosts for your server. To do this, click Enterprise Web Server *servername* > Server Preferences > Network Settings and type values for the MTA Host and NNTP Host fields.

WARNING: If either of these fields is not properly set an error message will appear, but the process will continue. Later when an agent is triggered, any functionality that relies on one of these server will fail.

Agent Management relies on Web Publishing. The server agents require information from Web Publishing to operate. If you want to use Agent Management you must turn Web Publishing on first.

To turn on the search state:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Agents & Search > Search State.
- 2** Selection On.
- 3** Click OK.

To turn on Web Publishing:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Web Publishing > Web Publishing State.
- 2** Select On.
- 3** Click OK.

To set up agent services for your server:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Agents & Search > Agent Management.
- 2** Select Yes from the Enable Agent Services list.
- 3** In the Agent Directory field, type the full path of the agent directory.

This is where files containing information about agents are kept. The default is *server_root\HTTPS-yourserver\AGENTS-DB*. If you want to specify a different directory, make sure you include a full path.

- 4** In the Agent File Prefix field, type the name of the file that you want to name the agent you are creating.
- 5** In the Agent File Prefix field, type the name of the file that you want to as the name of the agent you are creating.
- 6** In the Maximum Number of Agents per user field, type the maximum number of agents an individual user can have.

The number you specify must be an integer greater than 0.

- 7** In the Maximum life of an agent in days field, type the maximum number of days any agent can exist.

The number you specify must be an integer greater than 0.

This provides a calculated expiration date for your agents. Clients cannot input a longer agent life for a particular agent, although they may create an agent with a shorter lifespan.

- 8** In the Maximum Times That an Agent Can Trigger field, type in the maximum times an agent can be activated.

The number you specify must be an integer greater than 0.

Clients cannot input a larger amount of activations for a particular agent.

- 9** In the Agent Administrator's E-Mail Address field, type the agent's server administrator e-mail address.

This is the From address included on any e-mails that are sent from the server.

- 10** In the Organization Name field, type the name of your organization.

This identifies the organization associated with this server.

- 11** In the Minimum Timer Resolution in the Minutes field, type the minimum timer resolution in minutes.

This must be an integer and must be greater than or equal to 5.

This limits the period that clients can input as the interval at which periodic timer agents can activate.

- 12** Click OK or click Reset to delete all field information you have entered.

Agent Information in the Configuration Files

There are several configuration files that govern how agent services operate. In general, you don't access these files, but you may need to know about them.

The system configuration file is mapped to the AGENT_SYSTEM.INI file in the OBJ.CONF file. This file defines your system and data directories. This file in turn points to the AGENT_STRING.INI file that contains the text strings that are used to create the HTML agent services forms.

The information that you enter through the Enterprise Web Server *servername* button > Agents & Search > Agent Management form is reflected in the AGENT.CONF file. See "Configuring Agent Services" on page 151 for more information.

The access control list data is configured in the MAGNUS.CONF file, which points to another file,

server_root\HTTPACL\GENERATED.HTTPS-yourserver.ACL

Recovering Agent Files

Agent information is stored in a set of database files on the server. It is possible for the data in these files to become inconsistent or corrupted. If this happens, you can use the command-line utilities provided with the Web server to salvage and recover agent information from the corrupted files.

The next sections explain how agent information is stored and how to recover this information if file corruption occurs.

How Agent Information Is Stored

Agent information is stored in this database file:

*server_root\https-yourserver\agents-db\agent file
prefix_base*

In order to provide quicker access to information in this file, the Web server also creates and uses the following three index files.

- ♦ *prefix_BASE*.IDX is an index to the main agent database. It provides the other index files with quick access to the agent database.
- ♦ *prefix_USER*.IDX provides the server with a quick way to find agents based on username. Essentially, this file contains usernames and pointers to locations in the main database file. Given a username, the server can use this index file to look up that user's agents in the main database file.
- ♦ *prefix_CLASS*.IDX provides the server with a quick way to find agents based on class (or type). This file contains classes and pointers to locations in the main database file. Given a class, the server can use this index file to look up agents of that class in the main database file.

Fixing Inconsistencies and File Corruption

Agent information files are in one of the following three states:

- ♦ Valid: Nothing is wrong with the files.
- ♦ Inconsistent: The index files are not synchronized with the main database file.

- ♦ Corrupted: The main database file has become corrupted.

If the files are in the inconsistent state, you need to run the `agent_repair` command-line utility to fix the index files. For details, see “Recovering from Inconsistencies” on page 155.

If the files are in the Corrupted state, you need to run the `agent_salvage` command-line utility to recover the data. For details, see “Recovering from File Corruption” on page 156.

Recovering from Inconsistencies

In some cases, the index files develop inconsistencies and the server cannot use them to find entries in the main database file (NS_AGENT_BASE). When this situation occurs, the message “Agents database has become internally inconsistent” is logged to the AGENT.LOG file, and the server displays the following error: “Agent store files are out sync.”

Because both index files consist of data that is already stored in the main database file, you can recover the index files from the data in the main database file.

To recover the index files, run the AGENT_REPAIR.EXE utility, which is located in the `server_root\NETSCAPE\SUITESPOT\PLUGINS\AGENTS\BIN` directory. Unlike corrupted files, inconsistent files are repaired in place, with the “new” files overwriting the original files.

To run the AGENT_REPAIR.EXE utility:

1. Shut down your Web server.
2. On the command line, enter

```
agent_repair filepathname/fileprefix
```

Filepathname is the pathname for the agent file directory. The default is `server_root\HTTPS-yourserver\AGENTS-DB`.

Fileprefix is the prefix that is common to the names of the database and index files. Typically, this prefix is NS_AGENT.

For example, you could execute a command similar to the following:

```
agent_repair ns_agent
```

Recovering from File Corruption

In some cases, the main database file (NS_AGENT_BASE) might get corrupted. When this situation occurs, the server displays the following error message: “Agent store files are corrupted.”

The message “Agents database has become corrupted” is logged to the AGENT.LOG file. You can change this message by changing the text in the AGENT_STRINGS.INI file.

If the database is corrupted, you need to recover the data from the main database file. Be aware that corruption in the main database file may result in data loss.

To recover the data, run the AGENT_SALVAGE.EXE utility, which is located in the *server_root\NETSCAPE\SUITESPOT\PLUGINS\AGENTS\BIN* directory. This utility retrieves data from the corrupted files and creates new files for the data. Unlike inconsistent files, corrupted files cannot be repaired in place, so new files are created in addition to the original files rather than overwriting them.

To run the AGENT_SALVAGE.EXE utility:

1. Shut down your Web server.
2. On the command line, enter

```
agent_salvage filepathname/fileprefix newfileprefix
```

Filepathname is the pathname for the agent file directory. The default is *server_root\HTTPS-yourserver\AGENT_DB*.

Fileprefix is the prefix that is common to the names of the database and index files. Typically, this prefix is NS_AGENT.

Newfileprefix is a prefix that you assign to the newly created files that contain the recovered data. Typically, this prefix is NS_AGENT_RECOVERED.

For example, suppose you run the following command:

```
agent_salvage ns_agent ns_agent_recovered
```

The AGENT_SALVAGE utility retrieves data from the following corrupted files:

NS_AGENT_BASE

NS_AGENT_USER.IDX

NS_AGENT_CLASS.IDX

The utility then creates the following new files in the same directory as the original files and saves the data to these files:

NS_AGENT_RECOVERED_BASE

NS_AGENT_RECOVERED_USER.IDX

NS_AGENT_RECOVERED_CLASS.IDX

The utility also displays the following message to the user console when it is finished: “Agents database has been repaired.”

Accessing Agent Services

To access the agent services user interface access the following URL:

<http://yourserver/agents>

Yourserver can be either the IP address or the server name that was designated for your server when Enterprise was installed.

The NetWare Enterprise Web Server Agent Services page appears. The lower left frame contains the links you need to create and view agents. The lower right frame describes agent services and will be used to display information about a specific agent once you have selected one.

Standard and Advanced Options

Most users perform a standard set of agent tasks, and the standard options provide a simplified interface for these tasks. If you require additional capabilities, you can access a more complex interface by clicking Advanced Options at the bottom of the left frame.

Creating Agents

You can create agents with a standard set of options or with additional (advanced) options. Once you create an agent, you are considered the owner of that agent and only you can view or modify it.

In order to be able to create an agent, the Web Publishing state, search state, and agent services must all be on. See “Configuring Agent Services” on page 151 for directions on how to turn these states and services on.

Creating Document Agents

Standard Agents

To create a document agent using the standard set of options:

- 1 Access the following URL:

`http://yourserver/agents`

Yourserver can be either the IP address or the server name that was designated to your server when Enterprise was installed.

The NetWare Enterprise Web Server Agent Services page appears. The lower-left frame contains the links you need to create and view agents. The lower-right frame describes agent services and will be used to display information about a specific agent once you have selected one.

- 2 Click Standard Options at the bottom of the left frame.

If the button is labeled Advanced Options, you are already at the standard set of options.

- 3 Click Document Agent in the New Agent frame.

The New Document Agent form appears, with instructions to enter information as requested for each of the steps. See “Using the Standard Options” on page 158 for more information on these steps.

Advanced Agents

To create an agent with the advanced set of options:

- 1 Click Advanced Option at the bottom of the left frame.

If it is labeled Standard Options you are already at the advanced set of options.

- 2 Click Document Agent in the New Agent frame.

The New Document Agent form appears, with instructions to enter information as requested for each of the steps. See “Using the Advanced Options” on page 159 for more information on these steps.

Using the Standard Options

To create a document agent using the standard set of options:

- 1 Type the name you want to give the agent.

Select a name that helps you remember its function. Only 17 characters are visible at a time.

2 Type in the URI of the document (file) you want the agent to monitor.

Include a slash before the filename so the server can locate it correctly, such as /INDEX.HTML.

3 From the drop-down list select the Web Publisher event you want to activate the agent from.

If the document is modified or otherwise manipulated through some other file management program, the agent is not activated.

These are the actions you can monitor:

- ◆ Modify: The document is published.
- ◆ View: The document is viewed.
- ◆ Delete: The document is deleted.
- ◆ Move: The document is moved or renamed.
- ◆ Copy: The document is copied.
- ◆ View Document Attributes: The document's attributes are viewed.
- ◆ Modify Document Attributes: The document's attributes are modified.
- ◆ Document Locked: The document is locked.
- ◆ Document Unlocked: The document is unlocked.

4 Type the e-mail address of the users or newsgroup to notify when the event occurs for the document. You can enter more than one e-mail address or newsgroup, separated by commas.

5 Type any user-specified message contents you want added to the e-mail.

6 Type the e-mail address of the person or groups to notify when the agent is created.

You can enter more than one e-mail address, separated by commas, but only the first address is used in the standard Reply To field.

7 Click Create Agent to create the agent or Clear to clear the fields.

Using the Advanced Options

To create a document agent by using the advanced set of options:

1 Type the name you want to give the agent.

Select a name that helps you remember its function. Only the first 17 characters are visible at a time.

2 Type in the URI of the document (file) you want the agent to monitor.

Include a slash before the filename so the server can locate it correctly, such as /INDEX.HTML.

3 From the scrollable list, select the Web Publisher event you want to activate the agent from.

If the document is modified or otherwise manipulated through some other file management program, the agent is not activated. The actions you can monitor are listed here.

- ◆ Modify: The document is published.
- ◆ View: The document is viewed.
- ◆ Delete: The document is deleted.
- ◆ Move: The document is moved or renamed.
- ◆ Copy: The document is copied.
- ◆ View Document Attributes: The document's attributes are viewed.
- ◆ Modify Document Attributes: The document's attributes are modified.
- ◆ Document Locked: The document is locked.
- ◆ Document Unlocked: The document is unlocked.

4 Type the e-mail address of the users or newsgroups to notify when the event occurs for the document. You can enter more than one e-mail address or newsgroup, separated by commas.

5 You can define your message contents in the following ways:

- ◆ Type a user-specified message to be included as part of the notification.
- ◆ Identify a URL whose contents are to be included as part of the notification.
- ◆ Identify a URL to which you want to perform an HTTP post.
- ◆ You can define the URL-encoded message you want included as part of the HTTP post.

- 6** Type the e-mail address of the person or groups to notify when the agent is created.

You can enter more than one e-mail address, separated by commas, but only the first address is used in the standard Reply To field.

- 7** Change the maximum activations for an agent or its expiration date, as necessary.

You cannot extend beyond the limits set by your server administrator.

- 8** Click Create Agent to create the agent or Clear to clear the fields.

If you are using HTTP post operations, you may find it easiest to use some of the Agent API code provided as samples in the default installation of NetWare® Enterprise Web Server. You can access the sample code files in the /PLUGINS/AGENTS/EXAMPLES directory for your server. The API code builds an HTTP post method for you that you can revise and use as the post message contents. For example, the following code appears in the AGENTAPI.CPP file:

```
...
    if (urltopost) {
        //build the post message
        sprintf(header, "Content-Length:%d\n\
Content-Type: application/x-www-form-urlencoded'\
charset=US-ASCII\n\
Content-Transfer-Encoding: 7-bit",strlen(postmessage));
        sprintf(post_msg, "%s\n\n%s\n", header, postmessage);
    }...
```

Creating Directory Agents

To create a directory agent using the standard set of options, obtain the New Directory Agent form by accessing the following URL:

<http://yourserver/agents>

Yourserver can be either the IP address or the server name that was designated for your server when Enterprise was installed.

The NetWare Enterprise Web Server Agent Services page appears. The lower left frame contains the links you need to create and view agents. The lower right frame describes agent services and will be used to display information about a specific agent once you have selected one.

To create an agent with the standard set of options:

- 1 In the Agent Services page, click Standard Options at the bottom of the left frame.

If the button is labeled Advanced Options, you are already at the standard set of options.

- 2 Click Directory Agent in the New Agent frame.

The New Directory Agent form appears, with instructions to enter information as requested for each of the steps. See “Using the Standard Options” on page 158 for more information on these steps.

To create an agent with the advanced set of options:

- 1 In the Agent Services page click Advanced Options at the bottom of the left frame.

If the button is labeled Standard Options, you are already at the advanced set of options.

- 2 Click Directory Agent in the New Agent frame.

The New Directory Agent form appears, with instructions to enter information as requested for each of the steps. See “Using the Advanced Options” on page 159 for more information on these steps.

Using the Standard Options

To create a directory agent using the standard set of options:

- 1 Type the name you want to give the agent.

Select a name that helps you remember its function. Only 17 characters are visible at a time.

- 2 Type in the URI of the directory (folder) you want the agent to monitor.

Include a slash before the directory’s name so the server can locate it correctly, such as /PRODUCTS.

- 3 From the scrollable list, select the Web Publisher event you want to activate the agent from.

- ♦ Document Added/Deleted: A document in the agent directory is added or deleted.
- ♦ Directory Listed: The directory contents are listed, which occurs when a directory is opened in the Web Publisher file window.

- ♦ Delete: The directory is deleted.
- ♦ Move Directory: The directory is moved or renamed.
- ♦ Copy Directory: The directory is copied.
- ♦ View Directory Attributes: The directory's attributes are viewed.
- ♦ Modify Directory Attributes: The directory's attributes are modified.

If the directory is modified or otherwise manipulated through some other file management program, the agent is not activated. The actions you can monitor are described in "Using the Standard Options" on page 158.

- 4** Type the e-mail address of the users or newsgroup to notify when the event occurs. You can enter more than one e-mail address or newsgroup, separated by commas.
- 5** Type any user-specified message contents you want added to the e-mail sent as notification.
- 6** Type in the e-mail address of the person or groups to notify when the agent is created.

You can enter more than one e-mail address, separated by commas, but only the first address is used in the standard Reply To field.

- 7** Click Create Agent to create the agent or Clear to clear the fields.

Using the Advanced Options

To create a directory agent using the advanced set of options:

- 1** Type the name you want to give the agent.
Select a name that helps you remember its function. Only the first 17 characters are visible at a time.
- 2** Type the URI of the directory (folder) you want the agent to monitor.
Include a slash before the filename so the server can locate it correctly, such as /PRODUCTS.
- 3** Select the Web Publisher event you want to activate the agent from the drop-down list.
 - ♦ Document Added/Deleted: A document in the agent directory is added or deleted.
 - ♦ Directory Listed: The directory contents are listed, which occurs when a directory is opened in the Web Publisher file window.

- ♦ Delete: The directory is deleted.
- ♦ Move Directory: The directory is moved or renamed.
- ♦ Copy Directory: The directory is copied.
- ♦ View Directory Attributes: The directory's attributes are viewed.
- ♦ Modify Directory Attributes: The directory's attributes are modified.

If the document is modified or otherwise manipulated through some other file management program, the agent is not activated. The actions you can monitor are described in "Using the Standard Options" on page 158.

- 4** Type the e-mail address of the person or groups to notify when the event occurs for the document.

You can enter more than one e-mail address, separated by commas.

- 5** You can define your message contents in the following ways:

- ♦ Type a user-specified message to be included as part of the notification.
- ♦ Identify a URL whose contents are to be included as part of the notification.
- ♦ Identify a URL to which you want to perform an HTTP post.
- ♦ You can define the URL-encoded message you want included as part of the HTTP post.

- 6** Type the e-mail address of the person or groups to notify when the agent is created.

You can enter more than one e-mail address, separated by commas, but only the first address is used in the standard Reply To field.

- 7** Change the maximum activations for an agent or its expiration date.

You cannot extend beyond the limits set by your server administrator.

- 8** Click Create Agent to create the agent or Clear to clear the fields.

If you are using HTTP post operations, you may find it easiest to use some of the Agent API code provided as samples in the default installation of NetWare. You can access the sample code files in the /PLUGINS/AGENTS/EXAMPLES directory for your server. The API code builds an HTTP post method for you that you can revise and use as the post message contents. For example, the following code appears in the AGENTAPI.CPP file:

```

...
if (urltopost) {
//build the post message
sprintf(header, "Content-Length:%d\n\
Content-Type: application/x-www-form-urlencoded'\
charset=US-ASCII\n\
Content-Transfer-Encoding: 7-bit",strlen(postmessage));
sprintf(post_msg, "%s\n\n%s\n", header, postmessage);
}...

```

Creating Timer Agents

To create a timer agent with the standard set of options:

1 Access the following URL:

<http://yourserver/agents>

Yourserver can either be the IP address or server name that was designated to your server when Enterprise was installed.

The NetWare Enterprise Web Server Agent Services page appears. The lower-left frame contains the links you need to create and view agents. The lower-right frame describes agent services and will be used to display information about a specific agent once you have selected one.

2 Click Standard Options at the bottom of the left frame.

If the button is labeled Advanced Options, you are already at the standard set of options.

3 Click Timer Agent in the New Agent frame.

The New Timer Agent form appears, with instructions to enter information as requested for each of the steps.

To create an agent with the advanced set of options:

1 Click Advanced Options at the bottom of the left frame.

If the button is labeled Standard Options, you are already at the advanced set of options.

2 Click Timer Agent in the New Agent frame.

The New Timer Agent form appears, with instructions to enter information as requested for each of the steps.

Using the Standard Options

To create a timer agent using the standard set of options:

- 1** Type the name you want to give the agent.

Select a name that helps you remember its function. Only 17 characters are visible at a time.

- 2** Indicate when you want the agent to activate by making a selection and providing its associated data:

- ◆ **Once or Now:** Make the agent active as soon as it is created.

If you do not check Now as the activation time, type the exact date and time when you want the agent to activate.

- ◆ **Every:** Activate periodically on a regular basis.

Type a number and select the periodic interval from the drop-down list: minutes, hours, days, weeks, months, or years.

You can specify the beginning time for the agent's activation period, or check Now to make the agent active as soon as it is created.

Specify the time when the agent expires. You cannot create an agent that extends beyond your server's default expiration date for agents.

- ◆ **On the Same Day of the Week:** Specify the day or days on which the agent is to execute.

You can specify the beginning time for the agent's activation period, or check Now to make the agent active as soon as it is created.

Specify the time when the agent expires. You cannot create an agent that extends beyond your server's default expiration date for agents.

- 3** Type the e-mail address of the user or newsgroup to notify when the event occurs for the document.

You can enter more than one e-mail address or newsgroup, separated by commas.

- 4** Type any user-specified message contents you want added to the e-mail sent as notification.

- 5** Type the e-mail address of the person or groups to notify when the agent is created.

You can type more than one e-mail address, separated by commas, but only the first address is used in the standard Reply To field.

- 6 Click Create Agent to create the agent or Clear to clear the fields.

Using the Advanced Options

To create a timer agent using the advanced set of options:

- 1 Type in the name you want to give the agent.
Select a name that helps you remember its function. Only 17 characters are visible at a time.
- 2 Indicate when you want the agent to activate by making a selection and providing its associated data.
 - ◆ Once or Now: Make the agent active as soon as it is created.
If you do not check Now as the activation time, type the exact date and time when you want the agent to activate.
 - ◆ Every: Activate periodically on a regular basis.
Type a number and select the periodic interval from the drop-down list: minutes, hours, days, weeks, months, or years.
You can specify the beginning time for the agent's activation period, or check Now to make the agent active as soon as it is created.
Specify the time when the agent expires. You cannot create an agent that extends beyond your server's default expiration date for agents.
 - ◆ On the Same Day of the Week: Specify a day or days on which the agent is to execute.
You can specify the beginning time for the agent's activation period, or check Now to make the agent active as soon as it is created.
Specify the time when the agent expires. You cannot create an agent that extends beyond your server's default expiration date for agents.
- 3 Type the e-mail address of the user or newsgroup to notify when the event occurs for the document.
You can enter more than one e-mail address or newsgroup, separated by commas.
You can define your message contents in these ways:
 - ◆ Type in a user-specified message to be included as part of the notification.

- ♦ Identify a URL whose contents are to be included as part of the notification.
 - ♦ Identify a URL to which you want to perform an HTTP post.
 - ♦ You can define the URL-encoded message you want included as part of the HTTP post.
- 4** Type the e-mail address of the person or groups to notify when the agent is created.
- You can enter more than one e-mail address, separated by commas, but only the first address is used in the standard Reply To field.
- 5** Change the maximum activations for an agent.
- You cannot extend beyond the limit set by your server administrator.
- 6** Click Create Agent to create the agent or Clear to clear the fields.

If you are using HTTP post operations, you may find it easiest to use some of the Agent API code provided as samples in the default installation of the NetWare Enterprise Web Server. You can access the sample code files in the /PLUGINS/AGENTS/EXAMPLES directory for your server. The API code builds an HTTP post method for you that you can revise and use as the post message contents. For example, the following code appears in the AGENTAPI.CPP file.

```
...
    if (urltopost) {
        //build the post message
        sprintf(header, "Content-Length:%d\n\
Content-Type: application/x-www-form-urlencoded'\
charset=US-ASCII\n\
Content-Transfer-Encoding: 7-bit",strlen(postmessage));
        sprintf(post_msg, "%s\n\n%s\n", header, postmessage);
    }...
```

Creating Search Agents

To create a search agent using the standard set of options:

- 1** Access the following URL:

`http://yourserver/agents`

Yourserver can be either the IP address or server name that was designated to your server when Enterprise was installed.

The NetWare Enterprise Web Server Agent Services page appears. The lower left-frame contains the links you need to create and view agents. The lower-right frame describes agent services and will be used to display information about a specific agent once you have selected one.

2 Click Standard Options at the bottom of the left frame.

If it is labeled Advanced Options, you are already at the standard set of options.

3 Click Search Agent in the New Agent frame.

The New Search Agent form appears, with instructions to enter information as requested for each of the steps. See “Using the Standard Options” on page 158 for more information on these steps.

To create an agent with the advanced set of options:

1 Click Advanced Options at the bottom of the left frame.

If it is labeled Standard Options, you are already at the advanced set of options.

2 Click Search Agent in the New Agent frame.

The New Search Agent form appears, with instructions to enter information as requested for each of the steps. This is a Java* applet that provides additional searching capabilities. See “Using the Advanced Options” on page 159 for more information on these steps

Java must be enabled for Search to work. Click NetWare Enterprise Web Server *servername* button > Programs > Server-Side JavaScript > Yes > OK to enable Java.

Using the Standard Options

To create a search agent using the standard set of options:

1 Type the name you want to give the agent.

Select a name that helps you remember its function. Only 17 characters are visible at a time.

2 Indicate when you want the agent to activate by making a selection and providing its associated data. See “Using the Standard Options” on page 158 for more details on the options you can type into the field.

- ♦ Once or Now: Make the agent active as soon as it is created.

If you do not check Now as the activation time, type the exact date and time when you want the agent to activate.

- ♦ Every: Activate periodically on a regular basis.

Type a number and select the periodic interval from the drop-down list: minutes, hours, days, weeks, months, or years.

You can specify the beginning time for the agent's activation period, or check Now to make the agent active as soon as it is created.

Specify the time when the agents expires. You cannot create an agent that extends beyond your server's default expiration date for agents.

- ♦ On the Same Day of the Week:

Specify the day or days on which the agent is to execute.

You can specify the beginning time for the agent's activation period, or check Now to make the agent active as soon as it is created.

Specify the time when the agent expires. You cannot create an agent that extends beyond your server's default expiration date for agents.

- 3** In the When Activated field, type the e-mail address of the user or newsgroups to notify when the event occurs for the document.

You can enter more than one e-mail address or newsgroup, separated by commas.

- 4** You can define your message in the Message Content field in these ways:

- ♦ Select the scope of your search by clicking All Documents or Document Updated Since Last Activation.
- ♦ Select from the drop-down list the collection you want to search. The default is the Web Publishing collection. See "Using Search" on page 115 for more information about performing searches.
- ♦ In the Search For field, type the word or phrase you want to search for. Type an asterisk (*) to collect all documents that have been updated after the last search
- ♦ Type a user-specified message to be included as part of the notification.

- 5** Type the e-mail address of the person or groups to notify when the agent is created.

You can enter more than one e-mail address, separated by commas, but only the first address is used in the standard Reply To field.

- 6** Click Create Agent to create the agent or Clear to clear the fields.

Using the Advanced Options

When you choose to create an advanced search agent, you do so in a Java applet that guides you through making a search query

To use the applet, you must have Java enabled for your browser. See “Extending Your Server with Programs” on page 83 for more information about Java. Also, much of the search functionality is discussed at length in “Using Search” on page 115. The search agent uses the same query language, with the same set of rules and restrictions.

To create a search agent using the advanced set of options:

- 1** Type the name you want to give the agent.

Select a name that helps you remember its function. Only 17 characters are visible at a time.

- 2** Indicate when you want the agent to activate by making a selection and providing its associated data.

See “Using the Standard Options” on page 158 for more details on the options you can type into the field.

- ◆ Once or Now: Make the agent active as soon as it is created.

If you do not check Now as the activation time, type the exact date and time when you want the agent to activate.

- ◆ Every: Activate periodically on a regular basis.

Type a number and select the periodic interval from the drop-down list: minutes, hours, days, weeks, months, or years.

You can specify the beginning time for the agent’s activation period, or check Now to make the agent active as soon as it is created.

Specify the time when the agent expires. You cannot create an agent that extends beyond your server’s default expiration date for agents.

- ◆ On the Same Day of the Week: Specify a day or days on which the agent is to execute.

You can specify the beginning time for the agent’s activation period, or check Now to make the agent active as soon as it is created.

Specify the time when the agent expires. You cannot create an agent that extends beyond your server’s default expiration date for agents.

- 3 Type the e-mail address of the user or newsgroup to notify when the event occurs for the document.

You can enter more than one e-mail address or newsgroup, separated by commas.

- 4 You can define your message contents in these ways:

Java Search

You can use the embedded Java applet to guide you through constructing a search query. This is especially useful if you want to build a query that has several parts, such as searching for a word in the document's content as well as a specific attribute value. For more information on Java search, see "Guided Search" on page 133

To use Java search:

- 1 Type the e-mail address of the person or groups to notify when the agent is created.

You can enter more than one e-mail address, separated by commas, but only the first address is used in the standard Reply To field.

- 2 Change the maximum activations for an agent or its expiration date as necessary. (You cannot extend beyond the limits set by your server administrator.)

- 3 Click Create Agent to create the agent or Clear to clear the fields.

If you are using HTTP post operations, you may find it easiest to use some of the Agent API code provided as samples in the default installation of the NetWare Enterprise Web Server. You can access the sample code files in the /PLUGINS/AGENTS/EXAMPLES directory for your server. The API code builds an HTTP post method for you that you can revise and use as the post message contents. For example, the following code appears in the AGENTAPI.CPP file:

```
...
    if (urltopost) {
        //build the post message
        sprintf(header, "Content-Length:%d\n\
Content-Type: application/x-www-form-urlencoded'\
charset=US-ASCII\n\
Content-Transfer-Encoding: 7-bit",strlen(postmessage));
        sprintf(post_msg, "%s\n\n%s\n", header, postmessage);
    }...
```

Viewing and Managing Agents

You view only those agents that you have created. Once you are viewing an agent, you can perform some additional management tasks for it, such as modifying certain fields, deleting the agent, or disabling and re-enabling it.

To view an agent:

- 1 Access the following URL:

`http://yourserver/agents`

Yourserver can be either the IP address or the server name that was designated for your server when Enterprise was installed.

The NetWare Enterprise Web Server Agent Services page appears. The lower-left frame contains the links you need to create and view agents. The lower-right frame describes agent services and will be used to display information about a specific agent once you have selected one.

- 2 In the left frame, select an agent from the drop-down View Agent list. All of your agents are listed.
- 3 Click View.

The agent information is displayed in the right frame. This data differs according to the type of agent you are viewing. You can now modify, delete, or disable the agent.

Modifying an Agent

You can modify most of the agent options. For document and directory agents, you cannot change the document (or directory) and event that the agent is monitoring. For timer and search agents, you cannot change the timing of the agent's activation. If you want to change these options, you must delete the agent and create a new agent.

To modify an agent:

- 1 Access the following URL:

`http://yourserver/agents`

Yourserver can either be the IP address or server name that was designated to your server when Enterprise was installed.

The NetWare Enterprise Web Server Agent Services page appears. The lower-left frame contains the links you need to create and view agents.

The lower-right frame describes agent services and will be used to display information about a specific agent once you have selected one.

- 2** In the left frame, select an agent from the drop-down View Agent list. All of your agents are listed.

- 3** Click View.

You can modify the message contents, the maximum activations allowed for the agent, and, for document and directory agents, the agent's expiration date.

- 4** Click Modify Agent at the bottom of the frame.

Deleting an Agent

To delete an agent:

- 1** In the left frame of the Enterprise Web Server Agent Services page, select an agent from the drop-down View Agent list.
- 2** Click View.
- 3** Click Delete Agent at the bottom of the frame.
- 4** Click OK to delete the agent.

Disabling an Agent

To disable an agent:

- 1** In the left frame of the Enterprise Web Server Agent Services page, select an agent from the drop-down View Agent list.
- 2** Click View.
- 3** Click Disable Agent at the bottom of the frame.

This button changes to Enable Agent after the agent has been disabled.

Enabling a Disabled Agent

To enable an agent:

- 1** In the left frame of the Enterprise Web Server Agent Services page, select an agent from the View Agent drop-down list.
- 2** Click View.
- 3** Click Enable Agent at the bottom of the frame.

This button changes to Disable Agent after the agent has been disabled.

Managing Your Agents as a Group

You can enable, disable, or delete all of your agents at once. These actions only apply to agents that you created.

To manage all of your agents as a group:

- 1** Access the following URL:

`http://yourserver/agents`

Yourserver can be either the IP address or server name that was designated to your server when Enterprise was installed.

The NetWare Enterprise Web Server Agent Services page appears. The lower-left frame contains the links you need to create and view agents. The lower-right frame describes agent services and will be used to display information about a specific agent once you have selected one.

- 2** In the left frame, select an agent from the View Agent drop-down list.

- 3** Click Advanced Options at the bottom of the left frame.

If the button is labeled Standard Options, you are already at the advanced set of options.

- 4** To enable all of your agents simultaneously, click Enable All.

- 5** To disable all of your agents simultaneously, click Disable All.

- 6** To delete all of your agents simultaneously, click Delete All.

- 7** Click OK to delete all your agents.

10

Configuring Web Publishing

NetWare[®] Enterprise Web Server clients can use the Netscape* Web Publisher to collaborate on projects by directly accessing, editing, and managing files on remote servers. Web Publisher provides sophisticated features for server clients, such as file management, editing and publishing, document version control, search, agent services, access control, and link management.

As the server administrator, you can set many options that define how Web Publishing works for your server clients and how your server's Web Publishing data is maintained. One of the most important functions you can perform for your users is to create a database of searchable Web publishing data. This requires using the Index and Update Properties function to index a set of documents and directories so that when users start up Web Publisher, they can search on the contents and properties of these files.

Other Web Publishing setup and configuration functions include the following:

- ♦ Indexing and updating properties
- ♦ Setting the language that Web Publisher uses
- ♦ Turning off the link management component of Web publishing
- ♦ Setting the archive directory for version-controlled files
- ♦ Unlocking files that a client may have locked, thereby making them available again to other users
- ♦ Adding and managing custom Web publishing file properties
- ♦ Maintaining Web publishing data
- ♦ Turning Web publishing on and off for your server

For further information about Netscape Web Publisher, access the online help, by clicking the Help menu in Web Publisher, or you can click the Help button on one of the search interface forms, on the Agent Services page, or on the Web Publisher Services page.

Setting Access Control for an Owner

The access control (ACL) system supports a special user called owner. When an access control rule designates the user to be the owner, the permissions defined by this rule apply to the owner assigned by Web Publisher for each document. For example,

```
allow (write, delete) user = owner;
```

IMPORTANT: Do not create a user with the username of *owner*.

Ownership of Web publishing documents can be assigned either through the Enterprise Web Server *servername* > Web Publishing > Index and Update Properties form, or through Web Publisher. The Index and Update Properties form allows you to create a bulk assignment of ownership to a set of documents. Web Publisher performs individual assignments of file ownership to a user when the user publishes or uploads the file.

Only the owner can modify the access control rules for a file. These rules define the actions users can perform on the file, such as moving, copying, renaming, or deleting. An owner can reassign ownership of a file to another user, and if a file has no owner, anyone with a valid username can identify themselves as its owner. Because the username identified as the owner of a file can change, any access control that you place on a file should target the owner of a file rather than a specific username.

If the default ACL that governs your server is not restrictive or flexible enough for your Web publishing needs, you can use the Enterprise Web Server *servername* > Server Preferences > Restrict Access function to create an ACL that is more appropriate for Web publishing.

For example, you could create an ACL like the following:

```
acl "uri=/publisher/";  
    allow (read, execute, list, info) user = anyone;  
    allow (write, delete) user = owner;
```

This ACL sets a restriction such that only the owner of a file within the additional document directory of /PUBLISHER can modify or delete the file.

See “Controlling Access to Your Server” on page 53 for more information about setting access control.

Turning the Web Publishing On or Off

You can deactivate Web Publishing and you can turn it back on. If you turn off Web Publishing, you also turn off link management. Documents that are subsequently moved or renamed may have incorrect links, and the link status database may not be up to date. The solution is to use the Web Publishing > Link Management function to manually turn link management off and then turn it back on again. This starts the link management function again with an empty link status database. See “Changing the Link Management State” on page 183 for more details of link management.

If you turn Web Publishing off, all agents for the server are also turned off and clients cannot use Netscape Web Publisher to access agent services. When Web Publishing is turned back on, agents that were turned off solely because Web Publishing was turned off are turned back on. Agents that were disabled for other reasons are still disabled.

To turn the Web Publishing on or off:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Web Publishing > Web Publishing State.
- 2** To turn Web Publishing on, select On.
or
To turn it off, select Off.
The default value is Off.
- 3** Click OK.

Turning WebDAV On or Off

WebDAV stands for Web Distributed Authoring and Versioning. The WebDAV protocol makes it possible for Web users to write, edit and save shared documents without overwriting each others' work.

If you turn Web Publishing on WebDAV will be turned off. If you turn WebDAV on Web Publishing will be turned off.

IMPORTANT: You must be using NDS® as your directory service and Internet Explorer* as your browser in order for WebDAV to work.

For information on using WebDAV, refer to User Solutions for NetWare 5.1.

To turn WebDAV on or off:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Web Publishing > Web Publishing State.
- 2** To turn WebDAV on, select On.
or
To turn it off, select Off.
The default value is Off.
- 3** Click OK.

Enabling or Disabling My Network and NDSDAV

You must be using NDS as your directory service and WebDAV must be set to On in order for My Network and NDSDAV to work. These are the default settings on a new installation.

- 1** Using a text editor, open the configuration file
SYS:NOVONYX\SUITESPOT\BIN\WEBDAV\HTML\WEBDAV.
CONF.
- 2** To enable the My Network feature, edit the first line in the file to read
NDSDAV:True.
Or:
To disable the My Network feature, edit the first line to read
NDSDAV:False.
The default is False.
- 3** From the General Administration page, click the Enterprise Web Server *servername* button > Server Preferences > On/Off, to stop and restart the server so that your changes take effect.

Once NDSDAV is enabled, users can open the [https://servername/My Network](https://servername/MyNetwork) folder to browse their mapped drives, NDS Users and Groups, and User Home Directory using Web Folders or any other WebDAV enabled client. For more information, refer to the README contained in the My Network folder.

Setting the Web Publishing Language

You can change the Web Publishing language to any language supported by the user's installation. The different available languages are listed for the server administrator in a drop-down list on the form.

IMPORTANT: Be cautious when using this function. If you change the language of a collection, the system deletes all the existing data in the collection.

- 1 From the General Web Publishing > Web Publishing Language.
- 2 Select a language from the drop-down list. The default is English.
- 3 Click OK.

Your changes are reflected in the LANGUAGE.CONF file, which is located in the directory.

After you change the Web Publishing language, your server is automatically restarted to apply the change.

Maintaining Web Publishing Data

Web Publisher maintains multiple sets of data about the documents that are in the Web publishing collection. When all Web publishing data is synchronized, each file in the chosen document directory has a record in the Web publishing collection and each property record in the collection has a corresponding file in the document directory.

Although you can limit the scope of the Repair and Report functions to checking only the files in a particular document directory for collection records, each property record in the collection is checked for a corresponding source document, regardless of which directory the file might be in.

Occasionally, files can become unsynchronized. You can obtain a report on the state of your Web publishing files and then repair one or more directories as needed. For example, if a document that was indexed into a collection is deleted, there is a record in the collection that no longer has any corresponding source document. Repairing removes the collection records for these documents.

You can perform the following functions to maintain your Web publishing data:

- ♦ **Report:** You can produce a report on the current logical consistency of the Web publishing collection's data. This lists all the files in the selected

document directory and also lists all the records in the Web publishing collection, regardless of the directory the collection data corresponds to. The report indicates which files are not yet indexed (and therefore don't have records in the Web publishing collection) and which records have no source document (and therefore should be repaired). The report highlights errors and indicates what the result of the repair would be, for example, "Repair will delete properties record."

The report provides a short summary at the end of the log file, indicating how many directories and files have been checked, how many repairs are recommended, and how many errors have been encountered.

- ♦ **Repair:** You can repair the Web publishing collection's logical consistency. This function repairs the files in the selected document directory and produces a report similar to the one from the Report function. The report indicates which repairs have been completed and what the repair accomplished, for example, "Repair: removing properties record."
- ♦ **Optimize:** You can optimize the Web publishing collection to improve performance if you frequently add, delete, or update documents or directories in your collections. Optimizing is also done automatically when you reindex or update a collection, so this function is not often necessary. One situation when you might want to optimize a collection is just before Publishing it to another site or before putting it onto a read-only CD-ROM.

To access maintenance functions:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Web Publishing > Maintain Web Publishing Data.
- 2** Define the scope of the Repair and Report functions by selecting the document directory to check.

If you want to use a different directory, click View to see a list of directories. You can report on or repair any directory or subdirectory that is listed.

Click Back to return to the Maintain Web Publishing Data form. The directory name appears in the Document directory field.

- 3** To report on, repair, or optimize the subdirectories within the specified directory, check Include Subdirectories.
- 4** Click one of the following functions:
 - ♦ Report

- ♦ Repair
- ♦ Optimize

Changing the Link Management State

If you do not always need automatic link checking and updating, you can turn link management off to conserve resources and to improve searching and indexing performance. When you turn link management off, Web Publisher stops doing automatic link checking and you cannot use the Check Links function from the Web Publisher Services page.

You can also selectively turn the automatic link update feature on and off. When automatic link updating is on, Web Publisher changes the outgoing and incoming links in a file to keep them up to date as files are moved and renamed in Web Publisher. Because this revises the modification date for any file that has updated links, this feature is off by default.

The automatic link update feature effects only links outgoing to or incoming from moved or renamed files. It does not affect HTML files that are being uploaded or published. Provided that link management is on, the links in these files are always updated as part of the upload or publish operation.

Use a Netscape browser to access the Web Publisher and type:

`http://yourserver/publisher`

To change the link management state:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Web Publishing > Link Management.
- 2** To change the state of link management, select On or Off.

To deactivate link management, select Off. This clears the link status information so that when you try to check links in Web Publisher, you get an error message, and you cannot access any link status information.

To reactivate link management, select the On option. This starts link management up again, which creates a new empty link status database. To get link status information, you must again check links for all your files. Links that have changed status because you turned link management off may have to be manually fixed.
- 3** To turn on automatic link updating, select On. You can only turn this on when link management is on.

This starts automatic link updating, which revises links from or to files that are subsequently moved or renamed. It does not, however, affect the links in any files that were moved or renamed after automatic link updating was turned off.

- 4 Click OK.

Setting the Version Control Archive

The Web Publisher includes a version control system for keeping track of files and documents as they are updated and changed. Web Publisher manages version control, allowing you to compare different versions of a file, check version history for any file under version control, and use automatically incrementing version numbers for files edited under version control.

Files under version control are stored in an archive directory.

To specify which directory you want Web Publisher to use as the version control archive directory:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Web Publishing > Version Control.
- 2 Type the full path for the archive directory in the Archive Path field.

Web Publisher uses this archive to store all files under version control.

IMPORTANT: If you are changing the archive directory, but keeping the version history intact, you must have (a) already created the new directory, (b) moved the version history files to the new directory, and (c) deleted the old archive directory. If you don't want to keep the old version history, you don't need to move the files to the new directory, but you must do Steps a and b or this function will fail.

- 3 Click OK.

Unlocking Files

If a file that has been locked in Web Publisher is required by another user, you can unlock it. This is true for files that were locked manually by the client or automatically by Web Publisher as part of an edit or download operation.

Be cautious in using this function because if you unlock a file that was locked, you are allowing the file to be available for editing by other users.

To unlock a file:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Web Publishing > Unlock File.

The Choose field displays the currently selected file or directory.

- 2 To unlock a different file or a file from another directory, click View to see a list of resources, then browse to select another file.
- 3 Click the unlock link for a file to select it and return to the Unlock File form.

The filename appears in the Choose field.

- 4 Click OK.

After you unlock a file, your server is automatically restarted to incorporate the lock change.

If you want to unlock a file that begins with a period, as in .JSHRC, you cannot use this form to perform the unlocking. You must log in to Web Publisher as the user and unlock the file there.

Adding Custom Properties

As server administrator, you can add your own custom Web Publisher file properties. These properties are added to the default set of file properties stored in the Web publishing collection. Server clients can view visible custom properties in Web publisher and use them in their document searches.

If you want to add another custom property after creating the maximum number of custom properties for a given type, you cannot remove an existing custom property and re-use the property's slot in the collection by adding a new custom property of the same type. For example, if you want to add a numeric property after five have already been created, you cannot delete one of the existing five numeric properties and add another numeric property in its place. The only way to use the new property is to remove the entire collection and recreate it to include the new property.

To add a custom file property:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Web Publishing > Add Custom Property.
- 2 Type a name in the Property Name field.

The name has these restrictions:

- ♦ It cannot duplicate an existing Web Publisher property name
- ♦ It cannot exceed 128 characters
- ♦ It cannot contain spaces, quotes, or apostrophes
- ♦ It cannot begin with an underscore or have an underscore as the third-to-last character

3 Click the Property Type drop-down list > select the property's type.

This value is not modifiable. There is a limit to the number of each type you can have. These are the default settings:

- ♦ Text (a maximum of 30)
- ♦ Numeric (a maximum of 5)
- ♦ Date (a maximum of 5). Dates are formatted as month/day/year, and year can be two or four digits.

You can change the maximum settings for these property types in the WEBPUB.CONF file, although larger sets of attributes impact the performance of your server.

You cannot use the additional attributes in the existing Web publishing collection. If you want to use the new attributes in the Web publishing collection, you must use your file system to remove both the WEB_HTML and LINK_MGR collection files from the search collections directory and then restart your server. See "Configuring Manually" on page 121 for details on how to change the WEBPUB.CONF file.

4 Specify a permission by clicking either Read Only or Modifiable.

The default is Modifiable.

For modifiable custom properties defined as META-tagged attributes, the value in the document is extracted only the first time the document is indexed. Because users can input a different value in the attribute field through the Web Publisher Services Properties page, the server ignores the META-tagged value in subsequent indexing. In this way, the user's value is not overwritten.

5 Click one of the Visible to User buttons, either Invisible or Visible.

The default is set to Visible. This defines whether server clients can view the property through Web Publisher.

- 6 If the property you are adding is actually an HTML file attribute that has been tagged with the HTML META tag, you can check this check box.

From this point onward, when files containing this attribute are indexed, the contents of the META attribute is used as the value of the property, and you can search for files that contain this META-tagged property. The property must conform to the same conventions as property names

Managing Custom Properties

You can list all the file properties that are available for use. These include the default set plus any new custom properties you have created. You can remove or edit only those properties that you have created. These have active Remove and Edit links in the first two columns.

To remove a custom property:

- 1 From the General Administration page, click the Enterprise Web Server *servername* button > Web Publishing > Manage Properties.
- 2 Click Remove for the property you want to remove.
- 3 Click OK to remove the property.

or

Click Back to return to the Manage Properties page without removing the property.

To edit a custom property:

- 1 Click Edit for the property you want to edit.
- 2 Change the property as needed.

You can change only the property's name, permissions, visibility, and its option of whether to capture META-tagged attributes.

- 3 Click OK to update the property with your changes.

or

Click Back to return to the Manage Properties page without editing the property.

or

Click Reset to reset any property values you changed.

Indexing and Updating Properties

Before users can perform a search across a set of documents and directories, information about the documents and directories needs to be indexed into the Web publishing database. The Web publishing database is stored as a search collection and is created as part of the server installation process. Initially it contains no data and must be populated by indexing the documents in the document directories.

When a user starts Web Publisher, its window lists the files and folders that are in the selected document directory, but the data initially is not indexed and therefore is not available for searching. In addition, the files have no owners so anyone can define their username as the owner of a file, and thereby be able to set the access control for a file.

You can use the Index and Update Properties form to perform bulk indexing of documents to create searchable Web publishing data and you can also use it to do a bulk assignment of ownership for the files included in the collection. You can restrict or expand the scope of documents and directories to be indexed, and you can index just the file properties, called metadata. You can also index the documents' contents. If you choose to index the contents of the files, you can search on any word in the documents, although publishing and uploading files with Web Publisher may be slightly slower.

To index and update your properties:

- 1** From the General Administration page, click the Enterprise Web Server *servername* button > Web Publishing > Index and Update Properties.

The Document directory field displays the currently selected directory. You can index documents in the primary document directory, an additional document directory, or in a subdirectory.

If you want to index a different directory, click View to see a list of directories. You can index any directory that is listed or you can view the subdirectories in a listed directory and index one of those instead.

- 2** Click the index link for a directory and return to the Index and Update Properties form.
- 3** Select Include Subdirectories to also index the subdirectories within the specified directory.
- 4** Specify the files to index.

You can index all files in the chosen directory by leaving the default *.* pattern in the Include Files Matching Pattern field, or you can define your

own wildcard expression to restrict indexing to documents that match that pattern. For example, you could type ***.HTML** to only index the content in documents with the .HTML extension, or you could use this pattern (complete with parentheses) to index all HTML documents:

(*.HTM|*.HTML)

You can define multiple wildcards in an expression. See “Managing Your Server” on page 15 for details of the syntax for wildcard patterns.

- 5** If this is the first time you have indexed Web publishing documents, check the Index Unindexed Documents option.

In subsequent indexing operations, you can uncheck the option or you can leave it checked to index any new documents that have been added to the document directory.

- 6** Decide if you want to change files that have already been indexed. If you want to make a change to files that have already been indexed, use the Update Previously Indexed Documents option to create a bulk ownership assignment or to index the content of files that did not have this option set when they were first indexed.

These options are useful when you change many files at once. You can use the Web Publisher client to index and update individual files.

- 7** Decide if you want to create a bulk assignment of ownership. To create a bulk assignment of ownership to all files that match your criteria, check the Set Document Owner To check box > type a username in the field.

Be sure to type in a valid username because the server does not perform any validity checks on the name. This updates the owner property in each file’s collection entry.

- 8** To index the document content, check the Index Document Contents check box.

You can choose to index the documents’ contents as well as their file metadata.

- 9** Click OK to begin indexing and updating Web publishing.

A summary of the indexing operation is displayed in the Web browser window. The information is also logged to a local log file.

Once you have indexed documents into the Web publishing collection, you should not change any document directory’s URL mapping or the collection’s entries will target the URL mapping to the wrong physical file location. If you must change a document directory, you need to

reindex the documents in the new location. You can use the Repair function to remove the indexed data from the old directory mapping.



Novell Trademarks

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

Advanced NetWare is a trademark of Novell, Inc.

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

AppNotes is a registered service mark of Novell, Inc. in the United States and other countries.

AppNotes is a registered service mark of Novell, Inc. in the United States and other countries.

AppTester is a registered service mark of Novell, Inc. in the United States and other countries.

BrainShare is a registered service mark of Novell, Inc. in the United States and other countries.

C-Worthy is a trademark of Novell, Inc.

C3PO is a trademark of Novell, Inc.

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

Certified NetWare Administrator in Japanese and CNA-J are service marks of Novell, Inc.

Certified NetWare Engineer in Japanese and CNE-J are service marks of Novell, Inc.

Certified NetWare Instructor in Japanese and CNI-J are service marks of Novell, Inc.

Certified Novell Administrator and CNA are service marks of Novell, Inc.

Certified Novell Engineer is a trademark and CNE is a registered service mark of Novell, Inc. in the United States and other countries.

Certified Novell Salesperson is a trademark of Novell, Inc.

Client 32 is a trademark of Novell, Inc.

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

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

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

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

Custom 3rd-Party Object and C3PO are trademarks of Novell, Inc.

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

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

ElectroText is a trademark of Novell, Inc.

Enterprise Certified Novell Engineer and ECNE are service marks of Novell, Inc.

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

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

EXOS is a trademark of Novell, Inc.

Global MHS is a trademark of Novell, Inc.

Global Network Operations Center and GNOC are service marks of Novell, Inc.

Graphics Environment Manager and GEM are registered trademarks of Novell, Inc. in the United States and other countries.

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

GroupWise XTD is a trademark of Novell, Inc.

Hardware Specific Module is a trademark of Novell, Inc.

Hot Fix is a trademark of Novell, Inc.

InForms is a trademark of Novell, Inc.

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

Internetwork Packet Exchange and IPX are trademarks of Novell, Inc.

IPX/SPX is a trademark of Novell, Inc.

IPXODI is a trademark of Novell, Inc.

IPXWAN is a trademark of Novell, Inc.

LAN WorkGroup is a trademark of Novell, Inc.

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

LAN WorkShop is a trademark of Novell, Inc.

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

LANalyzer Agent is a trademark of Novell, Inc.

Link Support Layer and LSL are trademarks of Novell, Inc.

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

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

Media Support Module and MSM are trademarks of Novell, Inc.

Mirrored Server Link and MSL are trademarks of Novell, Inc.

Mobile IPX is a trademark of Novell, Inc.

Multiple Link Interface and MLI are trademarks of Novell, Inc.

Multiple Link Interface Driver and MLID are trademarks of Novell, Inc.

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

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

Natural Language Interface for Help is a trademark of Novell, Inc.

NDS Manager is a trademark of Novell, Inc.

NE/2 is a trademark of Novell, Inc.

NE/2-32 is a trademark of Novell, Inc.

NE/2T is a trademark of Novell, Inc.

NE1000 is a trademark of Novell, Inc.

NE1500T is a trademark of Novell, Inc.

NE2000 is a trademark of Novell, Inc.

NE2000T is a trademark of Novell, Inc.

NE2100 is a trademark of Novell, Inc.

NE3200 is a trademark of Novell, Inc.

NE32HUB is a trademark of Novell, Inc.

NEST Autoroute is a trademark of Novell, Inc.

NetExplorer is a trademark of Novell, Inc.

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

NetSync is a trademark of Novell, Inc.

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

NetWare 3270 CUT Workstation is a trademark of Novell, Inc.

NetWare 3270 LAN Workstation is a trademark of Novell, Inc.

NetWare 386 is a trademark of Novell, Inc.

NetWare Access Server is a trademark of Novell, Inc.

NetWare Access Services is a trademark of Novell, Inc.

NetWare Application Manager is a trademark of Novell, Inc.

NetWare Application Notes is a trademark of Novell, Inc.

NetWare Asynchronous Communication Services and NACS are trademarks of Novell, Inc.

NetWare Asynchronous Services Interface and NASI are trademarks of Novell, Inc.

NetWare Aware is a trademark of Novell, Inc.

NetWare Basic MHS is a trademark of Novell, Inc.

NetWare BranchLink Router is a trademark of Novell, Inc.

NetWare Care is a trademark of Novell, Inc.

NetWare Communication Services Manager is a trademark of Novell, Inc.

NetWare Connect is a registered trademark of Novell, Inc. in the United States.

NetWare Core Protocol and NCP are trademarks of Novell, Inc.

NetWare Distributed Management Services is a trademark of Novell, Inc.

NetWare Document Management Services is a trademark of Novell, Inc.

NetWare DOS Requester and NDR are trademarks of Novell, Inc.

NetWare Enterprise Router is a trademark of Novell, Inc.

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

NetWare Global Messaging and NGM are trademarks of Novell, Inc.

NetWare Global MHS is a trademark of Novell, Inc.

NetWare HostPrint is a registered trademark of Novell, Inc. in the United States.

NetWare IPX Router is a trademark of Novell, Inc.

NetWare LANalyzer Agent is a trademark of Novell, Inc.

NetWare Link Services Protocol and NLSP are trademarks of Novell, Inc.

NetWare Link/ATM is a trademark of Novell, Inc.
NetWare Link/Frame Relay is a trademark of Novell, Inc.
NetWare Link/PPP is a trademark of Novell, Inc.
NetWare Link/X.25 is a trademark of Novell, Inc.
NetWare Loadable Module and NLM are trademarks of Novell, Inc.
NetWare LU6.2 is trademark of Novell, Inc.
NetWare Management Agent is a trademark of Novell, Inc.
NetWare Management System and NMS are trademarks of Novell, Inc.
NetWare Message Handling Service and NetWare MHS are trademarks of Novell, Inc.
NetWare MHS Mailslots is a registered trademark of Novell, Inc. in the United States and other countries.
NetWare Mirrored Server Link and NMSL are trademarks of Novell, Inc.
NetWare Mobile is a trademark of Novell, Inc.
NetWare Mobile IPX is a trademark of Novell, Inc.
NetWare MultiProtocol Router and NetWare MPR are trademarks of Novell, Inc.
NetWare MultiProtocol Router Plus is a trademark of Novell, Inc.
NetWare Name Service is trademark of Novell, Inc.
NetWare Navigator is a trademark of Novell, Inc.
NetWare Peripheral Architecture is a trademark of Novell, Inc.
NetWare Print Server is a trademark of Novell, Inc.
NetWare Ready is a trademark of Novell, Inc.
NetWare Requester is a trademark of Novell, Inc.
NetWare Runtime is a trademark of Novell, Inc.
NetWare RX-Net is a trademark of Novell, Inc.
NetWare SFT is a trademark of Novell, Inc.
NetWare SFT III is a trademark of Novell, Inc.
NetWare SNA Gateway is a trademark of Novell, Inc.
NetWare SNA Links is a trademark of Novell, Inc.
NetWare SQL is a trademark of Novell, Inc.
NetWare Storage Management Services and NetWare SMS are trademarks of Novell, Inc.
NetWare Telephony Services is a trademark of Novell, Inc.
NetWare Tools is a trademark of Novell, Inc.
NetWare UAM is a trademark of Novell, Inc.
NetWare WAN Links is a trademark of Novell, Inc.

NetWare/IP is a trademark of Novell, Inc.

NetWire is a registered service mark of Novell, Inc. in the United States and other countries.

Network Navigator is a registered trademark of Novell, Inc. in the United States.

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

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

Network Support Encyclopedia and NSE are trademarks of Novell, Inc.

Network Support Encyclopedia Professional Volume and NSEPro are trademarks of Novell, Inc.

NetWorld is a registered service mark of Novell, Inc. in the United States and other countries.

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

Novell Alliance Partners Program is a collective mark of Novell, Inc.

Novell Application Launcher is a trademark of Novell, Inc.

Novell Authorized CNE is a trademark and service mark of Novell, Inc.

Novell Authorized Education Center and NAEC are service marks of Novell, Inc.

Novell Authorized Partner is a service mark of Novell, Inc.

Novell Authorized Reseller is a service mark of Novell, Inc.

Novell Authorized Service Center and NASC are service marks of Novell, Inc.

Novell BorderManager is a trademark of Novell, Inc.

Novell BorderManager FastCache is a trademark of Novell, Inc.

Novell Client is a trademark of Novell, Inc.

Novell Corporate Symbol is a trademark of Novell, Inc.

Novell Customer Connections is a registered trademark of Novell, Inc. in the United States.

Novell Directory Services and NDS are registered trademarks of Novell, Inc. in the United States and other countries.

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

Novell ElectroText is a trademark of Novell, Inc.

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

Novell Gold Authorized Reseller is a service mark of Novell, Inc.

Novell Gold Partner is a service mark of Novell, Inc.

Novell Labs is a trademark of Novell, Inc.

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

Novell NE/2 is a trademark of Novell, Inc.

Novell NE/2-32 is a trademark of Novell, Inc.

Novell NE3200 is a trademark of Novell, Inc.

Novell Network Registry is a service mark of Novell, Inc.

Novell Platinum Partner is a service mark of Novell, Inc.

Novell Press is a trademark of Novell, Inc.

Novell Press Logo (teeth logo) is a registered trademark of Novell, Inc. in the United States and other countries.

Novell Replication Services is a trademark of Novell, Inc.

Novell Research Reports is a trademark of Novell, Inc.

Novell RX-Net/2 is a trademark of Novell, Inc.

Novell Service Partner is a trademark of Novell, Inc.

Novell Storage Services is a trademark of Novell, Inc.

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

Novell Technical Services and NTS are service marks of Novell, Inc.

Novell Technology Institute and NTI are registered service marks of Novell, Inc. in the United States and other countries.

Novell Virtual Terminal and NVT are trademarks of Novell, Inc.

Novell Web Server is a trademark of Novell, Inc.

Novell World Wide is a trademark of Novell, Inc.

NSE Online is a service mark of Novell, Inc.

NTR2000 is a trademark of Novell, Inc.

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

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

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

Open Data-Link Interface and ODI are trademarks of Novell, Inc.

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

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

Open Socket is a registered trademark of Novell, Inc. in the United States.

Packet Burst is a trademark of Novell, Inc.

PartnerNet is a registered service mark of Novell, Inc. in the United States and other countries.

PC Navigator is a trademark of Novell, Inc.

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

Perform3 is a trademark of Novell, Inc.

Personal NetWare is a trademark of Novell, Inc.

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

Portable NetWare is a trademark of Novell, Inc.

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

Print Managing Agent is a trademark of Novell, Inc.

Printer Agent is a trademark of Novell, Inc.

QuickFinder is a trademark of Novell, Inc.

Red Box is a trademark of Novell, Inc.

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

Remote Console is a trademark of Novell, Inc.

Remote MHS is a trademark of Novell, Inc.

RX-Net is a trademark of Novell, Inc.

RX-Net/2 is a trademark of Novell, Inc.

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

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

Sequenced Packet Exchange and SPX are trademarks of Novell, Inc.

Service Response System is a trademark of Novell, Inc.

Serving FTP is a trademark of Novell, Inc.

SFT is a trademark of Novell, Inc.

SFT III is a trademark of Novell, Inc.

SoftSolutions is a registered trademark of SoftSolutions Technology Corporation, a wholly owned subsidiary of Novell, Inc.

Software Transformation, Inc. is a registered trademark of Software

Transformation, Inc., a wholly owned subsidiary of Novell, Inc.

SPX/IPX is a trademark of Novell, Inc.

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

Storage Management Services and SMS are trademarks of Novell, Inc.

Technical Support Alliance and TSA are collective marks of Novell, Inc.

The Fastest Way to Find the Right Word is a registered trademark of Novell, Inc. in the United States and other countries.

The Novell Network Symbol is a trademark of Novell, Inc.

Topology Specific Module and TSM are trademarks of Novell, Inc.

Transaction Tracking System and TTS are trademarks of Novell, Inc.

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

Virtual Loadable Module and VLM are trademarks of Novell, Inc.

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

Yes, It Runs with NetWare (logo) is a trademark of Novell, Inc.

Yes, NetWare Tested and Approved (logo) is a trademark of Novell, Inc.

ZENworks is a trademark of Novell, Inc.

