Novell QuickFinder™ Server

5.0

www.novell.com

ADMINISTRATION GUIDE

September 2007

Novell®

Legal Notices

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

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

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

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

Novell, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed on the Novell Legal Patents Web page (http://www.novell.com/company/legal/patents/) and one or more additional patents or pending patent applications in the U.S. and in other countries.

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

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

Novell Trademarks

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

Third-Party Materials

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

Contents

About This Guide 11

1 Overview of QuickFinder Server 13

1.1 Benefits . . . 14
1.2 Features . . . 14
1.2.1 QuickFinder Administrators . . . 14
1.2.2 End Users . . . 15
1.3 Key Components of QuickFinder Server . . . 16
1.3.1 Virtual Search Servers . . . 16
1.3.2 Indexes . . . 17
1.3.3 Templates . . . 17
1.4 Managing QuickFinder Server . . . 17
1.4.1 Accessing QuickFinder Server Manager on NetWare . . . 18
1.4.2 Accessing QuickFinder Server Manager on Linux . . . 18
1.5 Taking a Test Run: Performing a Search . . . 19
1.6 Getting Help . . . 19
1.7 What's Next . . . 20

2 What's New in QuickFinder Server 5.0 21

2.1 New and Updated Features . . . 21
2.2 File Reader Updates . . . 21

3 Setting Up QuickFinder Server 23

3.1 Installing QuickFinder Server . . . 23
3.1.1 Installing QuickFinder Server From the NetWare 6.5 SP7 Media . . . 23
3.1.2 Installing QuickFinder Server From the OES for Linux Media . . . 24
3.1.3 QuickFinder Languages . . . 24
3.2 Designing Your Search Solution . . . 24
3.2.1 Components of a Virtual Search Server . . . 25
3.2.2 Deciding If You Need More Than One Virtual Search Server . . . 26
3.2.3 Becoming a Search Service Host . . . 26

4 Migrating QuickFinder Server from OES 1 NetWare to OES 2 Linux 27

4.1 Coexistence . . . 27
4.1.1 Compatibility . . . 27
4.1.2 Coexistence Issues . . . 27
4.2 Migration Procedure . . . 27
4.3 Post-Migration Considerations . . . 29

5 Running QuickFinder Server in a Virtualized Environment 31

6 Configuring QuickFinder Server for Novell Cluster Services 33

7 Configuring QuickFinder Server Default and Services Settings 35

7.1 Differences between Default and Services Settings . . . 36
7.1.1 Default Settings . . . 36
7.1.2 Services Settings . . . 36
7.2 Configuring Default Settings . . . 36
7.2.1 Configuring Default General Settings . . . 37
7.2.2 Configuring Default Search Settings . . . 40
7.2.3 Configuring Default Print Settings . . . 44
7.2.4 Configuring Default Index Settings . . . 46
7.2.5 Configuring Default Security Settings . . . 49
7.3 Configuring Services Settings . . . 51
7.3.1 Configuring General Services Settings . . . 51
7.3.2 Configuring Search Services Settings . . . 54
7.3.3 Configuring Print Services Settings . . . 56
7.3.4 Configuring Synchronization . . . 58
7.4 Updating QuickFinder Server Software . . . 59

8 Creating and Managing Virtual Search Servers 61

8.1 About Virtual Search Servers . . . 61
8.2 Creating Virtual Search Servers . . . 61
8.2.1 Creating a Virtual Search Server . . . 62
8.2.2 Naming a Virtual Search Server . . . 63
8.2.3 Using the Virtual Search Server Alias . . . 63
8.2.4 Storing Virtual Search Server Files . . . 64
8.3 Enabling, Disabling, or Deleting a Virtual Search Server . . . 64
8.3.1 Enabling or Disabling a Virtual Search Server . . . 64
8.3.2 Deleting a Virtual Search Server . . . 65
8.4 Configuring Virtual Search Servers . . . 65
8.4.1 Configuring General Settings . . . 66
8.4.2 Configuring Search Settings . . . 71
8.4.3 Configuring Print Settings . . . 75
8.4.4 Configuring Index Settings . . . 77
8.4.5 Configuring Security Settings . . . 80
8.4.6 Configuring Synchronization Settings . . . 82
8.5 Creating Indexes . . . 83
8.5.1 Creating a Crawled Index . . . 83
8.5.2 Creating an Advanced Crawled Index . . . 84
8.5.3 Configuring Rights-Based Search Results for Crawled Indexes . . . 88
8.5.4 Creating a File System Index . . . 89
8.5.5 Creating an Advanced File System Index . . . 90
8.5.6 Configuring Rights-Based Search Results for File System Indexes . . . 93
8.5.7 Searching Across Multiple Indexes . . . 93
8.5.8 Indexing Content on a Password-Protected Web Site . . . 94
8.5.9 Indexing Volumes on Remote Servers . . . 94
8.5.10 Generating Indexes . . . 95
8.6 Managing Existing Index Files . . . 96
8.6.1 Editing an Index . . . 96
8.6.2 Deleting an Index . . . 97
8.6.3 Working with the Log File . . . 97
8.7 Indexing Dynamic Web Content . . . 98

8.8 Access Control to Search Results 99
8.9 Automating Index and Server Maintenance 100
8.9.1 Adding a Scheduled Event 100
8.9.2 Editing or Deleting an Event 101
8.10 Backing Up Your Virtual Search Server Files 101

9 Synchronizing Data Across Multiple QuickFinder Servers 103

9.1 Setting Up a Search Server to Receive Updates 104
9.1.1 Setting Up a QuickFinder Server to Receive Cluster Data 104
9.1.2 Setting Up a Virtual Search Server to Receive Updates 105
9.2 Setting Up a Search Master and Defining a Cluster 106
9.2.1 Setting Up a QuickFinder Server to Send Updates 106
9.2.2 Defining a Synchronization Cluster 108
9.3 Setting Up a Virtual Search Server to Send Updates 109
9.4 Sending Synchronization Updates Manually 109

10 Optimizing Search Results 111

10.1 Reviewing Query Log Reports 111
10.2 Improving Search Results through Intelligent Indexing 112
10.2.1 Excluding Documents from Being Indexed 113
10.3 Modifying Document Descriptions Returned in a Search Results List 114
10.4 Improving Search Results 114
10.4.1 Improving the Relevance of Search Results 115
10.4.2 Using Stop Words Processing to Improve Search Speed 116
10.4.3 Using Best Bets to Help Users Find the Right Information 117
10.4.4 Using Synonyms to Broaden Search Results 118
10.4.5 Redirecting Searches 121
10.4.6 Creating or Editing the Redirection URLs List 121
10.5 Helping Users Avoid Failed Searches 122
10.6 Weighted Queries 122
10.7 Ensuring Optimal Search Speed 123
10.8 Making Good Use of Document Fields 124
10.9 Searching XML Documents 124
10.10 Using the &filter Query Parameter 125

11 Understanding Templates 127

11.1 How Templates Work 127
11.2 Exploring the Default Search and Print Templates 129
11.2.1 Search Page Templates 130
11.2.2 Search Result Templates 130
11.2.3 Print Result Templates 130
11.2.4 Error and Response Message Templates 131
11.3 How Templates Use System Memory 131
11.4 Working with Additional Languages 131

12 Working with Template Variables and Search Parameters 133

12.1 Guidelines for Using Variables 133
12.1.1 Terminology Changes 134
12.2 Adding a Search Box to Your Corporate Web Site 134
12.3 Global Template Variables 134

12.4 Search Page Variables . . . 141
12.5 Search Result Variables . . . 141
12.6 Print Result Variables . . . 144
12.7 Highlighter Variables . . . 146
12.8 Error Message Variables . . . 147
12.9 Response Message Variables . . . 147
12.10 Query Report Template Variables . . . 148
12.10.1 General . . . 149
12.10.2 Report Summary . . . 149
12.10.3 Templates Used . . . 149
12.10.4 Load Details . . . 150
12.10.5 Search Details . . . 153
12.10.6 Search Results Pages . . . 154
12.10.7 Top X Queries . . . 154
12.10.8 Top Queries . . . 155
12.10.9 Error Searches . . . 156
12.10.10 Not Found Searches . . . 157
12.10.11 Redirected Searches . . . 158
12.10.12 Login Requests . . . 158
12.10.13 Print Details . . . 159
12.10.14 Print Errors . . . 159
12.10.15 Highlight Details . . . 160
12.10.16 Export Details . . . 160
12.11 Search Parameters . . . 162

13 Customizing Your Templates and Search Forms 173

13.1 Customizing Templates . . . 173
13.1.1 Customizing the Search Templates . . . 173
13.1.2 Customizing Search Result Templates . . . 174
13.1.3 Customizing Print Result Templates . . . 175
13.1.4 Customizing Error and Response Message Templates . . . 175
13.2 Testing Your Search and Print Solution . . . 176

14 Internationalizing Search Services 177

14.1 Working with Multiple Languages . . . 177
14.1.1 Specifying Locales within Template Filenames . . . 178
14.2 Understanding Character Set Encodings . . . 179
14.2.1 Unicode and UTF-8 . . . 179
14.2.2 Search Encodings . . . 180
14.2.3 Response Encodings . . . 180
14.2.4 HTML Encodings . . . 181
14.2.5 Template Encodings . . . 182
14.2.6 Encoding Issues When Printing . . . 182
14.3 Languages Included in the Default Templates . . . 183
14.4 Additional Resources . . . 183

15 Security Considerations for QuickFinder Server 185

15.1 Security Features . . . 185
15.2 Security Characteristics . . . 187
15.3 Security Configuration . . . 188
15.3.1 QuickFinder Configuration Settings . . . 188
15.3.2 Configuration Settings for Other Products . . . 192
15.4 Other Security Considerations . . . 192

A Troubleshooting QuickFinder **195**

B Combined Character Sets for Use with QuickFinder **197**

B.1 ASCII Character Set 198
B.2 Arabic Character Set 199
B.3 Chinese (Simplified) Character Set 199
B.4 Chinese (Traditional) Character Set 200
B.5 Cyrillic Character Set 201
B.6 European Character Set 202
B.7 Greek Character Set 205
B.8 Hebrew Character Set 206
B.9 Japanese Character Set 206
B.10 Korean Character Set 207
B.11 Thai Character Set 208
B.12 Turkish Character Set 209
B.13 Vietnamese Character Set 209

C Disabling QuickFinder Server **211**

C.1 Disabling QuickFinder Server 211
C.2 Consequences of Disabling QuickFinder Server 211

About This Guide

This guide describes how to use QuickFinder™ Server to add search functionality to your Internet or intranet Web site. It is divided into the following sections:

- Chapter 1, “Overview of QuickFinder Server,” on page 13
- Chapter 2, “What's New in QuickFinder Server 5.0,” on page 21
- Chapter 3, “Setting Up QuickFinder Server,” on page 23
- Chapter 4, “Migrating QuickFinder Server from OES 1 NetWare to OES 2 Linux,” on page 27
- Chapter 5, “Running QuickFinder Server in a Virtualized Environment,” on page 31
- Chapter 6, “Configuring QuickFinder Server for Novell Cluster Services,” on page 33
- Chapter 7, “Configuring QuickFinder Server Default and Services Settings,” on page 35
- Chapter 8, “Creating and Managing Virtual Search Servers,” on page 61
- Chapter 9, “Synchronizing Data Across Multiple QuickFinder Servers,” on page 103
- Chapter 10, “Optimizing Search Results,” on page 111
- Chapter 11, “Understanding Templates,” on page 127
- Chapter 12, “Working with Template Variables and Search Parameters,” on page 133
- Chapter 13, “Customizing Your Templates and Search Forms,” on page 173
- Chapter 14, “Internationalizing Search Services,” on page 177
- Chapter 15, “Security Considerations for QuickFinder Server,” on page 185
- Appendix A, “Troubleshooting QuickFinder,” on page 195
- Appendix B, “Combined Character Sets for Use with QuickFinder,” on page 197
- Appendix C, “Disabling QuickFinder Server,” on page 211

Audience

This guide is intended for anyone involved in installing, managing, and using QuickFinder Server to create search services.

Feedback

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

Documentation Updates

For the latest version of this documentation, see the QuickFinder documentation (http://www.novell.com/documentation/qfserver40) Web site.

Documentation Conventions

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

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

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

Overview of QuickFinder Server

1

QuickFinder™ Server allows your users to find the information they're looking for on any of your public and private Web sites, your partners' sites, any attached file systems or servers, and any number of additional Web sites across the Internet, all from a single search form on your Web page. You can easily modify the look and feel of any of the sample search results pages to match your corporate design.

You can create full-text indexes of HTML, XML, PDF, Word, OpenOffice.org, and many other document formats in almost any language with our indexing engine, which is base on Unicode*. You can configure and maintain your indexes remotely from anywhere on the Web with the QuickFinder Web-based administration module.

Whether it's simple search solutions for your own Web sites or complex, revenue-generating search services for hosted environments (ASPs/ISPs), QuickFinder provides the features, scalability, and flexibility you're looking for.

Figure 1-1 *How QuickFinder Server Works*

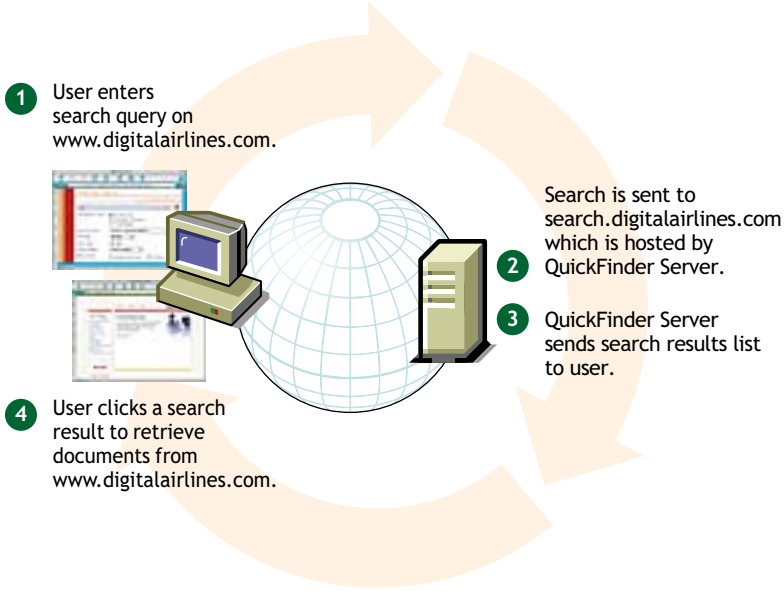

This overview includes the following topics:

- Section 1.1, "Benefits," on page 14
- Section 1.2, "Features," on page 14
- Section 1.3, "Key Components of QuickFinder Server," on page 16
- Section 1.4, "Managing QuickFinder Server," on page 17
- Section 1.5, "Taking a Test Run: Performing a Search," on page 19
- Section 1.6, "Getting Help," on page 19

1.1 Benefits

More than 70% of Web site visitors on the World Wide Web use the search functionality of the Web sites they visit, and over 40% of Web site owners feel that the search functionality is their most important feature for online shopping (even more important than customer service, wish lists, personalization, and product information). Yet more than 50% of large Web sites (over 500,000 pages) and 70% of moderately sized Web sites (up to 5,000 pages) do not offer search functionality to their visitors. And although Internet search engines index many of these Web sites, they either cannot or do not index all available content.

The top two reasons given by Web site owners for not offering search functionality are the difficulty in setting up a search service, and the length of time involved in doing so.

QuickFinder Server is easy to set up and manage, and is included free with NetWare® 6.5 SP7 and Novell® Open Linux Enterprise Server 2. It includes a browser-based administration utility with helpful examples and tips right in the interface. The help system offers additional information and provides links to relevant topics in this manual. Unlike other search engines available on the market today, there are no escalated fees or limits on the number of documents indexed, the number of users who use the service, or the number of user transactions.

Additionally, QuickFinder Server lets you do many things to increase your site's usability:

- Gather customer metrics by reviewing customer searches to identify what your customers look for the most
- Direct customers to the information you know they need, and away from the information you know they don’t want
- Spotlight products for marketing purposes by redirecting searches to specific Web pages
- Keep customers on your site longer
- Improve employee productivity by helping them find information more quickly
- Use the search page as an advertising portal (because the search page is often the most seen page on a Web site)
- Help customers navigate your Web site
- Provide an information portal for customer service departments to quickly answer customer questions

1.2 Features

QuickFinder Server offers many features for both the administrator and end user.

- Section 1.2.1, “QuickFinder Administrators,” on page 14
- Section 1.2.2, “End Users,” on page 15

1.2.1 QuickFinder Administrators

Many new features and enhancements have been added to this latest release of QuickFinder, making the creation and management of search services faster and easier. If you are a veteran of QuickFinder, see “What's New in QuickFinder Server 5.0” on page 21 for a list of the latest features. If you are new to QuickFinder, take a look at what QuickFinder lets you do:

- Host search services for one or more companies or organizations

- Quickly integrate search functionality throughout your corporate internal and external Web sites
- Define any number of indexes for any content on your Web or network servers
- Enforce access restrictions so only authorized personnel can search for sensitive documents
- Adjust the relevance of any documents or parts of your Web site to ensure only the best results
- Define special best bets search results of only the most popular, most recent, or most important information
- Generate reports of all search activity for a given period, so you can learn what customers are looking for and what's missing from your Web site and indexes
- Customize the look and feel of search and print results in all languages
- Create themes, which are defined collections of search and print result templates that allow you to deploy custom look and feel virtual search servers, each for a specific company or department
- Update your search indexes automatically by using the Scheduling page
- Synchronize indexes, templates, and configuration settings of two or more QuickFinder servers
- Redirect searches for common products or terms to a specific URL, similar to the AOL* keyword feature (http://www.amazon.com/exec/obidos/tg/detail/-/0764575023/103-5880842-6443046?v=glance&vi=excerpt)
- Improve search results by synonyms, redirection, stop words, best bets, relevance adjustments, search expansion, speller suggestions, and including additional search details with user queries
- Have QuickFinder notify you by e-mail when there are errors, warnings, and query reports

1.2.2 End Users

End users, or users who use your search services, benefit from fast and relevant searches and the ability to print formatted search results. Additionally, end users can do the following:

- Find the information they are looking for quickly and easily
- Search across multiple Web sites, servers, and file formats in any language, all from a single interface
- Print large collections of dispersed but related files as a single, coherently organized document
- Use the Highlighter feature to highlight search terms within the context of their original documents
- Focus their searches with wildcards and an easy-to-use query language
- Limit search results to documents created or modified recently (within any date range)
- Choose any level of detail for search results (QuickFinder supports up to 15 different fields)
- Select any number of indexes to search in
- Select the language QuickFinder should use when returning search results
- Sort search results by title, description, language, date, or URL (any of the 15 fields can be sorted, and you can perform multi-column sorts)

1.3 Key Components of QuickFinder Server

Four key components are at the heart of the QuickFinder Server: virtual search servers, indexes, templates, and configuration settings. To understand how QuickFinder operates, you should understand these components.

- Section 1.3.1, “Virtual Search Servers,” on page 16
- Section 1.3.2, “Indexes,” on page 17
- Section 1.3.3, “Templates,” on page 17

1.3.1 Virtual Search Servers

A virtual search server is a completely independent search solution established for a single Web site, department, or organization. QuickFinder Server can support any number of virtual search servers from a single machine, and consists of the following components:.

- A name and alias, which appear in your list of virtual search servers
- Index files containing key words and related URLs for use in generating search results
- Scheduled indexing events
- HTML-based templates for highlighter, search, and print results
- Log files
- Configuration settings, such as stop words, keyword-redirects, synonyms, and security settings
- Query reports
- Admin interface

Each of these components is managed through the QuickFinder Server Manager, which is accessed through a Web browser. For more information, see Section 3.2.1, “Components of a Virtual Search Server,” on page 25 and Section 1.4, “Managing QuickFinder Server,” on page 17.

Virtual search servers provide a method for categorizing similar information stores where each virtual search server is created for a specific purpose and audience. For example, you might create one virtual search server for your company's support organization, another for its public Web site, and yet another for your intranet. You might break these down even further by creating more focused virtual search servers for groups within these organizations.

A typical QuickFinder server might have a dozen virtual search servers, each containing several indexes of its own.

Figure 1-2 *A QuickFinder Server with Several Virtual Search Servers*

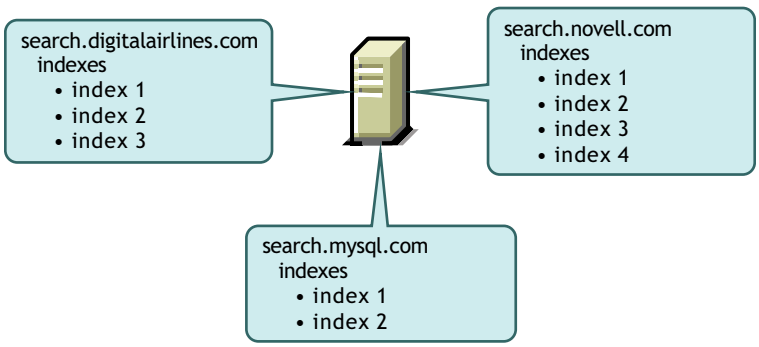

You can have two or more virtual search servers share an index by creating a duplicate index on each virtual search server that points to the same index directory. In this manner, all virtual search servers can search a shared index in addition to their own indexes.

To create virtual search servers and generate indexes for them, you use the QuickFinder Server Manager, a browser-based utility. For more information, see Chapter 8, “Creating and Managing Virtual Search Servers,” on page 61.

1.3.2 Indexes

An index is a file generated by QuickFinder Server that contains the key words and matching URLs gathered by QuickFinder Server from content located on file or Web servers. QuickFinder uses indexes to compile search results. You create one or more indexes for each virtual search server you create.

You can create two types of indexes: file system and crawled. When you want to index the contents of a file server, you create a file system index. When you want to index the contents of a Web server, you create a crawled index.

For more information about creating indexes, see Section 8.5, “Creating Indexes,” on page 83.

1.3.3 Templates

A template is an HTML document containing one or more QuickFinder Server variables. Template variables are used to produce dynamic results when a user performs a search on the virtual search server you have defined.

Templates can be shared across virtual search servers or each virtual search server can point to its own set of templates.

For more information about templates, see Chapter 11, “Understanding Templates,” on page 127.

1.4 Managing QuickFinder Server

QuickFinder Server Manager is a browser-based utility you use to create and manage all virtual search servers and their indexes on your QuickFinder server.

Figure 1-3 *QuickFinder Server Manager Interface*

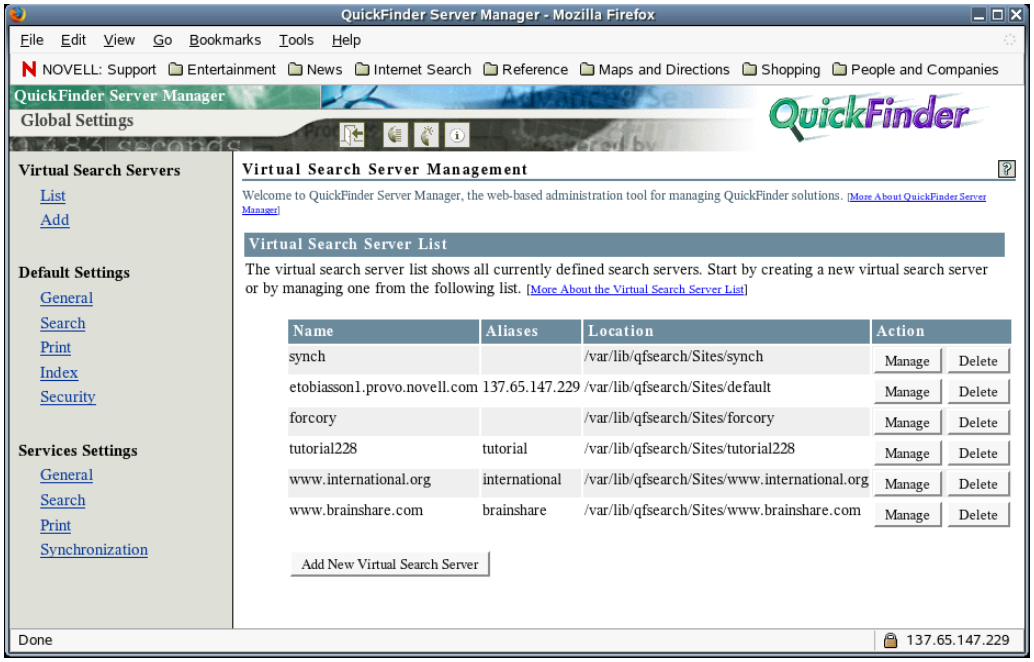

QuickFinder Server Manager lets you apply default (global) settings so that each time you create a virtual search server, you don't need to waste time configuring it. After you create a virtual search server, you can customize any of the default settings.

See Section 1.6, “Getting Help,” on page 19 for more information about where to get help using QuickFinder Server Manager.

- Section 1.4.1, “Accessing QuickFinder Server Manager on NetWare,” on page 18
- Section 1.4.2, “Accessing QuickFinder Server Manager on Linux,” on page 18

1.4.1 Accessing QuickFinder Server Manager on NetWare

1. Enter `http://`*`servername`*`/qfsearch/admin` (or Admin, AdminServlet, or adminservlet) in your Web browser's address field.

   You can also use the IP address of your server in place of *servername* as long as the IP address is included in the list of aliases for that virtual search server.
2. Enter the name and password of a valid eDirectory™ user (such as admin) that has write rights to the `/qfsearch` directory.
3. Click *Login*.

1.4.2 Accessing QuickFinder Server Manager on Linux

1. Enter `http://`*`servername`*`/qfsearch/admin` (or Admin, AdminServlet, or adminservlet) in your Web browser's address field.

   You can also use the IP address of your server in place of *servername* if the IP address is included in the list of aliases for that virtual search server.

2 Enter the name and password of a valid local Linux user that has write rights to the `/var/lib/qfsearch` directory.

You can log in as `root` (with the root password), the eDirectory administrator (with the admin password), or the user you created during the OES install (with that user's password).

You can also give any eDirectory user rights to administer QuickFinder by using Linux User Management (LUM):

2a Export the eDirectory user to the Linux machine.

2b Create a www group in eDirectory and give it the same GID as the www group on the Linux machine.

2c Assign the eDirectory user to the new www group.

The eDirectory user can now access QuickFinder Manager.

3 Click *Login*.

1.5 Taking a Test Run: Performing a Search

When you install QuickFinder Server, some of your server's content is automatically indexed and appears on the default search form as QuickFinder and DocRoot indexes. The first is a file system index of the QuickFinder help system and the second is a crawled index of the default Web server on your machine.

After the server has been started, you can open the search page through your Web browser and perform a search against the content that has been automatically indexed.

To test QuickFinder by using the default search page, do the following:

1 Type `http://`*`domainname`*`/qfsearch/quickfinder` in your Web browser's address field, then press Enter.

IMPORTANT: The URL is case sensitive. Use the exact case shown above. You can substitute your server's IP address for the domain name if the IP address is in list of aliases for the virtual search server.

---

2 Type `QuickFinder` in the *Search* field, then press Enter.

TIP: The Search form template, `SearchTemplate.html`, is stored on your server at `\`*`searchroot`*`\Templates` where *\searchroot* is the name of the folder where you installed QuickFinder. By default, this folder is found at `sys:\qfsearch` (found at the root of your sys volume) on NetWare and `/var/lib/qfsearch` on Linux. See "Customizing Your Templates and Search Forms" on page 173 for information.

1.6 Getting Help

The QuickFinder Server Manager interface includes tips, examples, and links into the QuickFinder help system. Whether you're an expert or a novice QuickFinder administrator, you can get as much or as little information as you need, when you need it.

When you need help using QuickFinder Server Manager, click one of the *More About* links found in some areas of the interface, or click to launch the help system. The help system contains information about each item on the current QuickFinder Server Manager page. At the bottom of

each help topic is the More Help On the Web section containing links to this *QuickFinder Server Administration Guide*.

1.7 What's Next

- To find out what's new, see Chapter 2, "What's New in QuickFinder Server 5.0," on page 21.
- To set up QuickFinder services on a physical server, see Chapter 7, "Configuring QuickFinder Server Default and Services Settings," on page 35.
- To set up QuickFinder services on a virtual server, see Chapter 5, "Running QuickFinder Server in a Virtualized Environment," on page 31.
- To set up QuickFinder services on a cluster, see Chapter 6, "Configuring QuickFinder Server for Novell Cluster Services," on page 33.
- To migrate QuickFinder services from NetWare to OES Linux, see Chapter 4, "Migrating QuickFinder Server from OES 1 NetWare to OES 2 Linux," on page 27.

What's New in QuickFinder Server 5.0

2

QuickFinder™ Server 5.0 has been enhanced with new features that speed up searching, improve the accuracy of search results, and enhance the experience users have while searching.

- Section 2.1, “New and Updated Features,” on page 21
- Section 2.2, “File Reader Updates,” on page 21

2.1 New and Updated Features

QuickFinder Server 5.0 gives you the ability to:

- Index all local and remote NCP™, Novell® Storage Services™, and NetWare® volumes.
- Filter search results based on NCP server rights.
- Properly encode file URLs so that they open directly through a file browser.

  For example, the following UNC path opens the specified file in OpenOffice.org: `file:////servername/path/filename.odt`
- Index large amounts of data per index (we have currently indexed over one million files).
- Dynamically load new external file readers, which makes it easier to update and add new file readers.
- Display arbitrary admin-specified metadata from indexed documents on the Search Results page.
- Display hits per index on the Search Results page (See sample ResultListTemplate.html for example).
- Search filenames by default with every query.

2.2 File Reader Updates

QuickFinder Server 5.0 includes the following file reader updates:

- **PDF:** QuickFinder Server 5.0 detects even more text in PDF files, and includes support for compressed Adobe* 6 and 7 files, a variety of unusual fonts (including Asian and double-byte fonts), and external CMAP files. The handling of descendant fonts has also been improved.
- **OpenOffice.org:** Support for Unknown field names was added, and specific OpenOffice.org types can now be detected using the mimetype file within a ZIP.
- **HTML:** Handling of entity references within a meta tag was added, as was improved handling of entity references. NoFollow to Robots comment tags were also added.
- **ASCII:** More encodings are now supported for ASCII files (such as UTF8, UNICODE, big-endian, and little-endian).
- **Microsoft Office:** Includes improved support for embedded documents, and support for additional metadata (including Unicode metadata).

Setting Up QuickFinder Server

3

You can install QuickFinder™ Server as a service on NetWare® 6.5 SP7 or Novell® Open Enterprise Server (OES) 2 for Linux, or you can dedicate a server to it using the *Novell QuickFinder* OES Services pattern. After installation, you can modify the global settings of your new QuickFinder server. Global settings make the creation of new virtual search servers easier by letting you specify the default settings to be used each time you create one. You can also enable or disable certain features, such as best bets, printing, highlighter, e-mailing, or synchronization.

- Section 3.1, "Installing QuickFinder Server," on page 23
- Section 3.2, "Designing Your Search Solution," on page 24

3.1 Installing QuickFinder Server

If you did not install QuickFinder Server during the NetWare 6.5 SP7 or OES 2 Linux installation, you can install it anytime using the NetWare 6.5 SP7 or the OES 2 for Linux installation media.

For more information about the install process, see the *OES 2: NetWare Installation Guide* (http://www.novell.com/documentation/oes2/inst_oes_nw/data/front.html#front) and the *OES 2: Linux Installation Guide* (http://www.novell.com/documentation/oes2/inst_oes_lx/data/front.html#front).

- Section 3.1.1, "Installing QuickFinder Server From the NetWare 6.5 SP7 Media," on page 23
- Section 3.1.2, "Installing QuickFinder Server From the OES for Linux Media," on page 24
- Section 3.1.3, "QuickFinder Languages," on page 24

3.1.1 Installing QuickFinder Server From the NetWare 6.5 SP7 Media

1. Insert the NetWare 6.5 SP7 Products CD/DVD into the drive of the server where you want to install QuickFinder.
2. Type `startx` at the system console, then press Enter to start the NetWare GUI.
3. Click *Novell* > *Install* > *Add*.
4. In the Source Path dialog box, specify the path to the CD, or click the *Browse* button and locate the CD.
5. Select the `postinst.ni` response file, then click *OK*.
6. From the products list, select *QuickFinder Server*.
7. Click *Next*.
8. When prompted, specify your administrator username and password, and your user context.
9. Click *OK*.
10. Follow the remaining screen prompts.

IMPORTANT: After the installation is complete, you must restart the NetWare server. Enter `restart server` at the system console. Also, if you are upgrading from a previous version of

QuickFinder Server or Web Search, you must regenerate any older indexes before they can be searched. For more information, see "Generating Indexes" on page 95.

For instructions about accessing the QuickFinder Server Manager, which is the interface for configuring and managing QuickFinder, see "Accessing QuickFinder Server Manager on NetWare" on page 18.

3.1.2 Installing QuickFinder Server From the OES for Linux Media

If you did not install QuickFinder when you installed OES 2, you can install it by using the YaST Control Center.

1. Launch the YaST Control Center.

   GNOME: Click *Computer* > *More Applications* > *System* > *YaST*.

   KDE: Click the menu button > *System* > *YaST*.
2. Click *Software* > *Add-on Product*.
3. Click *Open Enterprise Server* > *OES Install and Configuration*.
4. On the Software Selection page, select *Novell QuickFinder*.

   Services that are already installed have a blue check mark in the check box next to the service.
5. Click *Accept.*
6. Change the default configuration information as prompted, then click *Next*.
7. Follow the on-screen prompts to complete the installation.

3.1.3 QuickFinder Languages

QuickFinder is available in several languages. However, on NetWare, QuickFinder is only installed in the languages you selected during the initial operating system installation. You will have end-user search templates in all supported languages, but the QuickFinder Manager UI and both the admin and end-user help systems are only available in English and the initially selected languages.

To install the additional language pack files, simply extract the contents of the `QFindXX.zip` files into the `/QFSearch` directory.

NOTE: The language pack files are located in the `\Products\QFSearch` directory on the install CD and are also available from the QuickFinder product page (http://www.novell.com/products/quickfinder).

3.2 Designing Your Search Solution

QuickFinder Server is a search service solution designed for adding powerful search capabilities to individual Web sites. It is not intended to index the entire Internet. However, because many Web sites are comprised of multiple Web and file servers located across an enterprise, QuickFinder is designed to be able to index hundreds, even thousands, of Web sites and file servers as part of a single search solution.

This section includes the following information:

- Section 3.2.1, “Components of a Virtual Search Server,” on page 25
- Section 3.2.2, “Deciding If You Need More Than One Virtual Search Server,” on page 26
- Section 3.2.3, “Becoming a Search Service Host,” on page 26

3.2.1 Components of a Virtual Search Server

Providing search services involves creating one or more virtual search servers.

A virtual search server is a fully functioning, self-contained search service created for a particular audience, such as a department, organization, or a specific group of customers.

A virtual search server typically contains its own indexes, log files, administration interface, search and print templates, scheduled events, virtual search server name, and optional aliases. When you create a new virtual search server, a new directory is created in the `sys:/qfsearch/sites` subdirectory on NetWare and the `/var/lib/qfsearch/Sites` subdirectory on Linux. This subdirectory contains supporting configuration files and subfolders for storing indexes.

When you create a new virtual search server, you create an independent search service, meaning that it is self-contained and doesn't depend on, or interact with, other virtual search servers.

NOTE: Users cannot search more than one virtual search server at a time. However, a virtual search server can contain indexes created from content on multiple Web sites or file servers.

Taking the time to plan your search service strategy can save you time and money and improve the quality of your service.

In addition to an administrative interface, each virtual search server that you create typically contains one or more of the following components:

- **Indexes:** Files that hold key words and associated URLs of Web sites or file server content that have been indexed, or crawled.
- **Scheduled Events:** Index management, such as updating or regenerating, can be automated to occur at specific intervals using the Scheduling feature.
- **Search and Print Results Templates:** Templates that are populated with the results of a search and then displayed to the user. Depending on which templates are used, the level of detail displayed in search and print results varies.
- **Themes:** A collection of templates that share a common look and feel and that work together to provide a rich end-user search experience. A theme includes your search page, search and print results pages, and response and error message pages.
- **Configuration Settings:** Includes information such as stop words, keyword-redirects, synonyms, and security settings.
- **Query reports:** Summary reports, in HTML format, of the query log data.

3.2.2 Deciding If You Need More Than One Virtual Search Server

To determine if you need more than one virtual search server, answer the following questions:

- ❑ Do you want to host search services for multiple, independent organizations or Web sites?
- ❑ Do you want to consolidate multiple QuickFinder servers onto a single machine?
- ❑ Do you need to prevent users from being able to search across multiple indexes at the same time?

If you answered yes to any of these, you likely need to create more than one virtual search server. Otherwise, a single virtual search server might suffice. For more information about virtual search servers, see Chapter 8, “Creating and Managing Virtual Search Servers,” on page 61.

3.2.3 Becoming a Search Service Host

Companies who want customers to find information about their products typically outsource this function to search service companies. With QuickFinder Server, you can offer professional search services for other companies. Using a single installation of QuickFinder, you can host many virtual search servers simultaneously, which means that you could use the same server to host search services for several client or customer Web sites.

The following diagram shows how a search query on the Web site www.digitalairlines.com is sent to a QuickFinder server on the domain search.digitalairlines.com.

Figure 3-1 *How a Search Query is Sent to the QuickFinder Server and Processed*

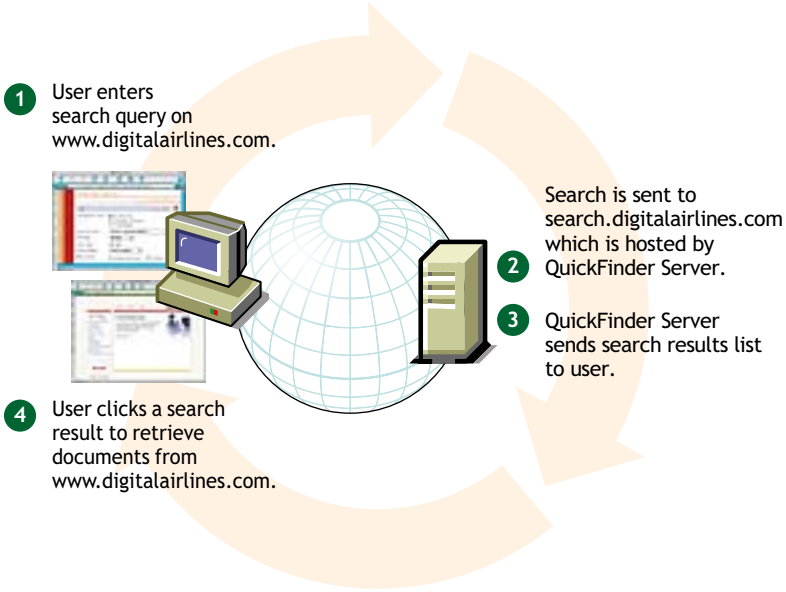

Migrating QuickFinder Server from OES 1 NetWare to OES 2 Linux

4

This section provides detailed instructions on how to migrate QuickFinder™ Server from NetWare® to Linux. For general information about the OES migration tools, see "Overview of the OES Migration Tools" in the *OES 2: Migration Tools Administration Guide*.

- Section 4.1, "Coexistence," on page 27
- Section 4.2, "Migration Procedure," on page 27
- Section 4.3, "Post-Migration Considerations," on page 29

4.1 Coexistence

This section provides information regarding the coexistence of QuickFinder Server 5.0 with existing NetWare or Linux networks, and with previous versions of the product.

- Section 4.1.1, "Compatibility," on page 27
- Section 4.1.2, "Coexistence Issues," on page 27

4.1.1 Compatibility

The following table summarizes the compatibility of QuickFinder with various network operating systems:

Table 4-1 *QuickFinder Compatibility*

Operating System	Compatible Versions
NetWare	OES NetWare
	NetWare 6.5 SP3 or later
Linux	SUSE® Linux Enterprise Server 8
	SUSE Linux Enterprise Server 9 SP1 or later
	SUSE Linux Enterprise Server 10 or later

4.1.2 Coexistence Issues

Web Search Server and QuickFinder Server cannot run on the same NetWare machine, but you can have as many versions of QuickFinder Server and Web Search Server as you want in your network.

4.2 Migration Procedure

Use the following steps to migrate QuickFinder Server from a NetWare server to an OES for Linux server.

NOTE: Migrating from QuickFinder Server running on OES NetWare to OES Linux replaces any indexes or configurations made on the Linux server. If you want to merge your changes from NetWare to Linux, you must manually re-create the changes using the QuickFinder Server Manager.

1. Install QuickFinder Server on Linux.

   For more information, see Section 3.1.2, "Installing QuickFinder Server From the OES for Linux Media," on page 24.

2. Access QuickFinder Server Manager on the NetWare server, then click Global Settings ( ) on the top toolbar.

   See Section 1.4.1, "Accessing QuickFinder Server Manager on NetWare," on page 18 for more information.

3. Write down the paths for each virtual search server displayed in the *Location* column.

4. On the OES for Linux server, mount the NetWare server (you can use ncpmount).

5. Make a backup of `/var/lib/qfsearch/SiteList.properties`.

   Make sure that you don't have a file with this name as a backup on the NetWare server.

6. Copy all `.properties` and `Cron.jobs` files from the root directory `sys:/qfsearch` on the NetWare server to `/var/lib/qfsearch` on the Linux server.

7. Copy `sys:/qfsearch/Sites` and all of its subdirectories to `/var/lib/qfsearch/Sites`.

8. Copy `sys:/qfsearch/Templates` and all of its subdirectories to `/var/lib/qfsearch/Templates`.

9. If any of the paths listed in Step 3 are not under `sys:/qfsearch` (for example, if you installed a virtual search server somewhere other than the default location), you must also copy those paths to Linux.

   For example, if you have the path `sys:/SearchSites/PartnerSite`, you must copy it to the Linux server. You could copy it to `/var/opt/SearchSites/PartnerSite` or `/var/lib/qfsearch/Sites/PartnerSite`.

10. Edit all NetWare paths in `/var/lib/qfsearch/SiteList.properties` to reflect the new Linux paths (such as `sys:/qfsearch to /var/lib/qfsearch`, or as in Step 9, `sys:/SearchSites/PartnerSite` to `/var/opt/SearchSites/PartnerSite`.

    NOTE: Some paths might have one or two backslashes (\) that must be replaced with one forward slash (/). For example, `sys:\\qfsearch\\docs` needs to be change to `/var/lib/qfsearch/docs`.

11. Update all NetWare paths in the properties and configuration files copied in the steps above to the Linux paths, and update any DNS names.

    The following is a list of files that must be updated:

    - `AdminServlet.properties`
    - `Cron.jobs`
    - `Sites/Highlighter.properties`
    - `Sites/Print.properties`
    - `Sites/Search.properties`

For each of the virtual search servers, modify the following:

- `qfind.cfg`
- Any of the above `.properties` files, if they exist.

NOTE: The names of most properties files are mixed case, so make sure the files copied from NetWare are the correct case. You can compare them to the `.properties.sample` files on Linux.

You might also need to update paths in templates. If you have problems such as a template not being found or some properties not being set properly, check the case of the filename. Linux filenames are case sensitive.

If you modified any "file" index paths to index directories on the Linux server, that index must be regenerated.

---

12 After all the files have been modified, run the following commands to set the access rights and owner/groups so that the QuickFinder engine has rights to access the files:

```
chown -R root:www /var/lib/qfsearch
chmod -R 770 /var/lib/qfsearch
```

4.3 Post-Migration Considerations

After upgrading from a previous version of Web Search or QuickFinder Server, be aware of the following issues:

- If you are upgrading QuickFinder Server (including the original OES version) on a NetWare server that is running Web Search Server, the QuickFinder Server install recognizes the older installation and migrates the configuration settings and indexes that QuickFinder Server uses. However, you must regenerate the older indexes before they can be searched (see "Generating Indexes" on page 95).
- If you are migrating from an existing NetWare server with Web Search to a new NetWare server with QuickFinder Server, the NetWare Migration Wizard handles the migration in such a way that everything is migrated from Web Search Server to QuickFinder Server on the new server.
- If you are migrating from an existing NetWare server with Web Search Server to a new Linux server with QuickFinder Server, you must manually copy files from NetWare to Linux so that everything doesn't have to be newly created.
- QuickFinder Server 5.0 indexes are not compatible with previous version of QuickFinder Server. The indexes must be regenerated, and you cannot synchronize QuickFinder Server 5.0 indexes with indexes from a previous version of QuickFinder Server (and visa-versa).
- The new template files are not placed in the `<qfsearch>/Templates` directory because this would overwrite any changes you might have made to them.

  Instead, the new template files are placed in the `<qfsearch>/Templates/Samples` directory. To take advantage of the functionality the new templates provide, copy the template files from the `<qfsearch>Templates/Samples` directory to the `<qfsearch>/Templates` directory.

Running QuickFinder Server in a Virtualized Environment

5

QuickFinder™ Server runs in a virtualized environment just as it does on a physical NetWare® server, or on a physical server running OES 2 Linux, and requires no special configuration or other changes. However, if you want to run QuickFinder within a virtual machine and intend to index large amounts of data, you need to assign at least 2 GB of memory to the virtual machine.

To get started with virtualization, see "Introduction to Xen Virtualization" in the *Virtualization: Getting Started* guide.

For information on setting up virtualized NetWare, see "Setting Up Virtual Machines" in the *Virtualization: Getting Started* guide and "NetWare Virtual Machines"in the *Virtualization: Guest Operating System Guide*.

For information on setting up virtualized OES 2 Linux, see "Setting Up Virtual Machines" in the *Virtualization: Getting Started* guide and "OES Linux Virtual Machines"in the *Virtualization: Guest Operating System Guide*.

Configuring QuickFinder Server for Novell Cluster Services

6

The QuickFinder™ Server Synchronization feature provides fault tolerance and load balancing. To include failover capabilities, you can enable QuickFinder for clustering.

To configure QuickFinder Server to work with Novell® Cluster Services™, you should install QuickFinder Server on each server in your cluster. This allows the installation program to correctly register QuickFinder Server with each of your server's Web and application servers.

With NetWare® 6.5, you cannot install QuickFinder Server to any volume other than sys:. However, you can use virtual search servers to enable QuickFinder Server to work in a cluster.

1. Make the cluster node active on one server.
2. If it is not already installed, install QuickFinder on that server.
3. In the QuickFinder Server Manager, click the *Global Settings* icon, then click *Add New Virtual Search Server*.
   - **3a** In the *Name* field, specify the DNS name of the cluster server.
   - **3b** In the *Location* field, specify the volume and path on the shared cluster volume where all the indexes and virtual search server settings will be located.
   - **3c** Click *Add*.
4. Set up and create any indexes that you want clustered in that virtual search server.

   For more information, see "Creating Indexes" on page 83.
5. Move the cluster node to the other server and install QuickFinder Server on that server.
6. Add a new virtual search server and give it the same name and location as the virtual search server you created in Step 3.

   The settings already created on the shared volume are displayed.

To do a search on the cluster volume, enter `http://DNS_CLUSTER/qfsearch/search`. QuickFinder Server sees the DNS and knows which virtual search server to send the request to.

IMPORTANT: All scheduling information is stored on the sys: volume, so you must set up scheduling on each server. When the cluster node is not active on one of the servers, the scheduled indexes on that server fail. However, the server that has the cluster node generates the indexes.

Configuring QuickFinder Server Default and Services Settings

7

Before you start creating virtual search servers and building indexes for them, you might want to modify the global settings of your search server, which affect all virtual search servers you create.

QuickFinder™ Server Manager's home page displays a list of all virtual search servers that exist on your QuickFinder Server (see Figure 7-1). This home page is called Global Settings because the changes you make from this page affect all new virtual search servers that you create, and they also affect the functionality of the servlets that provide the QuickFinder services.

Figure 7-1 *QuickFinder Server Virtual Search Global Settings Page*

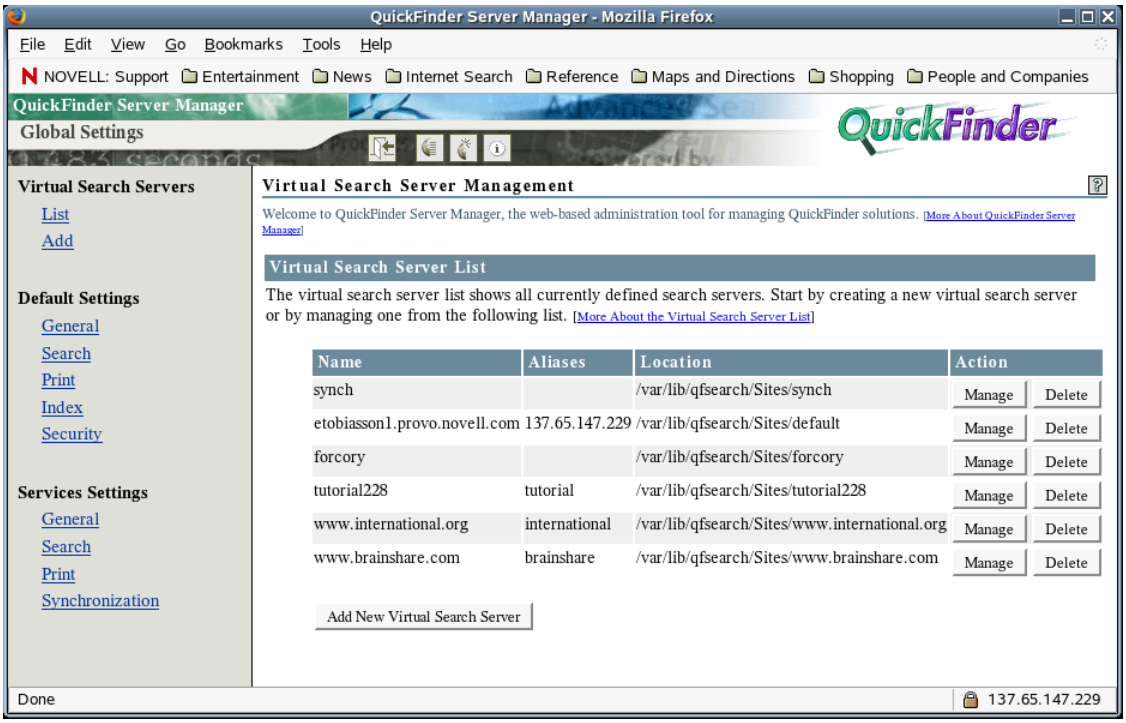

For instructions about accessing the QuickFinder Server Manager, which is the interface for configuring and managing QuickFinder, see “Accessing QuickFinder Server Manager on NetWare” on page 18.

This section contains the following topics:

- Section 7.1, “Differences between Default and Services Settings,” on page 36
- Section 7.2, “Configuring Default Settings,” on page 36
- Section 7.3, “Configuring Services Settings,” on page 51
- Section 7.4, “Updating QuickFinder Server Software,” on page 59

7.1 Differences between Default and Services Settings

The Global Settings page—also called the home page because it is the default page that appears when you first log in to QuickFinder Server Manager—is for use by the global search administrator, or the person who has administrative rights over the QuickFinder Server.

The Global Settings page has two groups of settings: Default Settings and Services Settings. The first lets you specify default settings for new virtual search servers, and the second gives you administrative control of all virtual search servers.

- Section 7.1.1, “Default Settings,” on page 36
- Section 7.1.2, “Services Settings,” on page 36

7.1.1 Default Settings

As the name implies, Default Settings lets you specify what the default settings should be each time you create a new virtual search server. For example, if you set the default query encoding from the General Settings page to Unicode (UTF-8), any new virtual search server you create after making this change defaults to the new setting. However, after you create a new virtual search server, you can then change its default settings.

For more information, see Section 7.2, “Configuring Default Settings,” on page 36 and Chapter 8, “Creating and Managing Virtual Search Servers,” on page 61.

7.1.2 Services Settings

Services Settings lets you control all new and existing virtual search servers. For example, when you disable a feature such as e-mail services, the e-mail services feature is removed from the virtual search server configuration pages, making it impossible for the virtual search server administrator to access it.

You can also perform other administrative tasks, such as checking for software updates and configuring QuickFinder synchronization.

For more information, see Section 7.3, “Configuring Services Settings,” on page 51.

7.2 Configuring Default Settings

Default settings are meant to make the creation of new virtual search servers easier. By letting you specify common settings for use by all new virtual search servers, you don't need to waste time making the same configuration changes each time you create a new virtual search server.

Before you begin making changes, start by exploring QuickFinder so that you can decide what the most common configuration settings should be for your situation. You might start by planning your search solution to get ideas about the purpose for it, who will be using it, what languages it must support, and so on. For additional information about planning your search solution before you begin creating it, see Section 3.2, “Designing Your Search Solution,” on page 24.

Existing virtual search servers are not affected by these changes.

- Section 7.2.1, “Configuring Default General Settings,” on page 37

- Section 7.2.2, “Configuring Default Search Settings,” on page 40
- Section 7.2.3, “Configuring Default Print Settings,” on page 44
- Section 7.2.4, “Configuring Default Index Settings,” on page 46
- Section 7.2.5, “Configuring Default Security Settings,” on page 49

7.2.1 Configuring Default General Settings

General Settings lets you modify general query, response, error log, and e-mail settings.

- “Modifying Query Settings” on page 37
- “Modifying Default Query Logging Settings” on page 38
- “Modifying Response Settings” on page 39
- “Modifying Error Log Settings” on page 40
- “Modifying E-Mail Settings” on page 40

Modifying Query Settings

1 On the QuickFinder Server Manager's Global Settings page, click *General* under *Default Settings*.

2 From the *Default query encoding* drop-down list, select an encoding that represents the character set encoding that the users of your search services are most likely to use.

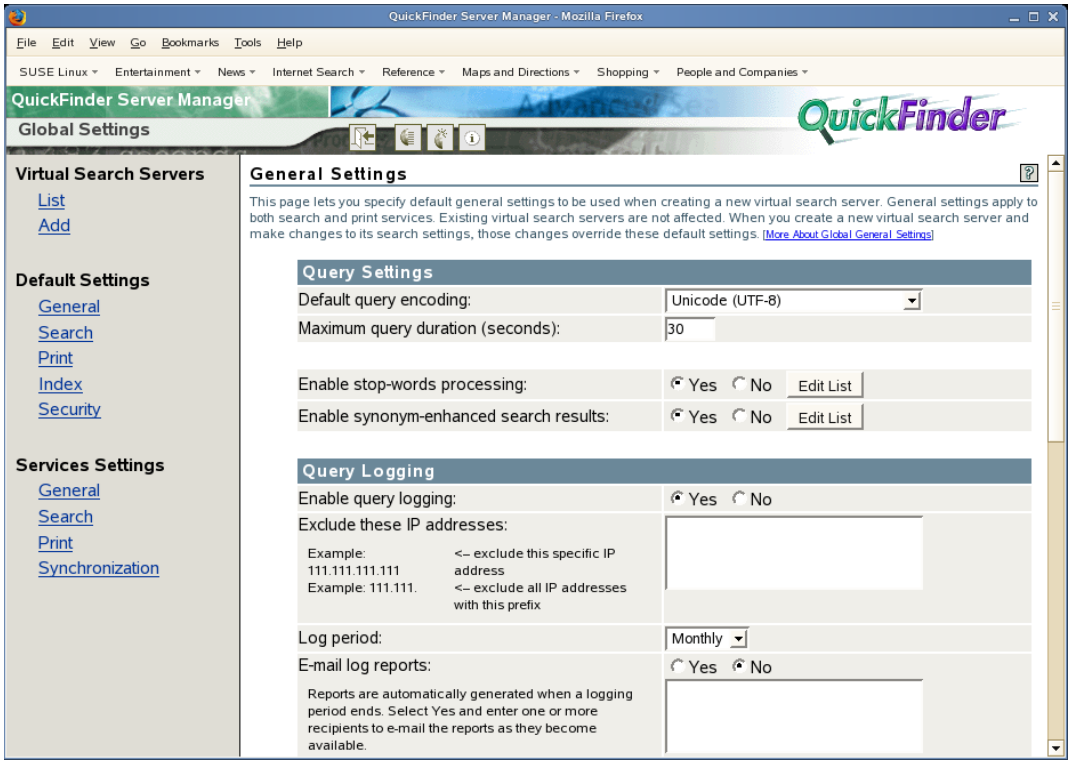

UTF-8 is a good choice for most situations because it works for all languages. However, there might be times you need to select an encoding other than UTF-8. For example, if you are hosting a Japanese Web site for a Japanese-only audience, you wouldn’t want to choose UTF-8

because it needs 3 bytes for most Japanese characters. In this case, a better choice might be Shift_JIS or ISO-2022-JP. For more information about character sets, see Appendix B, “Combined Character Sets for Use with QuickFinder,” on page 197.

3 In the *Maximum query duration* field, specify the maximum number of seconds before QuickFinder should end a query, regardless of whether a search has been completed.

This option is one of several methods that enable you to protect your server's resources from processing rogue searches, which are sometimes intended to harm your service by consuming server resources. For more information, see Appendix 15, “Security Considerations for QuickFinder Server,” on page 185.

4 (Optional) Click *Yes* next to *Enable stop-words processing* if you want QuickFinder to ignore insignificant words, such as adverbs, conjunctions, or prepositions, in a user's search query.

For information about stop words and how to edit the default list of stop words, see “Using Stop Words Processing to Improve Search Speed” on page 116.

5 (Optional) Click *Yes* next to *enable synonym enhanced search results* if you want QuickFinder to expand search queries by looking for synonyms of the original search terms.

Synonym enhancement is not performed within the FileFilter section of a search. For more information about synonym enhanced searches, see “Using Synonyms to Broaden Search Results” on page 118.

6 Click *Apply Settings*.

Modifying Default Query Logging Settings

When query logging is enabled (it is enabled by default), it captures many details about user queries, including total number of searches performed (for the current logging period), top 100 search terms, load statistics, most popular indexes, and the number of print jobs and highlights performed. It also provides information about queries that resulted in a “not found” error message. This information can be used to improve the performance of your search services, or to identify areas of your site that could be improved.

1 On the QuickFinder Server Manager's Global Settings page, click *General* under *Default Settings*.

2 Click *Yes* next to *Enable query logging*.

To disable query logging, click *No* next to *Enable query logging*.

3 In the *Exclude these IP addresses* field, specify one or more IP addresses of the computers that you want excluded from the query log, separating each additional address by a space or hard return.

This feature enables you to eliminate query logging for searches originating within your organization. Keeping query log information gathered from customer queries separate from log information related to employees of your own company can help you identify issues that directly affect your customers.

4 From the *Log period* drop-down list, select a log period of either *Monthly* or *Weekly*.

QuickFinder records search requests in the current log file until the end of the specified period. When QuickFinder switches to a new log file, a report for the old log file is automatically generated.

5 In the *E-mail log reports* field, specify one or more e-mail addresses where you want copies of each new log report sent, separating each additional address by a space or hard return.

If this field does not appear, it is because you have disabled e-mail services on your QuickFinder Server. To enable e-mail services, go to the QuickFinder Global Settings page and click *Services Settings* > *General* > *E-Mail Settings*.

6 To view an existing log report, select one from the *Available log reports* field, then click View *Report*.

If there are no logs listed, you have not yet generated a log report. Click *Generate Current* to create a report of the current (not yet completed) log period. Depending on the amount of information in the logs, this process could take anywhere from a few seconds up to several hours to complete. After a new report is generated, it appears in this field.

7 (Optional) If you want to specify a different logging template, specify the name of the template in the *Template to use when generating reports* field.

QuickFinder provides the following templates at install time:

- **ReportTemplate.html:** Produces an HTML summary report of the query log data.
- **ExportTemplate.xml:** Exports the raw query log data into an XML file format that can be easily consumed by other reporting tools such as Crystal Reports, databases, and executive information systems.

The templates are located in the virtual search server directory (for example `/qfsearch/sites/default`), and the logs and reports are located in `/qfsearch/sites/default/logs`. See Section 12.10, “Query Report Template Variables,” on page 148 for more information on editing or creating you own query report templates.

8 Click *Apply Settings*.

Modifying Response Settings

1 On the QuickFinder Server Manager's Global Settings page, click *General* under *Default Settings*.

2 Under *Response Settings*, select an output encoding from the *Default encoding for response pages* drop-down list.

This setting specifies the encoding QuickFinder should use when responding to user queries using the search and print results templates, and the error and response messages templates.

3 Specify the maximum number of potential search results (hits) in the *Refuse queries if potential hits exceed* field that QuickFinder allows before cancelling the search.

A *potential hit* refers to all initial occurrences of a user's search terms in an index. After all matches between search terms and their related documents are found in the index, filters are applied to narrow the list of search results. Determining a user's access rights to individual files when performing a rights-based search is an example of a filter.

If a search term is too vague, it can produce an excessively large list of potential hits, which can place an unnecessary burden on your server. By limiting the number of potential hits allowed, you can reduce the processing overhead that would otherwise be placed on your server. If a user's search query produces a number of potential hits that exceed the number you specify, a message is returned to them suggesting that they be more specific in their search query.

Although a hit might be found in the index, it might not yet be authorized by Novell® eDirectory™ and the file system to determine access rights. The number you specify here is critical in situations where a user might have access to only a few documents in an index that contains thousands of files.

4 Click *Apply Settings*.

Modifying Error Log Settings

1. On the QuickFinder Server Manager's Global Settings page, click *General* under *Default Settings*.
2. In the *Maximum log size* field, specify the maximum size (in bytes) that QuickFinder should allow the log file to grow to.

   Depending on the number of visitors that your virtual search servers hosts, log files can become large. This setting protects your system's hard drive resources.

   The number you specify here is divided evenly between two log files. For example, if you specify 30000, each log file allows up to 15000 bytes of logged data. This ensures that you always have at least 15000 bytes of logged synchronization data.
3. Click *Apply Settings*.

Modifying E-Mail Settings

1. On the QuickFinder Server Manager's Global Settings page, click *General* under *Default Settings*.
2. (Optional) Click *Yes* next to *Enable e-mail services* if you want QuickFinder to generate an e-mail message when errors occur during the generation (or regeneration) of your virtual search server's indexes, or when your indexes are being synchronized with other QuickFinder servers.

   IMPORTANT: If the E-mail Settings feature is not visible, you must enable it from the *Services Settings* > *General* page. See "Configuring General Services Settings" on page 51.

   ---
3. (Optional) If you enabled e-mail services, type one or more e-mail addresses indicating where error messages should be sent.

   The addresses you enter here become default addresses for each virtual search server. They can be removed from within each virtual search server, or more addresses can be added. Separate each address using a space, comma, semicolon, or return.
4. Click *Apply Settings*.

7.2.2 Configuring Default Search Settings

Default Search Settings lets you modify default search query, response, and template settings for all new virtual search servers.

- "Modifying Search Query Settings" on page 40
- "Modifying Search Response Settings" on page 42
- "Customizing (or Deleting) Highlighter Color Schemes" on page 42
- "Modifying Search Template Settings" on page 43

Modifying Search Query Settings

1. On the QuickFinder Server Manager's Global Settings page, click *Search* under *Default Settings*.
2. (Optional) Under *Query Settings*, click *Yes* next to *Enable search term redirection* if you want to redirect specific queries to specific URLs.

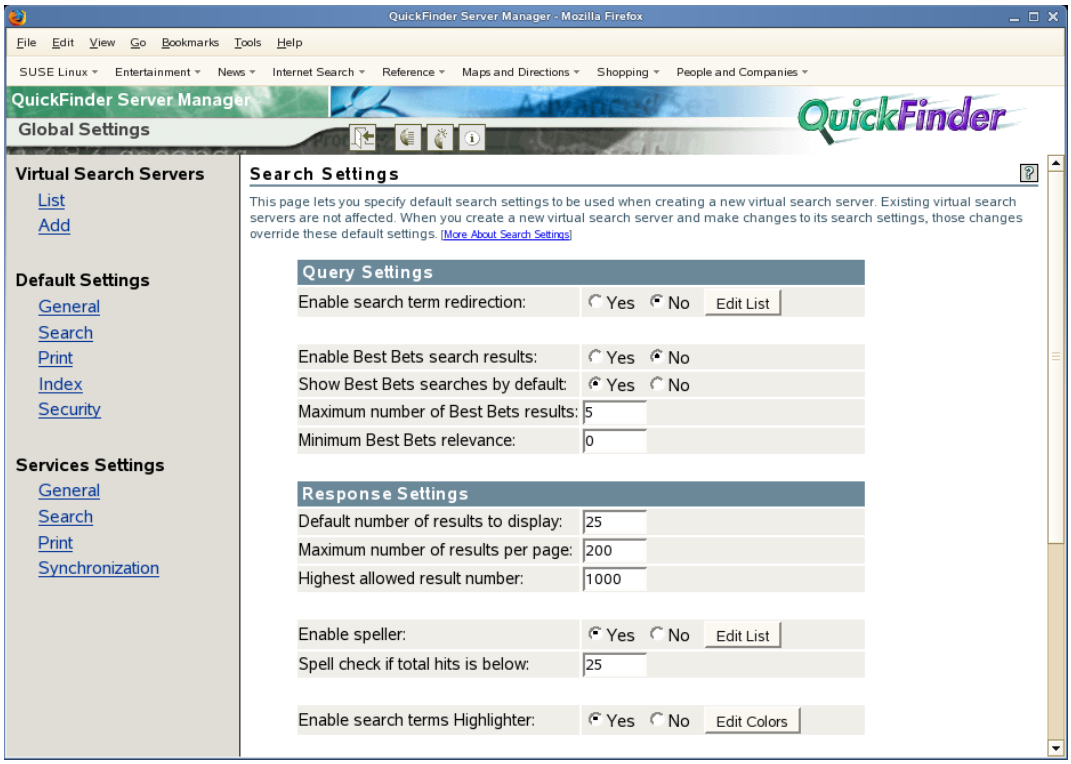

For example, searching for *one net* on www.novell.com redirects your browser to a special page designed to emphasize the Novell One Net strategy, rather than returning the search results page.

For more information about other QuickFinder features you can use to improve the search service experience for your users, see Chapter 10, “Optimizing Search Results,” on page 111.

3 (Optional) Click *Yes* next to *Enable Best Bets search results* if you want a specialized set of search results added to the top of the search results list.

Best bets is another method of helping users find what they are looking for more quickly. For more information, see “Using Best Bets to Help Users Find the Right Information” on page 117.

4 If you enabled best bets, click *Yes* next to *Show Best Bets searches by default* if you want the best bets results list to always appear on the search results page.

Click *No* if you want users to request the best bets results list at the time of their search request.

5 In the *Maximum number of Best Bets results* field, type the maximum number of best bets results to be included per search results page.

Consider keeping the list relatively short because the purpose of best bets is to narrow the search for users.

6 In the *Minimum Best Bets relevance* field, type the number of low relevance documents you want removed from the best bets display.

7 Click *Apply Settings*.

Modifying Search Response Settings

1. On the QuickFinder Server Manager's Global Settings page, click *Search* under *Default Settings*.
2. Under *Response Settings*, specify the number of search results in the *Default number of results to display* field that you want included on search results pages.

   For example, if you set this to 25 (the default setting) and the number of hits in a return was 200, QuickFinder would only return 25 hits per search results page.
3. Set a limit on the number of results allowed at one time on the results page by specifying a number in the *Maximum number of results to display* field.

   Use this setting (and the *Highest Allowed Result Number* setting) to protect your server from users who might attempt to set the number of results excessively high in order to intentionally abuse your computer. This can be done using available query parameters such as &NumHits=*number* and &StartHit=*number*. This setting blocks the use of such parameters.
4. Specify the highest number of search results that can be returned to a user query in the *Highest allowed result number* field.
5. (Optional) Click *Yes* next to *Enable speller* to have QuickFinder provide alternate spelling suggestions for search terms.

   Use the *Spell check if total hits is below* option to have QuickFinder spell-check search terms if the number of search results is lower than the number specified.
6. (Optional) Click *Yes* next to *Enable search terms highlighter* if you want users' search terms to be highlighted in the documents returned in a search.

   When the highlighter is enabled, a special icon appears on the search results page that users can click to show highlighted words in context. For instructions about how to customize the colors used by the highlighter, see "Customizing (or Deleting) Highlighter Color Schemes" on page 42.
7. Click *Apply Settings*.

Customizing (or Deleting) Highlighter Color Schemes

1. On the QuickFinder Server Manager's Global Settings page, click *Search* under *Default Settings*.
2. Click *Apply Settings* > *Back to Search Settings* to save any changes you have made.

   Clicking *Edit Colors* causes any unsaved changes to be lost
3. Under *Response Settings*, click *Edit Colors* next to *Enable search terms highlighter*.
4. Click the *Color* button in the *Foreground* and *Background* columns to select a color.

   or

   Type the color value (such as #RRGGBB) or color names. You can review your selections in the *Preview* column.

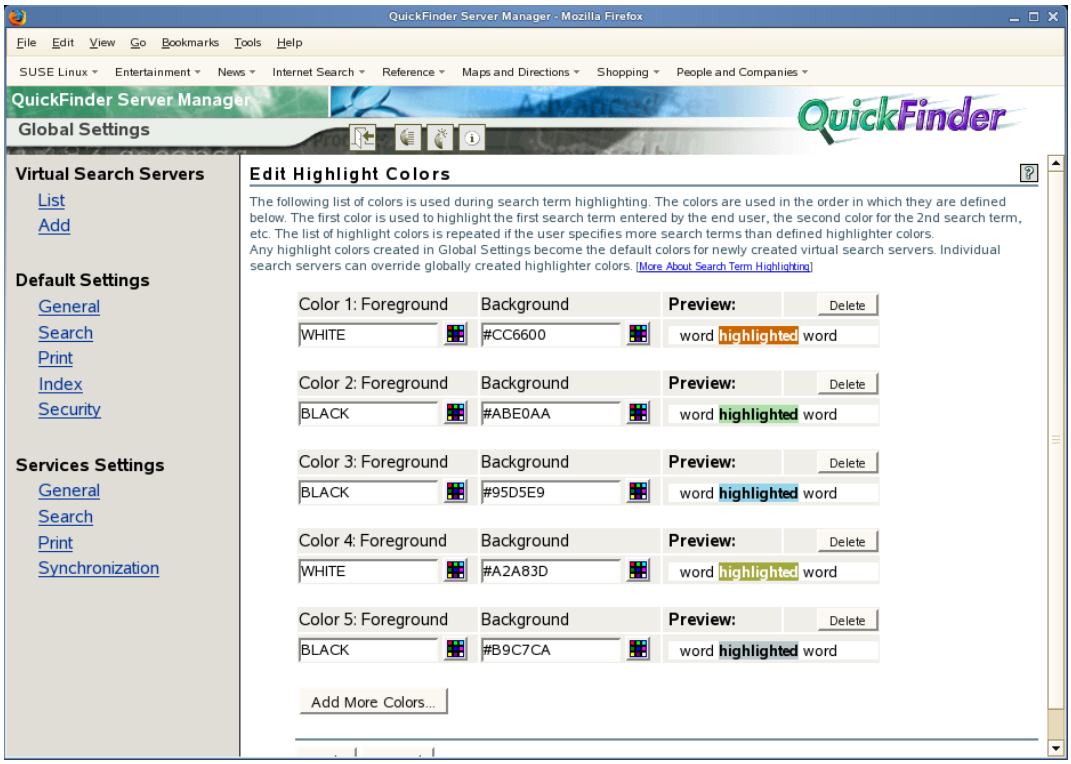

5 If you want to define additional colors, click *Add More Colors*.

Because a user can search for more than one search term at a time, you can define multiple colors so that each search term appears in a different color. For example, if a user searched for "HTTP AND Apache", the terms appear highlighted in the documents where they originated. The colors are used in the order in which they are defined here.

6 To remove a defined color set, click *Delete* in the row of the color set you want removed.

7 Click *Apply*.

Modifying Search Template Settings

1 On the QuickFinder Server Manager's Global Settings page, click *Search* under *Default Settings*.

2 Under *Template Settings*, specify a path to where your QuickFinder templates are stored in the *Templates Directory* field.

The default path is `volume:\searchroot\Templates` on NetWare and `/var/lib/qfsearch/Templates` on Linux. If you have created custom templates, or you want to keep your templates elsewhere, specify the path here so that QuickFinder knows where the templates are.

3 From the *Default encoding for templates* drop-down list, select the character set that your templates are written in.

This value is used with templates that do not specify an encoding. Encodings found in templates that do not match the encoding you specify here override this encoding.

4 In the *Default search page template* field, specify the filename of the search page template you want to use.

If you have created a custom template and want QuickFinder to use it as your search page, specify its name in this field.

5 In the *Default search results template* field, specify the filename of the search results template you want to use.

As with the default search page template, if you have created a custom search results template and want QuickFinder to use it as your default search results page, specify its name in this field.

6 In the *Default highlighter template* field, specify the filename of the highlighter template you want to use.

As with the default search and results page templates, if you have created a custom highlighter template and want QuickFinder to use it as your default highlighter page, specify its name in this field.

7 In the *Template to use if no results returned* field, specify the filename of the template that QuickFinder should return if no results are found.

8 In the *Template to use if error occurs* field, specify the filename of the template that QuickFinder should return if there are errors while processing a user's query.

9 Click *Apply Settings*.

For more information about QuickFinder templates, see Chapter 11, “Understanding Templates,” on page 127.

7.2.3 Configuring Default Print Settings

Default Print Settings lets you modify default print results and print template settings for all new virtual search servers.

- “Modifying Default Print Result Settings” on page 44
- “Modifying Default Print Template Settings” on page 45

Modifying Default Print Result Settings

1 On the QuickFinder Server Manager's Global Settings page, click *Print* under *Default Settings*.

2 Under *Print Results Settings*, specify the number of print results in the *Default number of results to print* field that you want displayed on each print results page.

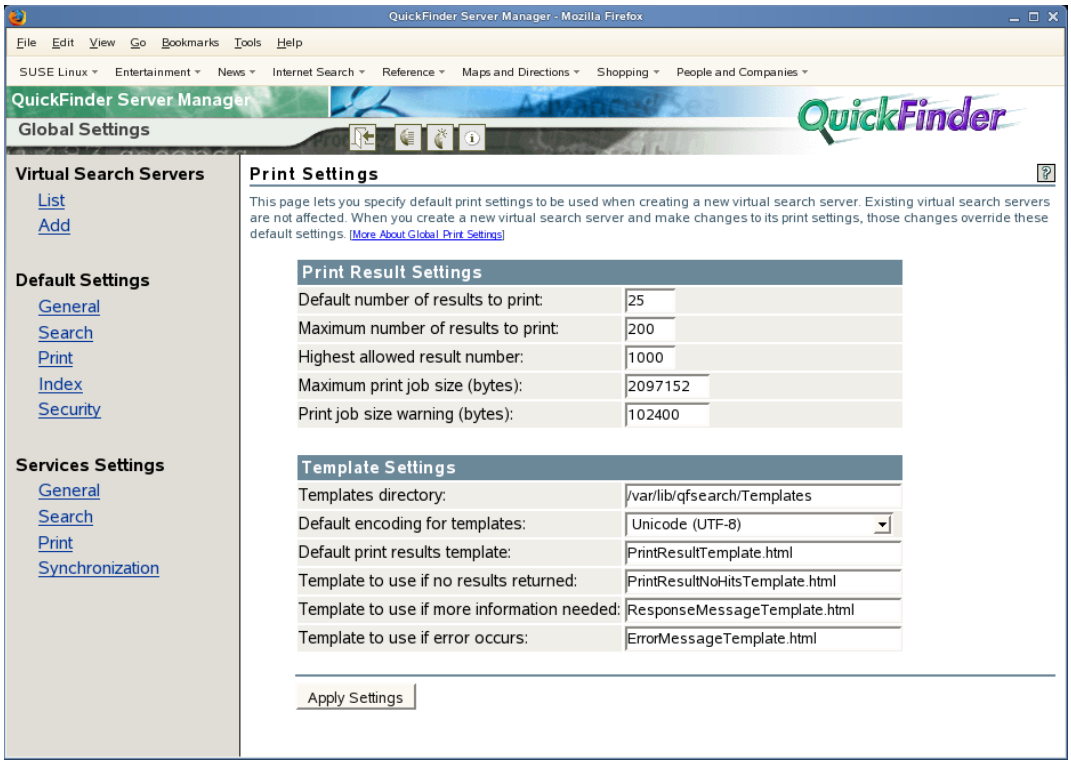

For example, if you set this to 25 (which is the default setting) and the number of hits in a return was 200, QuickFinder would only return 25 hits per print results page at a time.

3 Set a limit on the number of results allowed at one time on the results page by specifying a number in the *Maximum number of results to print* field.

4 Specify the highest number of search results that can be returned to a user query in the *Highest allowed result number* field.

5 To limit the size of a print job, specify the largest print job size that QuickFinder should allow in the *Maximum print job size* field.

Any users requesting a print job larger than this value receive a message informing them that the request was too large.

This is a useful feature for administrators who want to keep down the size of print jobs in their own companies, departments, or organizations.

6 To have users notified when a print job exceeds a certain size, specify the print job size in the *Print job size warning* field.

By default, this message is sent by using the `ResponseMessageTemplate.html` file and is intended as a warning to users that they are exceeding the allowed print job size. It then prompts the user to confirm the print job before continuing.

7 Click *Apply Settings*.

Modifying Default Print Template Settings

1 On the QuickFinder Server Manager's Global Settings page, click *Print* under *Default Settings*.

2 Under *Template Settings*, specify a path in the *Templates directory* field to where your QuickFinder templates are stored.

The default path is *volume*:*searchroot*\Templates on NetWare and /var/lib/qfsearch/Templates on Linux. If you have created custom templates, or you want to keep your templates elsewhere, specify the path here so that QuickFinder knows where the templates are.

3 From the *Default encoding for templates* drop-down list, select the character set that your templates are written in.

This value is used with templates that do not specify an encoding. Encodings found in templates that do not match the encoding you specify here override this encoding.

4 In the *Default print results template* field, specify the filename of the print results template you want to use.

If you have created a custom print results template and want QuickFinder to use it when returning print results, specify its name in this field.

5 In the *Template to use if no results returned* field, specify the filename of the template that QuickFinder should return if no print results match a user's query.

6 In the *Template to use if more information is needed* field, specify the filename of the template to be sent back to users whose print jobs exceed the size you specify in the *Print Job Size* field. (See Step 6 in "Modifying Default Print Result Settings" on page 44.)

7 In the *Template to use if error occurs* field, specify the filename of the template that QuickFinder should return if there are errors while processing a user's print query.

8 Click *Apply Settings*.

7.2.4 Configuring Default Index Settings

These settings are intended to make the process of creating indexes even easier by letting you configure common settings as default settings. This saves you time because you don't need to make the same selections each time you create a new index.

- "Modifying Default Index Settings" on page 46
- "Allowing Indexes to Be Used with QuickFinder Synchronization" on page 48
- "Enabling or Disabling User Authentication When Accessing Indexes" on page 48

Modifying Default Index Settings

1 On the QuickFinder Server Manager's Global Settings page, click *Index* under *Default Settings*.

2 Select the type of index that you want as the default index type on the Indexing Management page.

Each time you create a new index, the index type you select here is automatically selected. You can change the default index type when you create an index.

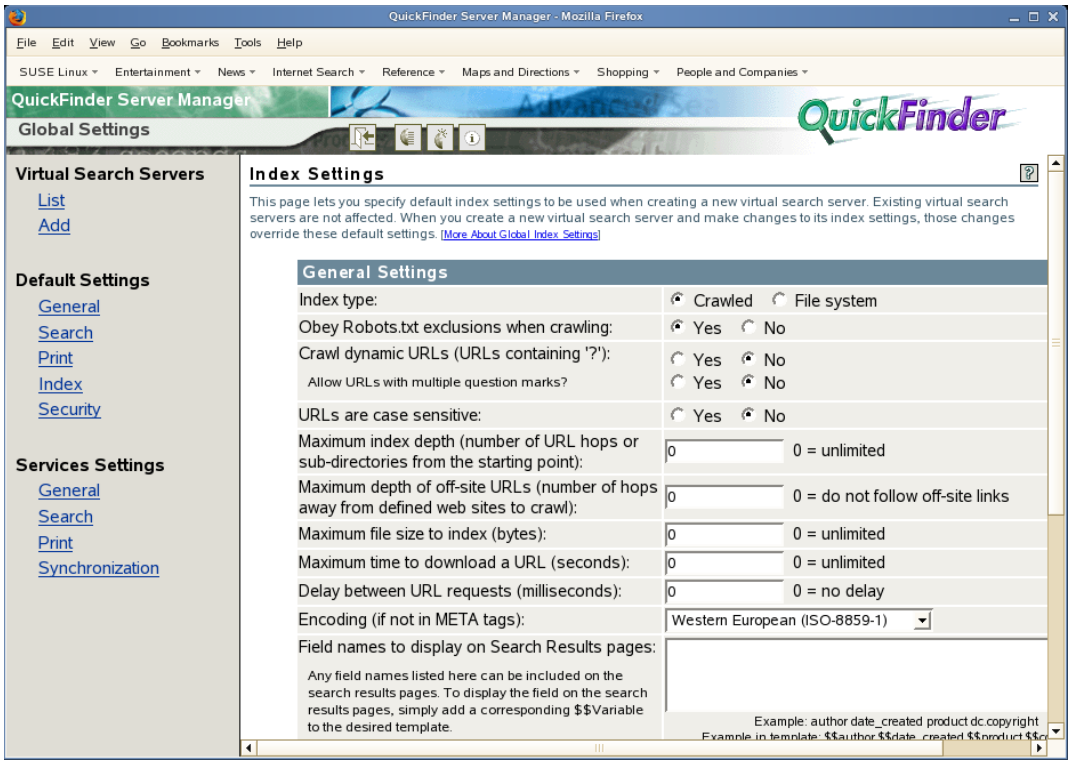

3 Click *Yes* next to *Obey Robots.txt exclusions when crawling* to instruct QuickFinder to obey the instructions in the `Robots.txt` file as it indexes Web sites.

For more information about `Robots.txt`, see "Using the Robots Meta Tag" on page 113.

4 Click *Yes* next to *Crawl dynamic URLs (URLs Containing '?')* if you want QuickFinder to index dynamic content, in addition to static content.

Some search engines do not crawl Web pages that are created dynamically, such as pages generated from forms submitted by a Web browser. The URLs of dynamic content typically contain a question mark (?) followed by additional parameters. (See Section 8.7, "Indexing Dynamic Web Content," on page 98.)

Because dynamic content can change at any time, you might want to schedule more frequent regeneration events for your indexes when enabling this feature. For information about scheduling automatic updates to your indexes, see Section 8.9, "Automating Index and Server Maintenance," on page 100.

5 Click *Yes* next to *Allow URLs with multiple question marks* if you want QuickFinder to crawl URLs with multiple query strings (parameters).

6 Select the *URLs are case sensitive* option if you want QuickFinder to recognize URLs that are different only in character case, but are otherwise identical. For example, www.digitalairlines.com versus www.DigitalAirlines.com.

IMPORTANT: Setting this option to *No* can help QuickFinder avoid indexing duplicate information, which can come from indexing URLs that are presented using different cases but actually point to the same information. However, if a Web server being indexed is configured to differentiate between cases, QuickFinder might leave out content that you want indexed.

---

7 In the *Maximum index depth* field, specify the maximum number of hypertext links from the starting URL that QuickFinder should follow before it stops indexing.

8 In the *Maximum depth of off-site URLs* field, specify the maximum number of hypertext links off-site from the starting URL that QuickFinder should follow before it stops indexing.

9 Specify a number (in bytes) in the *Maximum file size to index* field to keep QuickFinder from indexing files larger than the number you specify.

10 In the *Maximum time to download a URL* field, specify a number (in seconds) before QuickFinder automatically skips the indexing of the specified URL.

11 Type a number (in milliseconds) in the *Delay between URL requests* field that QuickFinder should pause between requests for URLs that it is trying to index.

12 From the *Encoding (if not in META tags)* drop-down list, select the encoding to be used when indexing files that do not contain an encoding specification.

For example, HTML files can specify their encoding with a Content-Type meta tag.

13 In the *Field names to display on Search Results pages* field, list any field names that you want to be included on the search results page (for example, author date_created product dc.copyright).

To display the field on the search results pages, add the corresponding $$Variable to the template (for example, $$author $$date_created $$product $$copyright).

The field data is stored in the index and causes the index size to increase.

14 Click *Apply Settings*.

Allowing Indexes to Be Used with QuickFinder Synchronization

1 On the QuickFinder Server Manager's Global Settings page, click *Index* under *Default Settings*.

2 Under *Synchronization Settings*, click *Yes* next to *Index may be copied to other clustered servers*.

For more information about QuickFinder Synchronization, see Chapter 9, “Synchronizing Data Across Multiple QuickFinder Servers,” on page 103.

3 Click *Apply Settings*.

Enabling or Disabling User Authentication When Accessing Indexes

1 On the QuickFinder Server Manager's Global Settings page, click *Index* under *Default Settings*.

2 (Optional) Under *Rights-based Search Results*, click *Off* to allow anyone to access the indexes.

No user authentication is required to search the indexes of the current virtual search server.

3 (Optional) Under *Rights-based Search Results*, click *by Index* to restrict access to the indexes only to those who have access rights to the filename you specify in the *by Index* field.

4 In the *by Index* field, specify the full path to a file on your server that already has rights protection assigned to it.

For example, on a NetWare 6.5 server where Apache is your Web server, you might type `sys:\\Apache2\htdocs\index.html`.

5 Select either *Search Engine* or *Templates* from the *Unauthorized hits filtered by* drop-down list to specify how QuickFinder should handle attempts to access indexes without authorization.

- Select *Engine* if you don't want any search results displayed when users attempt to search the index without first logging in. A No Results Found message is returned instead.

- Select *Templates* to have QuickFinder return search results to users who have not logged in. When they attempt to click a search result, they are then prompted for a username and password.

TIP: You can use the $$BeginUnAuthorized and $$EndUnAuthorized template variables to control how the template should display unauthorized hits. For more information about template variables, see Chapter 12, "Working with Template Variables and Search Parameters," on page 133.

---

6 Click *Apply Settings*.

7.2.5 Configuring Default Security Settings

Security settings let you manage access to indexed content by requiring users to authenticate to a server before seeing rights-protected search results. The Security Settings page lets you configure rights-based search results and specify secure connection settings, including enabling the HTTPS protocol for securing usernames and passwords as they are sent over the network.

- "Modifying Default Rights-based Search Results" on page 49
- "Modifying Default Connection Settings" on page 50

Modifying Default Rights-based Search Results

1 On the QuickFinder Server Manager's Global Settings page, click *Security* under *Default Settings*.

2 In the *Default search contexts* field, type the eDirectory context where user objects are stored.

If you provide a context here, users won't need to supply their fully qualified user IDs when logging in to QuickFinder. For example, if your user objects are all kept in a container named employees, type `employees` in the *Default search contexts* field so that QuickFinder knows where to look when a user attempts to log in. A user whose fully qualified name is sjones.employees.marketing.digitalairlines would only need to enter sjones.

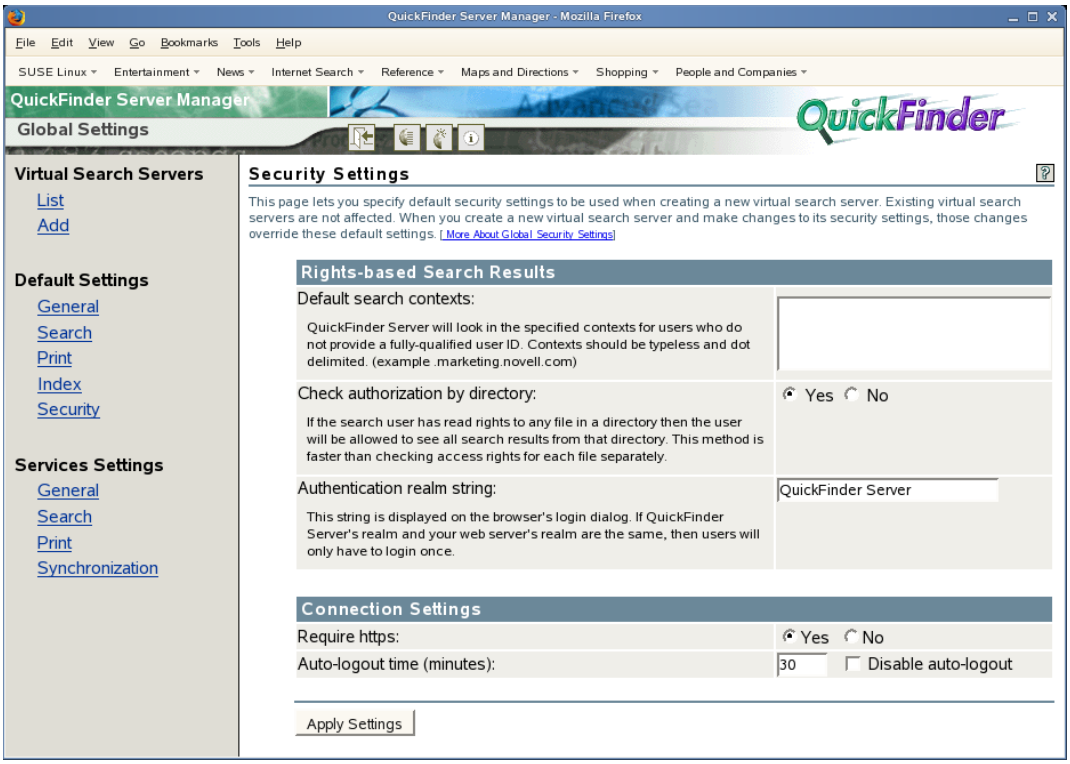

3 Click *Yes* next to *Check authorization by directory* if you want QuickFinder to validate user rights to files in a directory (search results) based on whether or not users have rights to any one file in the same directory.

4 In the *Authentication realm string* field, specify the correct authentication realm string, which is typically shown in your Web browser's login dialog box.

Specifying the Apache Web server's authentication realm string in this field makes it so that after users authenticate to it, they won't need to authenticate again when using QuickFinder to search and access protected information.

If Apache is your Web server, refer to http://httpd.apache.org/docs-2.0/howto/auth.html (http://httpd.apache.org/docs-2.0/howto/auth.html) for more information about authentication.

5 Click *Apply Settings*.

Modifying Default Connection Settings

1 On the QuickFinder Server Manager's Global Settings page, click *Security* under *Default Settings*.

2 Under *Connection Settings*, click *Yes* next to *Require https* if you want to protect usernames and passwords as they are sent across the network or Internet.

3 Specify a number (in minutes) in the *Auto-logout time* field to direct QuickFinder to log users out who have been idle for the specified period of time.

4 Click *Apply Settings*.

7.3 Configuring Services Settings

Services Settings are meant for global QuickFinder server administrators, giving them global control of the QuickFinder server and all virtual search servers, including the ability to completely disable a virtual search server. Services Settings also let the administrator control the overall performance of the QuickFinder server.

Unlike the Default Settings, changes made here affect both new and existing virtual search servers. For example, if the e-mail feature is enabled from the Default General Settings page (accessed from the Global Settings page), and on the General Settings page of a virtual search server, the global administrator can disable it from General Services Settings. When that is done, the e-mail feature is disabled and its settings are removed from all other configuration pages.

- Section 7.3.1, “Configuring General Services Settings,” on page 51
- Section 7.3.2, “Configuring Search Services Settings,” on page 54
- Section 7.3.3, “Configuring Print Services Settings,” on page 56
- Section 7.3.4, “Configuring Synchronization,” on page 58

7.3.1 Configuring General Services Settings

General Services Settings let you configure product update, error log, and e-mail settings. If you disable a feature here, it is disabled on all virtual search servers and removed from their configuration pages. If you re-enable the feature, it reappears on all configuration pages with the same settings used before it was disabled by the global administrator.

- “Modifying Product Update Settings” on page 51
- “Modifying Administrator Authentication Settings” on page 52
- “Modifying Error Log Settings” on page 52
- “Modifying E-Mail Settings” on page 53
- “Modifying Server Management Settings” on page 53

Modifying Product Update Settings

1. On the QuickFinder Server Manager's Global Settings page, click *General* under *Services Settings*.
2. (Optional) Under *Product Update Settings*, click *Yes* next *Check for product updates* if you want QuickFinder to regularly check for software updates and notify you when updates are available.

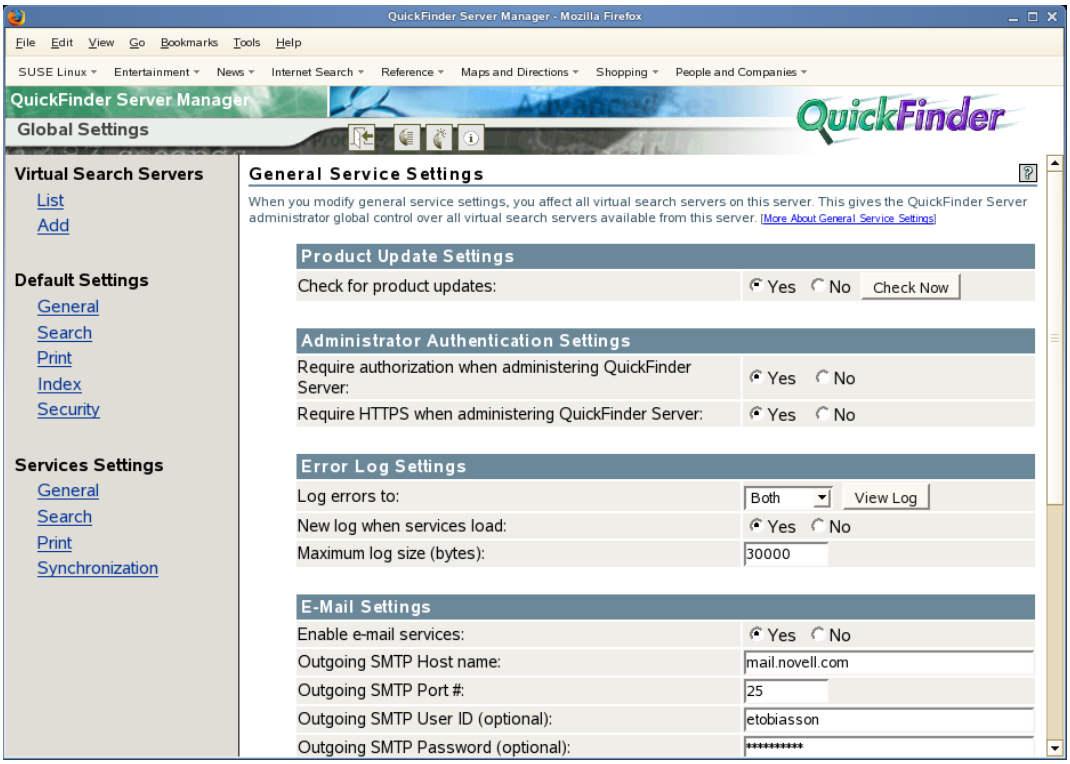

For more information, see Section 7.4, “Updating QuickFinder Server Software,” on page 59.

3 Click *Apply Settings* > *Back to General Service Settings* to save any changes you have made.

Clicking *Check Now* causes any unsaved changes to be lost.

4 To check for software updates immediately, click *Check Now*.

5 Click *Apply Settings*.

Modifying Administrator Authentication Settings

1 On the QuickFinder Server Manager's Global Settings page, click *General* under *Services Settings*.

2 Under *Administrator Authentication Settings*, choose from the following options:

- **Require authorization when administering QuickFinder Server:** Click *Yes* to require a login name and password when accessing QuickFinder Server Manager.
- **Require HTTPS when administering QuickFinder Server:** Click *Yes* to require the use of the secure HTTP protocol (HTTPS) when accessing QuickFinder Server Manager.

3 Click *Apply Settings*.

Modifying Error Log Settings

1 On the QuickFinder Server Manager's Global Settings page, click *General* under *Services Settings*.

2 Under *Error Log Settings*, select where you want log results displayed by choosing one of the following options from the *Log errors to* drop-down list:

- **File:** When this option is selected, you can click *View* next to the *Log Errors To* drop-down list to display the log results in your browser.
- **Console:** You can also view log results at the NetWare system console by selecting *Console*, pressing Ctrl+Esc on your server's keyboard, then pressing the number corresponding to the Tomcat servlet engine.
- **Both:** Displays results in both your browser and at the system console.

You can access the log file directly by going to `volume:\\searchroot\errors.log` on NetWare or `/var/lib/qfsearch/Errors.log` on Linux.

3 To start a new log file each time you restart the QuickFinder server, click *Yes* next to *New log when services load*.

You can also delete the log file at the path specified above. The log file is re-created on the first instance of a new error, statistics, etc.

4 To limit the size of the log file, specify a file size (in bytes) in the *Maximum log size* field.

The number you specify here is divided evenly between two log files. For example, if you specify 30000, each log file allows up to 15000 bytes of logged data. This ensures that you always have at least 15000 bytes of logged synchronization data.

5 Click *Apply Settings*.

Modifying E-Mail Settings

1 On the QuickFinder Server Manager's Global Settings page, click *General* under *Services Settings*.

2 Under *E-mail Settings*, click *Yes* next to *Enable e-mail services* if you want QuickFinder to generate an e-mail message whenever errors occur during the generation (or regeneration) of your virtual search server's indexes, or when your indexes are being synchronized with other QuickFinder servers.

3 In the *Outgoing SMTP Host name* field, specify the SMTP host name of the mail server that handles e-mail requests (for example, mail.novell.com).

4 In the *Outgoing SMTP Port #* field, specify the SMTP port number for the mail server specified above.

Typically, the port number for mail servers is 25.

5 (Optional) If your SMTP server requires authentication the, specify your user ID in the *Outgoing SMTP User ID* field.

6 (Optional) Specify the matching password for your SMTP user account in the *Outgoing SMTP Password* field.

7 (Optional) Specify an e-mail address in the *From Address* field.

8 Click *Apply Settings*.

Modifying Server Management Settings

1 On the QuickFinder Server Manager's Global Settings page, click *General* under *Services Settings*.

2 Under *Server Management Settings*, specify a number in the *Maximum number of Active Index Jobs* field to limit the number of indexing jobs that can run simultaneously.

   Any indexing jobs beyond the number you specify here are placed in a queue until one of the original indexing jobs is completed.

   Each server's performance varies because of hardware configuration and the amount of user traffic, so you should experiment with how many indexing jobs you can run simultaneously without negatively affecting the performance of your server. Also, you might consider using the *Index Scheduling* feature to regenerate indexes after business hours. See Section 8.9, "Automating Index and Server Maintenance," on page 100 for more information.

3 In the *Default location of virtual search servers* field, specify the path where you want all virtual search server files to be stored, including index and configuration files.

   Changing this setting won't move existing sites to the new default location, but all new virtual search servers will be placed here.

4 To direct QuickFinder to reload configuration files modified manually, outside of QuickFinder Server Manager, click *Yes* next to the *Detect manual search server changes* field.

   If you make changes outside of QuickFinder Server Manager, such as modifying configuration or properties files or any of the templates, QuickFinder re-reads those files as often as you indicate in the *Seconds between checking for changes* field.

   If you modify the `adminservlet.properties` file, you must reboot your server in order for the changes to take effect. This is the only file that cannot be re-read using the *Detect manual search server changes* feature.

5 In the *Seconds between checking for changes* field, specify how often QuickFinder should check for manual changes (changes made outside of QuickFinder Server Manager) to the configuration files.

6 To direct QuickFinder to reload QuickFinder templates that have been modified, click *Yes* next to the *Detect template changes* field.

   After making a change to a template from within your HTML editing tool and saving it on your server, QuickFinder re-reads the template as often as you specify in the *Seconds between template updates* field. This means you can test your changes almost immediately.

7 In the *Seconds between checking for template changes* field, specify how often QuickFinder should reload search, print, results, and error templates.

8 Click *Apply Settings*.

7.3.2 Configuring Search Services Settings

Search Services Settings let you turn search capabilities on or off and manage debugging and statistics settings.

- "Modifying General Search Services Settings" on page 54
- "Modifying Search Debug Settings" on page 55
- "Modifying Search Statistics Settings" on page 56

Modifying General Search Services Settings

1 On the QuickFinder Server Manager's Global Settings page, click *Search* under *Services Settings*.

2 Under *General Search Settings*, click *Yes* next to *Enable search service* to enable search services for all virtual search servers on your QuickFinder server.

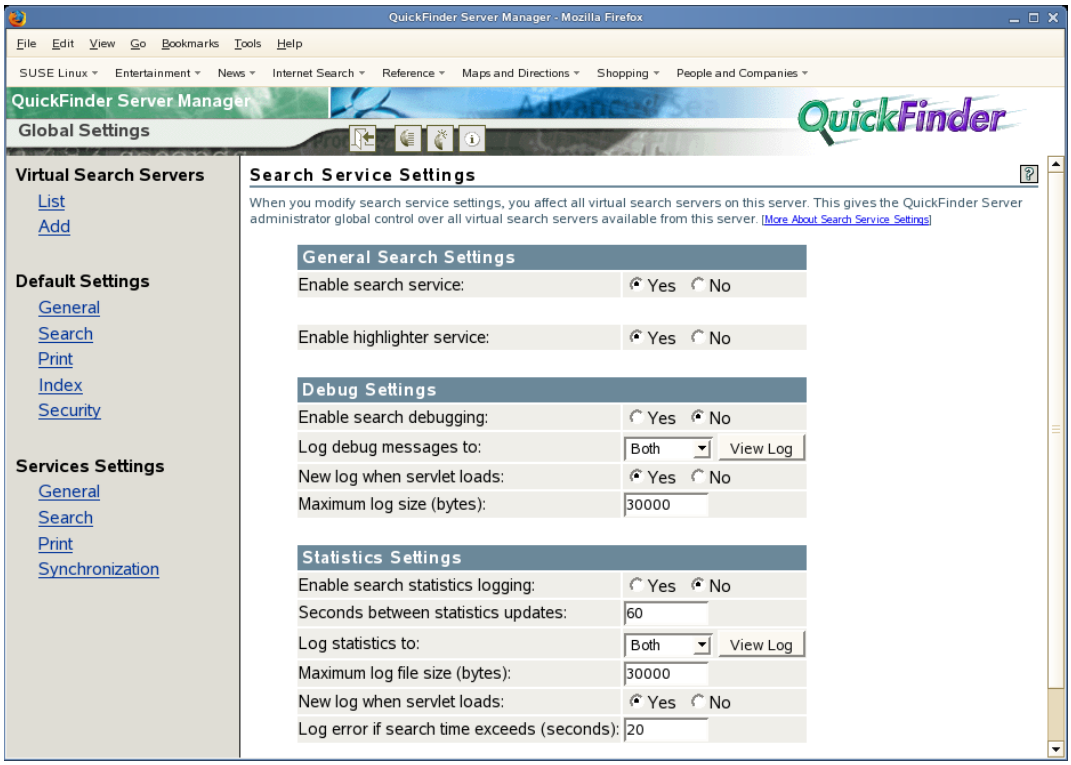

3 To enable the highlighter, click *Yes* next to *Enable highlighter service*.

4 Click *Apply Settings*.

Modifying Search Debug Settings

1 On the QuickFinder Server Manager's Global Settings page, click *Search* under *Services Settings*.

2 Under *Debug Settings*, click *Yes* next to *Enable search debugging* if you want to keep a log of all searches and query results going to all virtual search servers.

IMPORTANT: We recommend that you use this feature only while setting up or troubleshooting your search services because the log file can grow in size very quickly. In addition, logging query and response details can slow down searches.

---

3 Specify where you want log results displayed by selecting one of the following options from the *Log debug messages to* drop-down list:

- **File:** When this option is selected, you can click *View Log* next to the *Log debug messages To* drop-down list to display the log results in your browser.
- **Console:** You can also view log results at the system console by selecting *Console*, pressing Ctrl+Esc on your server's keyboard, then pressing the number corresponding to the Tomcat servlet engine.
- **Both:** Displays results in both your browser and at the system console.

4 To start a new log file each time you restart the QuickFinder server, click *Yes* next to *New log when servlet loads*.

5 To limit the size of the log file, specify a file size (in bytes) in the *Maximum log size* field.

The number you specify here is divided evenly between two log files. For example, if you specify 30000, each log file allows up to 15000 bytes of logged data. This ensures that you always have at least 15000 bytes of logged synchronization data.

6 Click *Apply Settings*.

Modifying Search Statistics Settings

1 On the QuickFinder Server Manager's Global Settings page, click *Search* under *Services Settings*.

2 Under *Statistics Settings*, click *Yes* next to *Enable search statistics logging* if you want an updated log file containing statistics about searches performed against all virtual search servers on your QuickFinder server.

3 In the *Seconds between statistics updates* field, specify a number (in seconds) that should elapse between updates of the statistics log file.

4 Specify where you want log results displayed by selecting one of the following options from the *Log statistics to* drop-down list:

- **File:** When this option is selected, you can click *View Log* next to the *Log statistics to* drop-down list to display the log results in your browser.
- **Console:** You can also view log results at the system console by selecting *Console*, pressing Ctrl+Esc on your server's keyboard, then pressing the number corresponding to the Tomcat servlet engine.
- **Both:** Displays results in both your browser and at the system console.

5 To limit the size of the log file, specify a file size (in bytes) in the *Maximum log file size* field.

The number you specify here is divided evenly between two log files. For example, if you specify 30000, each log file allows up to 15000 bytes of logged data. This ensures that you always have at least 15000 bytes of logged synchronization data.

6 To start a new log file each time you restart the QuickFinder server, click *Yes* next to *New log when servlet loads*.

7 In the *Log error if search time exceeds* field, specify a number (in seconds) before QuickFinder should record the current search as exceeding the specified time limit on the statistics display.

This appears as the *Limit* portion of the statistics display.

8 Click *Apply Settings*.

7.3.3 Configuring Print Services Settings

- "Modifying General Print Services Settings" on page 57
- "Modifying Print Debug Settings" on page 57
- "Modifying Print Statistics Settings" on page 58

Modifying General Print Services Settings

1 On the QuickFinder Server Manager's Global Settings page, click *Print* under *Services Settings*.

2 Under *General Print Settings*, click *Yes* next to *Enable print service* to enable print services for all virtual search servers on your QuickFinder Server.

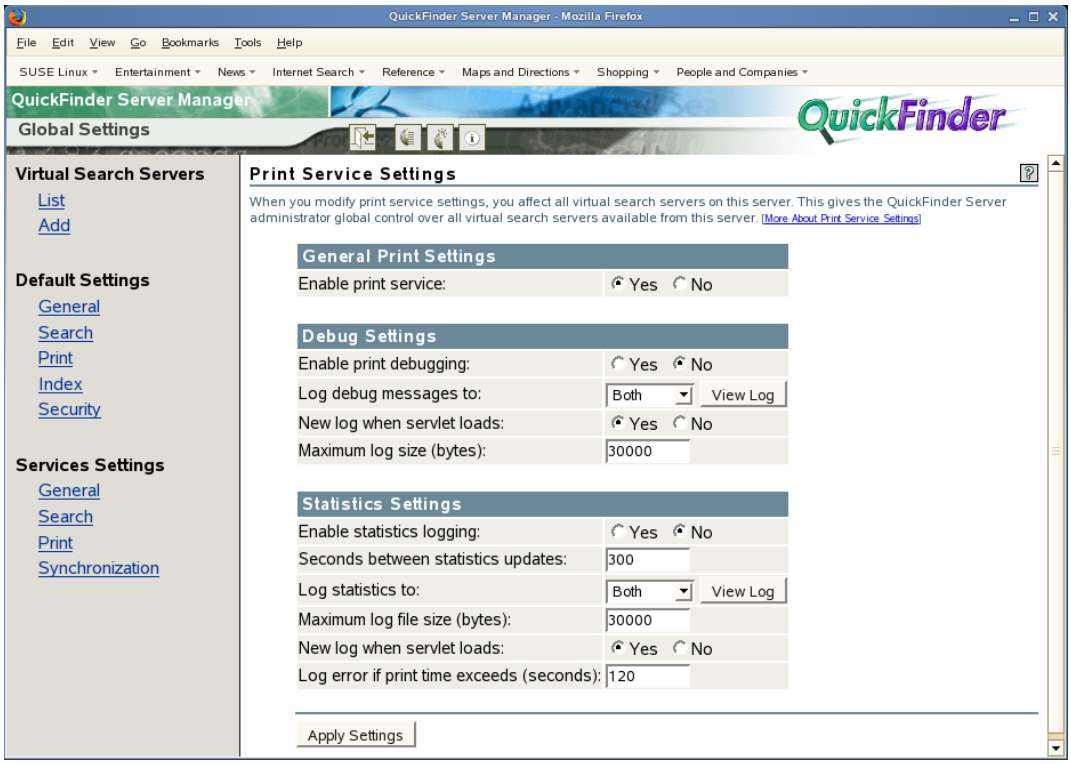

3 Click *Apply Settings*.

Modifying Print Debug Settings

1 On the QuickFinder Server Manager's Global Settings page, click *Print* under *Services Settings*.

2 Under *Debug Settings*, click *Yes* next to *Enable print debugging* if you want print debugging turned on.

IMPORTANT: We recommend that you use this feature only while setting up or troubleshooting your search services because the log file can grow in size very quickly. In addition, logging query and response details can slow down searches.

---

3 Specify where you want log results displayed by selecting one of the following options from the *Log debug messages to* drop-down list:

- **File:** When this option is selected, you can click *View Log* next to the *Log debug messages to* drop-down list to display the log results in your browser.
- **Console:** You can also view log results at the system console by selecting *Console*, pressing Ctrl+Esc on your server's keyboard, then pressing the number corresponding to the Tomcat servlet engine.

- **Both:** Displays results in both your browser and at the system console.

4 To start a new log file each time you restart the QuickFinder server, click *Yes* next to *New log when servlet loads*.

5 To limit the size of the log file, specify a file size (in bytes) in the *Maximum log size* field.

The number you specify here is divided evenly between two log files. For example, if you specify 30000, each log file allows up to 15000 bytes of logged data. This ensures that you always have at least 15000 bytes of logged synchronization data.

6 Click *Apply Settings*.

Modifying Print Statistics Settings

1 On the QuickFinder Server Manager's Global Settings page, click *Print* under *Services Settings*.

2 Under *Statistics Settings*, click *Yes* next to *Enable print statistics logging* if you want an updated log file containing statistics about print requests performed on your QuickFinder server.

3 In the *Seconds between statistics updates* field, specify a number (in seconds) that should elapse between updates of the statistics log file.

4 Specify where you want log results displayed by selecting one of the following options from the *Log statistics to* drop-down list:

- **File:** When this option is selected, you can click *View Log* next to the *Log statistics to* drop-down list to display the log results in your browser.
- **Console:** You can also view log results at the system console by selecting *Console*, pressing Ctrl+Esc on your server's keyboard, then pressing the number corresponding to the Tomcat servlet engine.
- **Both:** Displays results in both your browser and at the system console.

5 To limit the size of the log file, specify a file size (in bytes) in the *Maximum log file size* field.

The number you specify here is divided evenly between two log files. For example, if you specify 30000, each log file allows up to 15000 bytes of logged data. This ensures that you always have at least 15000 bytes of logged synchronization data.

6 To start a new log file each time you restart the QuickFinder server, click *Yes* next to *New log when servlet loads*.

7 In the *Log error if print time exceeds* field, specify a number (in seconds) before QuickFinder should record the current print job as exceeding the specified time limit on the statistics display.

This appears as the *Limit* portion of the statistics display.

8 Click *Apply Settings*.

7.3.4 Configuring Synchronization

QuickFinder Synchronization lets you designate one QuickFinder server as the search master from which updated indexes, templates, and configuration settings are systematically sent out to all other QuickFinder servers defined as part of a QuickFinder Synchronization cluster.

For more information about setting up and configuring QuickFinder Synchronization, see "Synchronizing Data Across Multiple QuickFinder Servers" on page 103.

7.4 Updating QuickFinder Server Software

As new QuickFinder software becomes available, you can be notified automatically from within QuickFinder Server Manager.

When new software is available, a message appears on the Global Settings page with a link indicating where you can download the latest version.

1 On the QuickFinder Server Manager's Global Settings page, click *General* under *Services Settings*.

2 (Optional) Under *Product Update Settings*, click *Yes* next *Check for product updates*.

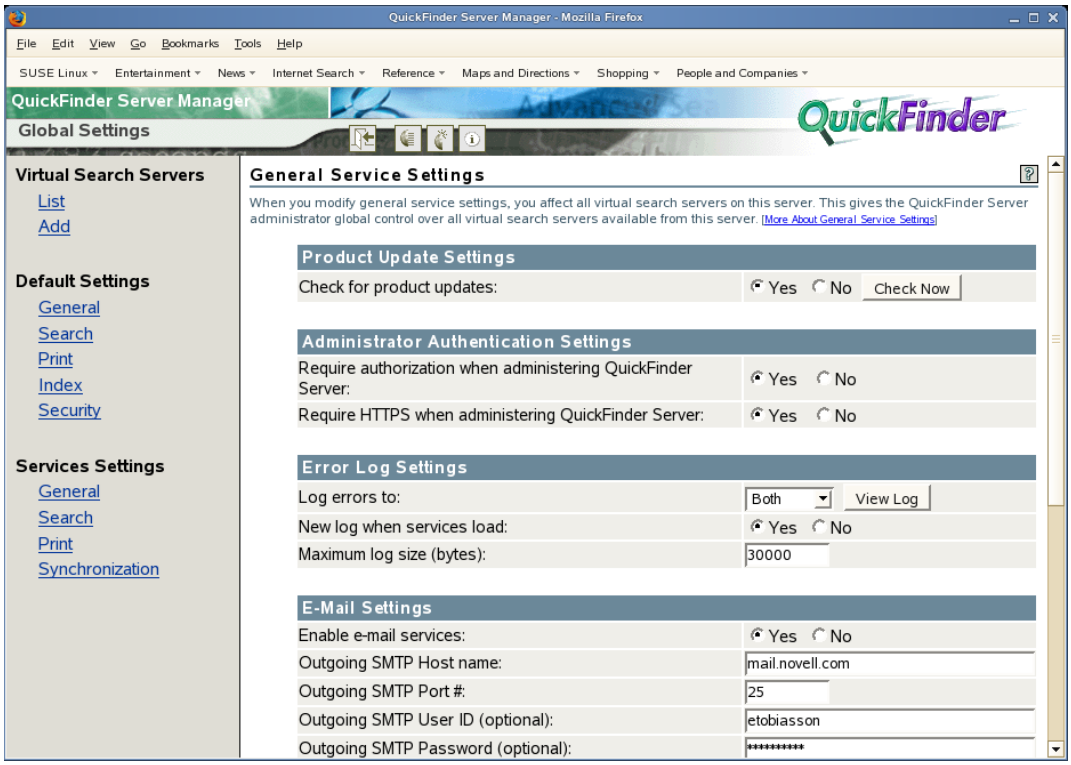

3 Click *Apply Settings* > *Back to General Service Settings* to save any changes you have made.

Clicking *Check Now* causes any unsaved changes to be lost.

4 To check for software updates immediately, click *Check Now*.

5 Click *Apply Settings*.

Creating and Managing Virtual Search Servers

8

This section provides detailed information about how to create and manage virtual search servers and their associated indexes using QuickFinder™ Server Manager. It contains the following topics:

- Section 8.1, “About Virtual Search Servers,” on page 61
- Section 8.2, “Creating Virtual Search Servers,” on page 61
- Section 8.3, “Enabling, Disabling, or Deleting a Virtual Search Server,” on page 64
- Section 8.4, “Configuring Virtual Search Servers,” on page 65
- Section 8.5, “Creating Indexes,” on page 83
- Section 8.6, “Managing Existing Index Files,” on page 96
- Section 8.7, “Indexing Dynamic Web Content,” on page 98
- Section 8.8, “Access Control to Search Results,” on page 99
- Section 8.9, “Automating Index and Server Maintenance,” on page 100
- Section 8.10, “Backing Up Your Virtual Search Server Files,” on page 101

8.1 About Virtual Search Servers

By definition, a virtual search server is a collection of one or more indexes and their related configuration files. Indexes and templates are at the heart of a virtual search server. An index is an optimized binary file that contains keywords found in documents hosted on a Web or file server. Indexes are used by QuickFinder to return search results to users' Web browsers. Templates are HTML documents that contain dynamic content. The dynamic content is generated by the use of search parameters and variables inserted in various places in the templates. Templates are customizable and are used for several things, including displaying search and print results to users.

For more information about templates, see Chapter 11, “Understanding Templates,” on page 127.

Before creating virtual search servers, particularly large or mission-critical ones, you should carefully plan how to best set up QuickFinder to meet requirements of your search services. A search service used by a small to medium department in a company requires different planning than a search service intended to serve thousands of customers on an enterprise wide Intranet Web site.

For information about how to plan an effective search service, see Section 3.2, “Designing Your Search Solution,” on page 24.

8.2 Creating Virtual Search Servers

After you have carefully planned your search service, you can start creating and configuring virtual search servers and begin adding indexes to them.

- Section 8.2.1, “Creating a Virtual Search Server,” on page 62
- Section 8.2.2, “Naming a Virtual Search Server,” on page 63
- Section 8.2.3, “Using the Virtual Search Server Alias,” on page 63

- Section 8.2.4, “Storing Virtual Search Server Files,” on page 64

8.2.1 Creating a Virtual Search Server

1 On the QuickFinder Server Manager Global Settings page, click *Add New Virtual Search Server*.

2 In the *Name* field, specify a new virtual search server name, which is typically the DNS or domain name of your server.

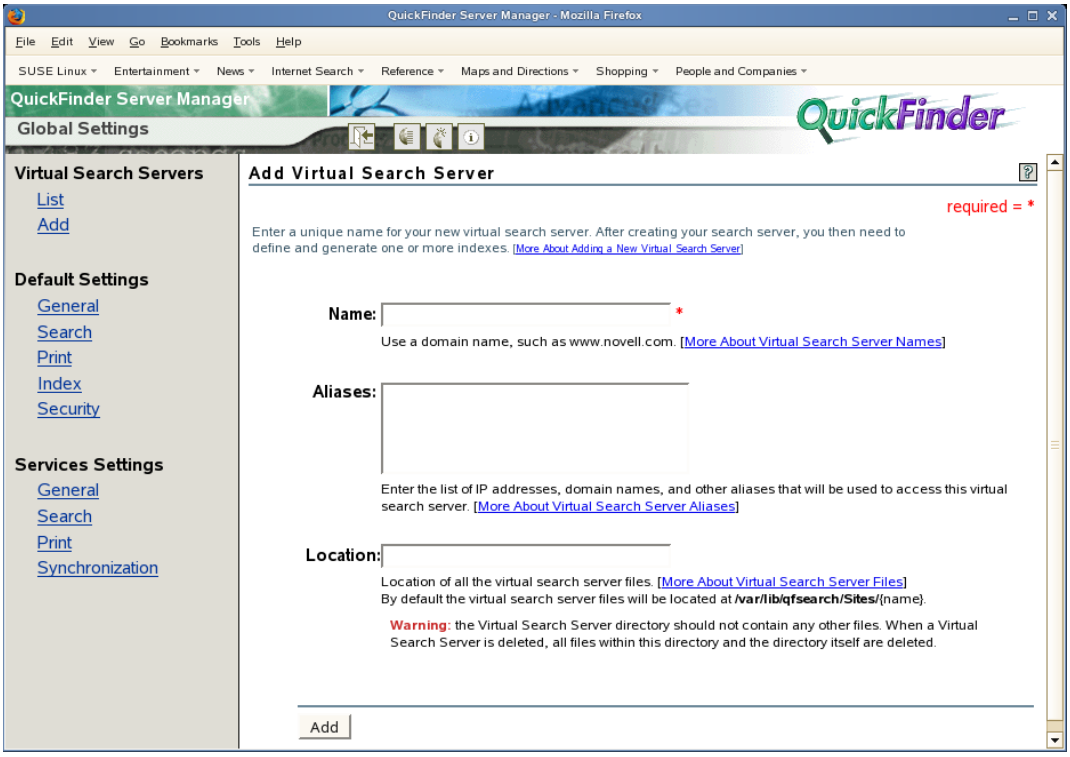

For more information about virtual search server names, see “Naming a Virtual Search Server” on page 63.

3 In the *Aliases* field, specify a virtual search server alias, which is typically the IP address of your server, the domain name, and any other alias that can be used to access this virtual search server.

See “Using the Virtual Search Server Alias” on page 63 for more information about aliases.

4 In the *Location* field, specify the path where you want the index and configuration files to be stored.

If this field is left blank, QuickFinder stores the virtual search server files in the `/searchroot/sites/sitename` directory. Also, you can store the files on any volume on the server where QuickFinder is installed, but not on other servers.

5 Click *Add*.

8.2.2 Naming a Virtual Search Server

When a user sends a search query to the QuickFinder Server, QuickFinder determines which of all of your virtual search servers it should use to handle the incoming search request.

QuickFinder uses two methods for determining this:

- Matching the domain name of the search query with the virtual search server names available in QuickFinder
- Using the server=*searchsitename* query parameter to find matching virtual search server names

For example, in the following search request, QuickFinder uses the domain name search.domainname1.com as the name of the virtual search server:

```
http://search.domainname1.com/qfsearch/
SearchServlet?query=find+something
```

This approach requires that your server be set up to recognize the domain name search.domainname1.com. Most servers can be set up to recognize and service multiple domain names in both software and hardware virtual server configurations.

You could also use an IP address to designate the virtual search server. For multiple virtual search servers, this approach works only in a hardware virtual server configuration where each virtual search server has its own unique IP address.

Hosting Multiple Virtual Search Servers Using One DNS Name

If you are hosting a search service for two or more customers, you can name each virtual search server according to the organization or company name of each customer and then use the &server query parameter when handling search queries. One of the advantages of using the &server query parameter is that it allows you to use a single DNS name.

For example, suppose your server's URL was searchit.novell.com. If you were setting up search services for a company called Digital Airlines and another company called DemoCity, you could host both services on your single server and then simply include the &server=digitalairlines and &server=democity query parameters within the search forms found on www.digitalairlines.com and www.democity.com.

Queries would be sent from the search forms on each Web site to the URLs corresponding to each virtual search server, as in the following example:

```
http://searchit.novell.com/qfsearch/
SearchServlet?server=digitalairlines
```

and

```
http://searchit.novell.com/qfsearch/SearchServlet?server=democity
```

8.2.3 Using the Virtual Search Server Alias

When defining a virtual search server, you are required to give it a name. However, you can also define an alias that can be used when identifying a specific virtual search server during a search request.

An alias name typically follows one of two conventions:

- An IP address.

This could be used either in the domain name portion of a URL or be included in a search query using the &server query parameter. Using an IP address in place of a domain name to select a virtual search server works only in a hardware virtual server configuration where each search server has its own unique IP address.

- Any other numeric or textual value that can be passed as the value of the &server query parameter.

For most virtual search servers, the best choice for a search server name and alias is the Web server's domain name and IP address.

8.2.4 Storing Virtual Search Server Files

Search server files include a set of index and configuration files for each virtual search server. When you create a new virtual search server, you can specify where you want virtual search server files to be stored, or you can accept the default path, which is determined by where you installed the QuickFinder Server.

Virtual search server files can be stored on any volume visible to the server that QuickFinder is installed on, regardless of which volume your QuickFinder Server is installed on. This includes SAN storage device volumes.

8.3 Enabling, Disabling, or Deleting a Virtual Search Server

You can disable a virtual search server or delete it from QuickFinder. When you disable a virtual search server, users receive an Unavailable Service message and cannot perform any searches. When you delete a virtual search server, all of its indexes and all related files are permanently deleted from your server.

- Section 8.3.1, “Enabling or Disabling a Virtual Search Server,” on page 64
- Section 8.3.2, “Deleting a Virtual Search Server,” on page 65

8.3.1 Enabling or Disabling a Virtual Search Server

1. On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to disable.
2. Under *Settings* (in menu on the left), click *On/Off*.
3. Click *Turn ON* or *Turn OFF* to enable or disable the current virtual search server.

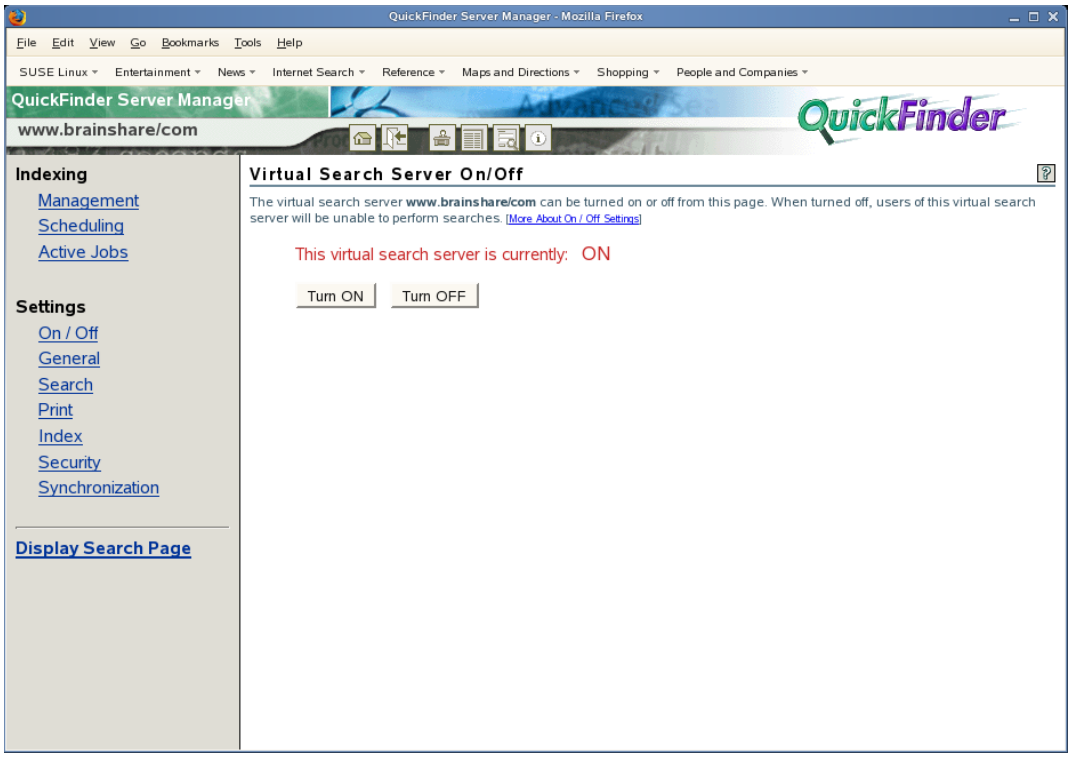

When QuickFinder turns off the current virtual search server, the status above the *Turn ON* and *Turn OFF* buttons changes to *OFF*.

8.3.2 Deleting a Virtual Search Server

1. On the Global Settings page of QuickFinder Server Manager, click *Delete* in the row of the virtual search server that you want deleted.
2. Click *Delete Virtual Search Server* to confirm the deletion.

   or

   Click *Cancel Deletion* to abort.

WARNING: There is no way to restore a deleted virtual search server. Exercise caution before proceeding.

8.4 Configuring Virtual Search Servers

After you have created a virtual search server, you can fine-tune it by modifying many of its default settings. This section describes how to configure an individual virtual search server.

TIP: When you create a new virtual search server, all of its default settings are inherited from the Global Settings pages of QuickFinder. If you are the global administrator, you can modify these settings so that when you create a new virtual search server, its default settings reflect your preferences. (See Section 7.1, “Differences between Default and Services Settings,” on page 36.)

- Section 8.4.1, “Configuring General Settings,” on page 66

- Section 8.4.2, “Configuring Search Settings,” on page 71
- Section 8.4.3, “Configuring Print Settings,” on page 75
- Section 8.4.4, “Configuring Index Settings,” on page 77
- Section 8.4.5, “Configuring Security Settings,” on page 80
- Section 8.4.6, “Configuring Synchronization Settings,” on page 82

8.4.1 Configuring General Settings

The General Settings page of QuickFinder Server Manager lets you configure the way that the virtual search server handles requests from clients for search services. You can modify general search query, search response, error log, and e-mail settings (if the global administrator has enabled the e-mail feature).

- “Modifying General Search Query Settings” on page 66
- “Modifying Query Logging Settings” on page 67
- “Modifying General Response Settings” on page 68
- “Modifying Error Log Settings” on page 69
- “Modifying E-Mail Settings” on page 69
- “Editing the Stop-Words List” on page 69

Modifying General Search Query Settings

1. On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.
2. Click *General* under *Settings*.
3. From the *Default query encoding* drop-down list, select an encoding that represents the character set encoding that the users of your search services are most likely to use.

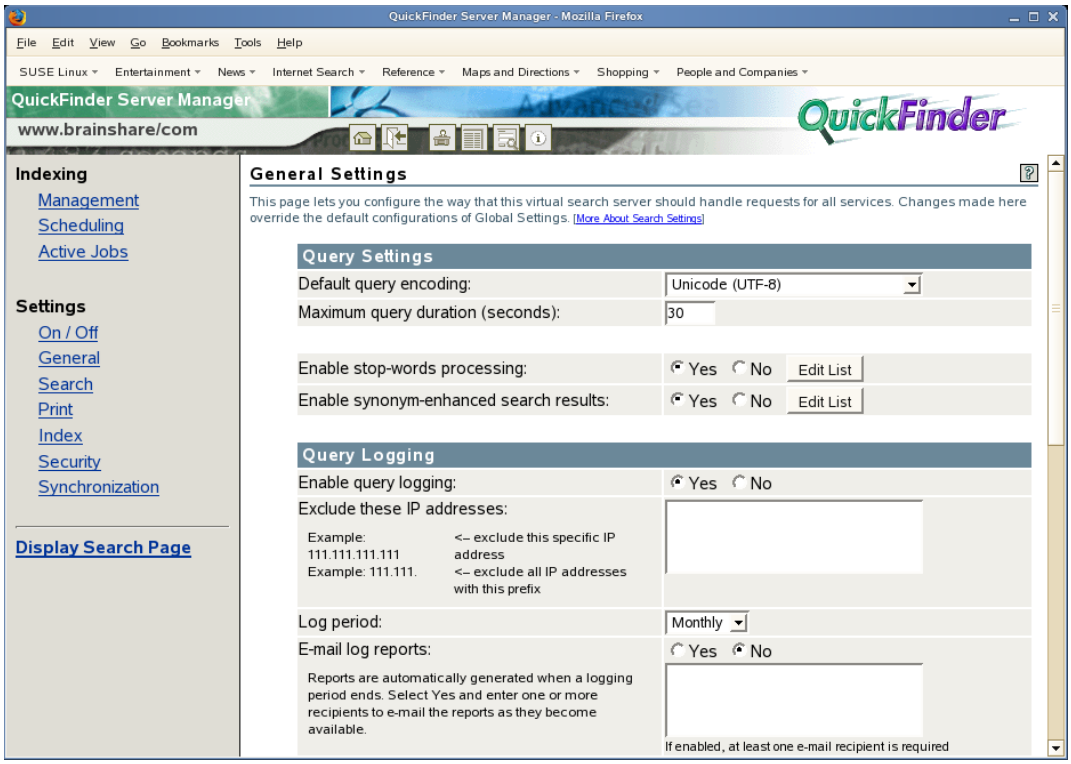

For more information about character sets, see Appendix B, “Combined Character Sets for Use with QuickFinder,” on page 197.

4 In the *Maximum query duration* field, specify the maximum number of seconds before QuickFinder should end a query, regardless of whether a search has been completed.

   This option is one of several methods for letting you protect your server's resources from processing potential rogue searches, which are sometimes intended to harm your service by consuming server resources.

5 (Optional) Click *Yes* next to *Enable stop-words processing* if you want QuickFinder to ignore insignificant words, such as adverbs, conjunctions, or prepositions, in a user's search query.

   For information about how to modify the list of stop-words used by QuickFinder, see “Editing the Stop-Words List” on page 69.

6 (Optional) Click *Yes* next to *Enable synonym enhanced search results* if you want QuickFinder to expand search queries by looking for synonyms of the original search terms.

   Synonym enhancement is not performed within the FileFilter section of a search. For more information about synonym enhanced searches, see “Using Synonyms to Broaden Search Results” on page 118.

7 Click *Apply Settings*.

Modifying Query Logging Settings

When query logging is enabled, it captures many details about user queries, including total number of searches performed (for the current logging period), top 100 search terms, load statistics, and the number of print jobs performed. It also provides information about queries that resulted in a “not found” error message. This information can be used to improve the performance of your search services, or to identify areas of your site that could be improved.

Query logging is enabled by default. However, it is optional and can be turned off by selecting *No* next to *Enable Query Logging*.

1. On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.
2. Click *General* under *Settings*.
3. Click *Yes* next to *Enable query logging*.

   or

   If Query Logging has been enabled, click *No* to turn it off.
4. In the *Exclude these IP addresses* field, specify one or more IP addresses of the computers that you want excluded from the query log, separating each additional address by a space or hard return.

   This feature enables you to eliminate query logging for searches originating within your organization. Keeping query log information gathered from customer queries separate from log information related to employees of your own company can help you identify issues that directly affect your customers.
5. From the *Log period* drop-down list, select a log period of either *Monthly* or *Weekly*.

   QuickFinder records search requests in the current log file until the end of the specified period. When QuickFinder switches to a new log file, a report for the old log file is automatically generated.
6. In the *E-mail log reports* field, specify one or more e-mail addresses where you want copies of each new log report sent, separating each additional address by a space or hard return.

   If this field does not appear, it is because you have disabled e-mail services on your QuickFinder Server. To re-enable e-mail services, go to the QuickFinder Global Settings page and click *Services Settings* > *General* > *E-Mail Settings*.
7. To view an existing log report, select one from the *Available log reports* field, then click *View Report*.

   If there are no logs listed, you have not yet generated a log report. Click *Generate Current* to create a report of the current (not yet completed) log period. Depending on the amount of information in the logs, this process could take anywhere from a few seconds up to several hours to complete. After a new report is generated, it appears in this field.
8. (Optional) If you want to specify a different logging template, specify the name of the template in the *Template to use when generating reports* field.
9. Click *Apply Settings*.

Modifying General Response Settings

1. On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.
2. Click *General* under *Settings*.
3. Under *Response Settings*, select an output encoding from the *Default encoding for response pages* drop-down list.

   This setting specifies the encoding QuickFinder should use when responding to user queries using the search and print results templates, and the error and response messages templates.
4. Specify the maximum number of queries in the *Refuse queries if potential hits exceed* field to cancel the processing of search results that might take a long time to complete.

5 Click *Apply Settings*.

Modifying Error Log Settings

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.

2 Click *General* under *Settings*.

3 In the *Maximum log size* field, specify the maximum size (in bytes) that QuickFinder should allow the log file to grow to.

Depending on the number of visitors that your virtual search server hosts, log files can become large. This setting protects your system's hard drive resources.

The number you specify here is divided evenly between two log files. For example, if you specify 30000, each log file allows up to 15000 bytes of logged data. This ensures that you always have at least 15000 bytes of logged synchronization data.

4 Click *Apply Settings*.

Modifying E-Mail Settings

If this section of the General Settings page is not visible, it is because the global administrator has disabled it. If you are the global administrator, you can enable this feature from the *Global Settings > Services Settings > General* page. In the *E-Mail Settings* section, enable e-mail services, then specify an outgoing SMTP host name, port number, user ID, and password. This makes the *E-Mail Settings* section appear on the General Settings page for a virtual search server. Follow the instructions below to enable or disable e-mail services and to specify recipients.

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.

2 Click *General* under *Settings*.

3 (Optional) Click *Yes* next to *Enable e-mail services* if you want QuickFinder to generate an e-mail message when errors occur during the generation (or regeneration) of your virtual search server's indexes, or when your indexes are being synchronized with other QuickFinder servers.

4 (Optional) If you enabled e-mail services, type one or more e-mail addresses indicating where error messages should be sent.

If the global administrator has specified default e-mail recipients, they appear here. You can remove them, or simply add additional addresses. Separate addresses by using a space, comma, semicolon, or carriage return.

5 Click *Apply Settings*.

Editing the Stop-Words List

Sometimes users include irrelevant words in their search strings, such as the conjunctions and, to, and of. These are referred to as stop words. The Stop Words feature of QuickFinder removes all occurrences of stop words from the search string before performing a search, except in the following situations:

- Stop words processing has been turned off.
- You prepend the word with a plus (+) sign (for example, +the).
- They are the only words in the query.

- They occur within a phrase.

You can add or remove words from the Stop-Words list. QuickFinder Server uses this list to determine which words it ignores in a user's search string. QuickFinder includes a predefined list of stop words, but you can edit or change these at any time.

1. On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.
2. Click *General* under *Settings*.
3. Click *Apply Settings* > *Back to General Settings* to save any changes you have made.

   Clicking *Edit List* causes any unsaved changes to be lost.
4. Under Query Settings, click *Edit List* next to *Enable stop-words processing*.

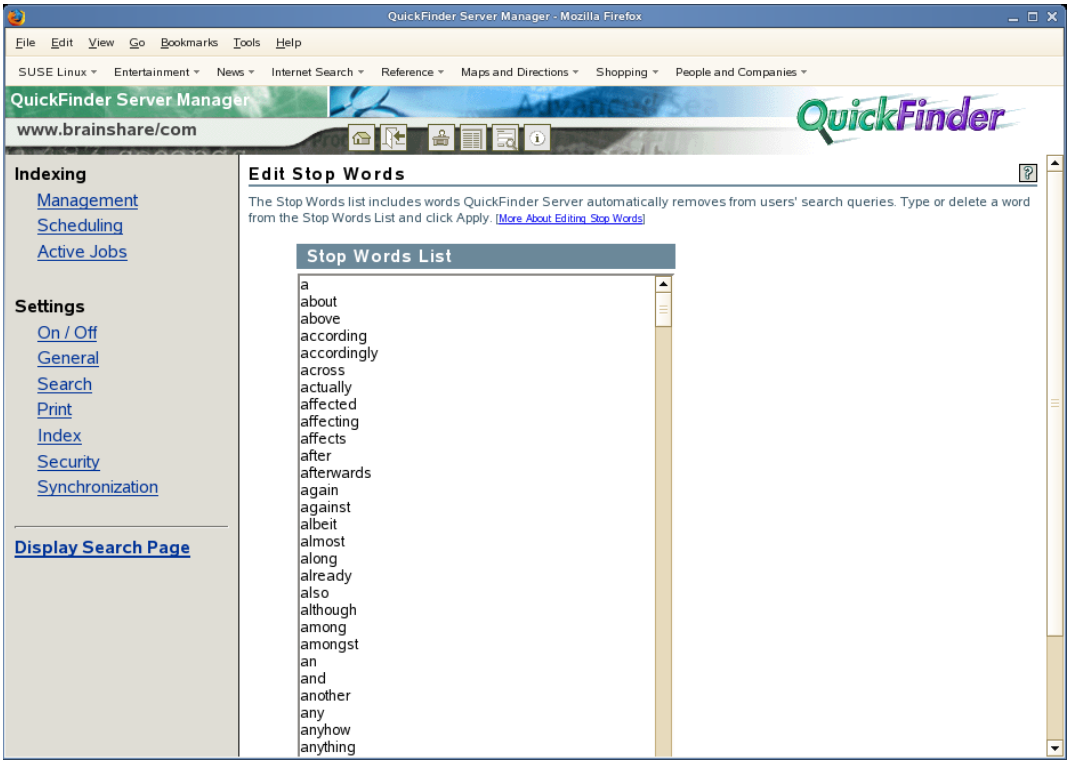

5. (Optional) To add a new word to the Stop-Words list, type a new word directly into the list.

   Separate each new word with a space or a carriage return. When you click *Apply*, the word is added to the list in alphabetical order, regardless of where you entered it in the list.
6. (Optional) To remove an existing stop word from the Stop-Words list, select it and press the Delete key.

   To removing a stop word that's on the *Global Stop Words* list, delete it from the global area.
7. Click *Apply*.

For information about enabling and editing the synonym enhanced search settings, see “Using Synonyms to Broaden Search Results” on page 118.

8.4.2 Configuring Search Settings

The Search Settings page lets you configure the way that the virtual search server handles end user's search queries. You can modify search query, search response, and template settings.

- “Modifying Search Query Settings” on page 71
- “Modifying Search Response Settings” on page 72
- “Modifying Template Settings” on page 73
- “Selecting Highlighter Colors” on page 74

Modifying Search Query Settings

1. On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.
2. Under *Settings*, click *Search*.
3. Under *Query Settings*, click *Edit* next to *Default indexes to search*.

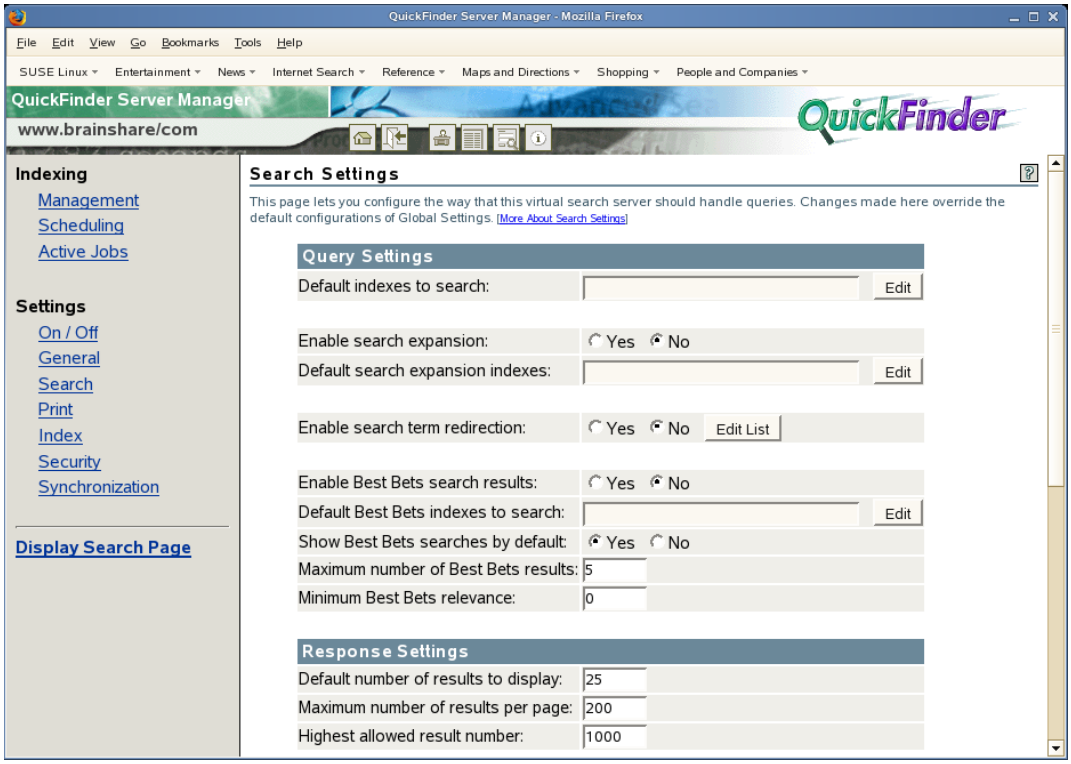

4. Select the indexes to be searched by default when a user does not select any indexes on the search form when performing a search, then click *Apply*.
5. (Optional) Click *Yes* next to *Enable search expansion* if you want QuickFinder to automatically perform a second search using additional indexes when a first search query fails to find results.

   For more information, see Section 10.5, “Helping Users Avoid Failed Searches,” on page 122.
6. (Optional) If you enabled search expansion, click *Edit* next to the *Default search expansion indexes* field to select the indexes to be used during an expanded search, then click *Apply*.

7 (Optional) Click *Yes* next to *Enable search term redirection* if you want QuickFinder to redirect common search terms to a specific URL.

For more information, see "Redirecting Searches" on page 121.

8 Click *Apply Settings* > *Back to Search Settings* to save any changes you have made.

Clicking *Edit List* causes any unsaved changes to be lost.

9 (Optional) If you enabled search term redirection, click *Edit List* to add common terms and their associated URLs, then click *Apply*.

For more information, see "Creating or Editing the Redirection URLs List" on page 121.

10 (Optional) Click *Yes* next to *Enable Best Bets search results* if you want QuickFinder to insert a secondary results list at the top of the search results page.

For more information, see "Using Best Bets to Help Users Find the Right Information" on page 117.

11 (Optional) Click *Edit* next to the *Default Best Bets indexes to search* field and select the indexes for use with best bets.

12 (Optional) If you want the best bets results to always appear on the search results page, select *Yes* next to *Show Best Bets searches By default*.

Select *No* if you want users to request the best bets results list at the time of their search request.

13 (Optional) In the *Maximum number of Best Bets results* field, type the maximum number of best bets results to be returned on each search results page.

14 (Optional) In the *Minimum Best Bets relevance field*, type the number of low relevance documents you want removed from the best bets display list.

15 Click *Apply Settings*.

Modifying Search Response Settings

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.

2 Under *Settings*, click *Search*.

3 Under Response Settings, specify the number of search results in the *Default number of results to display* field that you want displayed on each search results page.

For example, if you set this option to 25 (which is the default setting) and the number of hits in a return was 200, QuickFinder would only return 25 hits per search results page at a time.

4 Set a limit on the number of results allowed at one time on the results page by specifying a number in the *Maximum number of results per page* field.

5 Specify the highest number of search results that can be returned to a user query in the *Highest allowed result number* field.

6 (Optional) Click *Yes* next to *Enable speller* to have QuickFinder provide alternate spelling suggestions for search terms.

7 Click *Apply Settings* > *Back to Search Settings* to save any changes you have made.

Clicking *Edit List* causes any unsaved changes to be lost.

8 (Optional) Click *Edit List* to define your own words to ignore or misspelled words to replace.

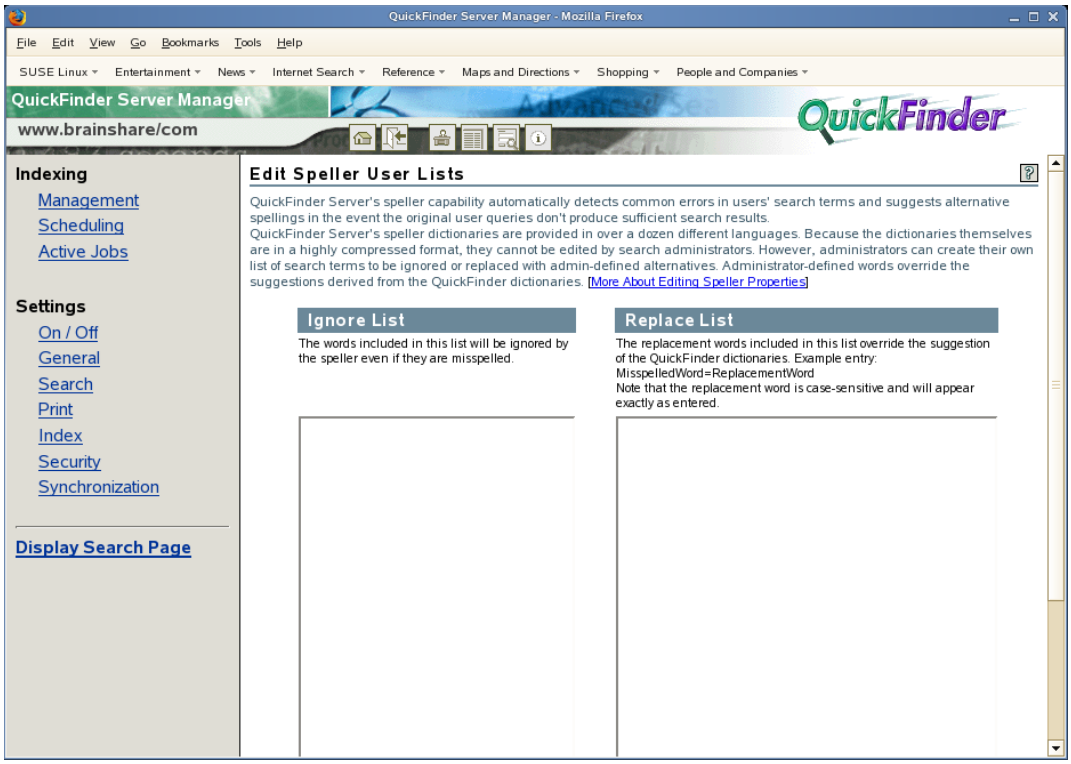

9 (Optional) Use the *Spell check if total hits is below* option to have QuickFinder spell-check search terms if the number of search results is lower than the number specified.

10 (Optional) Click *Yes* next to *Enable search terms Highlighter* if you want users' search terms to be highlighted in the documents returned in a search.

11 Click *Apply Settings* > *Back to Search Settings* to save any changes you have made.

Clicking *Edit Colors* causes any unsaved changes to be lost.

12 Click *Edit Colors* to modify the colors used by the Highlighter.

For information, see "Selecting Highlighter Colors" on page 74.

13 Click *Apply Settings*.

Modifying Template Settings

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.

2 Under *Settings*, click *Search*.

3 Under *Template Settings*, specify a path to where your QuickFinder templates are stored in the *Templates directory* field.

The default path is `volume:\searchroot\Templates` on NetWare, and `/var/lib/qfsearch/Templates` on Linux. If you have created custom templates, or if you want to keep your templates elsewhere, specify the path here so that QuickFinder knows where the templates are.

4 From the *Default encoding for templates* drop-down list, select the character set that your templates are written in.

This value is used even with templates that do not specify an encoding. Encodings found in templates that do not match the encoding you specify here override this encoding.

5 In the *Default search page template* field, specify the filename of the search page template you want to use.

If you have created a custom template and want QuickFinder to use it as your search page, specify its name in this field.

6 In the *Default search results template* field, specify the filename of the search results template you want to use.

If you have created a custom search results template and want QuickFinder to use it as your default search results page, specify its name in this field.

7 In the *Default highlighter template* field, specify the filename of the highlighter template you want to use.

If you have created a custom highlighter template and want QuickFinder to use it as your default search results page, specify its name in this field.

8 In the *Template to use if no results returned* field, specify the filename of the template that QuickFinder should return if no results are found.

9 In the *Template to use if error occurs* field, specify the filename of the template that QuickFinder should return if there are errors while processing a user's query.

10 Click *Apply Settings*.

Selecting Highlighter Colors

Because a user can search for more than one search term at a time, you can define multiple colors so that each search term appears in a different color. For example, if a user searched for “HTTP AND Apache,” the terms appear highlighted in the documents where they originated.

The colors are used in the order in which they are defined. If only two highlight colors are defined but a user searches for three search terms, the first color is repeated.

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.

2 Under *Settings*, click *Search*.

3 On the Search Settings page, click *Edit Colors* under *Response Settings*.

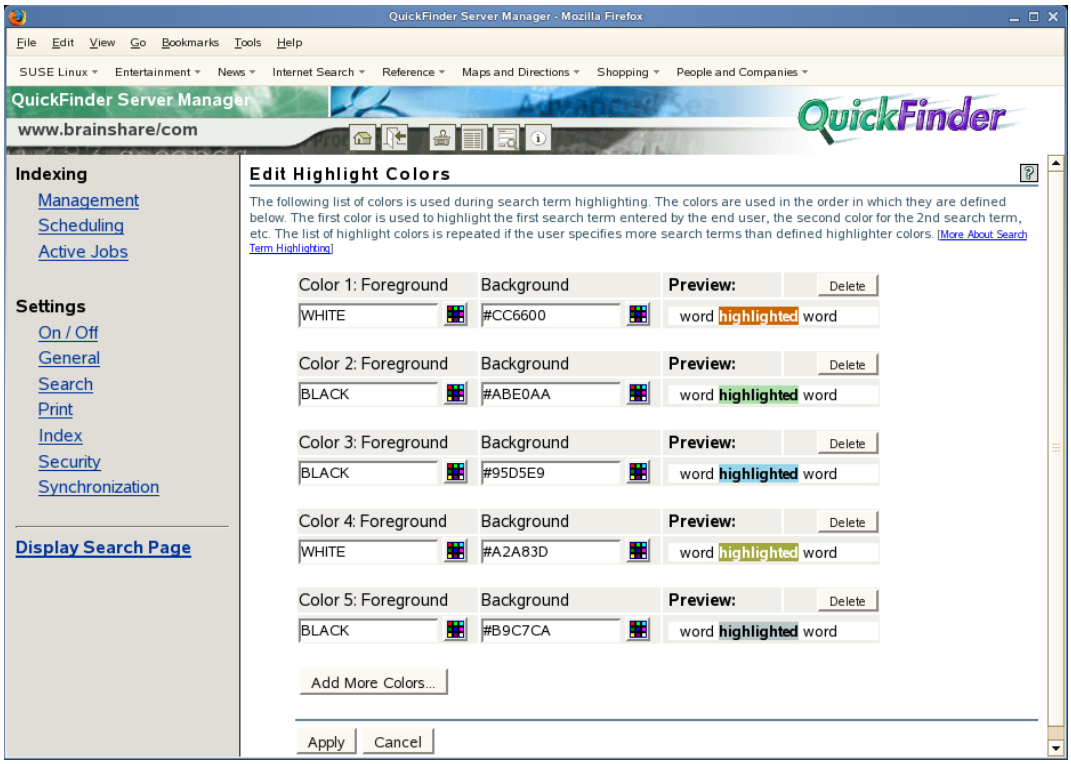

4 To define a new highlight color, click the *Color* button in the *Foreground* and *Background* columns to select a color.

   You can also enter the color value (such as #RRGGBB) or color names.

5 To define additional color sets, click *Add More Colors*.

6 To remove a defined color set, click *Delete* in the row of the color set you want removed.

7 Click *Apply*.

8.4.3 Configuring Print Settings

The Print Settings page lets you configure the way that the current virtual search server handles end user's print queries. You can modify print result and template settings.

- “Modifying Default Print Settings” on page 75
- “Modifying Default Print Template Settings” on page 76

Modifying Default Print Settings

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.

2 Click *Print* under *Settings*.

3 Under *Print Results Settings*, specify the number of print results in the *Default number of results to print* field that you want displayed on each print results page.

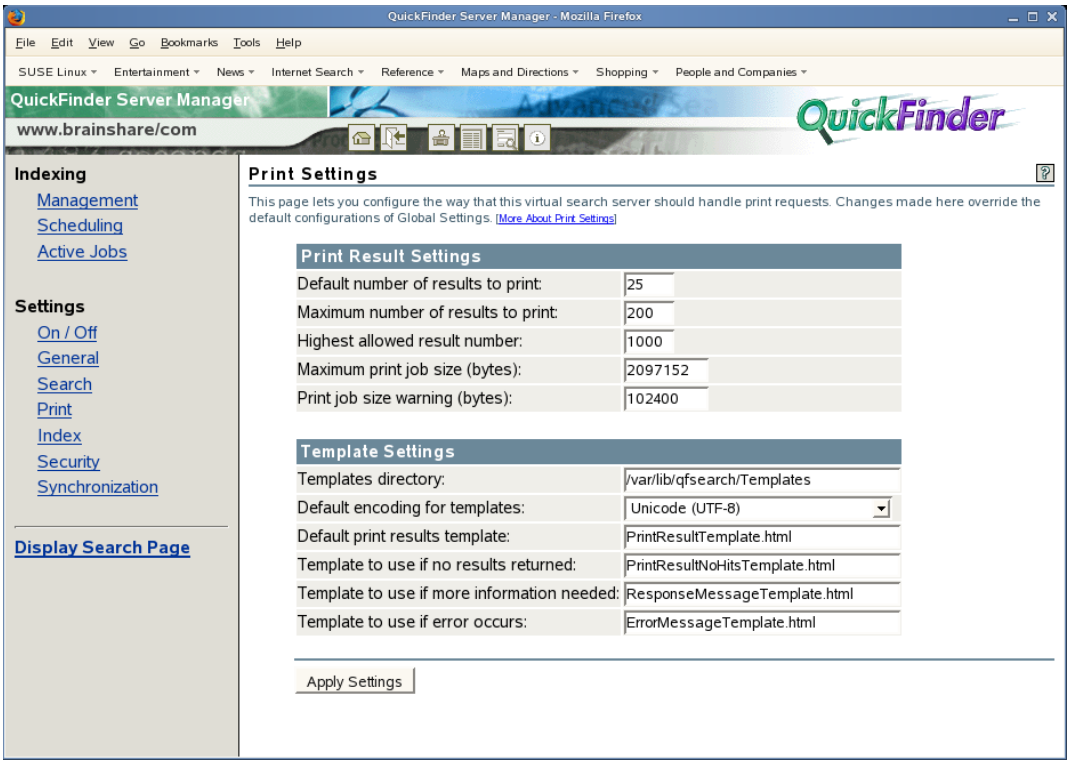

For example, if you set this to 25 (which is the default setting) and the number of hits in a return was 200, QuickFinder would only return 25 hits per print results page at a time.

4 Set a limit on the number of results allowed at one time on the results page by specifying a number in the *Maximum number of results to print* field.

5 Specify the highest number of search results that can be returned to a user query in the *Highest allowed result number* field.

6 To limit the size of a print job, specify the largest print job size that QuickFinder will allow in the *Maximum print job size* field.

Users requesting a print job larger than this value receive a message informing them that the request was too large.

This is a useful feature for administrators who want to keep down the size of print jobs in their own companies, departments, or organizations.

7 To be notified when a print job exceeds a certain size, specify the print job size in the *Print job size warning* field.

By default, this message is sent by using the `ResponseMessageTemplate.html` file and is intended as a warning to users that they are exceeding the allowed print job size. It then prompts the user to confirm the print job before continuing.

8 Click *Apply Settings*.

Modifying Default Print Template Settings

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.

2 Click *Print* under *Settings*.

3 Under *Template Settings*, specify a path in the *Templates directory* field to where your QuickFinder templates are stored.

   The default path is `volume:\searchroot\Templates` on NetWare, and `/var/lib/qfsearch/Templates` on Linux. If you have created custom templates, or want to keep your templates elsewhere, specify the path here so that QuickFinder knows where the templates are.

4 From the *Default encoding for templates* drop-down list, select the character set that your templates are written in.

   This value is used even with templates that do not specify an encoding. Encodings found in templates that do not match the encoding you specify here override this encoding.

5 In the *Default print results template* field, specify the filename of the print results template you want to use.

   If you have created a custom print results template and want QuickFinder to use it when returning print results, specify its name in this field.

6 In the *Template to use if no results returned* field, specify the filename of the template that QuickFinder should return if no print results match a user's query.

7 In the *Template to use if more information needed* field, specify the filename of the template to be sent back to users whose print jobs exceed the size you specify in the Print Job Size field. (See Step 6.)

8 In the *Template to use if error occurs* field, specify the filename of the template that QuickFinder should return if there are errors while processing a user's print query.

9 Click *Apply Settings*.

8.4.4 Configuring Index Settings

These settings are intended to make the process of creating indexes even easier by letting you configure common settings as default settings. This saves you time by not making you make the same selections each time you create a new index.

- "Modifying Default Index Settings" on page 77
- "Allowing Indexes to Be Used With QuickFinder Synchronization" on page 79
- "Enabling or Disabling User Authentication When Accessing Indexes" on page 79

Modifying Default Index Settings

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.

2 Click *Index* under *Settings*.

3 Select the type of index that you want as the default index type on the *Index Settings* page.

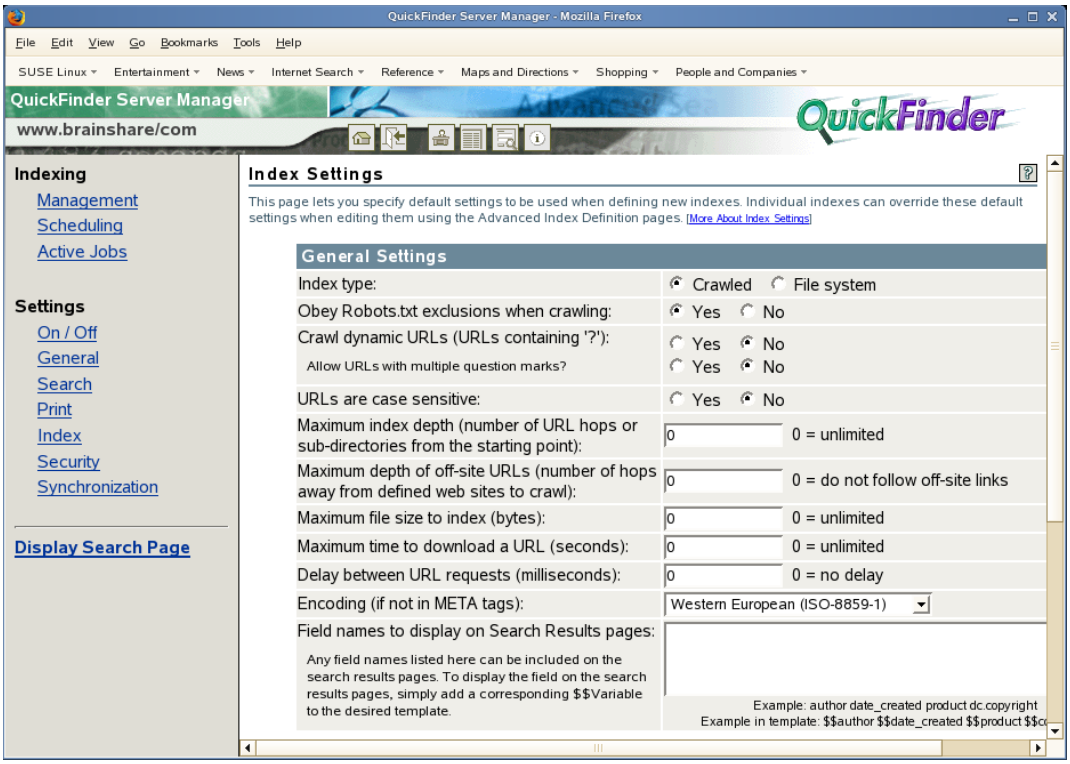

Each time you create a new index, the index type you select here is automatically selected. You can change the default index type when you create an index.

4 Click *Yes* next to *Obey Robots.txt exclusions when crawling* to instruct QuickFinder to obey the instructions in the `Robots.txt` file as it indexes Web sites.

For more information about `Robots.txt`, see "Using the Robots Meta Tag" on page 113.

5 Click *Yes* next to *Crawl dynamic URLs (URLs containing '?')* if you want QuickFinder to index dynamic content, in addition to static content.

Some search engines do not crawl Web pages that are created dynamically, such as pages generated from forms submitted by a Web browser. The URLs of dynamic content typically contain a question mark (?) followed by additional parameters. (See Section 8.7, "Indexing Dynamic Web Content," on page 98.)

Because dynamic content can change at any time, you might want to schedule more frequent regeneration events for your indexes when enabling this feature. For information about scheduling automatic updates to your indexes, see Section 8.9, "Automating Index and Server Maintenance," on page 100.

6 Click *Yes* next to *Allow URLs with multiple question marks* if you want QuickFinder to crawl URLs with multiple query strings (parameters).

7 Select the *URLs are case sensitive* option if you want QuickFinder to recognize URLs that are different only in character case, but are otherwise identical (for example, www.digitalairlines.com verses www.DigitalAirlines.com).

IMPORTANT: Setting this option to *No* can help QuickFinder avoid indexing duplicate information, which can come from indexing URLs that are presented using different cases but actually point to the same information. However, if a Web server being indexed is configured to differentiate between cases, QuickFinder might leave out content that you want indexed.

---

8. In the *Maximum index depth* field, specify the maximum number of hypertext links from the starting URL that QuickFinder should follow before it stops indexing.
9. In the *Maximum depth of off-site URLs* field, specify the maximum number of hypertext links off-site from the starting URL that QuickFinder should follow before it stops indexing.
10. Specify a number (in bytes) in the *Maximum file size to index* field to keep QuickFinder from indexing files larger than the number you specify.
11. In the *Maximum time to download a URL* field, specify a number (in seconds) before QuickFinder automatically skips the indexing of the specified URL.
12. Type a number (in milliseconds) in the *Delay between URL requests* field that QuickFinder should pause between requests for URLs that it is trying to index.
13. From the *Encoding (if not in META tags)* drop-down list, select the encoding to be used when indexing files that do not contain an encoding specification.

    For example, HTML files can specify their encoding with a Content-Type meta tag.
14. In the *Field names to display on Search Results pages* field, list any field names that you want to be included on the search results page (for example, author date_created product dc.copyright).

    To display the field on the search results pages, add the corresponding $$Variable to the template (for example, $$author $$date_created $$product $$copyright).

    The field data is stored in the index and causes the index size to increase.
15. Click *Apply Settings*.

Allowing Indexes to Be Used With QuickFinder Synchronization

1. On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.
2. Click *Index* under *Settings*.
3. Under *Synchronization Settings*, click *Yes* next to *Index may be copied to other clustered servers*.

    For more information about QuickFinder Synchronization, see Chapter 9, “Synchronizing Data Across Multiple QuickFinder Servers,” on page 103.
4. Click *Apply Settings*.

Enabling or Disabling User Authentication When Accessing Indexes

1. On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.
2. Click *Index* under *Settings*.
3. (Optional) Under *Rights-based Search Results*, click *Off* to allow anyone to access the indexes.

    No user authentication is required to search the indexes of the current virtual search server.
4. (Optional) Under *Rights-based Search Results*, click *by Index* to restrict access to the indexes only to those who have access rights to the filename you specify in the *by Index* field.
5. In the *by Index* field, specify the full path to a file on your server that already has rights protection assigned to it.

    For example, on a NetWare® 6.5 server where Apache is your Web server, you might type `sys:\\Apache2\htdocs\index.html`.

6 Select either *Search Engine* or *Templates* from the *Unauthorized hits filtered by* drop-down list to specify how QuickFinder should handle attempts to access indexes without authorization.

Select *Search Engine* if you don't want any search results displayed when users attempt to search the index without first logging in. A No Results Found message is returned instead.

Select *Template* to have QuickFinder return search results to users who have not logged in. When they attempt to click a search result, they are then prompted for a username and password.

TIP: You can use the $$BeginUnAuthorized and $$EndUnAuthorized template variables to control how the template should display unauthorized hits. For more information about template variables, see Chapter 12, "Working with Template Variables and Search Parameters," on page 133.

---

7 Click *Apply Settings*.

8.4.5 Configuring Security Settings

Security settings let you manage access to indexed content by requiring users to authenticate to a server before seeing rights-protected search results. The Security Settings page lets you configure rights-based search results and specify secure connection settings, including enabling the HTTPS protocol for securing usernames and passwords as they are sent over the network.

- "Modifying Default Rights-Based Search Results" on page 80
- "Modifying Default Connection Settings" on page 81

Modifying Default Rights-Based Search Results

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.

2 Under *Settings*, click *Security*.

3 In the *Default search contexts* box, type the Novell® eDirectory™ context where user objects are stored.

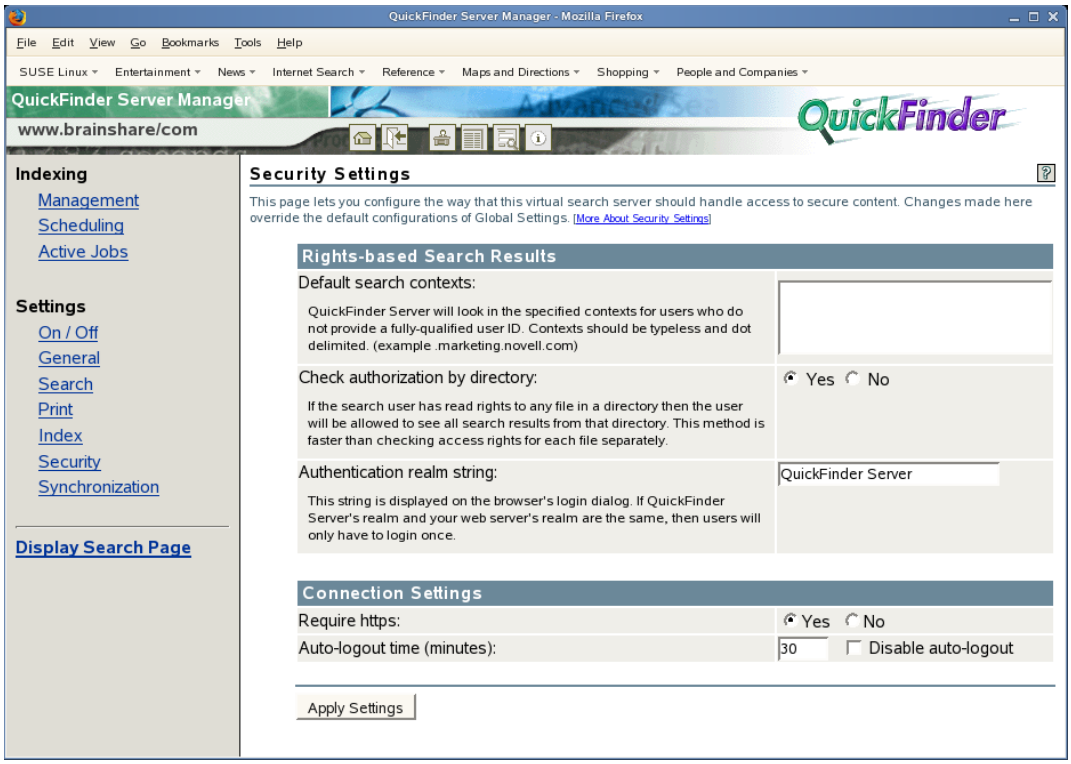

If you provide a context here, users won't need to supply their fully qualified user IDs when logging in to QuickFinder. For example, if your user objects are all kept in a container named employees, type `employees.marketing.digitalairlines` in the *Default Search Contexts* box so that QuickFinder knows where to look when a user attempts to log in. A user whose fully qualified name is sjones.employees.marketing.digitalairlines would only need to enter sjones.

4 Click *Yes* next to *Check authorization by directory* if you want QuickFinder to validate user rights to files in a directory (search results) based on whether or not users have rights to any one file in the same directory.

If a larger number of results are from the same directory, this can help to speed up the search.

5 (Optional) If you are using basic authentication (as apposed to using form-based authentication), then in the *Authentication realm string* field, specify the correct authentication realm string, which is typically shown in your Web browser's login dialog box.

Specifying the Apache Web server's authentication realm string in this field makes it so that after users authenticate to it, they won't need to authenticate again when using QuickFinder to search and access protected information.

If Apache is your Web server, refer to http://httpd.apache.org/docs-2.0/howto/auth.html (http://httpd.apache.org/docs-2.0/howto/auth.html) for more information about authentication.

6 Click *Apply Settings*.

Modifying Default Connection Settings

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.

2 Under *Settings*, click *Security*.

3 Under *Connection Settings*, click *Yes* next to *Require https* if you want to protect usernames and passwords as they are sent across the network or Internet.

4 Specify a number (in minutes) in the *Auto-logout time* field to direct QuickFinder to log users out who have been idle for the specified period of time.

5 Click *Apply Settings*.

8.4.6 Configuring Synchronization Settings

QuickFinder Synchronization lets you designate one QuickFinder server as the search master from which updated indexes, templates, and configuration settings are systematically sent out to all other QuickFinder servers defined as part of a QuickFinder Synchronization cluster.

For more information about setting up and configuring QuickFinder Synchronization, see Chapter 9, “Synchronizing Data Across Multiple QuickFinder Servers,” on page 103.

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.

2 Under *Settings*, click *Synchronization*.

3 Select *Yes* next to *Let Virtual Search Server receive cluster data*.

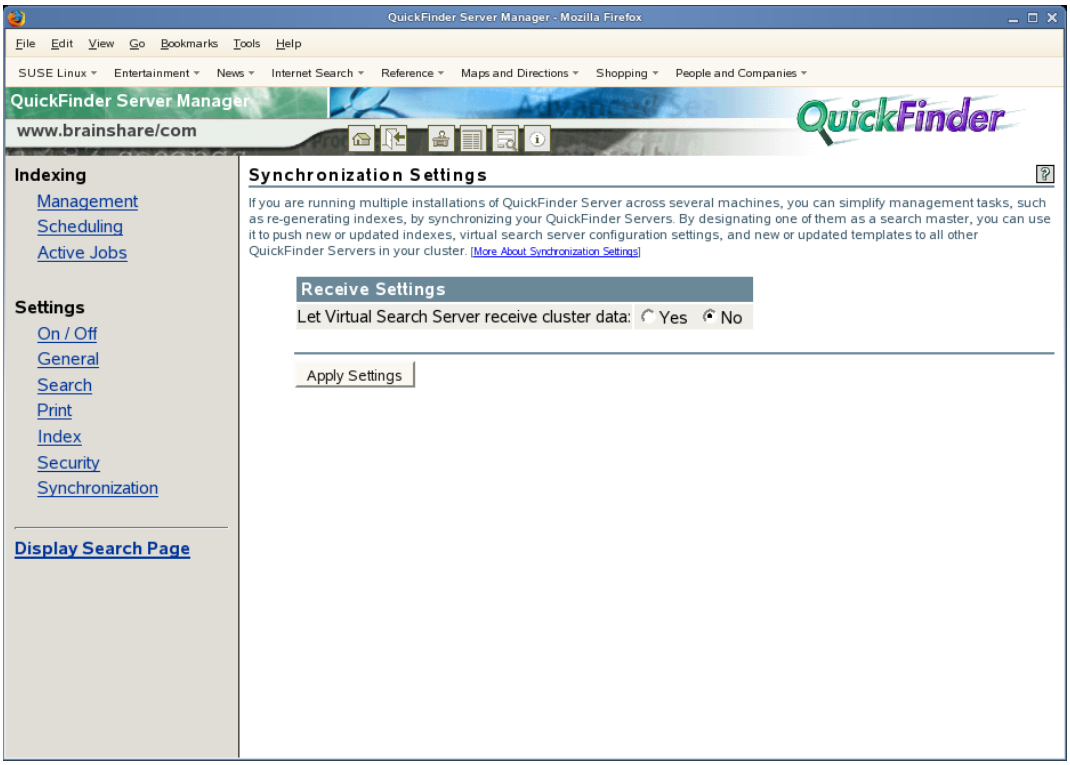

4 Click *Apply Settings*.

8.5 Creating Indexes

QuickFinder creates two types of indexes:

- **Crawled:** Created as QuickFinder follows (or crawls) hypertext links until it reaches a dead end. QuickFinder can crawl one or more Web sites, specific areas of a Web site, or specific URLs, even down to a specific filename.
- **File System:** Created as QuickFinder index content on a file server. QuickFinder can index one or more paths on multiple volumes, including Storage Area Network (SAN) storage devices.

There are two forms you can use to create each type of index: the standard form and the advanced form. The Define Crawled Index is the standard form for creating a crawled index, but the Define Crawled Index (Advanced) form offers more options than the standard form, including options that override default virtual search server settings. Both methods are described in the following sections.

- Section 8.5.1, “Creating a Crawled Index,” on page 83
- Section 8.5.2, “Creating an Advanced Crawled Index,” on page 84
- Section 8.5.3, “Configuring Rights-Based Search Results for Crawled Indexes,” on page 88
- Section 8.5.4, “Creating a File System Index,” on page 89
- Section 8.5.5, “Creating an Advanced File System Index,” on page 90
- Section 8.5.6, “Configuring Rights-Based Search Results for File System Indexes,” on page 93
- Section 8.5.7, “Searching Across Multiple Indexes,” on page 93
- Section 8.5.8, “Indexing Content on a Password-Protected Web Site,” on page 94
- Section 8.5.9, “Indexing Volumes on Remote Servers,” on page 94
- Section 8.5.10, “Generating Indexes,” on page 95

8.5.1 Creating a Crawled Index

1. On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.
2. Under *Define a New Index*, click *New crawled index*, then click *Define Index*.

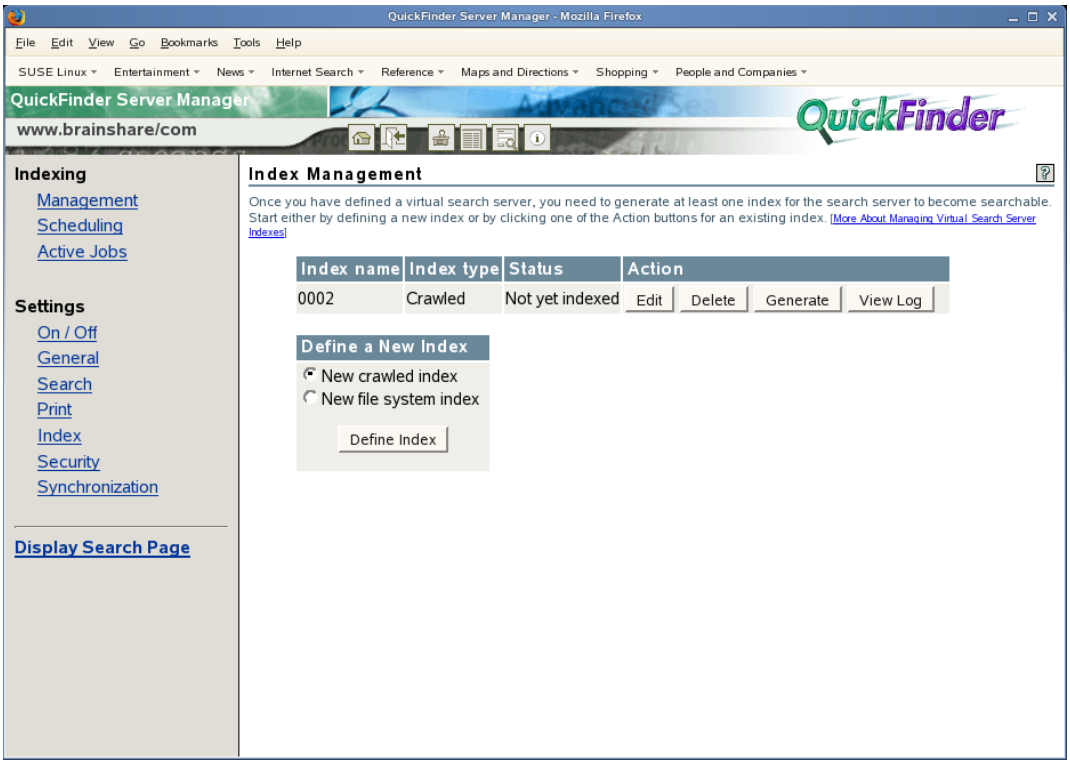

3 In the *Index Name* field, specify a name for your index.

   A name can be a word, phrase, or a numeric value. If the virtual search server you are working on contains, or will contain, a large number of indexes, you might want to use a numbering scheme to help you manage multiple indexes more effectively. However, the name you enter here appears on the default search page, so you might want to choose a name that can be understood by users of your search services.

4 Under *Web Sites to Crawl*, specify the URL of the Web site that you want indexed.

   You can specify just the URL, such as www.mycompany.com, or you can also append a complete path, down to the file level, such as www.mycompany.com/*path*/index.html.

5 If desired, add another URL.

6 To add additional URLs, click *Add More URLs*.

7 Click *Apply Settings*.

8.5.2 Creating an Advanced Crawled Index

The Define Crawled Index (Advanced) page offers some additional options beyond those available in the standard Define Crawled Index page. Changes made using this page override default virtual search server settings.

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.

2 Under *Define a New Index*, click *New crawled index*, then click *Define Index*.

3 On the Define Crawled Index page, click *Advanced Index Definition*.

4 In the *Index Name* field, specify a name for your new index.

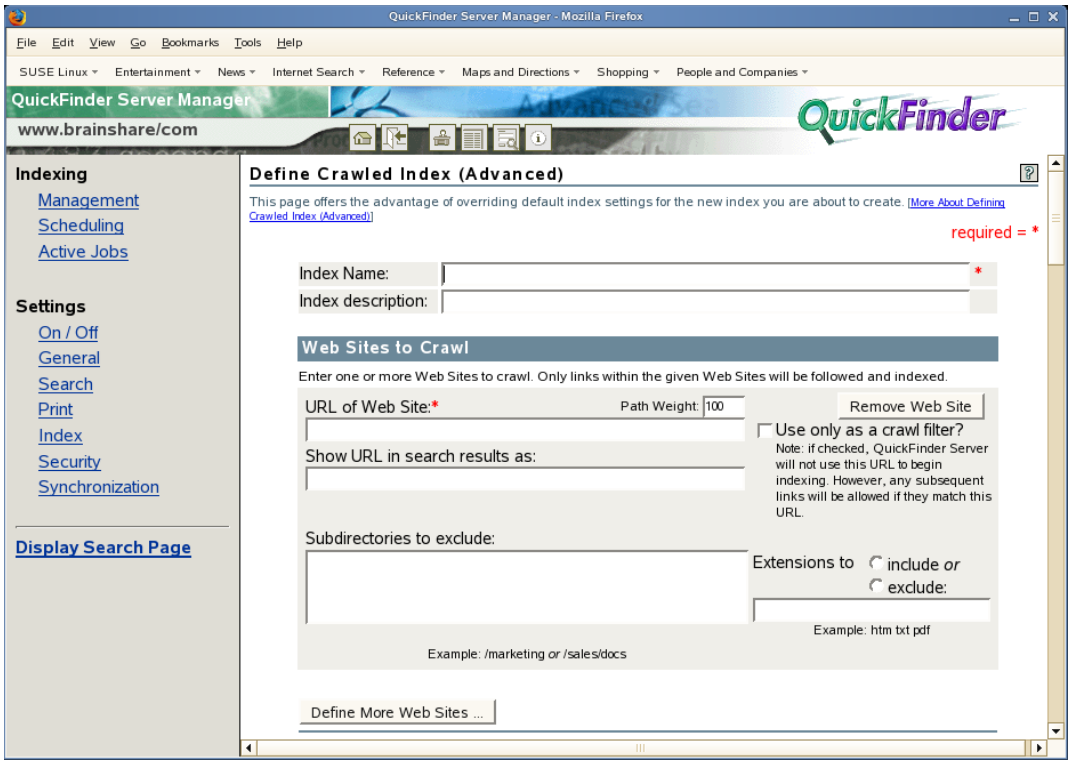

A name can be a word, phrase, or a numeric value. If the virtual search server you are working on contains, or will contain, a large number of indexes, you might want to use a numbering scheme to help you manage multiple indexes more effectively. However, the name you enter here appears on the default search page, so you might want to choose a name that can be understood by users of your search services.

5 In the *Index description* field, specify an optional description of the index to be created.

6 Under *Web Sites to Crawl*, specify the URL of the Web site to be indexed.

If you specify a filename at the end of the URL, then just that file is indexed.

7 (Optional) Use the *Path Weight* option to boost or degrade search results based on the path.

A weight of 100 makes the path's relevance normal. Increasing the weight makes the path more relevant, and lowering the weight makes the path less relevant.

8 (Optional) Select *Use only as a crawl filter* if you don't want QuickFinder to use the URL you specified in the *URL of Web Site* field to begin indexing.

Any subsequent links found that contain a URL matching the one you specified in the *URL of Web Site* field are followed and subsequently indexed.

9 (Optional) If you want to mask the actual URL displayed in the search results template, specify an alternate URL in the S*how URL in search results as* field.

For example, if you want to index a Web server that is used inside of your company but allow your customers access to some of the data, you could hide the actual internal URL with the URL of your public Web site.

10 In the *Subdirectories to exclude* text box, specify the directories that you want QuickFinder not to index.

For example, `/marketing` or `/sales/doc`.

11 To direct QuickFinder to include or exclude specific file types, click *Extensions to include* or *Extensions to exclude* and then specify the extensions, such as HTM PDF TXT, separating each one with a single space.

12 To add additional URLs, click *Define More Web Sites*.

13 To delete a URL, select it and click *Remove Web Site*.

14 In the *Additional URLs* text box, specify any other URLs that you want indexed (for example, www.mycompany.com/marketing).

This allows you to specify additional areas of information found on other Web sites, but not include all of the content of those sites to your searches.

When QuickFinder encounters links found in the pages of *Additional URLs* that point to pages specified in *Web Sites to Crawl*, QuickFinder follows those links. All other links that go outside of *Web Sites to Crawl* are not followed.

15 Use the *Off-Site URLs* option to determine the maximum number of off-site URLs (those URLs not located within any of the URLs specified in *Web Sites to Crawl*) that QuickFinder should index.

In the *URLs to Exclude* field, list the off-site URLs that you want to exclude from indexing.

16 Use the *Adjust Individual URL Relevance* option to adjust the relevance of individual items within the index.

Adjustment values can range from 1 to 200. Values higher than 100 increase the calculated relevance of the item on the search results page, and values lower than 100 decrease the calculated relevance of the item. The value specified here is combined with other values to determine the final relevance.

17 Under *Additional Settings*, specify the absolute path to where you want the index files stored in the *Location of Index Files* field.

For example, `volume:\searchroot\sites\mysites`.

By default, index files are stored at `volume:\searchroot\sites\default\indexes\` on Netware and `/var/lib/qfsearch/Sites/default/indexes\index_name` on Linux.

Changes made to *Additional Settings* override *Default Settings*.

18 From the *Level of detail in indexing logs* drop-down list, select the amount of information you want included in the index logs.

Option	Description
Disabled	Turns off index logging.
Terse	Lists only the URLs indexed.
Normal	Lists the URLs indexed and the results of the crawl.
Verbose	Lists the URLs indexed, the results of the crawl, and the links that were skipped during the crawl.
New Links	Lists the URLs indexed, the results of the crawl, the links that were skipped, and any new links found during the crawl.
All Links	Lists the URLs indexed, the results of the indexing, the links that were skipped, and all links found during the crawl.

19 From the *Encoding (if not in META tags)* drop-down list, select the encoding to be used by files being indexed that do not contain an encoding specification.

20 In the *Field names to display on Search Results pages* field, list any field names that you want to be included on the search results page (for example, author date_created product dc.copyright).

To display the field on the search results pages, add the corresponding $$Variable to the template (for example, $$author $$date_created $$product $$copyright).

21 Use the *Index Weight* option to boost or degrade search results based on the item's index.

A weight of 100 makes the item's relevance normal. Increasing the weight makes the item more relevant, and lowering the weight makes the item less relevant.

22 In the *Maximum index depth* field, enter the number of jumps (or links) from the starting URL that QuickFinder should crawl.

23 In the *Maximum file size to index* field, enter the maximum file size (in bytes) that QuickFinder should index.

Files exceeding this size are not indexed and are not included in search results.

24 In the *Maximum time to download a URL* field, enter a number (in seconds) before QuickFinder automatically skips the indexing of the specified URL.

25 In the *Delay between URL requests* field, specify the amount of time (in milliseconds) QuickFinder should delay before attempting to index a URL.

26 To direct QuickFinder to pay attention to case of filenames and directory names, click *Yes* next to *URLs are case sensitive*.

27 To direct QuickFinder to crawl dynamic content (URLs containing the question mark [?]), click *Yes* next to *Crawl dynamic URLs*.

For more information about indexing dynamic content, see Section 8.7, “Indexing Dynamic Web Content,” on page 98.

28 Click *Yes* next to *Obey Robots.txt exclusions when crawling* if you want QuickFinder to following instructions found in any Robots meta tags.

For more information, see “Using the Robots Meta Tag” on page 113.

29 Click *Yes* next to *Index may be copied to other clustered servers* if you want to allow this index to be copied to other servers in a QuickFinder Synchronization cluster.

For more information about QuickFinder Synchronization, see Chapter 9, “Synchronizing Data Across Multiple QuickFinder Servers,” on page 103.

30 If the URLs to be crawled require authentication, use the *Type of Authentication required to crawl web site* option to select the methods for providing the correct user credentials.

- **Basic:** If you know that the server to be indexed requires basic authentication, select *Basic*, then specify the username and password in the *Crawler Credentials* fields.

  For example, if you are indexing www.company1.com and it uses basic authentication, specify the username (user ID) and password in the *Crawler Credentials* fields. In this case, the credentials are sent using an HTTP authorization header with every request made to the server of the URL you have specified.

- **Form:** If the server to be indexed uses form-based authentication, type the correct user credentials in the *Form Fields* box. For example: UserIDField:$$UserID

In form-based authentication, the first time the Web site is indexed, the credentials are sent and a session cookie is returned. Thereafter, QuickFinder uses the session ID in the cookie for authentication and the credentials are no longer sent to the Web site.

If you are indexing more than one URL and each one requires a different set of credentials, we recommend that you create a separate index for each URL.

By default, QuickFinder Server sends the form-based credentials by using the HTTP Post protocol. If the Web sites being indexed require the HTTP Get protocol, deselect the *Send login data using HTTP post protocol* check box. When this option is not selected, QuickFinder Server sends the form-based credentials as query parameters to the URLs being indexed.

31 (Optional) If the Web sites you are indexing require users to log in at a specific URL (such as login.digitalairlines.com), specify the login URLs in the *Alternate Login URLs* field.

After the session cookies are returned, QuickFinder sends the appropriate cookies as needed to the Web sites being indexed.

32 Select *Yes* next to *Use Crawler Credentials when Highlighting* to use the *Crawler Credentials* specified Step 30 instead of the search user's credentials when requesting the specified documents.

33 In the *HTTP Headers* field, specify any additional headers and values you want included with each HTTP request, placing each header on a separate line.

Some Web sites require specific information in HTTP headers when attempts are made to access them. If your Web site uses form-based or cookie-based authentication, you can specify the information here.

34 Click *Apply Settings*.

After you define an index, you must generate it to make it searchable. See "Generating Indexes" on page 95.

8.5.3 Configuring Rights-Based Search Results for Crawled Indexes

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.

2 Under *Define a New Index*, select *New crawled index*, then click *Define Index*.

3 On the *Define Crawled Index* page, click *Advanced Index Definition*.

4 Under *Rights-based Search Results*, configure authorization checking by selecting one of the following options:

- **Use Default:** Select this option if you want this index to use the default authorization checking setting specified on the Index Settings page of your virtual search server.
- **Off:** If you want all users to have access to this index, select this option. No authorization checking is done.
- **by Index:** To enable rights checking for this index, specify a file that exists on your server that can be used in verifying user access. By creating a file and setting access rights to it, QuickFinder can verify access to this index based on the rights to the file. Click *Use default path* if a path was specified on the Index Settings page.

5 From the *Unauthorized hits filtered by* drop-down list, select one of the following filters:

- **Use Default:** Select this option if you want the current index to use the default setting found on the Index Settings page.
- **Search Engine:** When you select this option, users attempting to search the index without logging in do not see any of the unauthorized hits on the search results page. If the user doesn't have access to any search results, then the system returns a `No Results Found` message on the search results page.
- **Templates:** When you select this option, users attempting to search the index without logging in to the system receive results, but they are then required to provide a username and password before being allowed to see the contents.

6 Click *Apply Settings*.

After you define an index, you must generate it to make it searchable. See "Generating Indexes" on page 95.

8.5.4 Creating a File System Index

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.

2 Under *Define a New Index*, click *New file system index*, then click *Define Index*.

3 In the *Index Name* field, specify a name for your index.

A name can be a word, phrase, or a numeric value. If the virtual search server you are working on contains, or will contain, a large number of indexes, you might want to use a numbering scheme to help you manage multiple indexes more effectively. However, the name you enter here appears on the default search page, so you might want to choose a name that can be understood by users of your search services.

4 In the *Server Connection* field, select *Yes* if the files to be indexed are on an NCP™ server, then specify the NCP server name, a valid username, password, and the character set of the server.

The user must have at least read rights to all the files. If you want to do rights-based searches, the user must have administrator rights to the NCP server.

The *Server Charset* option must be set correctly so that the URLs can be properly encoded (according to the server encoding).

This option is useful if you have a local NSS volume on the same machine as your QuickFinder index and you want to create a rights-based search for your users, or if you have a local or remote NCP server (such as a NetWare server, another server with NCP on it, or a local indexing machine) and you want to centralize your indexing.

If you choose this option, make sure that the *Corresponding URL Prefix* option in the *Path Information* section contains a complete URL so that your users can access the indexed files from the NCP server. Also, if you are planning on indexing a large number of files (for example, over a million), your system should have at least 2 GB of memory (or more).

5 In the *Server path to be indexed* field, specify the absolute path to the folder containing the information that you want indexed (for example, `sys:\sales\reports`).

6 In the *Corresponding URL prefix* field, specify the URL that should be used by the search results page to access the individual files (for example, `/sales`).

You can also specify a file URL containing the UNC path of the server and path. The syntax is file://///server-dns-name/volume/path.

If the filename contains non-ASCII characters, set the return encoding of the Search Result page to match the encoding of the client's machine. You can set the return encoding on the General Settings page, or the client can set it by specifying the value of the encoding being sent to the search server in the retencoding search parameter. For example, the default English Windows* encoding is Windows-1252, and Japanese is Shift-Jis. In order for the page to open, the client must already be authenticated to the server that the path points to.

7 To add additional paths, click *Add More Paths*.

8 Click *Apply Settings*.

After you define an index, you must generate it to make it searchable. See "Generating Indexes" on page 95.

8.5.5 Creating an Advanced File System Index

1 On the QuickFinder Server Manager Global Settings page, click *Manage* in the row of the virtual search server that you want to work with.

2 Under *Define a New Index*, click *New file system index*, then click *Define Index*.

3 On the Define File System Index page, click *Advanced Index Definition*.

4 In the *Index Name* field, specify a name for your new index.

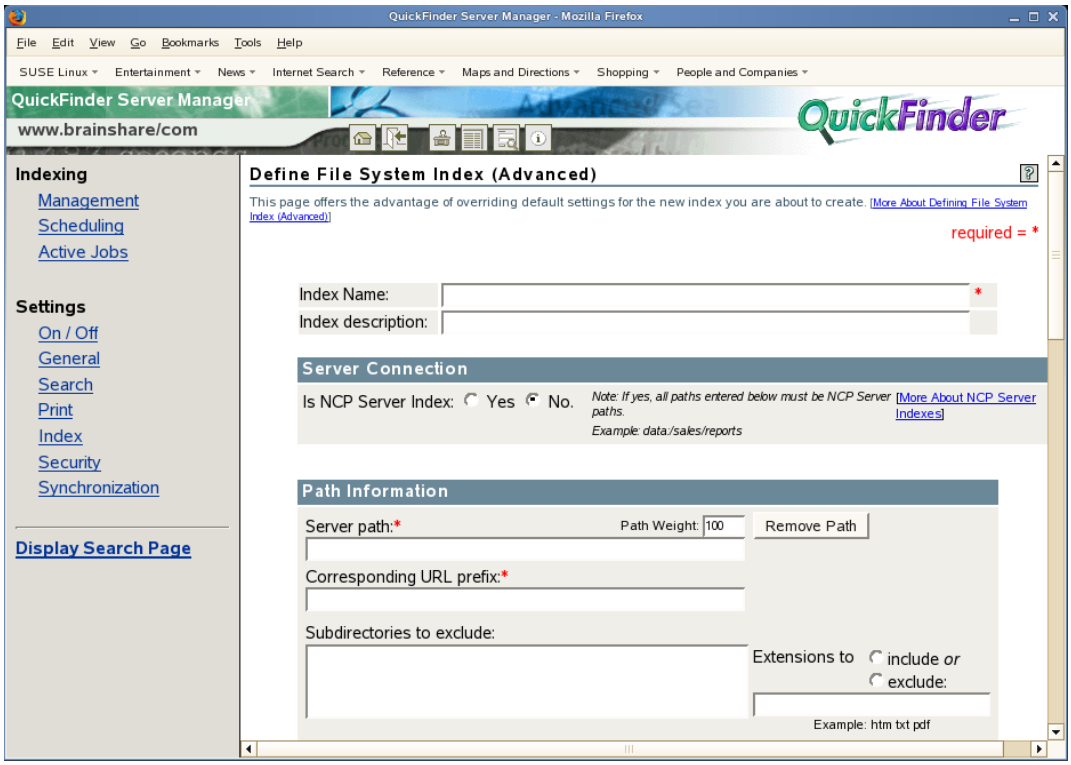

A name can be a word, phrase, or a numeric value. If the virtual search server you are working on contains, or will contain, a large number of indexes, you might want to use a numbering scheme to help you manage multiple indexes more effectively. However, the name you enter here appears on the default search page, so you might want to choose a name that can be understood by users of your search services.

5 In the *Index description* field, specify an optional description of the index to be created.

6 In the *Server Connection* field, select *Yes* if the files to be indexed are on an NCP server, then specify the NCP server name, a valid username, password, and the character set of the server.

   The user must have at least read rights to all the files. If you want to do rights based searches, the user must have administrator rights to the NCP server.

   The *Server Charset* option must be set correctly so that the URLs can be properly encoded (according to the server encoding).

   This option is useful if you have a local NSS volume on the same machine as your QuickFinder index and you want to create a rights-based search for your users, or if you have a local or remote NCP server (such as a NetWare server, another server with NCP on it, or a local indexing machine) and you want to centralize your indexing.

   If you choose this option, make sure that the *Corresponding URL Prefix* option in the *Path Information* section contains a complete URL so that your users can access the indexed files from the NCP server. Also, if you are planning on indexing a large number of files (for example, over a million), your system should have at least 2 GB of memory (or more)

7 Under *Path Information*, specify the absolute path to the folder containing the information that you want indexed in the *Server path* field (for example, `sys:\sales\reports`).

8 (Optional) Use the *Path Weight* option to boost or degrade search results based on the path.

   A weight of 100 makes the path's relevance normal. Increasing the weight makes the path more relevant, while lowering the weight makes the path less relevant.

9 In the *Corresponding URL prefix* field, enter the URL that should be used by the search results page to access the individual files (for example, /sales).

   You can also enter a file URL containing the UNC path of the server and path. The syntax is file:////server-dns-name/volume/path.

   If the filename contains non-ASCII characters, set the return encoding of the Search Result page to match the encoding of the client's machine. You can set the return encoding on the General Settings page, or the client can set it by specifying the value of the encoding being sent to the search server in the retencoding search parameter. For example, the default English Windows encoding is Windows-1252, and Japanese is Shift-Jis. In order for the page to open, the client must already be authenticated to the server that the path points to.

10 To exclude specific subdirectories from being indexed, specify their relative paths in the *Subdirectories to exclude* field.

11 To direct QuickFinder to include or exclude specific file types, click *Extensions to include* or *Extensions to exclude* and then type the extensions, separating each one with a single space, such as HTM PDF TXT.

12 (Optional) To add additional paths, click *Define More Paths*.

13 (Optional) To delete a path, select it and click *Remove Path*.

14 Use the *Adjust Individual File Relevance* option to adjust the relevance of individual items within the index.

   Adjustment values can range from 1 to 200. Values higher than 100 increase the calculated relevance of the item on the search results page, and values lower than 100 decrease the calculated relevance of the item. The value specified here is combined with other values to determine the final relevance.

15 In the *Location of index files* field, specify the absolute path to where you want the index files stored.

For example, `sys:\qfsearch\sites\mysites` on NetWare or `/var/lib/qfsearch/sites/mysites` on Linux.

By default, index files are stored at *`volume`*`:\`*`searchroot`*`\sites\`*`site_name`*`\indexes\`.on NetWare or `/var/lib/qfsearch/Sites/default/indexes/`*`index_name`* on Linux

16 From the *Level of detail in indexing logs* drop-down list, select the amount of information you want included in the index logs.

Option	Description
Disabled	Turns off index logging.
Terse	Lists only the files indexed.
Normal	Lists the files indexed and the results of the crawl.
Verbose	Lists the files indexed, the results of the crawl, and the links that were skipped during the crawl.
New Links	Lists the files indexed, the results of the crawl, the links that were skipped, and any new links found during the crawl.
All Links	Lists the files indexed, the results of the indexing, the links that were skipped, and all links found during the crawl.

17 From the *Encoding (if not in META tags)* drop-down list, select the encoding to be used when indexing files that do not contain an encoding specification.

For example, HTML files can specify their encoding with a Content-Type meta tag.

18 In the *Field names to display on Search Results pages* field, list any field names that you want to be included on the search results page (for example, author date_created product dc.copyright).

To display the field on the search results pages, add the corresponding $$Variable to the template (for example, $$author $$date_created $$product $$copyright).

The field data is stored in the index and causes the index size to increase.

19 Use the *Index Weight* option to boost or degrade search results based on the item's index.

A weight of 100 makes the item's relevance normal. Increasing the weight makes the item more relevant, and lowering the weight makes the item less relevant.

20 In the *Maximum index depth* field, specify the number of directories from the starting directory QuickFinder should search.

This lets you limit how far (or deep) into a file server QuickFinder should search.

21 In the *Maximum file size to index* field, specify the maximum file size (in bytes) that QuickFinder should index.

Files exceeding this size are not indexed and are not included in search results.

22 (Optional) Click *Yes* next to *Index may be copied to other clustered servers* if you want this index shared with other QuickFinder servers in a QuickFinder Synchronization cluster.

For more information about QuickFinder Synchronization, see Chapter 9, “Synchronizing Data Across Multiple QuickFinder Servers,” on page 103.

23 Click *Apply Settings*.

After you define an index, you must generate it to make it searchable. See "Generating Indexes" on page 95.

8.5.6 Configuring Rights-Based Search Results for File System Indexes

1. On the QuickFinder Server Manager Global Settings page, click *Manage* in the row of the virtual search server that you want to work with.
2. Under *Define a New Index*, select *New file system index*, then click *Define Index*.
3. On the Define File System Index page, click *Advanced Index Definition*.
4. Under *Rights-based Search Results*, configure authorization checking by selecting from one of the following options:
   - **Use Default:** Select this option if you want this index to use the default authorization checking setting specified on the Index Settings page of your virtual search server.
   - **Off:** If you want all users to have access to this index, select this option. No authorization checking is done.
   - **by Index:** To enable rights checking for this index, specify a file that exists on your server that can be used in verifying user access. By creating a file and setting access rights to it, QuickFinder can verify access to this index based on the rights to the file. Click *Use Default Path* if one was specified on the Index Settings page.
   - **by Result Item:** If checked, QuickFinder verifies the user's access rights to each hit. This is not recommended for high-traffic servers because checking every hit can slow down server performance.
5. From the *Unauthorized hits filtered by* drop-down list, select one of the following filters:
   - **Use Default:** Select this option if you want the current index to use the default setting found on the Index Settings page.
   - **Search Engine:** When you select this option, users attempting to search the index without logging in do not see any of the unauthorized hits on the search results page. If the user doesn't have access to any search results, then the system returns a `No Results Found` message on the search results page.
   - **Templates:** When you select this option, users attempting to search the index without logging in to the system receive results, but they are then required to provide a username and password before being allowed to see the contents.
6. Click *Apply Settings*.

After you define an index, you must generate it to make it searchable. See "Generating Indexes" on page 95.

8.5.7 Searching Across Multiple Indexes

QuickFinder can search across multiple indexes within a single virtual search server. However, searching a single index is generally faster than searching across multiple indexes.

Restricting Search Results to Specific Areas

You can restrict search results to specific areas of your file or Web server in the following ways:

- Using multiple indexes and using the &index=index_name query parameter.
- Using a single index and restricting results to certain URL paths by using the &filefilter=path query parameter.
- Using a single index and restricting results to certain values in document fields by including ^fieldname=value with either the query=value or filter=value search parameters.

TIP: Using the last option requires that indexed documents contain summary fields such as meta tags. This option works for almost any file format that contains document summary fields, including HTML, XML, PDF, Word, and WordPerfect*.

For information about preventing QuickFinder from indexing specific content, see "Excluding Documents from Being Indexed" on page 113.

8.5.8 Indexing Content on a Password-Protected Web Site

If the Web servers you want to index require authentication, there are two methods for providing the correct user credentials: basic authentication and form-based authentication. Which one you choose depends on how authentication is implemented on the Web sites you index. For example, if you are indexing www.company1.com and it uses basic authentication, then specify the username (user ID) and password in the *Crawler Credentials* fields. In this case, the credentials are sent using an HTTP authorization header with every request made to the server of the URL you have specified.

However, if www.company1.com uses a form-based authentication method, leave the *Crawler Credentials* fields blank and type the correct user credentials in the *Form Fields* text box. For example: UserIDField:$$UserID.

In Form-based authentication, the first time the Web site is indexed, the credentials are sent and a session cookie is returned. Thereafter, QuickFinder uses the session ID in the cookie for authentication and the credentials are no longer sent to the Web site.

TIP: If you are indexing more than one URL and each one requires a different set of credentials, we recommend that you create a separate index for each URL.

8.5.9 Indexing Volumes on Remote Servers

You can index file content stored on remote volumes by using the Novell NFS Gateway product, which is included with NetWare 6.5.

To be able to index remote servers, first install and configure the NFS Gateway. Then create a new file system index and specify the path to a volume on a remote server. QuickFinder does the rest.

For more information about installing and configuring the NFS Gateway, see the NFS Gateway for NetWare 6.5 (http://www.novell.com/documentation/nfsgynw65/index.html) documentation.

8.5.10 Generating Indexes

After you define an index, you must generate it before it can be used for searching. Generating an index is the actual process where QuickFinder Server examines file server or Web server content, gathers keywords, titles, and descriptions and then includes them in the index.

- "Generating an Index" on page 95
- "Generating an Index For a Linux-mounted NSS Volume" on page 96
- "Generating a File System Index" on page 96

Generating an Index

1. On the QuickFinder Server Manager Global Settings page, click *Manage* in the row of the virtual search server that you want to work with.
2. Click *Generate* in the *Action* column of the index that you want to work with.

   The Active Jobs page indicates the status of the current indexing jobs. When there is no current index job, the status page reads `No indexing jobs are currently running or defined`.

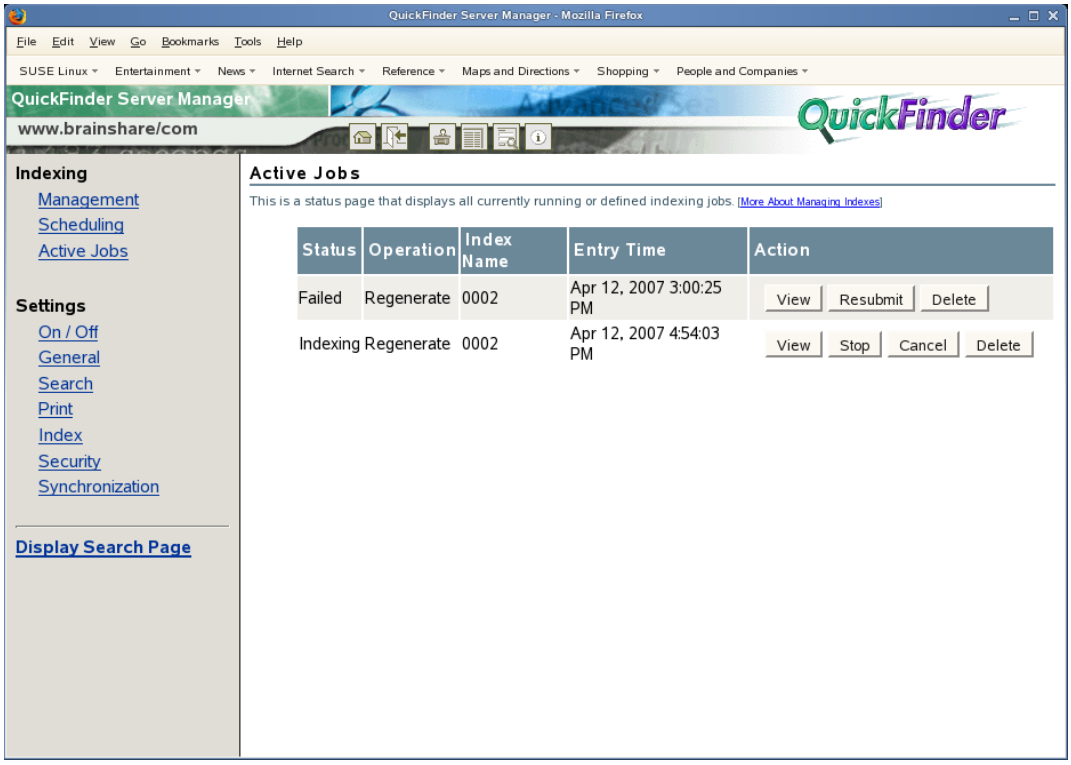

3. To cancel the current indexing jobs, click *Cancel* in the *Status* column.

You can direct QuickFinder to automatically update your indexes on specific dates and at specific times by scheduling events. For more information, see Section 8.9, "Automating Index and Server Maintenance," on page 100.

Generating an Index For a Linux-mounted NSS Volume

To generate an index for a Linux-mounted NSS volume, the novlwww user or www group must have read access to the NSS volume. To do this, verify that the novlwww user and the www group are LUM-enabled, then give read access to the NSS volume by running `rights` and assigning the user or group trustee rights to the volume.

NOTE: The novlwww user is the Open Enterprise Server defined Tomcat user. QuickFinder runs as a servlet so its access rights are the same as the Tomcat user.

Generating a File System Index

When generating a file system index and specifying a set of filename extensions to index, you could end up indexing files you don't want.

For example, you index your entire hard drive and look for only HTM and HTML files. There are about 10,000 properly matching files on your file system, but you end up with over 30,000 files in your index. This is because the file system scanner includes files with no filename extensions. In some cases, including files with no extension is better than not including them, but in this case, the index of all the HTML files on your hard drive is not helpful because it contains a large number of non-HTML files.

To avoid this kind of situation, manually modify the QuickFinder Server configuration files (`/var/lib/qfsearch/Sites/qfind.cfg` on Linux and `sys:/qfsearch/sites/qfind.cfg` on NetWare). Add `IncludeNoExtension N` to the `<Directory>` section of an index definition to prevent files with no filename extensions from being included (it defaults to Y). Put this line next to the `IncludeExtension` setting, which lists all of the filename extensions to be indexed.

TIP: QuickFinder Server can only index files accessible through local file system calls. If you mount a volume or map a drive to a remote server and the local system sees it as a local drive, QuickFinder Server can index it.

8.6 Managing Existing Index Files

After it is created, an index can then be edited or deleted. You can also view an index's log file.

- Section 8.6.1, “Editing an Index,” on page 96
- Section 8.6.2, “Deleting an Index,” on page 97
- Section 8.6.3, “Working with the Log File,” on page 97

8.6.1 Editing an Index

1. On the QuickFinder Server Manager Global Settings page, click *Manage* in the row of the virtual search server that you want to work with.
2. On the Index Management page, click *Edit* in the *Action* column of the index you want to work with.

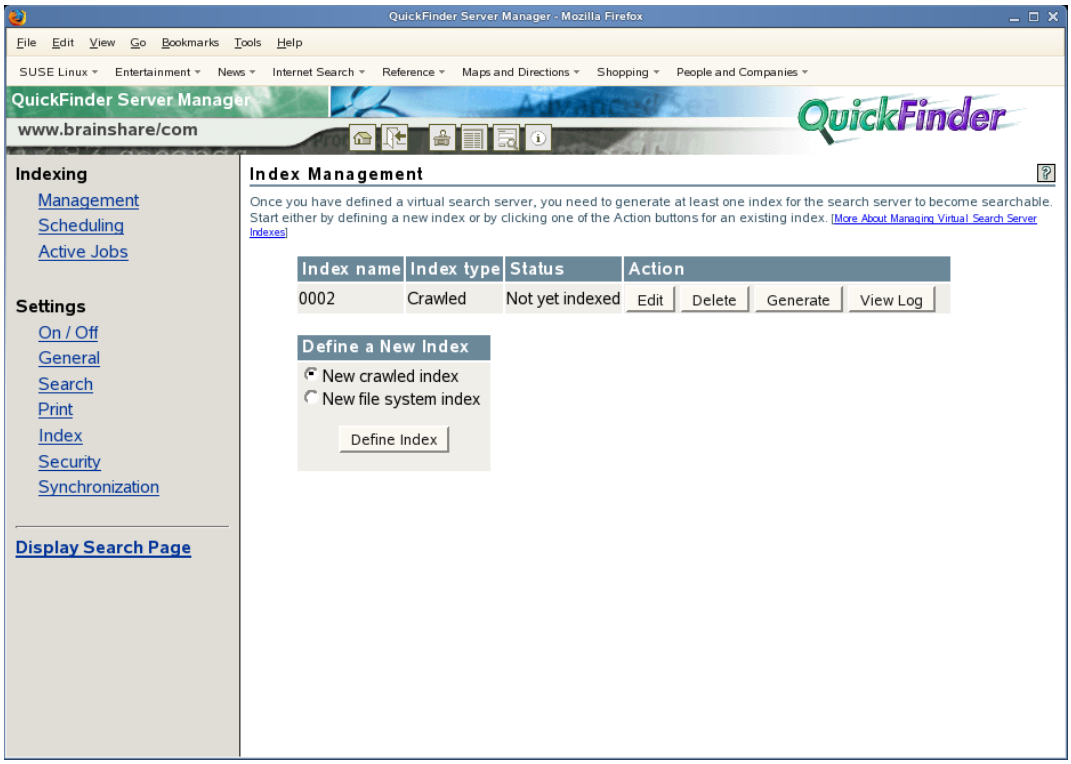

3 Make any of the changes you need to, then click *Apply Settings*.

If you used the Advanced page to create the index, it appears automatically. However, you can also click *Advanced Index Definition* to make advanced changes to an index you created using the standard Index Definition page.

4 If you added new paths or URLs or modified any of the existing ones, you should regenerate the index to include the new content.

8.6.2 Deleting an Index

1 On the QuickFinder Server Manager Global Settings page, click *Manage* in the row of the virtual search server that you want to work with.

2 On the Index Management page, click *Delete* in the *Action* column of the index you want to delete.

3 On the *Confirm Deletion of <indexname>* page, click *Delete Index* to proceed, or click *Cancel Deletion*.

WARNING: After an index has been deleted, it cannot be restored. You must generate a new index.

8.6.3 Working with the Log File

The purpose of the log file is to help you identify any errors (and their possible causes) during an indexing job.

In addition to reporting when the indexing job started and stopped, the log file also lists all files that were indexed, files that could not be found but were linked to, and even errors that might have occurred during the indexing process.

To view an index's log file:

1. On the QuickFinder Server Manager Global Settings page, click *Manage* in the row of the virtual search server that you want to work with.
2. Click *View Log* in the *Action* column of the index that you want to work with.
3. Review the contents of the log file.

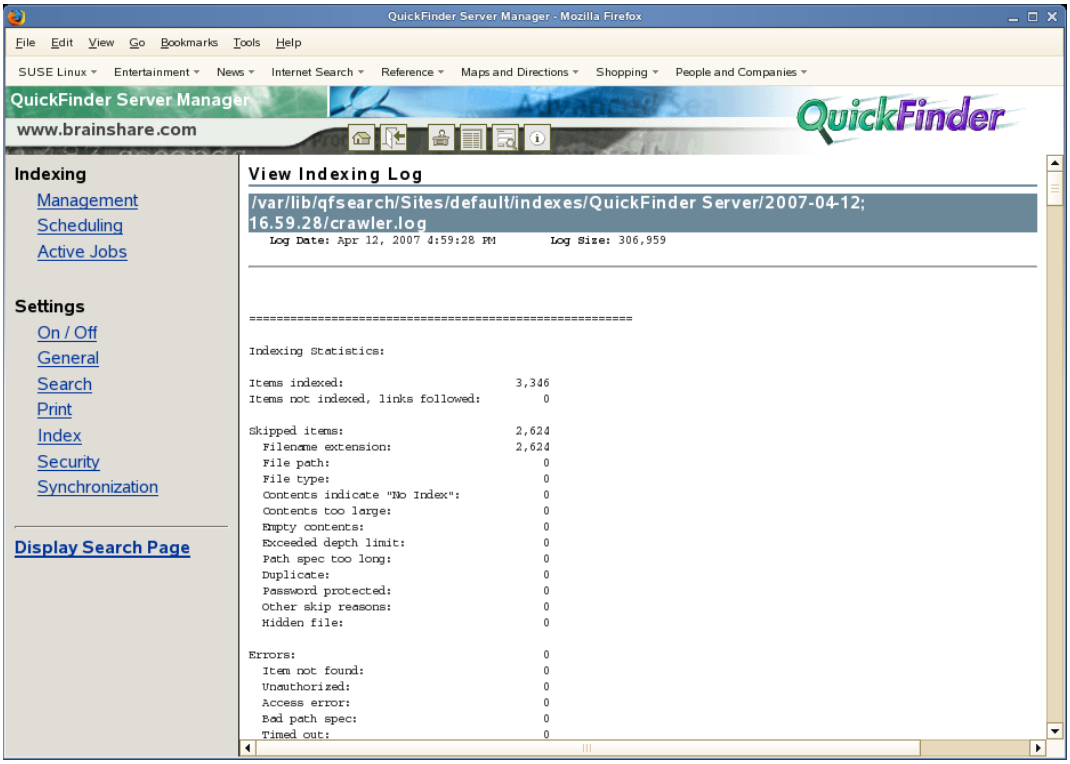

4. Click your browser's *Back* button to return to the Index Management page, or click *Management* in the left frame of the QuickFinder Server Manager.

8.7 Indexing Dynamic Web Content

Much of the content on the World Wide Web is static HTML, which means that after a static Web page is created, it remains the same until someone updates it. By contrast, many newer Web pages are created by Web applications, including servlets, Java* Server Pages (JSP*), Common Gateway Interfaces (CGI), and Perl scripts, and are usually created in response to user input.

An example of dynamic Web content is an eCommerce Web page where items to be purchased are stored in a virtual shopping cart and the total cost is updated as users add or remove items from the shopping cart.

Because the content changes regularly, many search engines don't index dynamic content.

QuickFinder includes the ability to index dynamic content. The URL of dynamic Web content typically includes a question mark (?). You can direct QuickFinder to index these URLs by setting the *Crawl dynamic URLs* option to *Yes*. You could then create a scheduled event that regenerates the specified indexes every few minutes.

8.8 Access Control to Search Results

Security settings let you manage access to indexed content by requiring users to authenticate to a server before seeing search result content.

In NetWare 6, Web Search depended on Enterprise Web Server NLM™ programs to authorize each search result for a particular user for those documents managed by eDirectory. QuickFinder Server now depends on Novell eDirectory. In addition to other enhancements, this allows user privileges to entire indexes rather than to each search result, which improves the overall speed at which search results requiring authentication are returned to the user.

To configure rights-based search results:

1. On the QuickFinder Server Manager Global Settings page, click *Manage* in the row of the virtual search server that you want to work with.
2. Under *Settings*, click *Security*.
3. In the *Default search contexts* field, specify the full context that your users share in the directory.

   Specifying the context here means that your users do not need to supply their fully qualified user IDs when logging in to QuickFinder.

   If User objects are stored in multiple containers, you can enter multiple contexts by separating them with carriage returns. For example:

   ```
   users.marketing.digitalairlines
   users.sales.digitalairlines
   users.support.digitalairlines
   ```
4. Click *Yes* next to *Check authorization by directory* if you want QuickFinder to verify user rights to all files in a directory.

   QuickFinder allows access to all other files in any directory where the user has rights to at least one file. This method speeds up access rights checking because QuickFinder doesn't need to check rights for each file returned in a set of search results.
5. In the *Authentication realm string* field, specify the correct authentication realm string, which is typically shown in the Web browser's login dialog box.

   If QuickFinder's realm string matches your Web server's realm string, users only need to log in once.

   If Apache is your Web server, refer to the Authentication, Authorization and Access Control (http://httpd.apache.org/docs-2.0/howto/auth.html) document on the Apache.org Web site for more information.
6. Select *Yes* next to *Require https* if you want to enable HTTPS for user connections.

   IMPORTANT: We recommend that you enable HTTPS by selecting *Yes*, which ensures that usernames and passwords are encrypted as they cross network wires.
7. In the *Auto-logout time* field, specify a number (in minutes) when QuickFinder should log out a user.

This option is a security feature. Sometimes, users walk away from their computers. Other users could then have access to their search results.

8 (Optional) Check *Disable auto-logout* if you don't want it used during rights-protected searches.

We recommend that you enable this feature as an added protection to secure content.

9 Click *Apply Settings*.

8.9 Automating Index and Server Maintenance

You can eliminate a lot of manual work in keeping indexes up-to-date by using the QuickFinder Server index scheduling feature. Because the Web and file content you have indexed will eventually change, you can direct QuickFinder to update your indexes on specific dates and at specific times or intervals.

- Section 8.9.1, “Adding a Scheduled Event,” on page 100
- Section 8.9.2, “Editing or Deleting an Event,” on page 101

8.9.1 Adding a Scheduled Event

1 On the QuickFinder Server Manager Global Settings page, click *Manage* in the row of the virtual search server that you want to work with.

2 Click *Scheduling* in the left frame of QuickFinder Server Manager.

3 Click *Add Event*.

4 Specify the month, days, days of the week, or time (in hours and minutes) when you want QuickFinder to run the event.

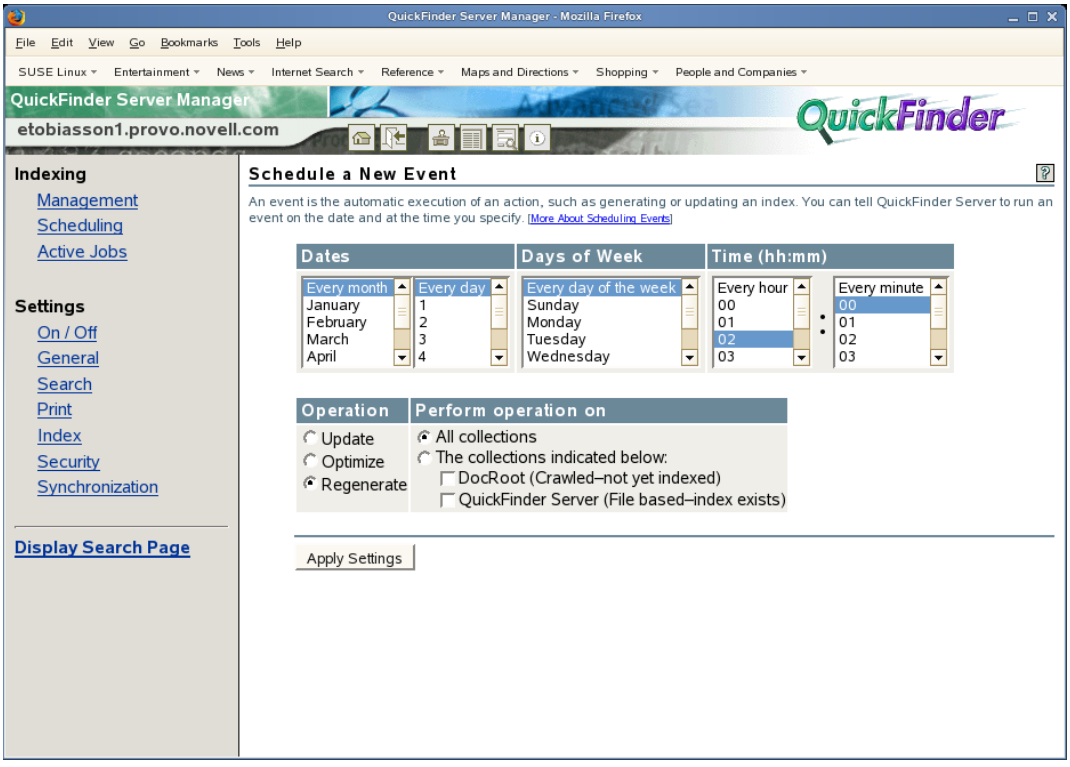

To select multiple dates and times, hold down the Ctrl key and click all of the items you want added. To select consecutive items, click the first item, hold down the Shift key, then click the last item.

5 Select the type of operation you want performed on your indexes.

- **Update:** QuickFinder identifies new content on Web or file servers and updates the index.
- **Optimize:** QuickFinder improves searching performance by removing unnecessary content and making the index file more compact.
- **Regenerate:** QuickFinder replaces the existing index with a newly generated one.

6 In the *Perform operations on* column, determine whether you want the chosen operation performed on all indexes or only on specified ones.

If you have large indexes, you might consider creating multiple events that update your indexes at varied times. Doing so minimizes CPU utilization. By default, QuickFinder supports up to 5 simultaneous indexing sessions. All other indexes wait until a previous index job has completed. You can control the number of simultaneous indexing jobs from Services Settings. (See Section 7.3, "Configuring Services Settings," on page 51.)

7 Click *Apply Settings*.

8.9.2 Editing or Deleting an Event

1 On the QuickFinder Server Manager Global Settings page, click *Manage* in the row of the virtual search server that you want to work with.

2 Click *Scheduling* in the left frame of QuickFinder Server Manager.

If no events have been scheduled, refer to Section 8.9.1, "Adding a Scheduled Event," on page 100 for instructions on adding a scheduled event.

3 To edit a scheduled event, click *Edit* in the row of the event you want to modify.

4 Make the desired changes, then click *Apply Settings*.

5 To delete a scheduled event, click *Delete* in the row of the event you want to delete.

6 Click *Delete Event* to confirm the deletion, or click *Cancel Deletion*.

8.10 Backing Up Your Virtual Search Server Files

As with any valuable data, you should make sure that your virtual search server files are backed up. At minimum, you should back up your index files, which by default are stored in `*volume*:\qfsearch\Sites\*sitename*\indexes` on NetWare and in `/var/lib/qfsearch/docs/index.html` on Linux.

If you have customized templates, you might also want to back them up. By default, they are stored in `*volume*:\qfsearch\Templates` on NetWare and in `/var/lib/qfsearch/Templates` on Linux.

Synchronizing Data Across Multiple QuickFinder Servers

9

If you are running multiple installations of QuickFinder™ Server across several servers, you can simplify management tasks, such as re-generating indexes and sharing configuration settings, by configuring all of your servers into a single QuickFinder Synchronization cluster.

QuickFinder Synchronization lets you designate one QuickFinder server as the search master from which updated indexes, templates, and configuration settings are systematically sent out to all other QuickFinder servers defined as part of a QuickFinder Synchronization cluster. Each of the other servers must be configured to receive updates from the search master.

QuickFinder Synchronization offers several key benefits:

- Saves system resources on all other machines because they do not need to regenerate indexes themselves.
- Lets an administrator manage all other QuickFinder servers from a single interface.
- Regenerates indexes nightly on the search master and then pushes them out to all other QuickFinder servers.
- Offers fault tolerance and load-balancing benefits by hosting QuickFinder indexes, templates, and configuration settings across multiple servers; if one server goes down, your users will not know the difference.

  You might want to use a Novell® iChain® server to front the QuickFinder Synchronization cluster servers. This would offer basic load-balancing capabilities using revolving IP addresses. For more information, see the *Novell iChain 2.3 Administration Guide* (http://www.novell.com/documentation/ichain23/index.html).
- Ensures that all QuickFinder servers within a cluster are using the same version of an index. This means that you never have a server with a newer version of an index than all other servers. Even the search master won't use the newly generated indexes unless all of the servers in the cluster can also do so.

IMPORTANT: QuickFinder Server Synchronization is different from Cluster Services. Novell Cluster Services™ groups two or more NetWare® servers together for the sake of failover and to ensure server availability. A QuickFinder Synchronization cluster synchronizes indexes, templates, and configuration settings between two or more QuickFinder servers, thus reducing the administrative work of updating the same index across several QuickFinder servers. (See Chapter 6, "Configuring QuickFinder Server for Novell Cluster Services," on page 33.)

To enable QuickFinder Synchronization, you need to do the following:

- Configure each QuickFinder server and virtual search servers that you want included in the cluster to receive updates. See Section 9.1, "Setting Up a Search Server to Receive Updates," on page 104.
- Designate one of the QuickFinder servers as the search master and configure it to send updates to other QuickFinder servers by identifying each of the other servers to be included in the cluster. A search master can continue to serve as a QuickFinder server, and you can have more

than one search master in your QuickFinder Synchronization cluster. See Section 9.2, “Setting Up a Search Master and Defining a Cluster,” on page 106.

- Configure which indexes on the search master will be included in the synchronized cluster. See Section 9.3, “Setting Up a Virtual Search Server to Send Updates,” on page 109

Although the order in which you perform these tasks is not critical, we recommend that you set up the Search servers intended to receive cluster data first, and then configure a QuickFinder server as the search master and use it to define the synchronization cluster.

- Section 9.1, “Setting Up a Search Server to Receive Updates,” on page 104
- Section 9.2, “Setting Up a Search Master and Defining a Cluster,” on page 106
- Section 9.3, “Setting Up a Virtual Search Server to Send Updates,” on page 109
- Section 9.4, “Sending Synchronization Updates Manually,” on page 109

9.1 Setting Up a Search Server to Receive Updates

For each QuickFinder server you want included in a synchronization cluster, you must configure it to receive updates. Also, each virtual search server must be configured to receive updates. You must make the same changes on every server to be included in the cluster.

- Section 9.1.1, “Setting Up a QuickFinder Server to Receive Cluster Data,” on page 104
- Section 9.1.2, “Setting Up a Virtual Search Server to Receive Updates,” on page 105

9.1.1 Setting Up a QuickFinder Server to Receive Cluster Data

1. On the Global Settings page of QuickFinder Server Manager, click *Synchronization* under *Services Settings*.
2. Under *General Synchronization Settings*, click *No* next to *Will this machine send cluster data*.

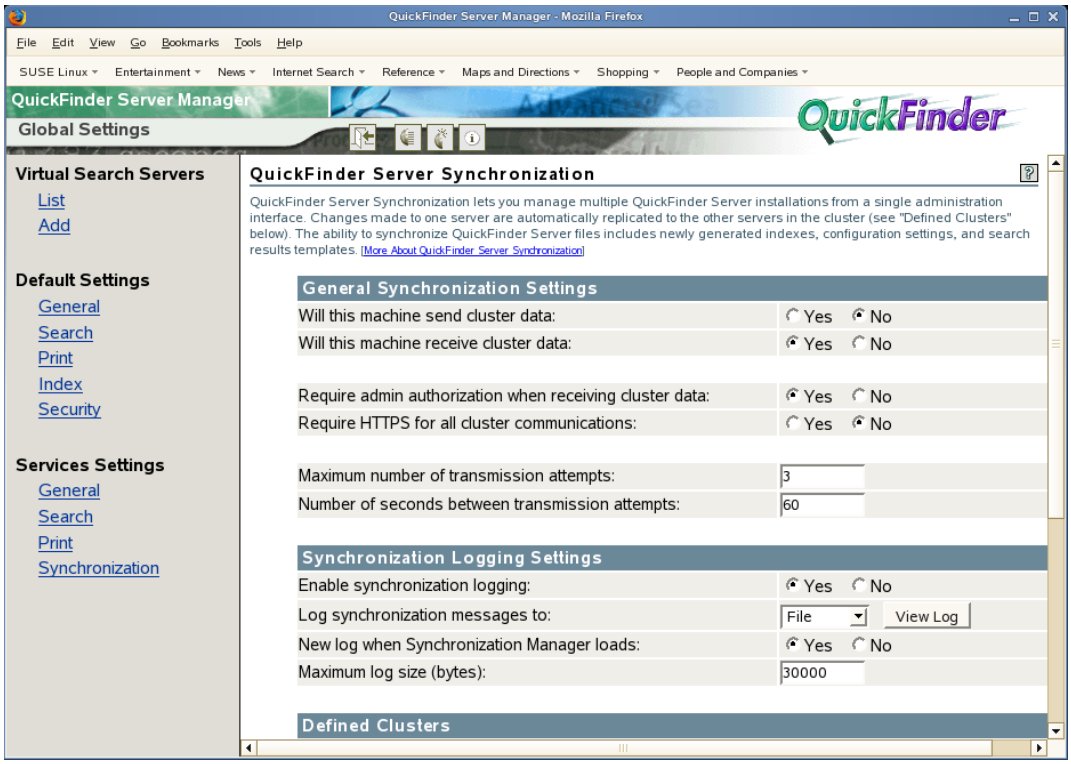

3 Click *Yes* next to *Will this machine receive cluster data*.

4 (Optional) Click *Yes* next to *Require admin authorization when receiving cluster data* if you want to require incoming requests to provide administrator credentials before proceeding.

   If the risk of intrusion to your server is low—for example, if your server is behind a firewall—it is better to disable this feature by clicking *No*. Enabling it can slow down system performance because of the rights checking that happens each time a request to send cluster data is received by the server.

5 (Optional) Click *Yes* next to *Require HTTPS for all cluster communications* to require a secure sockets layer (SSL) connection from a search master attempting to send updated data to this Search Server.

6 In the *Maximum number of transmission attempts* field, specify the number of times QuickFinder should attempt to connect to other servers in a cluster.

7 In the *Number of seconds between transmission attempts* field, type the number of seconds QuickFinder should wait between transmission attempts.

   You might want to set this number high enough to allow a server to fully restart after a failure.

8 Click *Apply*.

After you have configured a QuickFinder server to receive data, you must also configure at least one virtual search server to receive updates. See Section 9.1.2, “Setting Up a Virtual Search Server to Receive Updates,” on page 105 for more information.

9.1.2 Setting Up a Virtual Search Server to Receive Updates

1 On the Global Settings page of QuickFinder Server Manager, select a virtual search server from the list by clicking *Manage*.

2 Under *Settings*, click *Synchronization*.

3 Click *Yes* next to *Let virtual search server receive cluster data*.

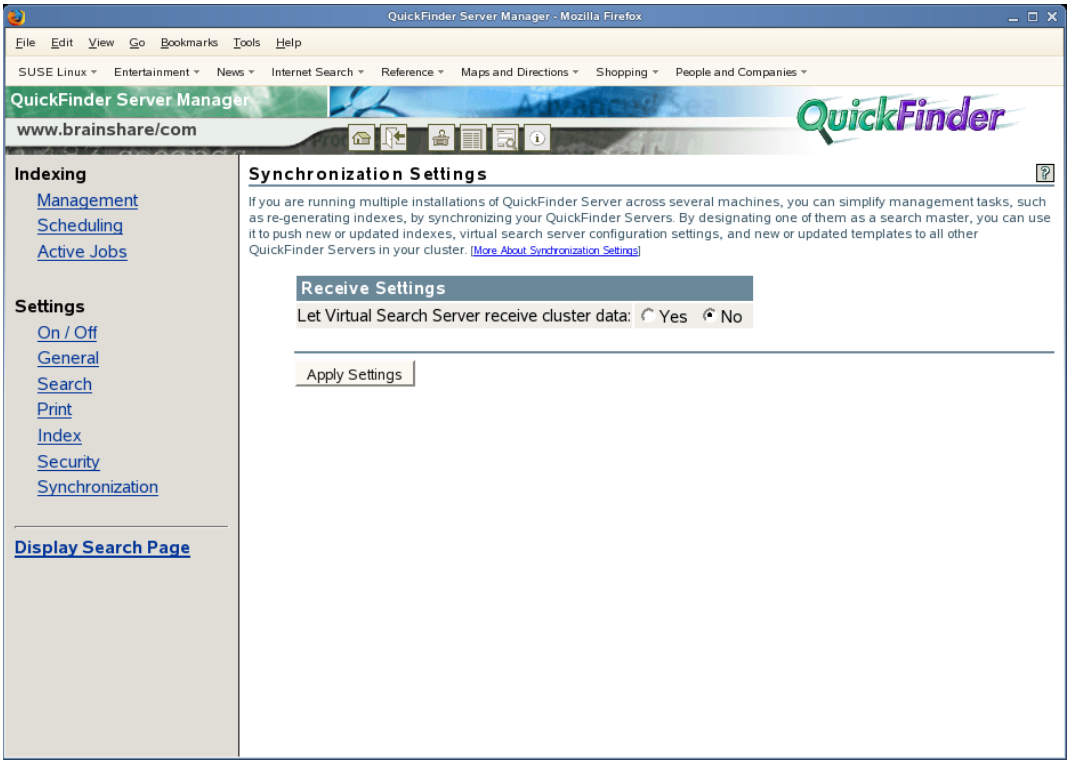

4 Click *Apply Settings*.

9.2 Setting Up a Search Master and Defining a Cluster

After you have configured both the QuickFinder servers and virtual search servers for each server to be included in the QuickFinder Synchronization cluster, you now need to configure a search master and define one or more clusters.

- Section 9.2.1, “Setting Up a QuickFinder Server to Send Updates,” on page 106
- Section 9.2.2, “Defining a Synchronization Cluster,” on page 108

9.2.1 Setting Up a QuickFinder Server to Send Updates

1 On the Global Settings page of QuickFinder Server Manager, click *Synchronization* under *Services Settings*.

2 Under *General Synchronization Settings*, click *Yes* next to *Will this machine send cluster data*.

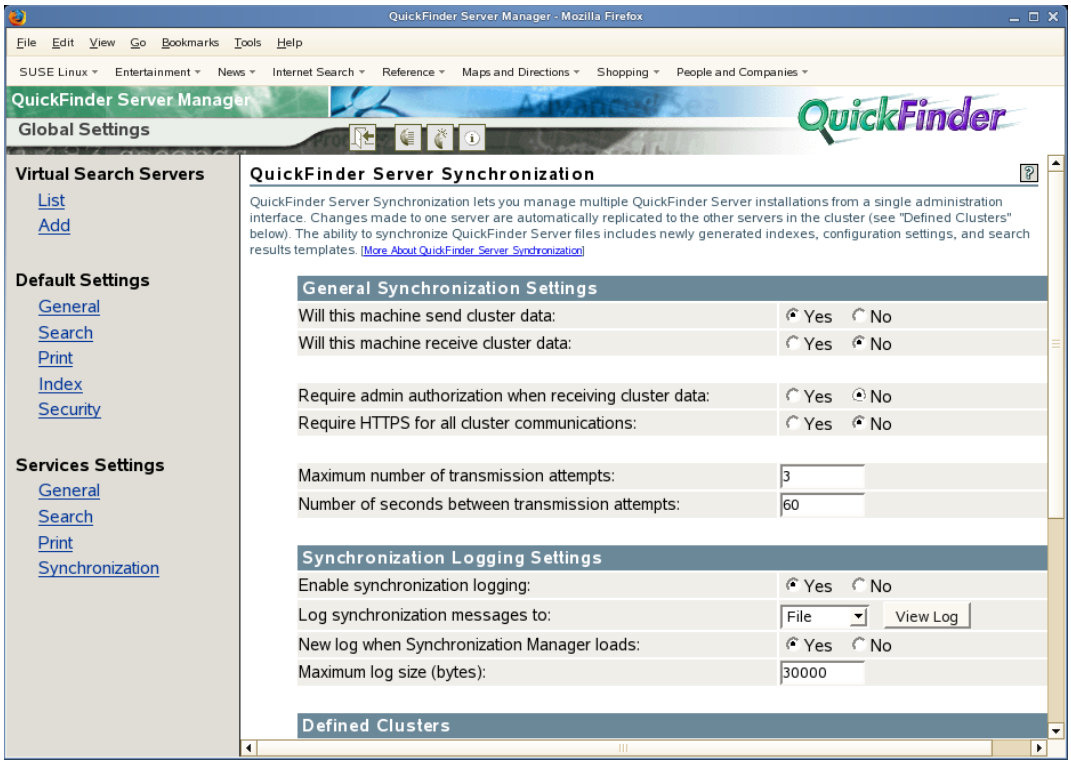

3 Click *No* next to *Will this machine receive cluster data*.

If you want the search master to receive updates from another server in the cluster, select *Yes*. You can configure a QuickFinder server to both send and receive updates.

4 Because this is the search master and will probably not be receiving cluster data, click *No* next to *Require admin authorization when receiving cluster data*.

5 (Optional) To ensure a secure connection with the other QuickFinder Servers, click *Yes* next to *Require HTTPS for all cluster communications*.

If the risk of intrusion to your server is low—for example, if your server is behind a firewall—it is better to disable this feature by clicking *No*. Enabling it can slow down system performance because of the rights checking that happens each time a request to send cluster data is received by the server.

6 In the *Maximum number of transmission attempts*, specify the number of times the search master should attempt to transmit data to an unresponsive QuickFinder Server in the cluster before quitting.

After the number of transmission attempts specified here is reached, QuickFinder stops trying to transmit data to other servers in a cluster. If the e-mail notification capability has been enabled, then QuickFinder also sends an e-mail identifying the cause of the problem. For information about enabling the e-mail feature, see Section 7.3, “Configuring Services Settings,” on page 51.

7 In the *Number of seconds between transmission attempts*, type the number of seconds QuickFinder should wait between each transmission attempt specified in the *Maximum Number of Transmission Attempts*.

8 (Optional) Under *Synchronization Logging Settings*, click *Yes* next to *Enable synchronization logging* if you want QuickFinder to log cluster-related errors.

9 (Optional) Specify where you want log results displayed by selecting one of the following options from the *Log synchronization messages to* drop-down list:

- **File:** When this option is selected, you can click *View Log* next to the *Log synchronization messages to* drop-down list to display the log results in your browser.
- **Console:** You can also view log results at the system console by selecting *Console*, pressing Ctrl+Esc on your server's keyboard, then pressing the number corresponding to the Tomcat servlet engine.
- **Both:** Displays results in both your browser and at the system console.

10 (Optional) Click *Yes* next to *New log when synchronization manager loads* to start a new logging session when the synchronization manager servlet is loaded (or restarted).

The contents of the current log are replaced with the new data.

11 (Optional) In the *Maximum log size* field, specify the maximum file size of the log file (in bytes) before a new one is started.

The number you specify here is divided evenly between two log files. For example, if you specify 30000, each log file allows up to 15000 bytes of logged data. This ensures that you always have at least 15000 bytes of logged synchronization data.

12 Click *Apply*.

9.2.2 Defining a Synchronization Cluster

You must identify each of the QuickFinder Servers (and their virtual search servers) that you have configured to receive cluster data. To do so, you must know the URL of each server to be included, as well as the names of each virtual search server on each of the QuickFinder servers that you want included.

Also, if you have configured any of the participating servers to require administrator authorization when being accessed by the search master, you must provide their administrator usernames and passwords.

1 On the Global Settings page of QuickFinder Server Manager, click *Synchronization* under *Services Settings*.

2 Under *Defined Clusters*, specify a name for the new cluster in the *Cluster <number> Name* field.

3 In the *Server <number> URL* field, specify the URL to a QuickFinder Server that you want included in the cluster, and specify the name of the virtual search server to be included.

4 In the *Virtual Search Server* field, specify the name of a virtual search server.

5 (Optional) If you enabled (or will enable) the *Require admin authorization when receiving cluster data* feature on this QuickFinder Server, specify the admin username and password in the appropriate fields.

6 (Optional) To add another server to the cluster, click *Add Server* and repeat Step 1 through Step 5.

7 Click *Apply*.

8 To add additional clusters, click *Add New Cluster*.

9.3 Setting Up a Virtual Search Server to Send Updates

The last step in setting up QuickFinder Synchronization is to configure one or more virtual search servers on your search master to send updates.

1. On the Global Settings page of QuickFinder Server Manager, select a virtual search server from the list by clicking *Manage*.
2. Under *Settings*, click *Synchronization*.
3. Under *Send Settings*, click *Yes* next to *Let virtual search server send cluster data*.
4. Click *Edit* next to the *Under name of clusters to send to* field.
5. Select one or more clusters that you have defined, then click *Apply*.
6. Select one or more indexes from *Indexes to send to clusters*.

   The indexes you select are sent to the designated clusters every time they are regenerated (manually or scheduled). Also, you can click the *Indexes* button under the *Send indexes* section to send these indexes immediately to the clusters you've selected.
7. Click *Apply Settings*.

9.4 Sending Synchronization Updates Manually

After QuickFinder Synchronization is configured properly, you can go to a virtual search server and send index, template, and configuration updates at any time. For example, if you need to make a correction or addition to a template, you can open the Synchronization Settings page and click *Templates* under *Send Now* and the templates are then pushed out to all other virtual search servers in the clusters you selected in the *Name of Clusters to Send To* field.

1. On the Global Settings page of QuickFinder Server Manager, select a virtual search server from the list by clicking *Manage*.
2. Under *Settings*, click *Synchronization*.
3. Under *Send Now*, click the *Indexes*, *Templates*, or *Settings* button to send updates from the current virtual search server to all others in the cluster.

When the process is underway, a confirmation page is returned. You don't need to wait for the process to complete before continuing to work in QuickFinder Server Manager.

Optimizing Search Results

10

There are a number of ways you can optimize the performance of your virtual search servers. Some of the optimization strategies are put into place as you plan and create your search solution. Others are done after your search solution has been created.

This section discusses things you can do after your search solution is already in place. It includes the following topics:

- Section 10.1, “Reviewing Query Log Reports,” on page 111
- Section 10.2, “Improving Search Results through Intelligent Indexing,” on page 112
- Section 10.3, “Modifying Document Descriptions Returned in a Search Results List,” on page 114
- Section 10.4, “Improving Search Results,” on page 114
- Section 10.5, “Helping Users Avoid Failed Searches,” on page 122
- Section 10.6, “Weighted Queries,” on page 122
- Section 10.7, “Ensuring Optimal Search Speed,” on page 123
- Section 10.8, “Making Good Use of Document Fields,” on page 124
- Section 10.9, “Searching XML Documents,” on page 124
- Section 10.10, “Using the &filter Query Parameter,” on page 125

10.1 Reviewing Query Log Reports

Before you begin changing various QuickFinder™ settings, set up and review query log files for clues about where your search service needs improvements.

Query logging captures many details about user queries, including total number of searches performed (for the current logging period), top 100 search terms, load statistics, and the number of print jobs performed. It also provides information about queries that resulted in a “not found.”

1. On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to work with.
2. Click *General* under *Settings*.
3. To view an existing log report, select a report from the *Available log reports* list, then click *View Report*.

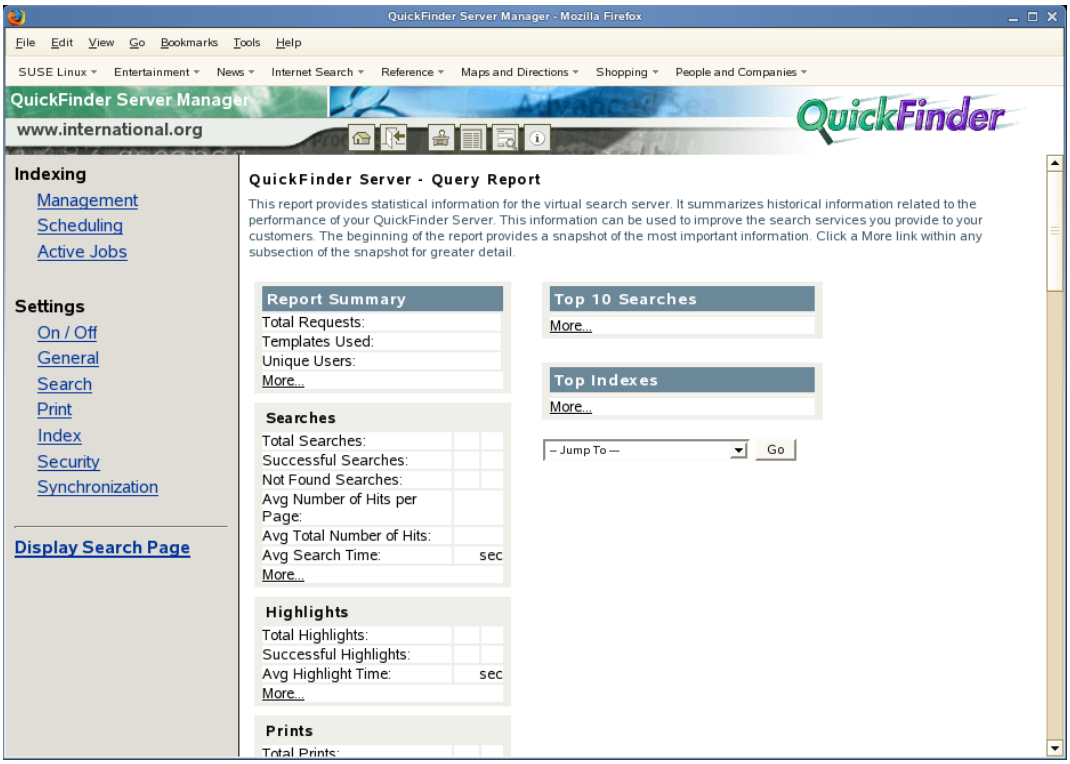

If there are no logs listed, you have not yet generated a log report. Click *Generate Current* to create a report of the current (not yet completed) log period. Depending on the amount of information in the logs, this process could take anywhere from a few seconds up to several hours to complete. After a new report is generated, it appears in this list.

10.2 Improving Search Results through Intelligent Indexing

You can improve the accuracy of your search results by following these indexing guidelines:

- When defining and creating your indexes, start with the highest possible Web site URLs and file system paths.
- If content is showing up in your search results that you don't want included, try removing some paths or URLs from your defined indexes. Also, try excluding specific subdirectories that you know or suspect might contain content that you don't want searched.
- If you've indexed too many file types and cluttered your search results, try removing file types that you don't want indexed by using the *Extensions to Exclude* option on the Define Index page. See "Creating an Advanced Crawled Index" on page 84 and "Creating an Advanced File System Index" on page 90 for more information.
- Use the Robots meta tag (http://www.robotstxt.org/wc/exclusion.html) in your Web site's content.
- Exclude documents or specific sections of documents, including headers, footers, and navigation bars.
- Use the Robots Exclusions standard (http://www.robotstxt.org/wc/exclusion.html) (the `/robots.txt` files found on almost all Web sites). The Index Definition's excludePaths

settings are automatically combined with the `robots.txt` settings when crawling and indexing. You can turn `robots.txt` support on or off.

10.2.1 Excluding Documents from Being Indexed

One way to improve search results is to guard what content is actually indexed, thus clearing a path for relevant information.

- "Using the Extensions to Exclude Option" on page 113
- "Using the Extensions to Include Option" on page 113
- "Using the Robots Meta Tag" on page 113
- "Using the Robots Comment Tag" on page 114

Using the Extensions to Exclude Option

You can use the *Extensions to Exclude* option to direct QuickFinder to ignore specific file types. For example, if you don't want Word or PowerPoint* documents to be included in search results, you would enter DOC and PPT in the *Extensions to Exclude* field. When these document types are encountered during an indexing job, QuickFinder skips them.

For more information on the *Extensions to Exclude* option, see "Creating an Advanced Crawled Index" on page 84 and "Creating an Advanced File System Index" on page 90.

Using the Extensions to Include Option

As mentioned above, you can use the *Extensions to Exclude* option to direct QuickFinder to ignore specific file types. However, if you can't specify all of the extensions to exclude, use the *Extensions to Include* option and specify all acceptable file extensions. A typical list would specify HTM, HTML, PDF, TXT, and DOC.

TIP: When entering extensions in the *Extensions to Exclude* box, separate each extension by a space or a hard return. Don't use commas. For example:

htm html pdf txt doc

Using the Robots Meta Tag

Another effective way of controlling what QuickFinder indexes is using the Robots meta tag. This tag is inserted into header section of a document and contains instructions about what should (or should not) be indexed.

When a Web-based search engine encounters a document containing the Robots meta tag, the search engine does as the meta tag instructs.

There are several values you can specify in the Robots meta tag:

- NOINDEX: Indicates that the document is not to be indexed.
- NOFOLLOW: Indicates that hypertext links in the document are not to be crawled.
- FOLLOWINDEX: Indicates that hypertext links in the document should be crawled.
- ALL: Indicates that the document can be indexed and all links can be crawled.

- NONE: Indicates that the document is not to be indexed and that hypertext links are not to be crawled.

To include the Robots meta tag, use this syntax:

```
<META name="Robots" content="value, optional_value">
```

Using the Robots Comment Tag

You can also use the Robots Comment tag to exclude specific sections of HTML documents from your search results. For example, you might not want such sections as repetitive headers, footers, navigation bars, and server-side includes to be indexed.

TIP: You can also place these tags at the top and bottom of all include files so these sections never get indexed when they are part of a larger document.

To direct QuickFinder where to begin skipping content while indexing:

1. At the point in your HTML document where you want QuickFinder to begin skipping content while indexing, enter the following tag:

   ```
   <!--*Robots NoIndex- >
   ```

2. Just after the content you want skipped, enter the following tag:

   ```
   <!--*Robots Index- >
   ```

3. Save your changes and index (or reindex) the content.

10.3 Modifying Document Descriptions Returned in a Search Results List

QuickFinder returns descriptions of each hit that is listed on the search results page. By default, the following information is returned for each result:

- *Description* field
- *Summary* field
- *Abstract* field
- The first 255 characters of the document (beginning with first heading and skipping links)

The first three fields are taken from the content of meta tags in HTML documents or from document summary fields in other document types such as Word or PDF files. If these tags or fields are not defined, QuickFinder tries to find the first heading and begin selecting words. If it can't find a heading, then it begins at the top of the document and selects the first 255 relevant display bytes as the description.

10.4 Improving Search Results

There are many things you can do to improve the speed, accuracy, and relevance of search results for end users.

- Section 10.4.1, "Improving the Relevance of Search Results," on page 115
- Section 10.4.2, "Using Stop Words Processing to Improve Search Speed," on page 116
- Section 10.4.3, "Using Best Bets to Help Users Find the Right Information," on page 117

- Section 10.4.4, “Using Synonyms to Broaden Search Results,” on page 118
- Section 10.4.5, “Redirecting Searches,” on page 121
- Section 10.4.6, “Creating or Editing the Redirection URLs List,” on page 121

10.4.1 Improving the Relevance of Search Results

QuickFinder uses a sophisticated relevance-ranking algorithm. During a search, QuickFinder considers

- The number of times words appear in a document
- The proximity of words in a multiple word search (the closer the words appear, the more relevant the document will be)
- The order of words in a multiple word search (the exact order of words is more relevant)
- The location of words in a document (specifically words that appear in a meta tag, title, body, header, footer, etc.)
- The formatting of words in a document (such as bold, font type and size, etc.)
- Query weighting in a multiple query scenario
- The number of times words occur within an entire index (for example, the word “the” has low relevance)
- The date of the document.
- The depth of the document.
- The index, path, and URL weights.
- Off-site links.
- Synonyms (synonym-derived hits are slightly less relevant by default).

To illustrate how these criteria work, consider the following examples:

- Words in bold are more relevant than regular words.
- Words contained in the `<Title>` tag are more relevant than words contained within the `<body>` tag.
- Words contained in the Keywords and Description meta tags are more relevant than content words.
- Words contained within the `<A HREF=>` tag used for creating links are less relevant than words outside of this tag.
- A document containing a specified search term multiple times is more relevant than a document that contains the search term only once.
- A word within a 36-point body text is more relevant than within 4-point footer text.
- Documents returned from a query that is weighted at 100% are more relevant than those returned from a 50% weighted query. This is normally used in multi-query searches where each query has a specified weight. For example:

  ```
  query0=netware&weight0=100&query1=groupwise&weight1=100
  ```

In addition to QuickFinder's inherent capabilities, you can further enhance the speed and accuracy of search results by enabling and configuring additional search features.

10.4.2 Using Stop Words Processing to Improve Search Speed

Sometimes users include irrelevant words in their search strings, such as the conjunctions and, to, and of. These are referred to as stop words. The Stop Words feature of QuickFinder removes all occurrences of stop words from the search string before performing a search, except in the following situations.

- Stop words processing has been turned off.
- You prepend the word with a plus (+) sign (for example, +the).
- They are the only words in the query.
- They occur within a phrase.

NOTE: The Boolean operators AND, OR, and NOT are treated as regular words if they are lowercase or in quotes and as Boolean operators if they appear uppercase. For example:

jack and jill (searches for three words)
“jack AND jill” (also searches for three words)
jack AND jill (searches for two words)

A set of common stop words is included in QuickFinder, but you can easily add your own, or remove any of the ones we have included. See “Editing the Stop-Words List” on page 69 for more information.

Before a virtual search server can use stop words processing, it must first be enabled on the QuickFinder server on the Global Settings page of QuickFinder Server Manager.

- “Enabling Stop Words Processing on a QuickFinder Server” on page 116
- “Enabling Stop Words Processing on a Virtual Search Server” on page 116

Enabling Stop Words Processing on a QuickFinder Server

1. On the Global Settings page of QuickFinder Server Manager, click *General* under *Default Settings*.
2. Under *Query Settings*, click *Yes* next to *Enable stop-words processing*.
3. (Optional) Click *Edit List* to modify the default list of stop words, modify the list, then click *Apply*.
4. Click *Apply Settings*.

Enabling Stop Words Processing on a Virtual Search Server

1. On the QuickFinder Server Manager Global Settings page, click *Edit* in the row of the virtual search server that you want to enable stop words processing on.
2. Under *Settings*, click *General*.
3. Click *Yes* next to *Enable stop-words processing*.
4. (Optional) Click *Edit List* to modify the default list of stop words, modify the list, then click *Apply*.

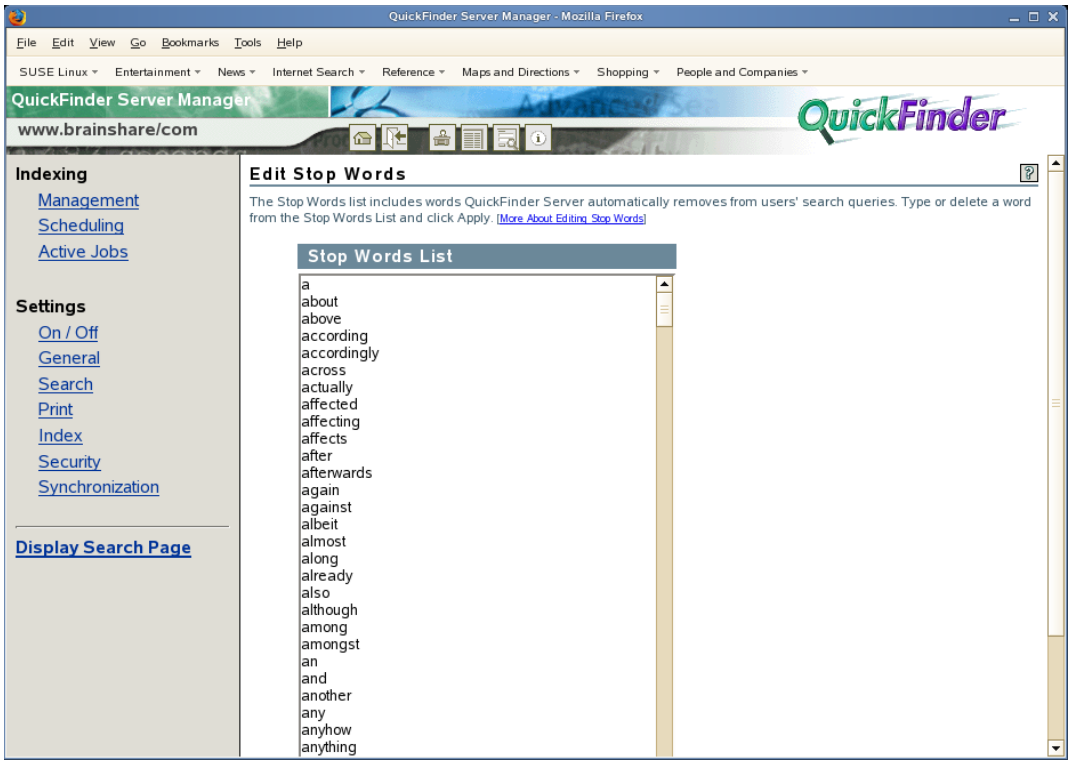

5 Click *Apply Settings*.

10.4.3 Using Best Bets to Help Users Find the Right Information

Best bets is a secondary results list that appears at the top of the search results page. It is generated from a special-purpose index, or an index created for the express purpose of generating a best bets results list.

A best bets list can help users find what they are looking for more quickly by bringing the most popular or most important things to the top of the results page.

A special-purpose index typically contains information about the most popular, most recent, or most important documents.

- “Enabling Best Bets on a QuickFinder Server” on page 117
- “Enabling Best Bets on a Virtual Search Server” on page 118

Enabling Best Bets on a QuickFinder Server

1 On the Global Settings page of QuickFinder Server Manager, click *Search* under *Default Settings*.

2 Click *Yes* next to *Enable Best Bets search results*.

3 (Optional) If you want the best bets results to show automatically on the search results page, click *No* next to *Show Best Bets searches by default*.

4 (Optional) In the *Maximum number of Best Bets results* field, type the maximum number of best bets results to be returned on each search results page.

5 Click *Apply Settings*.

Enabling Best Bets on a Virtual Search Server

1 On the Global Settings page of QuickFinder Server Manager, click *Manage* in the row of the virtual search server that you want to modify.

2 Under *Settings*, click *Search*.

3 Click *Yes* next to *Enable Best Bets search results*.

4 Click *Edit* next to the *Default Best Bets indexes to search* field and select the indexes for use with best bets.

5 If you want the best bets results to show automatically on the search results page, click *Yes* next to *Show Best Bets searches by default*.

6 In the *Maximum number of Best Bets results* field, specify the maximum number of best bets results to be returned on each search results page.

7 (Optional) In the *Minimum Best Bets relevance field*, type the number of low relevance documents you want removed from the best bets display list.

8 Click *Apply Settings*.

10.4.4 Using Synonyms to Broaden Search Results

When enabled, a synonym-enhanced search returns documents in the search results that contain synonyms of the user's original search terms. This can lead users to documents they might not otherwise find because they don't know the exact terminology.

If search results contain the original term, the relevance does not change. If results contain the original term and a synonym, relevance is boosted. If results contain only synonyms, relevance is adjusted by the value entered for the Relevance of Synonym-derived Search Results setting, which is usually slightly lower than the normal relevance.

Common synonyms are included in the QuickFinder Synonym List, but you can add new ones and delete existing ones.

When enabling synonym enhanced searches, keep the following things in mind:

- A synonym-derived search result can be slightly less relevant than the original term.
- A document containing both the original search term and one of its synonyms is more relevant than a document with only the original term.
- Administrators control the Synonym List.
- Synonyms enhancement is not performed within the FileFilter section of a search.

This section contains the following information:

- "Enabling Synonym-Enhanced Searches on a Virtual Search Server" on page 119
- "Editing the Synonyms List" on page 120

Enabling Synonym-Enhanced Searches on a Virtual Search Server

1 On the Global Settings page of QuickFinder Server Manager, select a virtual search server from the Virtual Search Server List, then click *Manage*.

2 Click *General* under *Settings*.

3 Click *Yes* next to *Enable synonym-enhanced search results*.

4 Click *Apply Settings* > *Back to General Settings* to save any changes you have made.

Clicking *Edit List* causes any unsaved changes to be lost.

5 (Optional) To reduce the relevance factor, click *Edit List* to the right of *Enable synonym-enhanced search results*, specify a number in the *Relevance of synonym-derived search results* field, then click *Apply*.

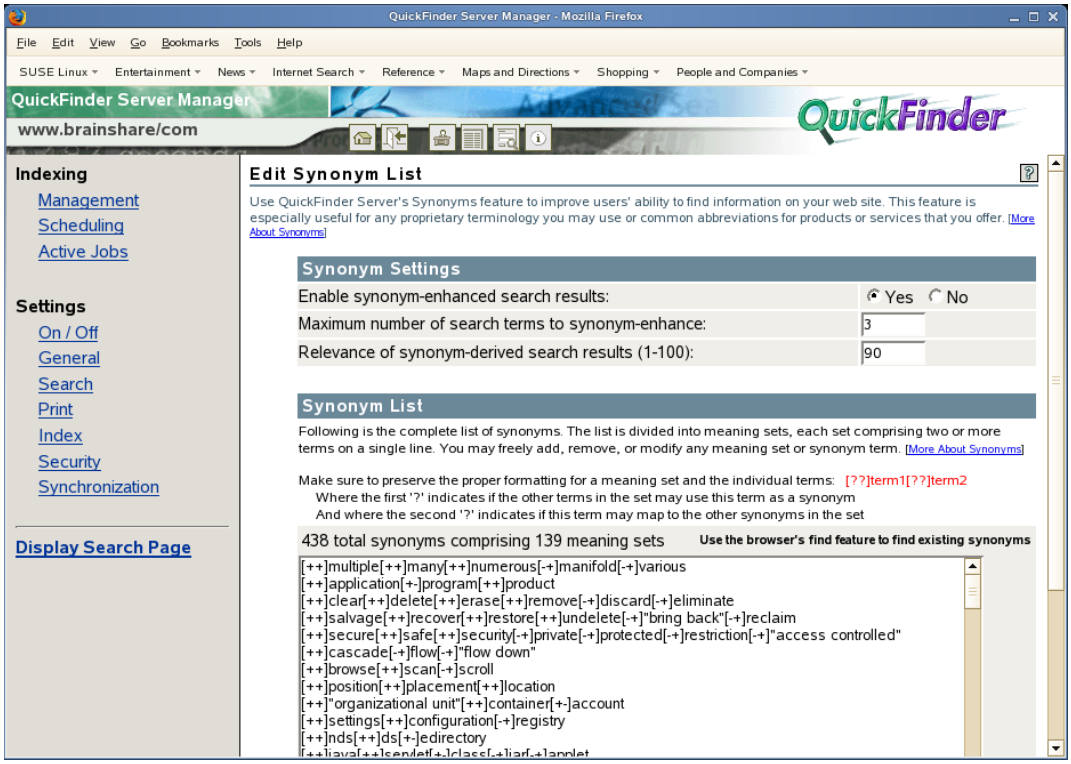

The purpose of this feature is to keep long queries from timing out. Because a single search term could potentially have many synonyms, this feature lets you reduce the relevance factor, which reduces the total number of hits.

A value below 100 generates a second query containing the synonyms where the query weight is slightly lower than the original query weight. A value of 100 simply adds the synonyms to the user's original query and gives all words equal weight.

The most effective relevance numbers are in the 90s. Searches are fastest when you set the relevance to 100. However, you l ose the benefit of synonym-derived search results when you set this number too high.

6 Click *Apply Settings*.

Editing the Synonyms List

1 On the Global Settings page of QuickFinder Server Manager, select a virtual search server from the Virtual Search Server List, then click *Manage*.

2 Click *General* under *Settings*.

3 Click *Yes* next to *Enable synonym-enhanced search results*.

4 Click *Apply Settings* > *Back to General Settings* to save any changes you have made.

Clicking *Edit List* causes any unsaved changes to be lost.

5 Click *Edit List* to the right of *Enable synonym-enhanced search results*.

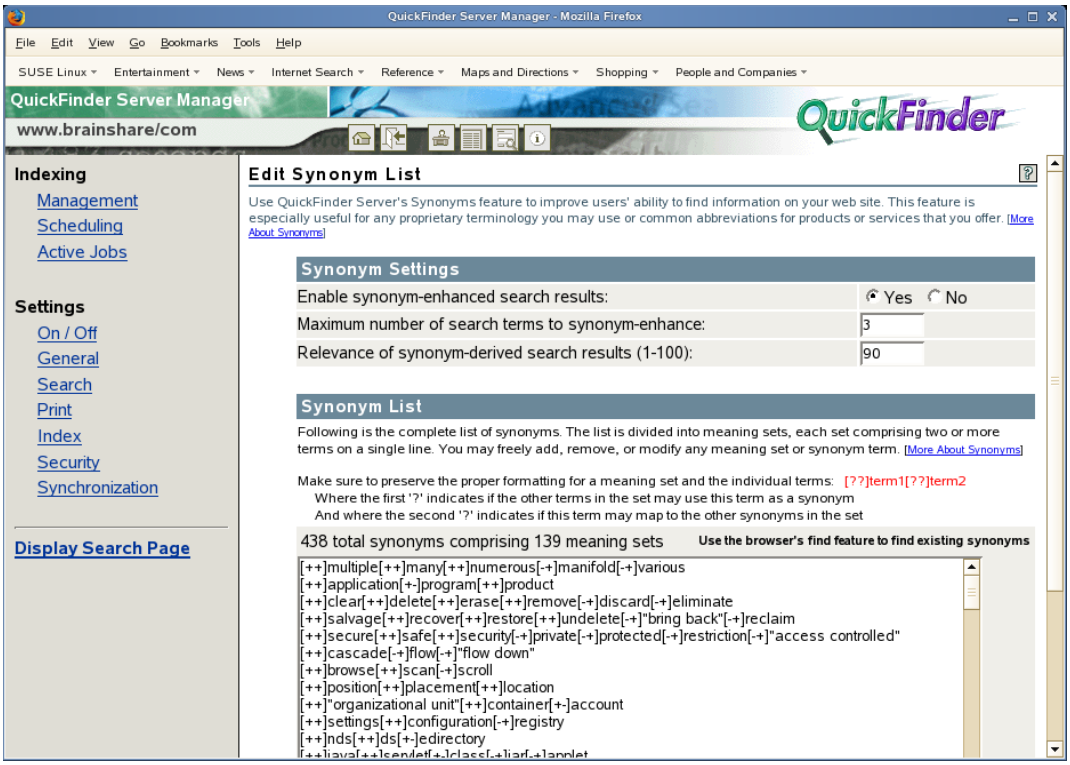

The Synonym List contains a list of default synonyms included with QuickFinder. The synonyms are grouped by sets. A set is one row of synonyms in the list and contains words that are common synonyms of each other word in the same set.

Each word in a set is preceded by a pair of brackets containing a combination of two plus (+) or minus (-) signs, which identifies how the word they precede is to be interpreted.

The first + (or -) before a word specifies whether the other words in the set can use this word as a synonym. For example, if the word adjust was preceded by a +, any other word in the same set could use adjust as a synonym.

The second + (or -) before a word specifies whether that word can use other words in the set as its synonym. For example, if the word adjust was preceded by a -, then adjust could use any other word in the same set as its synonym.

6 When you have finished making changes by adding or deleting words and assigning the proper + or - symbols, click *Apply*.

10.4.5 Redirecting Searches

The redirection feature lets you specify key words that redirect the user's Web browser to a specific URL.

To enable search term redirection on a QuickFinder server:

1. On the Global Settings page of QuickFinder Server Manager, click *Search* under *Default Settings*.
2. Under *Query Settings*, click *Yes* next to *Enable search term redirection*.
3. (Optional) Click *Edit List* to create (or modify) the list of redirection terms and their associated URLs.

   For more information, see "Creating or Editing the Redirection URLs List" on page 121.
4. Click *Apply Settings*.

10.4.6 Creating or Editing the Redirection URLs List

1. On the Global Settings page of QuickFinder Server Manager, click *Search* under *Default Settings*.
2. Click *Apply Settings* > *Back to Search Settings* to save any changes you have made.

   Clicking *Edit List* causes any unsaved changes to be lost.
3. Click *Edit List* to create (or modify) the list of redirection terms and their associated URLs.

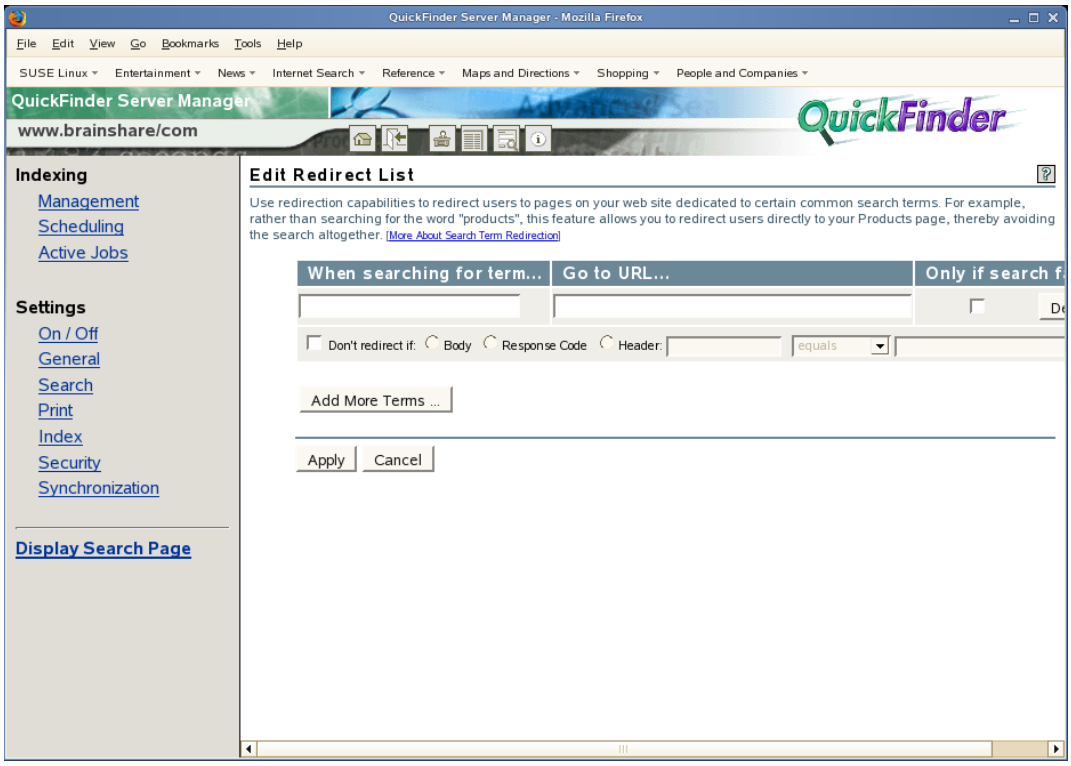

4. In the *When searching for term* field, specify a search term.
5. In the *Go to URL* field, specify the associated URL.

6 (Optional) Select the *Only if search fails* option if you want the search query to be redirected if a search fails to return any results.

7 (Optional) If you want additional control over when a search term is redirected to an alternate URL, check the *Don't redirect if* check box and specify one or more of the following conditions:

- Select *Body* if you want redirection disabled when a specific word or phrase is found in the body of a document returned during a search. From the *Conditions* drop-down list, select a condition (equals, begins with, contains, or ends with) and specify a word or phrase in the text field at the right.
- Select *Response Code* if you want redirection disabled when a specific response code is returned. For example, if a user enters a search term that you have identified for redirection and a 404 Not Found error is returned for any reason, QuickFinder won't redirect the request to the alternate URL. Specify the response code to be used in the text field to the right of the *Conditions* drop-down list.
- Select *Header* if you want redirection disabled when a specific word or phrase is found in the header section of a document returned during a search. From the *Conditions* drop-down list, select a condition (equals, begins with, contains, or ends with) and specify a word or phrase in the text field at the right.

8 Click *Apply*, then click *Apply Settings*.

10.5 Helping Users Avoid Failed Searches

If QuickFinder cannot locate any documents matching a user's search query, a Not Found message is returned. This is sometimes called a failed search. When you enable the search expansion feature, QuickFinder directs failed searches to additional indexes that you create for use with expanded searches.

The best way to implement expanded searches is to create special indexes that are targeted to very specific purposes. For example, you might create an index dedicated to sales reports information, which could help users who are trying to pinpoint sales reports for specific fiscal quarters. (For more information about how to create and manage indexes, see Section 8.5, “Creating Indexes,” on page 83.)

To enable search expansion on a virtual search server:

1 On the Global Settings page of QuickFinder Server Manager, select a virtual search server from the *Virtual Search Server List*, then click *Manage*.

2 Under *Settings*, click *Search*.

3 Under *Query Settings*, click *Yes* next to *Enable search expansion*.

4 Click *Edit* next to the *Default search expansion indexes* and select the indexes to be used.

5 Click *Apply*, then click *Apply Settings*.

10.6 Weighted Queries

A weighted query is used anytime you want to modify the order or relevance of certain hits in a user's normal search results list or when you want to add additional search results users might not have identified in their queries.

QuickFinder allows users to submit more than one query item as part of a single search request.

The following query parameters are combined to identify a single search query item:

&filter#=
&filteroperator#=
&operator#=
&query#=
&weight#=

One use of this feature could be to provide profile-enhanced search requests. For example, the following query returns French product downloads higher up in the search results list but does not eliminate results of any other language downloads:

```
&query0=product+downloads&weight0=100&query1=^language=french&weight1=
90
```

This example directs QuickFinder to perform two completely separate searches. The search results from the two queries are then merged based on the relevance of the individual search results and the weighting of the respective query that produced them.

Another example might be to give the search results from one index more or less relevance than the search results of another index when performing a multiple-index search. For example, the search results from Novell might be more relevant than the search results from Novonyx.

To send multiple query items, these parameters must be grouped using a number (#) at the end of the parameter name so they are interpreted properly. The numbering should begin at 0 or 1 and increment sequentially for each additional query item.

10.7 Ensuring Optimal Search Speed

After a virtual search server has been accessed, all of its configuration files are read into memory. For speed reasons, the virtual search server remains cached in memory until a period of inactivity has elapsed. The virtual search server is then dynamically removed from memory until its next use. Because of this, the first time a virtual search server is accessed is usually the slowest.

However, there are other factors that can affect the performance of your QuickFinder services. As with any software, the amount of available system resources (CPU, RAM, and hard drive) affects QuickFinder Server performance. QuickFinder speed depends on the following factors:

- System processor speed
- Number of processors
- Amount of system memory (RAM)
- Number of hosted virtual search servers
- Number of indexes within each virtual search server
- Number of files included within each index
- Number of indexes included within each query
- Number of queries performed at one time
- Complexity of users' queries
- Number of search results returned with each results page
- Number of concurrent active indexing jobs
- Other functions being performed by your server

Adjusting any of these values can have a significant impact on the performance of your search services.

As a general guideline, use the fastest CPU possible and include as much RAM as possible. Although the duration of each user query is very short, while it is active it consumes an average of 500 KB of memory. Memory consumption varies widely while the indexer is calculating the final search results list, depending on the number of possible search results.

Also, try to schedule the regeneration of your indexes during off-peak hours so that it won't interfere with normal user searches. (See Section 8.9, “Automating Index and Server Maintenance,” on page 100.)

10.8 Making Good Use of Document Fields

A document field is any meta tag or document summary field that helps to identify the document's contents. A document summary field might be a title, heading, or paragraph contained in a title or meta tag within an HTML document.

QuickFinder is designed to take advantage of document fields in order to improve the accuracy, relevance, display information, and speed of search results.

By design, QuickFinder always indexes all document fields in many document types, including HTML, PDF, Word, WordPerfect, XML, etc. Users can then constrain searches to the contents of any document field.

As a QuickFinder administrator, you can also use document fields to further restrict search results to certain products, categories, authors, titles, keywords, or any other content belonging to a document field.

To perform a field-restricted search, use the `^fieldname=`*search_criteria* search operator.

TIP: You might consider sending this information as hidden data using the &filter= query parameter. For example:

```
&filter=^product=netware
```

10.9 Searching XML Documents

XML documents provide a tremendous advantage to narrowing search results because of their hierarchical structure and use of multiple document summary fields.

QuickFinder provides complete hierarchical searching using the fieldname=search_criteria operator. For example, you can find information anywhere in the XML document, within any of the title tags, or limit it to within the title tag that is part of the `<DOCUMENT><SUMMARY>` hierarchy.

The following table shows example uses of the fieldname=*search_criteria* operator when performing a search in XML documents.

Table 10-1 *Example Uses of the fieldname=search_criteria Operator*

Example Values	Result
search_criteria	Finds *search_criteria* anywhere in the document.

Example Values	Result
^Document*=*search_criteria*	Finds *search_criteria* anywhere within any tag that is part of the <DOCUMENT> hierarchy.
^<Document<Summary*=*search_criteria*	Finds *search_criteria* within any tag that is part of the <DOCUMENT> or <SUMMARY> hierarchy.
^<Document<Summary<Title=*search_criteria*	Finds *search_criteria* only within the <DOCUMENT><SUMMARY><TITLE> hierarchy.
^<Document*<Title=*search_criteria*	Finds *search_criteria* within any TITLE tag, located at any level within the <DOCUMENT> hierarchy.
^<*<Title=*search_criteria*	Finds the *search_criteria* within any TITLE tag in the document.

10.10 Using the &filter Query Parameter

The &filter query parameter lets you enhance searches by adding hidden, additional query details when users submit a search query. This is an enhancement over previous versions of QuickFinder, which required that you use JavaScript* to add additional details to search queries.

The &filter query parameter works just like the &query= parameter and can be used together by using the optional number (#) value. For example, if the query parameter was &query0=*search_criteria*, the matching filter parameter would be &filter0=*additional_hidden_search_criteria*. This allows the multiple weighted queries feature to work as designed while allowing administrators to add additional query details to each query.

Unlike the &query parameter, the &filter parameter can be sent multiple times. For example, if users search for software patches, you could include the various products to be searched, which could then improve search time and accuracy:

```
query=software
patchesfilter=^Products=Product257filter=^Products=Product16filter=^Pr
oducts=Product302
```

The resultant URL might appear as follows, but with the HTTP and domain name prefix:

```
&query=software+patches&filter=%2FProducts%3DProduct257&filter=%2FProd
ucts%3DProduct16&filter=%2FProducts%3DProduct302
```

NOTE: All &filter operators are combined using the default &operator=*value*, AND. Also, the default Boolean conjunction joining the various filter operators is an OR search. You can change the default Boolean conjunction by using the &filteroperator=# query parameter. The pound sign (#) here acts just like the one used in the #operator=# query parameter.

Understanding Templates

11

QuickFinder™ Server uses templates to generate search forms and search and print results as well as user feedback such as error or response messages.

A template is an HTML document containing one or more QuickFinder Server variables. Template variables are used to produce dynamic results when a user performs a search on the virtual search server you have defined.

Templates can be shared across virtual search servers or each virtual search server can point to its own set of templates.

This section describes how templates work and discusses the default QuickFinder Server templates that are included with QuickFinder. It includes the following topics:

- Section 11.1, "How Templates Work," on page 127
- Section 11.2, "Exploring the Default Search and Print Templates," on page 129
- Section 11.3, "How Templates Use System Memory," on page 131
- Section 11.4, "Working with Additional Languages," on page 131

11.1 How Templates Work

As defined above, a template is an HTML document containing one or more QuickFinder Server variables. When users search your virtual search server, they use a Web browser to access the search form template. See Figure 11-1, "The QuickFinder Form as It Appears in a Web Browser," on page 128.

The Search form template, `SearchTemplate.html`, is stored by default in *`volume`*`:\`*`searchroot`*`\templates` on NetWare® and in `/var/lib/qfsearch/Templates` on Linux. This path might be different if you chose to install QuickFinder in another directory.

Figure 11-1 *The QuickFinder Form as It Appears in a Web Browser*

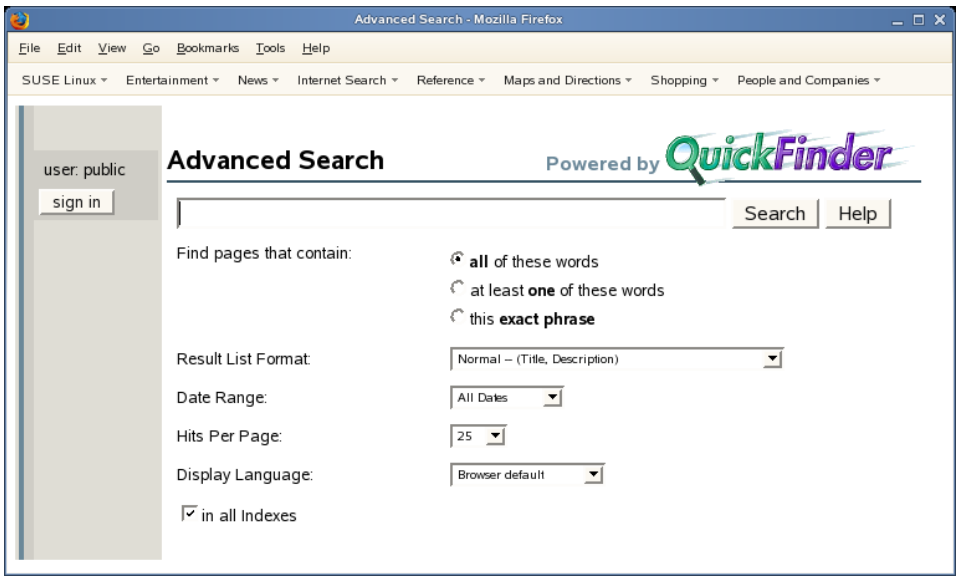

The QuickFinder form is used to capture user input, select available indexes, and then return the results in either a search or print results template, which appears to the user in a dynamically updated HTML document.

Search result templates display hits according to user selections on the search form. For more information about these search result templates, see "Search Page Templates" on page 130.

There are also search and print templates for several different languages. For a discussion about creating templates for international languages, see Chapter 14, "Internationalizing Search Services," on page 177.

After a query is submitted and results are found, QuickFinder populates a results template with all relevant information for each search result. (See Figure 11-2 on page 129.)

Figure 11-2 *A Search Results Page Produced by the ResultListTemplate.html Search Results Template*

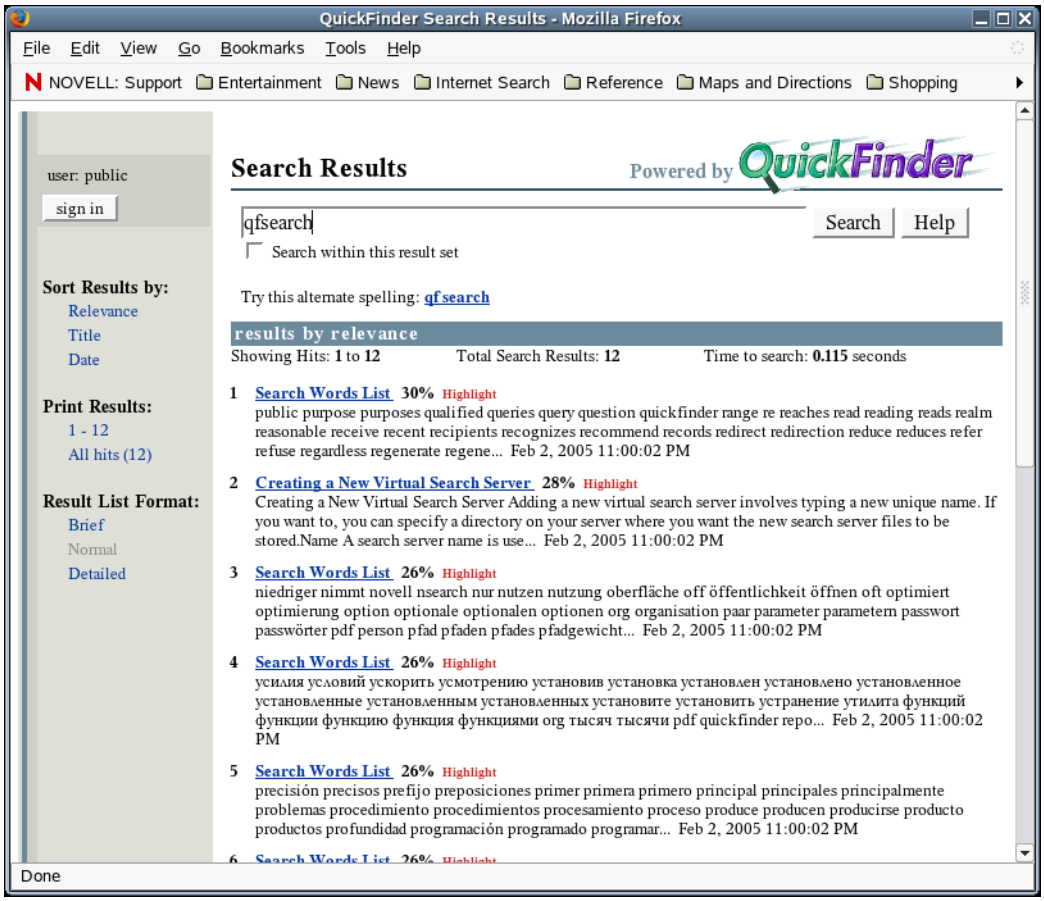

You can also customize the search form to include additional parameters that allow you to offer more options to your users for more accurate searching.

11.2 Exploring the Default Search and Print Templates

QuickFinder Server includes several default templates used to create search forms and to format, display, and print search results for users.

You can use the templates as they are or you can modify them to look and feel how you want them to. You can also create as many additional templates as you need or replace the default templates with your own templates.

QuickFinder includes the following template categories:

- Section 11.2.1, “Search Page Templates,” on page 130
- Section 11.2.2, “Search Result Templates,” on page 130
- Section 11.2.3, “Print Result Templates,” on page 130
- Section 11.2.4, “Error and Response Message Templates,” on page 131

The templates are stored in *volume*:*searchroot*\templates on NetWare and in /var/lib/qfsearch/Templates on Linux.

11.2.1 Search Page Templates

QuickFinder includes two search page templates that are used to generate a search page, as described in the following table:

Table 11-1 *Search Page Templates*

Template Name	Purpose
SearchTemplate.html	Lets users select a variety of options when performing searches. This is the default search template used by QuickFinder Server.
SearchTemplate.Simple	Similar to SearchTemplate.html, except that this template contains no dynamic indexes.

11.2.2 Search Result Templates

QuickFinder includes several ready-made result templates, as described in the following table:

Table 11-2 *Search Results Templates*

Template Name	Purpose
ResultListTemplate.html	Formats and organizes search results and offers additional sorting functions to the user.
ResultListNoHitsTemplate.html	Indicates when no hits are found during a search and offers users a chance to refine their search.
ResultListTerseTemplate.html	Similar to ResultListTemplate but returns less information, such as dates and titles only.
ResultListVerboseTemplate.html	Similar to ResultListTemplate, but returns more information, such as file date, time, and language. Additional sort options are also provided.

11.2.3 Print Result Templates

From the search results page, users have the option of printing all files matching their search or only those files displayed on the current search results page. When one of these options is selected, the print result templates described in the following table are displayed:

Table 11-3 *Print Results Templates*

Template Name	Purpose
`PrintResultTemplate.html`	Combines the full contents of each of the files in the print request into a single document which is then displayed in the user's Web browser. A dynamic table of contents is also created. When the entire content is downloaded to the browser, the browser's print dialog box appears.
`PrintResultNoHitsTemplate.html`	Indicates when no documents are found during a print request.

11.2.4 Error and Response Message Templates

QuickFinder Server includes error and response message templates that are returned when an error occurs or when information is needed from the user.

The default response message template is returned to convey a specific message to the user such as "Print job exceeds recommended size limits," typically returned when a user attempts to print more content than the Web administrator has allowed.

The error and response message templates are found at *volume*`:\`*searchroot*`\TEMPLATES\ErrorMessageTemplate.html` and `ResponseMessageTemplate.html` on NetWare, and `/var/lib/qfsearch/Templates/ErrorMessageTemplate.html` and `ResponseMessageTemplate.html` on Linux.

11.3 How Templates Use System Memory

Templates are cached in memory for quick rendering speed. Each template consumes approximately 10 KB.

Similar to the virtual search server cache, templates remain cached in memory until a period of inactivity has elapsed. The template is then dynamically removed from memory until its next use. The first time a template is accessed, therefore, is normally the slowest.

TIP: Too many templates in the template cache can consume a great deal of memory. Try to share templates across sites to minimize the impact on system memory resources.

11.4 Working with Additional Languages

QuickFinder includes each of the templates described above in each of several languages. Using standard encoding practices, you can internationalize your templates.

Any changes made to the default templates should also be made to the language templates you use. For a more complete discussion about creating a multilingual search solution, see Chapter 14, "Internationalizing Search Services," on page 177.

Working with Template Variables and Search Parameters

12

If you have a basic understanding of HTML, you can create an advanced search solution that your users can use to perform complex searches.

Building an advanced search solution involves the use of QuickFinder™ template variables and search parameters to create or customize QuickFinder templates, and to create or customize one or more search forms.

This section contains the following topics:

- Section 12.1, “Guidelines for Using Variables,” on page 133
- Section 12.2, “Adding a Search Box to Your Corporate Web Site,” on page 134
- Section 12.3, “Global Template Variables,” on page 134
- Section 12.4, “Search Page Variables,” on page 141
- Section 12.5, “Search Result Variables,” on page 141
- Section 12.6, “Print Result Variables,” on page 144
- Section 12.7, “Highlighter Variables,” on page 146
- Section 12.8, “Error Message Variables,” on page 147
- Section 12.9, “Response Message Variables,” on page 147
- Section 12.10, “Query Report Template Variables,” on page 148
- Section 12.11, “Search Parameters,” on page 162

12.1 Guidelines for Using Variables

Use the following guidelines when you are using variables to either customize the default templates or to create new templates:

- **Case Sensitivity:** All variables are case sensitive. Changing case in a variable causes QuickFinder to ignore the variable.
- **Variable Formatting:** All variables must be used exactly as they appear in the tables in this document. Variables always begin with two dollar signs ($$).
- **Success of a Variable:** The inclusion of a variable does not guarantee that information is returned after a search is performed. For example, using the $$Author variable might not return the name of a document's author if that information is not available in the meta tag of the document.
- **Internationalizing Templates:** If you want to internationalize your templates, you must create a template for each language you want to support in your search solution. For more information about languages, see Chapter 14, “Internationalizing Search Services,” on page 177.

For more information about how to implement variables in a search or print template, or how to implement search parameters in an HTML document to create a search form, see Chapter 13, “Customizing Your Templates and Search Forms,” on page 173.

12.1.1 Terminology Changes

In prior versions of QuickFinder Server, the term “search site” was defined as a collection of one or more indexes and related configuration files. To avoid confusion with the term Web site, the term was changed wherever it appeared in the documentation and in the variables and parameters. A search site is now referred to as a “virtual search server.”

New variables and parameters that parallel the term virtual search server have been added. They function identically to the previous variables and parameters, and the old variables and parameters can still be used.

Similarly, the term “collection” has been changed to “index.”

We recommend that you start using the newer variables and parameters in order to avoid confusion.

12.2 Adding a Search Box to Your Corporate Web Site

To add a basic search box to a Web page on your Web site, add the following code at the location of the page where you want the search box to appear:

```
<form name=”SearchForm” method=”get” action=”/qfsearch/SearchServlet”>
   <input type=”text” name=”query”>
   <input type=”submit” value=”Search”>
</form>
```

12.3 Global Template Variables

Global template variables can be used in any of the QuickFinder templates.

For more information about how to implement variables in a template, see Chapter 13, “Customizing Your Templates and Search Forms,” on page 173.

Table 12-1 *Global Template Variables*

Name	Description
$$Authenticated	Indicates whether or not the user is authenticated by returning either a 1 (True) or 0 (False).
$$BeginAuthenticated	Begins a section for a valid, logged-in user. If a user is authenticated, the text between the begin and end tags is processed and appears in the output. If a user is not authenticated, the text is removed from the search results page. To control the appearance of unauthenticated search results, see “$$BeginUnAuthenticated” on page 136. See also “$$EndAuthenticated” on page 136.
$$BeginBestBetIndexesLoop	Begins a repeating section that lists Best Bet indexes used to produce the best bets search results list. This section is repeated for each Best Bet index. See also “$$EndBestBetIndexesLoop” on page 136.

Name	Description
$$BestBetIndex[*number*]	Names of the best bets indexes the user specified in the search query. If the optional number is not provided, the value of $$BestBetIndexesCurrent is used. See also “$$BeginBestBetIndexesLoop” on page 134.
$$BestBetIndexesCount	Number of best bets indexes associated with the user query. See also “$$BeginBestBetIndexesLoop” on page 134.
$$BestBetIndexesCurrent	Number of the current best bets index. See also “$$BeginBestBetIndexesLoop” on page 134.
$$BestBetNum	Total number of best bets results to display at a time.
$$BestBetShow	Returns True or False, depending on QuickFinder Server Manager’s Search Settings page. To override the default, use the &bbshow= true | false query parameter.
$$BeginExpandIndexesLoop	Begins a repeating section that shows the search expansion indexes used if the primary indexes do not find any search results. This section is repeated for each search expansion index.
$$BeginFiltersLoop	Begins a repetitive section that is processed for each filter parameter associated with the current query item. Multiple query items can be sent as part of a single query. See also “$$EndFiltersLoop” on page 136.
$$BeginIndexesLoop	Begins a repetitive section that is processed for each index the user specified in the search query. See also “$$EndIndexesLoop” on page 136 and “$$QueryIndex[number]” on page 139.
$$BeginLogin	Begins a section of code used to define a login form that the user can use to log in. This section is processed only if the &login=true query parameter is included with the QuickFinder request.
$$BeginLoop	End of the header section. Beginning of the repeating body section. This section is repeatedly parsed until there are no further result items to process. See also “$$EndLoop” on page 137.
$$BeginNoLogin	Informs users that the current search form does not require a login. Typically, this section is followed by a login button that users can click to receive a login form, which comes from the $$BeginLogin variable. If no login query parameter is passed to the server, then everything between this variable and $$EndNoLogin is returned to the user. See also “$$EndNoLogin” on page 137.
$$BeginQueryLoop	Begins a repetitive section that is processed for each query item associated with the current search query. See “$$NumQueryItems” on page 138 for more information. See also “$$EndQueryLoop” on page 137.

Name	Description
$$BeginReturnFieldsLoop	Begins a repetitive section that is reprocessed for each return field the user specified in the search query. See also “$$QueryReturnField[number]” on page 139.
$$BeginSortKeysLoop	Begins a repetitive section that is processed for each sort key the user specified in the search query. See also “$$EndSortKeysLoop” on page 137 and “$$SortKeysCurrent” on page 140.
$$BeginUnAuthenticated	Begins a section for an unrecognized or logged-out user. Used in conjunction with $$EndUnAuthenticated. If a user is not recognized, the text between these two tags is processed and appears in the output. If the user is recognized as a valid, logged-in user, this text does not appear in the output. To control the appearance of authenticated search results, see “$$BeginAuthenticated” on page 134.
$$Counter[*variable_number*, *increment_number*]	Inserts the value of the specified *variable_number* counter into the search result page. All counters initialize to zero. The optional second parameter specifies the amount to increment or decrement the current value. A maximum of 10 counters is supported. For example: $$Counter[1] = insert value of counter #1 $$Counter[1,1] = increment counter #1 by 1 and display the new value $$Counter[5,-3] = decrement counter #5 by 3 and display the new value
$$DefaultQueryEncoding	Default encoding of the user query if not specified by using the &encoding= query parameter.
$$EndAuthenticated	Ends a section for a valid, logged-in user. See also “$$BeginAuthenticated” on page 134.
$$EndBestBetIndexesLoop	Closes the $$BeginBestBetIndexesLoop, which lists the best bets indexes used to produce the best bets results list. See also “$$BeginBestBetIndexesLoop” on page 134.
$$EndExpandIndexesLoop	Ends the $$BeginExpandedIndexesLoop section, which is used to get a list of search expansion indexes used by the current query. See also “$$BeginExpandIndexesLoop” on page 135.
$$EndFiltersLoop	Ends a repetitive section that is processed for each filter parameter associated with the current query item. See also “$$BeginFiltersLoop” on page 135.
$$EndIndexesLoop	Ends a repetitive section that is processed for each index the user specified in the search query. See also “$$BeginIndexesLoop” on page 135 and “$$QueryIndex[number]” on page 139.

Name	Description
$$EndLogin	Closes the $$BeginLogin section, which is used to produce a login form. See also “$$BeginLogin” on page 135.
$$EndLoop	End of the repeating body section. Beginning of the footer section.
$$EndNoLogin	Closes the $$BeginNoLogin section, which is used to specify that the user is not currently logged in. This section appears when no login parameter is passed to the server. See also “$$BeginNoLogin” on page 135.
$$EndQueryLoop	Ends a repetitive section that is processed for each query item associated with the current search query. See “$$NumQueryItems” on page 138 for more information. See also “$$BeginQueryLoop” on page 135.
$$EndReturnFieldsLoop	End of a repetitive section that is reprocessed for each return field the user specified in the search query. See also “$$QueryReturnField[number]” on page 139.
$$EndSortKeysLoop	Ends a repetitive section that is processed for each sort key the user specified in the search query. See also “$$BeginSortKeysLoop” on page 136 and “$$SortKeysCurrent” on page 140.
$$EndUnAuthenticated	Ends a section for an unrecognized or logged-out user. For more information, see “$$BeginUnAuthenticated” on page 136.
$$ExpandIndexesCount	Number of search expansion indexes associated with the user query. See also “$$BeginExpandIndexesLoop” on page 135.
$$ExpandIndexesCurrent	Number of the current search expansion index. See also “$$BeginExpandIndexesLoop” on page 135.
$$FilterCount	Number of filters associated with the current query item. Multiple query items can be associated with a single query.
$$FilterCurrent	Number of the current filter associated with the current query item.
$$FilterOperator	Boolean operator used to join the filters associated with the current query item. The full set of filters is always joined to the current query item using the Boolean AND.
$$FilterValue[*number*]	Pulls the value of the specified filter associated with the current query item. If the optional # parameter is not provided, the current filter loop value ($$FilterCurrent) is used.
$$HostURL	Generates the current URL up to, and including, the optional port number.

Name	Description
$$IncludeFile[*template_name*]	Automatically pulls in the designated template at the location of this variable. The included template can contain other template variables, which are processed as though they were a part of the original template. The template name parameter can be either a full FILE:// URL based on the file system of the server or a relative path based on the location of the parent template. The template name parameter can be located within quotation marks. See the `ResultListTemplate.html` file for an example use of this variable.
$$IndexesCount	Number of indexes associated with the user query. See also "$$BeginIndexesLoop" on page 135.
$$IndexesCurrent	Number of the current index. See also "$$BeginIndexesLoop" on page 135.
$$IsInvalidLogin[*text*]	If a user's login is invalid, the text contained between the brackets is returned to the user. A typical message might be `Invalid login. Please try again`.
$$NumQueryItems	Number of query items contained within the current query. Although most queries use only one query item, it is possible to construct a query with multiple search criteria, each weighted with a value between 1 and 100. The resulting search contains hits from each of the queries, but the search results are organized with the most relevant hits first (from any of the individual queries).
$$Query[*number*]	Query specified by the client in the search field. The optional number identifies the corresponding query item. The value of $$QueryCurrent is used if the optional number is not provided. See also "$$NumQueryItems" on page 138 for more information.
$$QueryCount	Number of query items associated with the search query. See also "$$NumQueryItems" on page 138.
$$QueryCountry	The country requested by the client. This must be an uppercase, two-character value as specified in ISO 3166-1.
$$QueryCurrent	Number of the current query item. See "$$NumQueryItems" on page 138 for more information. See also "$$BeginQueryLoop" on page 135.
$$QueryDate	Begin date requested by the client. Only those documents dated on or after the specified date are returned in the search results. See the query parameter "date" on page 163 for more information.
$$QueryEncoding	Actual encoding used to interpret the query. This can be the same as the $$DefaultQueryEncoding, the value of the &encoding= query parameter, or UTF-8.
$$QueryExpandIndex[*number*]	Names of the search expansion indexes the user specified in the search query. If the optional number is not provided, the current value of $$ExpandIndexesCurrent is used. See also "$$BeginExpandIndexesLoop" on page 135.

Name	Description
$$QueryFileFilter	Returns the filename filter associated with the user query.
$$QueryIndex[*number*]	Names of the indexes the user specified in the search query. If the optional number is not provided, the current value of $$IndexesCurrent is used. See also "$$BeginIndexesLoop" on page 135.
$$QueryLanguage	Language requested by the client. This must be a lowercase, two-character value as specified in ISO 639.
$$QueryNumHits	Number of search results requested by the client.
$$QueryOperator	The type of the current search: 0 = Boolean AND search 1 = Boolean OR search 2 = phrase search
$$QueryReturnField[*number*]	Name of the return fields the user specified in the search query. If the optional number is not provided, the current value of $$ReturnFieldsCurrent is used. See also "$$BeginReturnFieldsLoop" on page 136 and "$$EndReturnFieldsLoop" on page 137.
$$QueryServerName[*text*]	Identifies the name of the Virtual Search Server provided with the &server= query parameter. The optional text parameter can be provided in the following formats: $$QueryServerName = NameOfServer $$QueryServerName[*text*] = *text* NameOfServer $$QueryServerName[%*text*] = *text* URLEncodedNameOfServer $$QueryServerName[*text* $$QueryServerName *text*] = *text* NameOfServer *text* $$QueryServerName[%*text* $$QueryServerName *text*] = *text* URLEncodedNameOfServer *text* See also "$$ServerName" on page 140.
$$QuerySynonymFlag	Set to either True or False based on the current setting under QuickFinder Server Manager's General Settings page. See "Using Synonyms to Broaden Search Results" on page 118 for information about enabling synonym-enhanced searching.
$$QueryTemplate	Template name requested by the client. See also "$$TemplateName" on page 140.
$$QueryTemplateTheme	Template theme requested by the client. This is not necessarily the theme of the search result because the specified theme might not exist. See also "$$TemplateTheme" on page 141.
$$QueryVersion	Version number of the current query format.

Name	Description
$$QueryWeight[*number*]	Weighting of the current query item, 1 to 100. If *number* is not specified, the value of $$QueryCurrent is used. See "$$NumQueryItems" on page 138 for more information.
$$ResultEncoding	Encoding used to return the current search results page. This is either the value of the valid &retencoding=query parameter or the default specified by the search administrator in the QuickFinder Server Manager.
$$ReturnField[*number*]	Name of the return fields the user specified in the search query. If the optional number is not provided, the value of $$ReturnFieldsCurrent is used. See also "$$BeginReturnFieldsLoop" on page 136 and "$$EndReturnFieldsLoop" on page 137.
$$ReturnFieldsCount	Number of return fields specified in the search query. See also "$$BeginReturnFieldsLoop" on page 136.
$$ReturnFieldsCurrent	Number of the current iteration of the $$BeginReturnFieldsLoop.
$$SearchFor[*number*]	Query entered by the client into the search field. If the optional number is not provided, the value of $$QueryCurrent is used. See "$$Query[number]" on page 138 for more information.
$$ServerName	Name of the virtual search server that produced the current output. See also "$$QueryServerName[text]" on page 139.
$$ServerLocation	Path on the network server to the virtual search server configuration files and indexes.
$$SortField[*number*]	The name of the field to sort on. If the optional number is not provided, the value of $$SortKeysCurrent is used. See "$$SortByURL[sortfield.sortorder ...]" on page 144 and the query parameter "sortorderquery_number" on page 169 for more information.
$$SortKeysCount	Number of sort keys associated with the current query.
$$SortKeysCurrent	Current sort keys number. See "$$BeginSortKeysLoop" on page 136 for more information.
$$SortOrder[*number*]	Method of sorting the field (ascending, descending, and default). If the optional number is not provided, the value of $$SortKeysCurrent is used. See "$$SortOrder[number]" on page 140 for more Information.
$$TemplateExt	Filename extension of the template currently displayed in the browser. When an error occurs, QuickFinder automatically attempts to find an error message template that matches the filename extension of the requested search template.
$$TemplateLocale	Locale of the template, such as zh_TW. The locale information is taken from the template filename.
$$TemplateName	Filename of the template currently displayed in the browser. See also "$$QueryTemplate" on page 139.

Name	Description
$$TemplateTheme	Theme (or theme directory) that the current template belongs to. See also “$$QueryTemplateTheme” on page 139.
$$UserID	The User ID corresponding to the user who is currently logged in.

12.4 Search Page Variables

The following table lists all available search page variables that can be used in addition to the global template variables to extend the functionality of the default search templates (`SearchTemplate.html` or `SearchTemplate.Simple`) or to create new templates.

Table 12-2 *Search Page Variables*

Name	Description
$$BeginServerIndexesLoop	Begins a repeating section in the search template where information for each of the defined indexes is written.
$$EndServerIndexesLoop	Ends a repeating section in the search template where information for each of the defined indexes is written.
$$ServerIndexDescription	Description of the virtual search server defined in the QuickFinder Server Manager.
$$ServerIndexName	Name of the virtual search server defined in the QuickFinder Server Manager.

12.5 Search Result Variables

The following table lists all available search result variables that can be used in addition to the global template variables to extend the functionality of the default search result templates or to create new templates.

For more information about how to customize templates, see Chapter 11, “Understanding Templates,” on page 127 and Chapter 13, “Customizing Your Templates and Search Forms,” on page 173.

Table 12-3 *Search Result Variables*

Variable Name	Description
$$Author	Name of the original author of a document returned in a hit.
$$BeginAuthorized	Begins a section for a search result that the user has rights to see. Used in conjunction with $$EndAuthorized. If a search result is authorized, this section of text and template variables is processed. If it is unauthorized, this section is removed from the output. See also “$$BeginUnAuthorized” on page 142.

Variable Name	Description
$$BeginBestBet	Begins a repetitive section that is processed for each Best Bet search result item. Result detail variables, such as $$Title, $$Description, $$Index, and $$Relevance, can be used within this section. See also “$$EndBestBet” on page 142.
$$BeginUnAuthorized	Begins a section for a search result that the user does not have rights to see. Used in conjunction with $$EndUnAuthorized. If a search result is not authorized, this section of text and template variables is processed. If the search result is authorized, this section is removed from the output. See also “$$BeginAuthorized” on page 141.
$$DateTime[*date_format*]	The date and time of a hit. This is automatically written in the Java “medium” date format using the client's locale (all calendars, translations, and date and time formats are observed). $$DateTime[ ] can use an optional date and time format provided within the brackets [ ]. The text should conform to the Java *DateFormat* syntax.
$$Description	Abstract, description, or first 255 display bytes of the result item.
$$Encoding	Character set encoding of the current search result item.
$$EndAuthorized	Ends a section for a search result that the user has rights to see. For more information, see “$$BeginAuthorized” on page 141.
$$EndBestBet	Ends the $$BeginBestBet section, which is used to display best bets results. See “$$BeginBestBet” on page 142.
$$EndUnAuthorized	Ends a section for a search result item that the user does not have rights to see. For more information, see “$$BeginUnAuthorized” on page 142.
$$FileFormat	Indicates a specific document type. For example, HTML, Word 97, or PowerPoint. **TIP:** You can control the text output for each file format by using the optional `FileFormats.properties` file. For more information, see the Novell® QuickFinder Server File Formats (http://www.novell.com/products/quickfinder/fileformats.html) page.
$$FirstHit	Hit number of the first item in the current result page. Displayed using the client's locale.
$$HighlightURL[*text*]	If the source document of a search result is capable of being highlighted (HTML or text files), then the text contained within the brackets is displayed. This variable can be used as a marker to identify which of the search results can display highlighted search terms. Typically, $$HighlightURL is also included within the brackets and thereby outputs the URL to be used to highlight the current document.
$$Index	Name of the index in which a particular search result item was found.

Variable Name	Description
$$IsExpandedSearch[*text*]	Displays a message indicating when a set of results was produced from expanded search indexes. The message is placed between the brackets [ ] and might read something like `The query you were looking for was not found in the selected indexes. The displayed search results were produced from additional search indexes.`
$$Language	Language of the result item. $$Language is displayed in the language of the client's locale.
$$LastHit	Hit number of the last item in the current result page. Displayed using the client's locale.
$$MoreHits[*page#*, *text*]	Conditional text section to be included only if there are additional hits in the search results that can be retrieved. If the first section of the conditional text contains a number followed by a comma (for example: $$MoreHits[3, text to be included]), the server first determines if the designated search results page exists. If *page#* is missing, 1 (the next page) is assumed. If the designated page is available, the remaining text after the comma and up to the closing bracket is written to the result page. The initial number is relative to the current page. That is, -1 references the page immediately before the current page and 1 references the page immediately after. Zero (0) refers to either the previous page or the next page.
$$MoreHitsURL[*page_number*]	URL needed to display another page of search results. The optional parameter identifies the desired search result page number. If it is not provided, 1 is assumed. The page number is relative to the current page. That is, -1 refers to the page immediately before the current page and 1 references the page immediately after. Zero (0) refers to the current page. The URL is inserted only if the designated page exists.
$$Number	Hit number of the current result item. Possible numbers begin with 1 and end with $$TotalHits. Displayed using the client's locale.
$$PageNum[*page#*]	Page number corresponding to the current search results page. This variable is used when more search results are available than can fit on one page (such as when $$TotalHits is greater than the number of hits specified per page in the &numhits=# query parameter). The optional *page#* identifies the relative page from the current result page. That is, -1 refers to the page immediately before the current page and 1 references the page immediately after. Zero (0) refers to the current page. The page number is inserted only if the designated page exists.
$$PrintURL[first_hit_number, number_of_hits]	URL used to print the hits listed on the current search result page. The optional parameters can be specified to define the beginning search result number and the number of search results to include in the print job. The *number_of_hits* parameter can use the $$TotalHits template variable.

Variable Name	Description
$$Relevance	How closely the result matches the user's query, indicated by percentages (1% to 100%).
$$SearchFor[*number*]	Query entered by the client into the search field. See "$$Query[number]" on page 138 for more information.
$$SearchTime	Amount of time used to process the current search request. $$SearchTime is displayed using the client's locale.
$$Size	Size of the data pointed to by the result item's URL. Displayed using the client's locale.
$$SortByDate	Returns True or False depending on the &sortbydate=*true_false* query parameter.
$$SortByURL[*sortfield.sortorder* ...]	The URL used to show the current result page sorted by one or more search result fields. Sort field names include title, author, changedate, filelength, language, summary, relevance, url, index, format, and document_number. Optional sort orders include ascending and descending. Sort field and sort order names are separated by a period. Multiple sort fields are separated by a space.
$$Title	If a title is not available in documents being searched, $$URL is used instead; if the URL is unavailable, < *title unavailable* > is used.
$$TotalHits	Total number of hits that match the search query. This is not the same as the number of hits displayed in any particular result page. Displayed using the client's locale.
$$URL	URL of the result item.

12.6 Print Result Variables

The following table lists all available print result variables that can be used in addition to the global template variables to extend the functionality of the default print result templates or to create new templates.

For more information about how to implement variables in a template (HTML) page, see Chapter 13, "Customizing Your Templates and Search Forms," on page 173.

Table 12-4 *Print Result Variables*

Name	Description
$$BeginAuthorized	Begins a section for a print result that the user has rights to see. Used in conjunction with $$EndAuthorized. If a search result is authorized, this section of text and template variables is processed. If the result is unauthorized, this section is removed from the output. See also "$$BeginUnAuthorized" on page 142.

Name	Description

$$BeginMimeType[*text/plain*]	Identifies special formatting to apply to individual print result items based on their Mime type. For example, text formats should generally be output using HTML's `<pre>` tag, such as in the following:
	`$$BeginMimeType[text/plain]`
	`<pre>`
	`$$EndMimeType`
	See also "$$EndMimeType" on page 145.

$$BeginTOCList[*text*]	Beginning of the table of contents repeating section. This section is repeatedly parsed until there are no further TOC result items to process. This is a conditional text section. The items within the brackets ([ ]) are processed only if the current item represents a change in the depth of the hierarchy. If $$Product appears within the conditional text, it is replaced only if the current item also represents a new product.
$$BeginUnAuthorized	Begins a section for a print result that the user does not have rights to see. Used in conjunction with $$EndUnAuthorized. If a print result is not authorized, this section of text and template variables is processed. If the print result is authorized, this section is removed from the output. See also "$$BeginAuthorized" on page 141.
$$Bookmark	HTML anchor name of the current result item. This can be used to jump from a TOC entry to the corresponding section within the print job. All bookmark entries begin with "novell_print_toc_" and are followed by the number of the current result item, as in novell_print_toc_1.
$$Description	Abstract, description, or first 255 display bytes of the result item.
$$EndAuthorized	Ends a section for a print result that the user has rights to see. For more information, see "$$BeginAuthorized" on page 141.
$$EndMimeType	Ends a $$BeginMimeType section. For more information, see "$$BeginMimeType[text/plain]" on page 145.
$$EndTOCList[*text*]	End of the table of contents section. This is a conditional text section. The items within the brackets ([ ]) are written out each time a result item occurs that decreases the depth of the hierarchy. If the depth of the current item is several levels fewer than the previous item, the text within the conditional text block is written out that many times.
$$EndUnAuthorized	Ends a section for a print result item that the user does not have rights to see. For more information, see "$$BeginUnAuthorized" on page 142.
$$Number	Hit number of the current result item. Possible numbers begin with 1 and end with $$TotalHits. Is displayed using the client's locale.
$$NumIndents	Number of indentations required for the current Table of Contents entry.
$$Product	Name of the product associated with the current item in the table of contents. This displays only if this is the first result item within that product. See also "$$BeginTOCList[text]" on page 145.

Name	Description
$$Title	Title of the result item. For empty titles, <title unavailable> is displayed. Displayed using the client's locale.
$$TotalHits	Total number of hits that match the search query. This is not the same as the number of hits displayed in any particular result page. Displayed using the client's locale.
$$URL	URL of the result item.
$$URLContent	The entire contents of the URL are placed into the template at this location. The URL contents are not parsed to validate their data type, formatting, or functionality. Only text/plain and text/html files are printed. All other files are inserted into the print job as an error message.

12.7 Highlighter Variables

The following table lists all available highlighter variables that can be used in addition to the global template variables to extend the functionality of the default highlighter template or to create a new template.

For more information about how to implement variables in a template, see Chapter 13, "Customizing Your Templates and Search Forms," on page 173.

Table 12-5 *Highlighter Variables*

Name	Description
$$BeginMetaDataLoop	Begins a loop used for formatting the meta tags associated with the currently highlighted file. For example, if three meta tags were found, this loop could produce three rows in a table, placing the data from one meta tag in each of the three rows. See also "$$MetaName" on page 146 and "$$MetaValue" on page 146.
$$BeginWordsLoop	Begins a loop that outputs the search terms used to find the current document. Each search term is produced using the colors specified in QuickFinder Server Manager. See also "$$Word" on page 146.
$$EndMetaDataLoop	Ends the $$BeginMetaDataLoop section.
$$EndWordsLoop	Ends the $$BeginWordsLoop section.
$$MetaName	Name of the metadata tag, such as Title or Description.
$$MetaValue	Lists the values of each of the metadata tags found in a file.
$$Word	Lists all of the words to be highlighted. The words are produced using the colors specified in QuickFinder Server Manager. See the `HighlighterTemplate.html` file for an example.

12.8 Error Message Variables

The following table lists all available error message variables that can be used in addition to the global template variables to enhance the organization of the default error message template, or to create new templates.

For more information about how to implement variables in a template, see Chapter 13, "Customizing Your Templates and Search Forms," on page 173.

Table 12-6 *Error Message Variables*

Name	Description
$$ErrorDescription	Longer version of the message. This might include additional error details or problem resolution information.
$$ErrorMessage	Text version of the error. Generally quite short.
$$ErrorNumber	Numeric version of the error.

12.9 Response Message Variables

The following table lists all available response message variables that can be used in addition to the global template variables to enhance the organization of the default response message templates or to create new templates.

For more information about how to implement variables in a template, see Chapter 13, "Customizing Your Templates and Search Forms," on page 173.

NOTE: The repeating variables $$BeginLoop and $$EndLoop should not be used in a response message and are ignored if used.

Table 12-7 *Response Message Variables*

Name	Description
$$Cancel[*text*]	If the *Cancel* button is specified by Server logic, this parses and inserts the conditional text into the response page. Currently used by PrintServlet when a print job exceeds the print job size warning limit.
$$Continue[*text*]	If the *Continue* button is specified by Server logic, this parses and inserts the conditional text into the response page. Currently used by PrintServlet when a print job exceeds the print job size warning limit.
$$Help[*text*]	If the *Next* button is specified by Server logic, this parses and inserts the conditional text into the response page.
$$Ignore[*text*]	If the *Ignore* button is specified by Server logic, this parses and inserts the conditional text into the response page.
$$Next[*text*]	Compare to $$Prev.
$$No[*text*]	If the *No* button is specified by Server logic, this parses and inserts the conditional text into the response page.

Name	Description
$$OK[*text*]	If the *OK* button is specified by Server logic, this parses and inserts the conditional text into the response page. Currently used by PrintServlet when a print job exceeds the maximum print job size.
$$Prev[*text*]	If the *Previous* button is specified by Server logic, this parses and inserts the conditional text into the response page.
$$ResponseNumber	Numeric version of the response required of the user.
$$ResponseMessage	Text version of the response required of the user. Generally quite terse. Can often be used as a title.
$$ResponseDescription	Longer version of the message. This might include additional details or “see also” information.
$$Retry[*text*]	If the *Retry* button is specified by server logic, this parses and inserts the conditional text into the response page.
$$URL	URL to use when the *Parent* button is clicked. This must appear within the brackets of a button's conditional text section. The URL logic is generated by the server.
$$Yes[*text*]	If the *Yes* button is specified by server logic, this parses and inserts the conditional text into the response page.

12.10 Query Report Template Variables

The following tables list all available query report template variables that can be used in addition to the global template variables to extend the functionality of the default query report templates or to create new templates.

All numbers, dates, times, and day names throughout the report use either the locale of the template or server to properly display this information according to the rules defined by each language and country.

- Section 12.10.1, “General,” on page 149
- Section 12.10.2, “Report Summary,” on page 149
- Section 12.10.3, “Templates Used,” on page 149
- Section 12.10.4, “Load Details,” on page 150
- Section 12.10.5, “Search Details,” on page 153
- Section 12.10.6, “Search Results Pages,” on page 154
- Section 12.10.7, “Top X Queries,” on page 154
- Section 12.10.8, “Top Queries,” on page 155
- Section 12.10.9, “Error Searches,” on page 156
- Section 12.10.10, “Not Found Searches,” on page 157
- Section 12.10.11, “Redirected Searches,” on page 158
- Section 12.10.12, “Login Requests,” on page 158
- Section 12.10.13, “Print Details,” on page 159
- Section 12.10.14, “Print Errors,” on page 159

- Section 12.10.15, “Highlight Details,” on page 160
- Section 12.10.16, “Export Details,” on page 160

12.10.1 General

Table 12-8 *General Query Report Template Variables*

Name	Description
$$ReportName	Displays the filename of the current report without the terminating filename extension. A query report’s filename usually corresponds to the date range represented within the report.
$$SiteName	The name of the virtual search server to which the query report belongs.

12.10.2 Report Summary

Table 12-9 *Report Summary Query Report Template Variables*

Name	Description
$$AccessTotal	Displays the total number of requests received during the current report period.
$$UniqueUsers	Displays the number of unique users that were successfully authenticated through QuickFinder.

12.10.3 Templates Used

Table 12-10 *Templates Used Query Report Template Variables*

Name	Description
$$BeginRptTemplatesLoop	Begins a repeating section that displays the list of templates used to return search results information to the end user. This includes all template types.
$$EndRptTemplatesLoop	Ends a repeating section that displays the list of templates used to return search results information to the end user. This includes all template types.
$$RptTemplate	The name of the current QuickFinder template.
$$RptTemplateTotal	The number of times the current template was used.
$$RptTemplateCurrent	The number of the current loop. 1-based.
$$RptTemplateCount	The total number of unique templates requested by all users during the logging period.

12.10.4 Load Details

Table 12-11 *Load Details Query Report Template Variables*

Name	Description
$$DayOfWeek[#]	Displays the name of the current day of the week. # - optional day of week number (1 - 7)
By Week	
$$BeginChartWeek[#]	Begins a repeating section that displays a load chart for the specified week. # - optional week value (1 - 6). If the optional week value is provided, but the week has no load information, this entire section is skipped. If the optional week value is not provided and the “current” week has no load information, the section is skipped for the current week, but repeated for any remaining weeks that follow.
$$EndChartWeek	Ends a repeating section that displays a load chart for the specified week.
$$BeginChartWeekDay[#2] or [#1,#2]	Begins a repeating section that displays information for each day of the week. If the specified day and week do not have load information available, the current section is not skipped. Use $$BeginSkipIfEmptyDay to skip the current section. #1 - optional week value (1 - 6). #2 - optional day value (1 - 7)
$$EndChartWeekDay	Ends a repeating section that displays information for each day of the week.
$$BeginSkipIfEmptyDay[#2] or [#1,#2]	Skips the current section (up to $$EndSkipIfEmptyDay) if the specified day of the week does not have any load information available. #1 - optional week value (1 - 6); defaults to the current value of the $$BeginChartWeek loop. #2 - optional day value (1 - 7); defaults to the current value of $$BeginChartWeekDay loop.
$$EndSkipIfEmptyDay	Ends the current section if the specified day of the week does not have any load information available.
$$WeekDayTotal[#2] or [#1,#2]	Displays the total number of requests handled on the specified day and week. #1 - optional week value (1 - 6); defaults to the current value of the $$BeginChartWeek loop. #2 - optional day value (1 - 7); defaults to the current value of $$BeginChartWeekDay loop.

Name	Description
$$WeekDayPeek[#2] or [#1,#2]	Displays the peak number of requests handled in a single hour on the specified day and week. #1 - optional week value (1 - 6); defaults to the current value of the $$BeginChartWeek loop. #2 - optional day value (1 - 7); defaults to the current value of $$BeginChartWeekDay loop.
$$WeekDayPeekTime[#2] or [#1,#2]	Displays the time of day when the peak number of requests occurred for the specified day and week. #1 - optional week value (1 - 6); defaults to the current value of the $$BeginChartWeek loop. #2 - optional day value (1 - 7); defaults to the current value of the $$BeginChartWeekDay loop.
$$ChartWeekDay[#3] or [#3,#4] or [#2,#3,#4] or [#1,#2,#3,#4]	Displays a chart showing the relative number of requests received on the specified day and week. The chart is produced using a single character (#4) repeated #3 times. #1 - optional week (1 - 6); defaults to $$BeginChartWeek. #2 - optional day value (1 - 7); defaults to $$BeginChartWeekDay. #3 - optional repeat value; defaults to 100. #4 - optional character; defaults to capital I
$$CurrentChartWeek	Displays the number of the current week (1 - 6).
$$WeekDateRange	Displays the date range represented by the current week. Dates are displayed using the Java “medium” date format.
By Day of Week	
$$BeginChartDayAverages[#]	Begins a repeating section that displays a load chart for the specified day of week. # - optional day of week value (1 - 7). If the optional day of week value is provided, but the day has no load information, this entire section is skipped. If the optional day of week value is not provided and the “current” day has no load information, the section is skipped for the current day of week, but repeated for any remaining days that follow.
$$EndChartDayAverages	Ends a repeating section that displays a load chart for the specified day of week.
$$BeginChartHour[#2] or [#1,#2]	Begins a repeating section that displays information for each hour of the current day of week. If the specified hour and day do not have load information available, the current section is not skipped. Use $$BeginSkipIfEmptyHour to skip the current section. #1 - optional day value (1 - 7). #2 - optional hour value (1 - 24).

Name	Description
$$EndChartHour	Ends a repeating section that displays information for each hour of the current day of week.
$$IsAverageDay[averageText, actualText]	Displays either averageText or actualText based on the number of days represented by the current day of week.
$$BeginSkipIfEmptyHour[#2] or [#1,#2]	Skips the current section (up to $$EndSkipIfEmptyHour) if the specified hour does not have any load information available. #1 - optional day of week value (1 - 7); defaults to the current value of the $$BeginChartDayAverages loop. #2 - optional hour value (1 - 24); defaults to the current value of the $$BeginChartHour loop.
$$EndSkipIfEmptyHour	Ends the current section if the specified hour does not have any load information available.
$$CurrentChartHour	Displays the time of day corresponding to the current hour.
$$HourTotal[#2] or [#1,#2]	Displays the average number of requests handled during the specified hour and day. #1 - optional day of week value (1 - 7); defaults to the current value of the $$BeginChartDayAverages loop. #2 - optional hour value (1 - 24); defaults to the current value of the $$BeginChartHour loop.
$$HourPeek[#2] or [#1,#2]	Displays the peak number of requests handled during the specified hour. This # matches $$HourTotal if only 1 day is represented for the day of week. #1 - optional day of week value (1 - 7); defaults to the current value of the $$BeginChartDayAverages loop. #2 - optional hour value (1 - 24); defaults to the current value of the $$BeginChartHour loop.
$$HourPeekTime[#2] or [#1,#2]	Displays the time of day when the peak number of requests occurred for the specified hour and day of week. #1 - optional day of week value (1 - 7); defaults to current value of $$BeginChartDayAverages loop. #2 - optional hour value (1 - 24); defaults to current value of $$BeginChartHour loop.

Name	Description
$$ChartHour[#3] or [#3,#4] or [#2,#3,#4] or [#1,#2,#3,#4]	Displays a chart showing the relative number of requests received on the specified hour and day of week. The chart is produced by using a single character (#4) repeated #3 times. #1 - optional day of week (1 - 7); defaults to $$BeginChartDayAverages. #2 - optional hour value (1 - 24); defaults to $$BeginChartHour. #3 - optional repeat value; defaults to 100. #4 - optional character; defaults to capital I

12.10.5 Search Details

Table 12-12 *Search Details Query Report Template Variables*

Name	Description
$$QueryTotal	Displays the total number of search requests received by QuickFinder. This includes all types of search, including failed searches.
$$SuccessfulQueries	Displays the total number of regular search requests that were successfully completed.
$$NotFoundQueries	Displays the total number of search requests that could not find any hits. This includes the searches that were successfully completed by looking in other indexes using the Search Expansion feature. See “$$ExpandedQueries” on page 153 for more information.
$$ErrorQueries	Displays the total number of search requests that resulted in an error condition. Errors can arise if the server is currently unable to complete the request, the request is too complex or takes too long to process, the query syntax is incorrect, the server load is too great, or because of a number of hardware-related or software-related issues that can prevent a search from processing properly.
$$ExpandedQueries	Displays the total number of “Not Found” search requests that were successfully completed by looking in other indexes as a result of the Search Expansion feature.
$$RedirectedQueries	Displays the total number of search requests that were redirected to a specific, admin-defined URL.
$$ShowFirstQueries	Displays the total number of search requests that resulted in a successful search, but were redirected to the first hit in the search results list.
$$SearchPageQueries	Displays the total number of search requests that resulted in showing the Advanced Search page to the end user.
$$LoginQueries	Displays the total number of search requests that represent an attempt by the search user to log in.
$$LogoutQueries	Displays the total number of search requests that represent an attempt by the search user to log out.

Name	Description
$$SwitchProtocolQueries	Displays the total number of search requests that result in a switch of protocols. If specified by the administrator, QuickFinder can redirect login attempts from the unsecured HTTP protocol to the more secure HTTPS.
$$AvgNumHitsPerPage	Displays the average number of hits listed on a single search results page.
$$AvgTotalNumHits	Displays the average total number of hits for user queries.
$$AvgQueryTime	Displays the average amount of time required to complete user queries.

12.10.6 Search Results Pages

Table 12-13 *Search Results Pages Query Report Template Variables*

Name	Description
$$BeginResultPageLoop[#]	Begins a repeating section that displays information regarding the search results page on which users found their information. # - optional page number to begin the loop. 1-based. The loop terminates as soon as a search results page returns a use percentage less than 1%.
$$EndResultPageLoop	Ends a repeating section that displays information regarding the search results page on which users found their information.
$$ResultPageNum	The number of the current search result page. 1-based.
$$ResultPageTotal	The total number of times user found their hit on the current search results page.
$$ResultPageCount	The maximum search results page number includes all of the <1% search results pages.
$$ResultTotalPages	The maximum search results page number includes all of the <1% search results pages.
$$ResultPagePercentage	Percent of users who found a hit on the current search results page.
$$ResultPageCurrent	The number of the current loop; see the optional # in $$BeginResultPageLoop. 1-based.

12.10.7 Top X Queries

Table 12-14 *Top X Queries Query Report Template Variables*

Name	Description
$$BeginTopQueriesLoop[#]	Begins a repeating section that displays the most popular user searches. The loop is repeated for the top number of queries. The optional number defaults to 100 if not provided.
$$EndTopQueriesLoop	Ends a repeating section that displays the most popular user searches.
$$TopQuery	The current most-popular query

Name	Description
$$TopSubQuery	All additional queries (separated by beginning and ending parentheses).
$$TopFilterQuery	All filter items (separated by beginning and ending parentheses).
$$TopFileFilterQuery	File filter string.
$$TopQueryTotal	The number of times the current query was requested by all users during the logging period.
$$TopQuerySubTotal	The number of searches represented by the top number of queries listed.
$$TopQueryCurrent	The number of the current loop; see the optional # in $$BeginResultPageLoop[#]. 1-based
$$TopQueryCount	The total number of unique queries requested by all users during the logging period.
$$BeginHasSubQuery	Everything between BeginHasSubQuery and EndHasSubQuery is displayed if any subqueries or filters were sent in the query.
$$EndHasSubQuery	Everything between BeginHasSubQuery and EndHasSubQuery is displayed if any subqueries or filters were sent in the query.

12.10.8 Top Queries

Table 12-15 *Top Queries Query Report Template Variables*

Name	Description
$$BeginIndexQueryLoop[#]	Begins a repeating section that displays the queries associated with the most popular indexes or the most popular index combinations. The loop is repeated for the top number of queries. The optional number defaults to 25 if not provided.
$$EndIndexQueryLoop	Ends a repeating section that displays the queries associated with the most popular indexes or the most popular index combinations.
$$IndexQuery	The current query associated with the index or index combination.
$$IndexQueryTotal	The number of times the current query was requested by all users during the logging period.
$$IndexQuerySubTotal	The number of searches represented by the number of queries listed.
$$IndexQueryCurrent	The number of the current loop; see the optional # in $$BeginIndexQueryLoop[#]. 1-based.
$$IndexQueryCount	The number of unique queries associated with the current index or index combination.
By Index	
$$BeginRptIndexesLoop[#]	Begins a repeating section that displays the most popular indexes. The loop is repeated for the top number of indexes. The optional number defaults to 1000 (that is, all indexes)
$$EndRptIndexesLoop	Ends a repeating section that displays the most popular indexes.

Name	Description
$$RptIndex	The name of the current index.
$$RptIndexQueriesTotal	The total number of times the current index was used in a query.
$$RptIndexCurrent	The number of the current loop; see the optional # in $$BeginRptIndexesLoop[#]. 1-based.
$$RptIndexCount	The total number of times the current index searches unique indexes available for searching (including those that were never used in a user query).
By Index Combinations	
$$BeginRptFullIndexesLoop	Begins a repeating section that displays the most popular index combinations. All index combinations are displayed.
$$EndRptFullIndexesLoop	Ends a repeating section that displays the most popular index combinations.
$$RptFullIndex	Displays the list of indexes corresponding to the current loop.
$$RptFullIndexQueriesTotal	The total number of times the current index combination was used in a query.
$$RptFullIndexCurrent	The number of the current loop. 1-based.
$$RptFullIndexCount	The total number of unique index combinations available for searching (including those that were never used in a user query).

12.10.9 Error Searches

Table 12-16 *Error Searches Query Report Template Variables*

Name	Description
$$BeginErrorQueriesLoop[#]	Begins a repeating section that displays the searches that produced an error. The loop is repeated for the top number of error queries. The optional number defaults to 50.
$$EndErrorQueriesLoop	Ends a repeating section that displays the searches that produced an error.
$$ErrorQuery	The current error query.
$$ErrorSubQuery	All additional queries (separated by beginning and ending parentheses).
$$ErrorFilterQuery	All filter items (separated by beginning and ending parentheses).
$$ErrorFileFilterQuery	File filter string.
$$ErrorQueryTotal	The number of times the current error query was requested by all users during the logging period.
$$ErrorQuerySubTotal	The number of searches represented by the top number of error queries listed.
$$ErrorQueryCurrent	The number of the current loop; see the optional # in $$BeginErrorQueriesLoop[#]. 1-based.

Name	Description
$$ErrorQueryCount	The total number of unique error queries requested by all users during the logging period.
$$BeginHasErrorSubQuery	Everything between BeginHasErrorSubQuery and EndHasErrorSubQuery is displayed if any subqueries or filters were sent in the query.
$$EndHasErrorSubQuery	Everything between BeginHasErrorSubQuery and EndHasErrorSubQuery is displayed if any subqueries or filters were sent in the query.

12.10.10 Not Found Searches

Table 12-17 *Not Found Searches Query Report Template Variables*

Name	Description
$$BeginNotFoundQueriesLoop[#]	Begins a repeating section that displays the searches that produced a Not Found response. This includes the searches that were successfully resolved by looking elsewhere using the Search Expansion feature. The loop is repeated for the top number of Not Found queries. The optional number defaults to 50.
$$EndNotFoundQueriesLoop	Ends a repeating section that displays the searches that produced a Not Found response.
$$NotFoundQuery	The current Not Found query.
$$NotFoundSubQuery	All additional queries (separated by beginning and ending parentheses).
$$NotFoundFilterQuery	All filter items (separated by beginning and ending parentheses).
$$NotFoundFileFilterQuery	File filter string.
$$NotFoundQueryTotal	The number of times the current Not Found query was requested by all users during the logging period.
$$NotFoundQuerySubTotal	The number of searches represented by the top number of Not Found queries listed.
$$NotFoundQueryCurrent	The number of the current loop; see the optional # in $$BeginNotFoundQueriesLoop[#]. 1-based.
$$NotFoundQueryCount	The total number of unique Not Found queries requested by all users during the logging period.
$$BeginHasNotFoundSubQuery	Everything between BeginHasNotFoundSubQuery and EndHasNotFoundSubQuery is displayed if any subqueries or filters were sent in the query.
$$EndHasNotFoundSubQuery	Everything between BeginHasNotFoundSubQuery and EndHasNotFoundSubQuery is displayed if any subqueries or filters were sent in the query

Name	Description
$$BeginNotFoundIndexesLoop	Begins a repeating section that displays the indexes associated with the current Not Found query. The loop is repeated for all associated indexes.
$$EndNotFoundIndexesLoop	Ends a repeating section that displays the indexes associated with the current Not Found query.
$$NotFoundIndex	The name of the current index.
$$NotFoundIndexQueriesTotal	The number of Not Found queries for a particular index.
$$NotFoundIndexCurrent	The number of the current NotFoundIndex loop. 1-based.
$$NotFoundIndexCount	The total number of indexes associated with the current Not Found query.

12.10.11 Redirected Searches

Table 12-18 *Redirected Searches Query Report Template Variables*

Name	Description
$$BeginRedirectsLoop[#]	Begins a repeating section that displays the searches that were redirected to a specific, admin-defined URL. The loop is repeated for the top number of redirected queries. The optional # defaults to 50.
$$EndRedirectsLoop	Ends a repeating section that displays the searches that were redirected to a specific, admin-defined URL.
$$RedirectedQuery	The current redirected query.
$$RedirectedQueryTotal	The number of times the current redirected query was requested by all users during the logging period.
$$RedirectedQuerySubTotal	The number of searches represented by the top number of redirected queries listed.
$$RedirectedQueryCurrent	The number of the current loop. 1-based.
$$RedirectedQueryCount	The total number of unique queries that were redirected.

12.10.12 Login Requests

Table 12-19 *Login Requests Query Report Template Variables*

Name	Description
$$BeginLoginUsersLoop	Begins a repeating section that displays the individual users who were authenticated through QuickFinder.
$$EndLoginUsersLoop	Ends a repeating section that displays the individual users who were authenticated through QuickFinder.
$$LoginUser	Displays the name of the current logged-in user.

Name	Description
$$LoginUserTotal	Displays the total number of times the current user logged in through QuickFinder.
$$LoginUserCurrent	The number of the current loop. 1-based.
$$LoginUserCount	The total number of unique users that successfully logged-in through QuickFinder.

12.10.13 Print Details

Table 12-20 *Print Details Query Report Template Variables*

Name	Description
$$PrintTotal	Displays the total number of print requests handled by QuickFinder during the log period. This includes all print request types (successful and unsuccessful).
$$SuccessfulPrints	Displays the total number of print requests that resulted in successful printed output.
$$NotFoundPrints	Displays the total number of print requests that resulted in a Not Found response.
$$ErrorPrints	Displays the total number of print requests that resulted in an error condition.
$$AvgNumHitsPerPrint	Displays the average number of results per print request.
$$AvgTotalPrints	Displays the average number of total results per print request.
$$AvgPrintTime	Displays the average amount of time required to process a print request. Print requests take much longer than searches because the individual documents must be retrieved from their source location, converted into QuickFinder's internal Unicode* format, then streamed back to the user using the specified output encoding.

12.10.14 Print Errors

Table 12-21 *Print Errors Query Report Template Variables*

Name	Description
$$BeginPrintErrorsLoop	Begins a repeating section that displays the print requests that produced an error condition. The loop is repeated for the top number of print requests. The optional number of defaults to 50.
$$EndPrintErrorsLoop	Ends a repeating section that displays the print requests that produced an error condition.
$$PrintErrorQuery	The current print request that resulted in an error response.
$$PrintErrorSubQuery	All additional queries (separated by beginning and ending parentheses).

Name	Description
$$PrintErrorFilterQuery	All filter items (separated by beginning and ending parentheses).
$$PrintErrorFileFilterQuery	File filter string.
$$PrintErrorQueryTotal	The number of times the current print request was issued by all users during the logging period.
$$PrintErrorQuerySubTotal	The total number of print requests represented by the current list of print errors.
$$PrintErrorQueryCurrent	The number of the current loop. 1-based.
$$PrintErrorQueryCount	The total number of unique print requests that produced an error response.
$$BeginHasPrintErrorSubQuery	Everything between BeginHasPrintErrorSubQuery and EndHasPrintErrorSubQuery is displayed if any subqueries or filters were sent in the query.
$$EndHasPrintErrorSubQuery	Everything between BeginHasPrintErrorSubQuery and EndHasPrintErrorSubQuery is displayed if any subqueries or filters were sent in the query.

12.10.15 Highlight Details

Table 12-22 *Highlight Details Query Report Template Variables*

Name	Description
$$HighlightTotal	Displays the total number of requests for QuickFinder to highlight the user's search terms within the original document. This includes both successful and unsuccessful highlight attempts.
$$SuccessfulHighlights	Displays the total number of successful requests for QuickFinder to highlight the user's search terms within the original document.
$$ErrorHighlights	Displays the total number of highlight requests that resulted in an error condition.
$$AvgHighlightTime	Displays the average amount of time required to highlight a user's search terms within the original document. The majority of the time is spent retrieving the original document.

12.10.16 Export Details

These template variables are used to export a raw version of the entire contents of the query log. The `ReportTemplate.html` file summarizes the queries for a certain date, and the `ExportTemplate.xml` file contains a complete copy of the query log formatted so that other query reporting engines can consume the results.

Table 12-23 *Export Details Query Report Template Variables*

Name	Description
$$BeginEXPORTLoop	Begins a repeating section that exports the query details and subsequent response information associated with all QuickFinder Search, Print, and Highlight requests.
$$EndEXPORTLoop	Ends a repeating section that exports the query details and subsequent response information associated with all QuickFinder search, print, and highlight requests.
$$EXPORTDate	Exports the time and date of the request.
$$EXPORTType	Exports the type of request: search, print, or highlight.
$$EXPORTUser	Exports the IP address making the request.
$$EXPORTUserID	Exports the logged-in UserID, if available.
$$EXPORTLocale	Exports the first locale provided by the request. For example, en or en_US.
$$BeginEXPORTIndexes	Begins a repeating section that exports the indexes associated with a request.
$$EndEXPORTIndexes	Ends a repeating section that exports the indexes associated with a request.
$$EXPORTIndex	Exports the name of the current index associated with the request.
$$BeginEXPORTQueries	Begins a repeating section that exports the user queries associated with each request. Users can send multiple weighted queries with each QuickFinder request.
$$EndEXPORTQueries	Ends a repeating section that exports the user queries associated with each request.
$$EXPORTQuery	Exports the current user query associated with the QuickFinder request. Users can send multiple weighted queries with each request.
$$BeginEXPORTFilters	Exports the filters associated with each QuickFinder request. A filter is used to send additional hidden details with each query. Multiple filters can be associated with each query. Multiple queries can be associated with each request.
$$EndEXPORTFilters	Exports the filters associated with each QuickFinder request. A filter is used to send additional hidden details with each query. Multiple filters can be associated with each query. Multiple queries can be associated with each request.
$$EXPORTFilter	Exports the value of the current filter.
$$EXPORTFileFilter	Exports the value of the file filter associated with the current request.
$$EXPORTResult	Exports the type of the response returned from processing the request: Normal | Expanded | Redirected | ShowFirst | SearchPage | Login | Logout | NotFound | Error | SwitchProtocols.
$$EXPORTTemplate	Exports the name of the template associated with the QuickFinder response.

Name	Description
$$EXPORTPage	Exports the number of the search results page. Begins with 1. Used only for Normal, Expanded, and ShowFirst result types.
$$EXPORTHits	Exports the number of hits associated with the current Search Results page.
$$EXPORTTotalHits	Exports the total number of hits associated with the current request.
$$EXPORTTime	Exports the number of milliseconds required to fully process the request.

12.11 Search Parameters

The following table lists and describes all available search parameters. Each of these parameters can be used to extend or enhance the functionality of the search page templates or to create new search page templates.

For more information about how to implement parameters in an HTML document, see Chapter 13, "Customizing Your Templates and Search Forms," on page 173.

NOTE: If you use a parameter but leave its value blank, the default value for that parameter is used.

Table 12-24 *Search Parameters*

Name	Value	Description
basiclogin	Boolean	Displays the browser's basic authentication dialog box. **Syntax:** `basiclogin=true|false` **Options:** True, False **Example:** `basiclogin=True`
bbindex	String	Specifies which best bets index to use. Can be used multiple times to specify additional indexes. Overrides the default setting. **Syntax:** `bbindex=variable` **Alternate Syntax:** `bbindex=indexname;indexname` **Example:** `bbindex=myindex` **Default:** None (specified using QuickFinder Server Manager)
bbshow	Boolean	Specifies whether or not the best bets list should be shown on the search results page. Overrides the default setting. **Syntax:** `bbshow=true|false` **Options:** True, False **Example:** `bbshow=true` **Default:** True (specified using QuickFinder Server Manager)

Name	Value	Description
bbnum	Integer	Specifies the number of best bets to show on the results page. Overrides the default setting. **Syntax:** `bbnum=variable` **Options:** Any positive number less than or equal to the maximum number specified in QuickFinder Server Manager. **Example:** `bbnum=5` **Default:** None (specified using QuickFinder Server Manager)
buttonpressed	String	A button clicked by the user. If this value is part of the query, then a response message should not be sent to the client. **Syntax:** `buttonpressed=button_name` **Options:** Yes, No, OK, Cancel, Continue, Ignore, Retry, Prev, Next, Help. **Example:** `buttonpressed=Cancel`
context	String	Overrides the user's default context for the current login. **Syntax:** `context=variable` **Example:** `context=sales.novell` **Default:** None (specified using QuickFinder Server Manager)
country	String	Specifies your country using the two-character, uppercase country value derived from ISO 3166 (http://www.iso.org/iso/en/prods-services/iso3166ma/02iso-3166-code-lists/list-en1.html). **Syntax:** `country=country code` **Example:** `country=TW` **Default:** As specified by Web browser language preferences.
date	Integer	Specifies a date range to be searched in milliseconds. If you pass a positive number such as 940457147873, then QuickFinder creates a date and time based on the number of milliseconds elapsed since January 1, 1970, 12:00 a.m. For example, the number 940457147873 produces the search start date of October 20, 1999, at 4:05:47 p.m. The example shows the number of milliseconds spanning a three-month time frame. The minus sign (-) before the number indicates three months back in time. **Syntax:** `date=number` **Example:** date=-7905600000

Name	Value	Description
encoding	String	Specifies the character set encoding used to encode the search request itself. **Syntax:** `encoding=value` **Example:** `encoding=Shift_JIS` **Default:** UTF-8 (specified using QuickFinder Server Manager)
expandindex	String	Specifies an expanded search index. You can specify more than one index at a time. Overrides the default setting. **Syntax:** `expandindex=variable` **Alternate Syntax:** `expandindex=indexname1;indexname2;indexname3` **Example:** `expandindex=developerindex` **Default:** None (specified using QuickFinder Server Manager)
filefilter	String	Filters search results based on their path, domain, filename, or extension. Uses the same query syntax as the &query= parameter. **NOTE:** If a primary &query= parameter is not provided, then QuickFinder automatically performs a filename-only search. **Syntax:** `filefilter=search_criteria` **Example:** `filefilter=.exe` **Example:** `filefilter=.com`
filter*query_number*	String	The &filter#= query parameter is used to send additional query details not specified by the user in order to help limit the scope of a search. Normally, these are included as hidden fields on an HTML form. This parameter supports all of the same features and functionality as the &query= parameter. However, unlike the &query= parameter, this parameter can be sent multiple times for a single query item. The individual filters associated with a single query item are joined using the value of the filteroperator parameter. The set of filters is logically joined with the rest of the query item by using the Boolean AND operator. **Syntax:** `filterquery_number=searchcriteria` **Example:** `filter0=^product=GroupWise` See also “filteroperatorquery_number” on page 165.

Name	Value	Description
filteroperator*query_number*	Number	Boolean conjunction to be used between multiple filters (several filters can be associated with a single query item). The complete set of filters is always associated with the corresponding query item using the Boolean AND operator. **Syntax:** `filteroperator`*`query_number=number`* **Options:** 0 (AND), 1 (OR), 2 (PHRASE) **Example:** `filteroperator0=1` **Default:** 1 (OR) See also the query parameters “operatorquery_number” on page 166 and “filterquery_number” on page 164.
gettotalhits	Boolean	Enables or disables the total number of hits calculation. For example, if you set the gettotalhits parameter to False, the Total Number of Hits label on the results page displays 0 (zero). Setting this parameter to True shows the total number of hits found during the search. In some complex situations, this can save valuable processing time. **Syntax:** `gettotalhits=`*`True|False`* **Options:** True, False **Example:** `gettotalhits=False` **Default:** True
id*query_number*	String	Document ID that is used to narrow a search. You can specify more than one ID by using the same field name more than once. **Syntax:** `id`*`query_number=documentID`* **Example:** `&id0=z1.0010.&id0=z1.0020` **Default:** None
index	String	Restricts a search to one or more specified indexes. The index name you specify with this parameter must exactly match the name of an index defined at the server. You can specify more than one index by either sending this parameter more than once or by separating the list of indexes with a semicolon (;). **Syntax:** `index=`*`index_name1[;index_name2]`* **Example:** `index=GroupWise&index=NetWare` **Example:** `index=GroupWise;NetWare` **Default:** All indexes or the list of default indexes specified in QuickFinder Server Manager.

Name	Value	Description
lang	String	Specifies a language using the two-character, lowercase language value derived from ISO 639-1 (http://www.geo-guide.de/info/tools/languagecode.html). **Syntax:** `lang=language_code` **Example:** `lang=ja` **Default:** Browser language preference.
login	Boolean	When set to True, displays a form-based login rather than the standard browser-based login dialog box. **Syntax:** `login=True|False` **Options:** True, False **Example:** `login=True`
logout	Boolean	When set to True, logs out the current user. **Syntax:** `logout=True|False` **Options:** True, False **Example:** `logout=True`
noexpand	Boolean	When set to True, turns off search expansion. **Syntax:** `noexpand=variable` **Options:** True, False **Example:** `noexpand=True`
noredirect	Boolean	When set to True, turns off URL redirection. **Syntax:** `noredirect=True|False` **Options:** True, False **Example:** `noredirect=True`
numhits	Integer	Number of hits you want returned at one time in the search results page. **Syntax:** `numhits=number` **Example:** `numhits=25` **Default:** 25 (specified using QuickFinder Server Manager)
operator*query_number*	Integer	Which operator to use between two or more words in a search. **Syntax:** `operatorquery_number=number` **Options:** 0 (AND), 1 (OR), 2 (PHRASE) **Example:** `operator0=1` **Default:** 0 (specified using QuickFinder Server Manager)

Name	Value	Description
password	String	User's password for form-based logins. **Syntax:** `password=variable` **Example:** `password=novell`
query*query_number*	String	Actual search criteria passed to the QuickFinder Server. Because multiple queries can be defined in a single search request, the number variable lets you group various parameters together that apply to a specific query. Adding the same number to two or more parameters groups them together so they can be applied to the same query. For example, &query0=product+downloads&weight0=100 groups the &weight query parameter together with &query0. This parameter is used most often with operator*number*, filter*number*, filteroperator*number*, id*number*, weight*number*, and type*number*. **Syntax:** `query*query_number*=*searchcriteria*` **Example:** `query0=novell+AND+groupwise` For more information, see Section 10.6, “Weighted Queries,” on page 122.
relevance	String	Whether or not to sort the search results by relevance. Turning this feature off is a potential speed gain because the sort algorithm won't run. **Syntax:** `relevance=*True*|*False*` **Options:** True, False **Example:** `relevance=False` **Default:** True
retencoding	String	Specifies the character set encoding to be used by the next results page returned to the user. **Syntax:** `retencoding=*character_set_encoding*` **Example:** `retencoding=iso-8859-1` **Default:** UTF8 (specified using QuickFinder Server Manager)

Name	Value	Description
retfield	String	Level of detail given about each result item. The fewer the details, the faster a search is returned to a user. Field names include title, author, URL, changedate, language, summary, relevance, index, format, and filelength. **IMPORTANT:** Type these fields exactly as they appear above. To specify more than one field, separate them with an ampersand (&). **Syntax:** `retfield=`*`field_name`* **Example:** `retfield=title` **Example:** `retfield=title&retfield=author`
server	String	Name of the virtual search server that is to receive this request. This query parameter is optional if the domain name of the request matches the name or alias of a registered virtual search server. **Syntax:** `server=`*`virtual_search_server_name`* **Example:** `server=digitalairlines` **Default:** Domain name portion of search request
showfirsthit	Boolean	If set to True, this parameter automatically goes to the URL of the first hit on the current page rather than displaying the search results page. **Syntax:** `showfirsthit=`*`True`*`|`*`False`* **Options:** True, False **Example:** `showfirsthit=True` **Default:** False
sortbydate	String	Sorts the Total Search Results list by date, ignoring the normal relevance ordering. **Syntax:** `sortbydate=`*`True`*`|`*`False`* **Options:** True, False **Example:** `sortbydate=True` **Default:** False

Name	Value	Description
sortfield*query_number*	String	Specifies the fields on which to sort the search results returned in a results page. Grouped together with the sortorder query parameter by adding a *number* to the end of the parameter name. Field names include title, author, URL, changedate, language, summary, relevance, index, format, and filelength. **IMPORTANT:** Type these fields exactly as they appear above. **Syntax:** `sortfield`*`query_number`*`=`*`field_name`* **Example:** `sortfield1=title`
sortkeys	Integer	Specifies the number of sort fields that should be used to sort the search results. **Syntax:** `sortkeys=`*`number`* **Example:** `sortkeys=1`
sortorder*query_number*	Integer	Specifies the alphanumeric ordering of search result items (hits). Grouped together with the sortfield query parameter by adding a *number* to the end of the parameter name. **Syntax:** `sortorder`*`query_number`*`=`*`number`* **Options:** 0 (Ascending), 1 (Descending), 2 (Default for each field) **Example:** `sortorder1=0`
starthit	Integer	Hit number you want QuickFinder to begin searching from. If you specify a value of 35, QuickFinder returns hits beginning with hit number 35. **Syntax:** `starthit=`*`number`* **Example:** `starthit=35` **Default:** 1
synflag	String	When set to False, turns off synonym-enhanced searching. Also overrides the default specified in QuickFinder Server Manager. **Syntax:** `synflag=`*`True`*`|`*`False`* **Options:** True, False **Example:** `synflag=False` **Default:** True

Name	Value	Description
template	String	Specifies the results template you want your search results returned in. The following list includes the default templates included with the QuickFinder Server. However, your QuickFinder Server administrator might have created custom templates using different names. Check with your administrator if the default templates do not work for you. You must type the names of these templates exactly as they appear in this list: ◆ `ResultListTemplate.html` ◆ `ResultListTerseTemplate.html` ◆ `ResultListVerbose.html` ◆ `PrintResultTemplate.html` ◆ `SearchTemplate.html` ◆ `HighlighterTemplate.html` Localized versions for multiple languages can also be used. See Section 14.1, “Working with Multiple Languages,” on page 177. **TIP:** A number of different file formats can be used as QuickFinder templates. For example, QuickFinder ships both HTML and XML templates. QuickFinder uses the current template filename extension to determine which template filename to use next. For example, if a user specifies the `ResultListTerseTemplate.xml` file for the current query, but the search produces a Not Found condition, QuickFinder automatically tries to find a template named `ResultListNoHitsTemplate.xml`. **Syntax:** `template=`*`filename`* **Example:** `template=ResultList.html` **Default:** Search Page (`SearchTemplate.html`), Search Results (`ResultListTemplate.html`), and Print (`PrintResultTemplate.html`) Default templates are specified using QuickFinder Server Manager.
theme	String	Name of the theme (or directory within the templates directory) where a complete set of search and print templates are stored. **Syntax:** `theme=`*`theme_name`* **Example:** `theme=Intranet`

Name	Value	Description
type*query_number*	Integer	Type of search. 0 = Normal search; this is the default. 1 = Searches only the given document numbers. 2 = Root search used by the search tree control to get the top tree nodes. 3 = Used to get the children of the given document number. 4 = Searches the descendants of the given document numbers. Used to narrow a search or a print request, including all of its children. **Syntax:** `type*query_number*=*number*` **Options:** 0 = Normal search; this is the default. 1 = Searches only the given document numbers. 2 = Root search used by the search tree control to retrieve the top tree nodes. 3 = Retrieves the children of the given document number. 4 = Searches the descendants of the given document numbers. Used to narrow a search or a print request, including all of its children. **Example:** `type0=2` **Default:** 0 (zero)
url	String	Specified as a parameter to the Highlight servlet to determine the URL to highlight. The URL is highlighted only if the current user has rights to the file. The value must be properly URL-escaped. **Syntax:** `url=*value*` **Example:** `url=http%3A%2F%2Fwww.novell.com%2Findex.html`
userid	String	User ID for form-based logins. **Syntax:** `userid=*variable*` **Example:** `userid=jsmith`
weight*query_number*	Integer	Assigns a level of importance to the current query item as it relates to the other query items that are part of the search query. QuickFinder Server uses this number along with the relevance number to determine a search result's final relevance and then orders the results accordingly. **Range:** 0 to 100 **Syntax:** `weight*query_number*=*number*` **Example:** `weight0=75` **Default:** 100

Customizing Your Templates and Search Forms

13

You can quickly create a custom search solution by modifying the default QuickFinder™ templates. Templates include some fundamental options for users, but you can add or remove options and modify the form layout and design to give the search form the look, feel, and function you need. If you are creating a hosted search service for another company's Web site, you can modify the templates to match the look and feel of their Web site.

If you are confident in coding with HTML, you can start with the default search page template to get a feel for the available parameters and then begin coding completely new search and print templates from scratch.

This section includes the following topics:

- Section 13.1, "Customizing Templates," on page 173
- Section 13.2, "Testing Your Search and Print Solution," on page 176

13.1 Customizing Templates

You can extend the capabilities of QuickFinder Server by customizing the templates.

The first step is to determine which components of QuickFinder you want to customize. For example, if you only want to add a few additional search features to the search page template and modify its background color and table size, you would modify the `SearchTemplate.html` or `SearchTemplate.Simple` files.

This section discusses how to customize the search, print, and result templates and how to use available parameters and variables to create a customized search solution.

- Section 13.1.1, "Customizing the Search Templates," on page 173
- Section 13.1.2, "Customizing Search Result Templates," on page 174
- Section 13.1.3, "Customizing Print Result Templates," on page 175
- Section 13.1.4, "Customizing Error and Response Message Templates," on page 175

13.1.1 Customizing the Search Templates

You can customize the design and functionality of the static or dynamic search templates by

- Modifying HTML code
- Adding or removing search parameters

If you are familiar with HTML, you can quickly modify the design of the default (dynamic) search template or the static search template. For example, you can change the colors of the search page or add new custom graphics.

To modify the functionality of the default search template, you can add or remove search parameters. Search parameters are used to communicate with QuickFinder. By embedding them in

the correct places in your HTML source, you can extend or limit the functionality of the default search template.

For example, if you wanted your users to use a specific set of templates found in a themes directory, you would add the following HTML code, including the theme parameter, to the `SearchTemplate.html` file:

```
<INPUT TYPE="Hidden" NAME="theme" VALUE="$$QueryTemplateTheme">
```

This sample HTML code tells QuickFinder to look for templates only in the specified template directory. All themes are located within the templates directory specified in QuickFinder Server Manager.

For a complete list of available search parameters, see Section 12.11, "Search Parameters," on page 162.

13.1.2 Customizing Search Result Templates

QuickFinder Server includes several default search result templates that are used to display hits, provide feedback to a user, or request information from a user after a search is performed. For more information about the default search result templates, see Chapter 11, "Understanding Templates," on page 127.

You can customize the design and functionality of the default search result template, which is the template used when a user selects *Normal* from the *Result List Format* drop-down list in the QuickFinder form. For information about how to access the QuickFinder form, see Section 1.5, "Taking a Test Run: Performing a Search," on page 19.

Customizing the default search result template involves

- Modifying the HTML code
- Adding or removing search result variables

If you are familiar with HTML, you can quickly modify the design of the default search result template. For example, you can change the colors of the search page or add new graphics.

To modify the functionality of the default search result template, you can add or remove search result variables. Search result variables are placed in the template where you want search results to be displayed.

For example, if you want to display the total number of hits returned when a user performs a search and you want the information to appear in the upper left corner of the search results page, you would add the following HTML code to the search result template file:

```
Total Search Results: $$TotalHits
```

After a user performs a search, the $$TotalHits variable is replaced by the actual total number of hits found during the search.

The $$TotalHits variable is used to retrieve the total number of hits found during a search. You can place this variable anywhere in the results list template to organize the display of information in the way you want.

Default search result templates are located in *volume*`:\`*searchroot*`\TEMPLATES\`. For a complete list of search result variables that you can use to customize default search result templates or to create new ones, see Section 12.5, "Search Result Variables," on page 141.

13.1.3 Customizing Print Result Templates

Print result templates are used to organize and format search results for printing and to provide feedback to a user when no hits are found. QuickFinder Server includes two default print result templates: the default print result template and a “no hits” template.For more information about the default print result templates, see Chapter 11, “Understanding Templates,” on page 127.

You can customize the design and functionality of the default print result template in the same way you customize the search result template:

- Modifying the HTML code
- Adding or removing print result variables

If you are familiar with HTML, you can quickly modify the design of the default print result template. For example, you can change the colors of the print results page or add new graphics to it.

To modify the functionality of the default print results template, you can add or remove print result variables. Variables are placed in the template where you want search results to be displayed.

For example, if you wanted to remove the table of contents from the default print results template, you would remove, or comment out, the following HTML code in the `PRINTRESULTLIST.HTML` template, which would include the $$BeginTOCList variable:

```
<CENTER><H2>Table of Contents</H2></CENTER><p><!--           TABLE OF
CONTENTS                    - >$$BeginTOCList[<BIG><B>$$Product</B></
BIG><DL>]<DT><A HREF=”#$$Bookmark”><BIG>$$Title</BIG></A><SPACER
TYPE=HORIZONTAL SIZE=20><I><SMALL>[$$URL]</SMALL></I>$$EndTOCList[</
DL>]
```

You could either save your changes in the default print result list template or you could save it with a new name, thereby creating an alternative template for users who don't want a table of contents in the print results. For the template to be effective, you would then need to add a hypertext link in the search result template that would include the &template=*new_template_name* query parameter.

Default print result templates can be found at *`volume`*`:\`*`searchroot`*`\templates`. For a complete list of print result variables that you can use to customize default search result templates or to create new ones, see Section 12.6, “Print Result Variables,” on page 144.

13.1.4 Customizing Error and Response Message Templates

Error and response messages are used to either provide feedback to the user or to request information from the user.

Error and response message templates are used to display the content of error and response messages sent by the QuickFinder Server in response to search or print errors. Similar to search and print templates, error and response templates can be customized. However, because the contents of error and response messages are built into QuickFinder Server, you cannot modify the contents of the messages or the button objects that might appear, depending on the type of response being generated.

- “Customizing Error Messages” on page 176
- “Customizing Response Messages” on page 176

Customizing Error Messages

There are several error messages that can be returned to a user. For example, when users incorrectly use a search operator in a search form, they might get the message, “Search Error: Incorrect use of Boolean operator.” An error number might also appear.

Although you can use HTML tags to format an error message, add or remove variables to determine what information is shown to the user, or even reorganize where the messages appear in the template, you cannot modify the message itself.

Customizing Response Messages

The same concepts apply to response messages, but response messages return buttons that a user can click. Which buttons appear are determined by the QuickFinder Server. Although you can modify the labels of these buttons, you cannot determine which buttons appear, or when.

13.2 Testing Your Search and Print Solution

After you customize the templates and the search form, you can test them in your Web browser by pointing to the search form URL and entering a search string. See Section 1.5, “Taking a Test Run: Performing a Search,” on page 19 for information about how to access the QuickFinder form.

Remember that a search cannot be performed until you have defined at least one index and generated it by using QuickFinder Server Manager. See Section 8.5, “Creating Indexes,” on page 83 and “Generating Indexes” on page 95 for more information.

Internationalizing Search Services 14

QuickFinder™ Server is capable of handling search queries, search results, templates, and Web content in many languages and character sets. QuickFinder can auto-detect languages and character sets, but to ensure a complete international search solution, you must identify language, country, and character information throughout your QuickFinder implementation.

This section discusses key issues related to supporting multiple languages from a single search solution and includes the following topics:

- Section 14.1, “Working with Multiple Languages,” on page 177
- Section 14.2, “Understanding Character Set Encodings,” on page 179
- Section 14.3, “Languages Included in the Default Templates,” on page 183
- Section 14.4, “Additional Resources,” on page 183

14.1 Working with Multiple Languages

Customizing your search solution is important only if you want to let your users conduct language-specific searches. You specify the language of a template by inserting a language identifier in the meta tag of your templates or HTML files. The language identifier can also be used in Search Results pages to let users quickly recognize the search results that interest them.

QuickFinder Server also lets Web clients specify their locale at the time the search query is entered. The default Search page illustrates this feature by auto-detecting a user's locale and selecting the appropriate language from the Display Language drop-down list. This selection sends two parameters to the QuickFinder Server: language and country. The country parameter is almost always blank. The search engine uses this information to find locale-specific versions of the templates used to return search results.

To specify the language of a template or of any HTML content that is indexed as part of your virtual search server, you must enter a language identifier within an HTML file's header section. For example, if you wanted to identify a Russian template, you would add the following meta tag:

```
<meta http-equiv=”Content-Language” content=”ru”>
```

In some cases, such as Traditional and Simplified Chinese, you need to use the two-character, uppercase country codes. For example:

```
<meta http-equiv=”Content-Language” content=”zh-TW”><meta http-
equiv=”Content-Language” content=”zh-CN”>
```

The first line of the example indicates the Chinese language (ZH) and the geographic location as Taiwan. The second line of the example indicates the Chinese language (ZH) but China as the geographic location.

This combination of language and country codes is called a locale. For more information about locales, refer to Section 14.4, “Additional Resources,” on page 183.

14.1.1 Specifying Locales within Template Filenames

QuickFinder Server consists of three primary servlets: SearchServlet, PrintServlet, and AdminServlet. Each servlet returns information to the Web client using server-side templates. Templates are stored at *volume*:*searchroot*\TEMPLATES. For more information about templates, see Chapter 11, "Understanding Templates," on page 127.

After determining a Web client's locale, QuickFinder attempts to locate a matching search result template. That is, each of the QuickFinder services automatically attempts to locate a version of the requested template that most closely matches the Web client's locale.

IMPORTANT: QuickFinder cannot find locale-specific templates without the two-character language code and the optional two-character country code. See Section 14.4, "Additional Resources," on page 183 for more information about language code syntax.

For example, if a Web client requests to see search results using the `ResultListTemplate.html` file and the client is a Chinese language user from Taiwan and the server is Russian, then QuickFinder will try to find a Chinese-Taiwan version of the template first (`ResultListTemplate_zh_TW.html`) because that exactly matches the client's language and country. The following table lists the template names the system looks up in this example in order of priority.

Table 14-1 *Template Name Priority (Chinese-Taiwan)*

Template Name	What QuickFinder Concludes
1. `ResultListTemplate_zh_TW.html`	Specific client locale
2, `ResultListTemplate_zh.html`	Simplified client locale
3. `ResultListTemplate.html`	Client requested name
4. `ResultListTemplate_ru.html`	Specific server locale (no simplified versions)
5. `ResultListTemplate_en.html`	English language version
6. `ResultListTemplate.html`	Up to the first underscore (_)

If this scenario were reversed so that the search client was Russian and the server was Chinese (Taiwan), and the client requested the `ResultListTemplate_ja.html` template, then the lookup order follows the order shown in the following table.

Table 14-2 *Template Name Priority (Japanese)*

Template Name	What QuickFinder Concludes
1. `ResultListTemplate_ja_ru.html`	Specific client locale (no simplified versions)
2. `ResultListTemplate_ja.html`	Client requested name
3. `ResultListTemplate_ja_zh_TW.html`	Specific server locale
4. `ResultListTemplate_ja_zh.html`	Simplified server locale

Template Name	What QuickFinder Concludes
5. `ResultListTemplate_ja_en.html`	English language version
6. `ResultListTemplate.html`	Up to the first underscore (_)

All templates undergo this rigorous lookup system. When a template is located, its name is stored and associated with the original client locale so that all subsequent requests for that template from the same locale automatically find the template without performing the same rigorous lookup.

No further lookups are attempted for that combination of client locale and template name until the QuickFinder Server is restarted. If all template lookups fail, then an error message is returned to the client performing the search.

14.2 Understanding Character Set Encodings

A character set is a grouping of alphabetic, numeric, and other characters that have some relationship in common. For example, the standard ASCII character set includes letters, numbers, symbols, and control codes that make up the ASCII coding scheme. A character set encoding is the mapping of a character set to a value that can be understood and processed by a computer.

QuickFinder relies on character set encodings to identify the characters used when performing a search, reading a template, posting results to a Web browser, or indexing Web-based content. If the encoding information is missing in any of these areas, QuickFinder uses the default encodings identified in the SearchServlet and PrintServlet properties files. You can modify these settings by using QuickFinder Server Manager.

Because most languages have several encodings that identify their character sets, QuickFinder Server supports a wide variety of character set encodings and encoding aliases.

Some examples of character set encodings include iso-8859-1, shift_jis, big5, and latin2. The official list of registered encodings is available from the Internet Assigned Numbers Authority (see Section 14.4, “Additional Resources,” on page 183). These are the official names for character sets that can be used in the Internet and can be referred to in Internet documentation. However, not all IANA-registered character set encodings are supported by QuickFinder Server. Refer to Section 14.4, “Additional Resources,” on page 183 for a list of encodings and encoding aliases that are supported by QuickFinder Server.

- Section 14.2.1, “Unicode and UTF-8,” on page 179
- Section 14.2.2, “Search Encodings,” on page 180
- Section 14.2.3, “Response Encodings,” on page 180
- Section 14.2.4, “HTML Encodings,” on page 181
- Section 14.2.5, “Template Encodings,” on page 182
- Section 14.2.6, “Encoding Issues When Printing,” on page 182

14.2.1 Unicode and UTF-8

Unicode is a 16-bit character encoding standard developed by the Unicode Consortium. By using two bytes to represent each character, Unicode enables almost all of the written languages of the world to be represented using a single character set. Unicode does not require any special processing to access any character in any language.

This makes Unicode very easy to use when processing text from multiple languages and scripts. This is the reason QuickFinder converts all external files into Unicode for processing.

As already mentioned, Unicode is two bytes wide for all characters. Although this is ideal for computer processing, it doubles the size of all single-byte languages. This has a significant impact on Internet performance. For this reason, QuickFinder also supports an alternate representation of Unicode known as UTF-8. UTF-8 is a Unicode Transformation Format that uses sequences of 1 to 6 bytes to represent all the characters in the Unicode standard. Most notably, ASCII characters are transmitted without any conversion at all. This means that most Internet content is already in the UTF-8 representation. Many Asian languages, however, require three bytes per character in the UTF-8 format. Other languages can require up to six bytes to represent each of their characters.

You need to decide if Unicode or UTF-8 best meets your needs when creating HTML content, QuickFinder templates, or search pages.

14.2.2 Search Encodings

The only encodings QuickFinder currently supports when performing a search are Unicode and UTF-8. Therefore, any page that allows Web users to enter a search must ensure that the results are passed to the server in one of these two formats. See “Template Encodings” on page 182 for more information.

To pass Unicode characters to QuickFinder, use the syntax %uHHHH, where

- The percent sign (%) is used as the CGI escape character
- The lowercase letter U (u) indicates that the subsequent 4 characters represent a Unicode value.
- Four uppercase H letters (HHHH) indicate four hexadecimal characters (0-9, A-F)

To pass UTF-8 characters to QuickFinder, just use normal ASCII characters or the syntax %HH... for all other characters, where

- % is the CGI escape character
- HH indicates two he xi decimal characters (0-9, A-F)
- %HH indicates additional %HH groupings that might be required to properly transmit a character

---

TIP: If the encoding of the page containing a search form is already set to UTF-8 or Unicode, most browsers automatically transmit the entered search text correctly by using the designated encoding.

By default, QuickFinder uses UTF-8 in its sample search pages.

14.2.3 Response Encodings

One of the many parameters that can be sent when conducting a search is the encoding that should be used when returning the results to the browser. All QuickFinder encodings listed in Appendix B, “Combined Character Sets for Use with QuickFinder,” on page 197 can be used.

If the search result page contains the ability to refine or redo the search, then the response encoding can significantly impact the possible characters that can be entered when conducting the next search from this page. For example, if the user requests results in the iso-8859-1 encoding (HTML's default), then only iso-8859-1 characters can be entered in the subsequent search from that page.

Other characters can still be sent to the QuickFinder services using the %uHHHH and %HH formats, but the browser does not allow users to enter normal text characters other than those supported by iso-8859-1.

Although QuickFinder can return search results from many languages, some characters found in titles and descriptions might be returned as question marks (?) indicating that these characters are not available in the current response encoding. If a character can be represented in the current encoding but a font is not available, many browsers substitute an alternate character such as an empty box character. After the appropriate fonts have been installed, these characters will then display properly.

By default, QuickFinder returns all search, print, and administration pages in UTF-8.

14.2.4 HTML Encodings

Because HTML content can contain text written in many character sets, all HTML files need to include a tag that identifies the character set encoding. To identify the encoding of an HTML file (or search template), use the following meta tag at the top of the file's header section:

```
<meta http-equiv="Content-Type" content="text/html;
charset=Shift_JIS">
```

In this example, you would replace Shift_JIS with the appropriate Internet Assigned Numbers Authority (IANA)-assigned encoding value.

It is very important that the CHARSET value accurately represent the character set encoding that was actually used when the HTML Web content or QuickFinder template was created. A correct entry allows QuickFinder to accurately interpret and convert the characters in the document. An incorrect entry prevents QuickFinder from being able to read the characters as valid data in the authored language.

IMPORTANT: Improperly identified characters result in garbled text. In some cases, the Web-based content cannot be properly indexed or printed. In the most severe cases, the document being read might produce a server-side exception, which will ultimately discontinue processing the document and perhaps the entire operation.

Because QuickFinder is Unicode-based, when reading templates or when indexing or printing HTML content, all character encodings are converted from their source encoding to Unicode for internal processing.

During indexing, if a document contains characters not supported by the designated encoding, if the document doesn't have an encoding designation, or if the designation is inaccurate, the indexer attempts to recover. But if it cannot, it might index the information incorrectly or quit indexing that page entirely.

When reading a template file, QuickFinder might automatically cease processing the file if it contains any characters not supported by the current encoding. It tries to ignore the invalid text and continue, but this might not be possible.

When displaying search results or when printing HTML content, any character that does not match the specified response encoding receives a question mark (?) in its place when rendered at the browser. Although some characters are properly supported by the current encoding, the browser might not have the required fonts to display the characters. In this case, users might see square boxes representing these characters. This is an indication that the valid character reached the browser, but

the operating system could not provide a font to properly render the character. The user then needs to either change fonts or install the correct fonts in order to properly display the characters.

TIP: If a document does not contain a CHARSET encoding value, the default encoding for HTML documents is ISO-8859-1, also known as Latin1. The default encoding for plain text documents is US-ASCII.

QuickFinder also allows administrators to define the default encodings for templates, HTML content when printing, and search and print responses. Refer to the QuickFinder Server Manager Help for information about changing the default encodings.

14.2.5 Template Encodings

All HTML documents should include a Content-Type meta tag identifying their character set encodings. The character set encoding allows HTML Web clients (or browsers) to understand the contents of the file. This tag is also used by browsers to automatically switch their display system and fonts to correctly show the Web page's contents. This lets users surf the World Wide Web without constantly changing their display system as they encounter content from various languages and character sets.

However, because QuickFinder lets administrators specify both template encodings and response encodings, browsers might get confused when presented with the valid response encoding in the HTTP header and one or more alternate encodings from the Content-Type meta tags within the file that was part of the original QuickFinder template.

NOTE: $$IncludeFile[ ] templates can also contain their own Content-Type meta tags.

To solve this problem, QuickFinder allows placing the Content-Type meta tag specifying the template's encoding within an HTML comment. This effectively obscures the original template encoding from the browser, but still allows QuickFinder to read the encoding when the template file is processed.

A sample QuickFinder template is given below. The Content-Type meta tag has been hidden inside of an HTML comment. This template can be embedded within other templates using the $$IncludeFile[ ] template variable without affecting QuickFinder's ability to distinguish between the various encodings. This file can also be processed and then sent to a user's Web browser without conflicting with the response encoding provided by QuickFinder in the HTTP response headers.

```
<html>
<head><!-- Note that the HTML encoding command (meta tag) is hidden
within a comment so that it does not affect a user's browser display. -
><!-- The actual encoding used when sending this file to the user is
controlled by the response encoding - ><!-- <META HTTP-EQUIV="Content-
Type" CONTENT="text/html; charset=iso-8859-1"> - ></head>
<body>
Template data here.</body>
</html>
```

14.2.6 Encoding Issues When Printing

When QuickFinder processes a print request, it gathers the entire contents of each file and builds an appended print job page, one file after another. Each file can contain its own Content-Type meta tag

identifying its encoding. Each file's encoding is used by QuickFinder to convert that file into Unicode before being sent out using the response encoding.

Unfortunately, all of these encoding meta tags might confuse the browser's display system. Although QuickFinder has already properly converted the files into a single response encoding, the browser sees the Content-Type meta tags that direct it to do something else, and becomes confused.

To solve this problem, you can create a print results template that contains a Content-Type meta tag encoding at both the top and bottom of the file, before and after the various documents are printed. All current browsers take either the first Content-Type meta tag that they encounter or the last. Constructing a print template with both satisfies all browsers.

14.3 Languages Included in the Default Templates

There are additional search and print templates for each of the following languages:

- Chinese (Traditional and Simplified)
- English
- French
- German
- Italian
- Japanese
- Korean
- Portuguese
- Russian
- Spanish

Templates are stored in *volume*:*searchroot*\TEMPLATES.

14.4 Additional Resources

The following table lists additional resources for learning more about locales, country and language codes, and encodings.

Table 14-3 *Additional Internationalization Resources*

Component	Resource Location
Language and country codes (locale)	RFC 1766 (http://www.ietf.org/rfc/rfc1766.txt)
	NOTE: Although RFC 1766 uses the hyphen character (-) to separate language and country information, QuickFinder uses the underscore character (_) in order to conform to the Java convention.
	ISO 639 (http://www.ics.uci.edu/pub/ietf/http/related/iso639.txt)
	ISO 3166 (http://www.chemie.fu-berlin.de/diverse/doc/ISO_3166.html)

Component	Resource Location
Character sets	Internet Assigned Numbers Authority (IANA) Character Set registry (http://www.iana.org/assignments/character-sets)
Unicode	Unicode Consortium home page (http://www.unicode.org)
UTF-8	UTF-8: A Transformation Format of ISO10646 (http://nis.nsf.net/internet/documents/rfc/rfc2279.txt)

Security Considerations for QuickFinder Server

15

This section provides specific instructions on how to configure and maintain QuickFinder™ Server in the most secure way possible.

- Section 15.1, “Security Features,” on page 185
- Section 15.2, “Security Characteristics,” on page 187
- Section 15.3, “Security Configuration,” on page 188
- Section 15.4, “Other Security Considerations,” on page 192

15.1 Security Features

The following table contains a summary of the security features of QuickFinder:

Table 15-1 *QuickFinder Security Features*

Feature	Yes/No	Details
Users are authenticated	Yes	Administrative users are authenticated via eDirectory™ (or PAM on Linux) and authorized access if they have write rights to the configuration file in the product directory (`SYS:\qfsearch` on NetWare® and `/var/lib/qfsearch` on Linux).
Servers, devices, and/or services are authenticated	No	

Feature	Yes/No	Details
Access to information is controlled	Yes	Access to the administrative interface is restricted to valid users that have write rights to the configuration file in the product directory. Rights-based search results can be restricted to those that have rights to view them based on the following: ◆ The files or index are identified as public. or They are a valid user. ◆ The index is specified as rights-controlled at the index level and the user has rights to read the index control file. or The index is specified as rights-controlled at the file or path level and the user has rights to read the file or the path that contains that file.
Roles are used to control access	No	
Logging and/or security auditing is done	Yes	QuickFinder keeps log files containing the logged-in users' UserIDs and the incoming IP address. However, the UserIDs are not exposed in the summary reports that are generated. Administrators can create their own exports that expose the UserIDs and IP addresses.
Data on the wire is encrypted by default	Yes	The following data is encrypted on the wire: ◆ QuickFinder administration via browser UI. ◆ When logging in (if the administrator specified switching to the HTTPS protocol). ◆ When crawling HTTPS-based Web sites. ◆ When synchronizing QuickFinder indexes, configuration settings, and templates to other QuickFinder servers in a server farm (if the administrator specified to use the HTTPS protocol). ◆ Any time the user switches the browser's URL to use the HTTPS protocol. ◆ If crawling a password-protected Web site, but not using the HTTPS protocol, then UserIDs and passwords might be passed in the clear.
Data stored is encrypted	No	

Feature	Yes/No	Details
Passwords, keys, and any other authentication materials are stored encrypted	No	QuickFinder stores the credentials needed to crawl password-protected Web sites in its configuration files. These files are stored in the product directory (which should be protected). Both the UserID and the Password are visible if using the Form-based login method when crawling a Web site. The password is not visible in the UI when using the Basic Authentication method to access password-protected Web sites.
Security is on by default	Yes	

15.2 Security Characteristics

QuickFinder Server communicates using port 80 for normal searches, port 443 for rights-based searches and to log in (controllable by the administrator), and port 2200 on NetWare for administration (also controllable by the administrator). On Linux, QuickFinder uses port 443 for administration (also controllable by the administrator). QuickFinder Server's Highlighter and Print servlets can use whatever port a URL was originally crawled on.

When crawling a Web site, QuickFinder Engine uses port 80 for most Web sites and port 443 for most HTTPS-based Web sites. However, the actual ports are controlled by each Web site administrator. If a Web site is password-protected, user credentials can be configured by the search administrator when he or she defines the indexes. These credentials are sent with the URLs requested.

When indexing a File System, the QuickFinder Server engine only indexes what it has rights to see. On NetWare, it has full access to all mounted volumes. On Linux, it has rights to only the files that the novelwww user (within the www group) has rights to see. QuickFinder cannot control what user is used to run QuickFinder Server; it simply runs with whatever user the Tomcat servlet engine was launched with. QuickFinder also adds the novelwww user to the shadow group, which allows QuickFinder Server and QuickFinder Engine to determine if a user is a valid user through PAM.

When synchronizing indexes, configuration settings, and search templates between QuickFinder servers, QuickFinder Server uses either port 80, port 443, or port 2200 when communicating, controlled by the administrator. Administrators can also optionally configure the synchronization to require administrator credentials and HTTPS communications.

Anyone that logs in as a valid user via eDirectory on NetWare and PAM (possibly eDirectory) on Linux and has write rights to the specified `qfind.cfg` file (`/var/lib/qfsearch/Sites/qfind.cfg` on Linux and `sys:/qfsearch/sites/qfind.cfg` on NetWare) can administer QuickFinder.

QuickFinder Server was originally designed to be capable of hosting search services for multiple independent organizations (possibly completely different enterprises). After an administrator has gained entry to QuickFinder Manager (through the specific virtual search server he or she has rights to), he or she then has administrative rights to all of the other virtual search servers.

On the searching side, QuickFinder Server does not usually perform rights-based searching. However, using QuickFinder Manager, administrators can configure any index to restrict access to

the search results within it (only the ability to see results on the search results pages, not access to them). After an administrator has configured an index to perform rights-based searching, approximately the same logic as for administering the product is performed: users are authenticated by eDirectory on NetWare and PAM (possibly eDirectory) on Linux and individual files are authorized based on users' individual read rights to various files in the file system.

15.3 Security Configuration

The following sections provides a summary of security-related configuration settings for QuickFinder:

- Section 15.3.1, "QuickFinder Configuration Settings," on page 188
- Section 15.3.2, "Configuration Settings for Other Products," on page 192

15.3.1 QuickFinder Configuration Settings

The following table lists the QuickFinder configuration settings that are security-related or impact the security of QuickFinder.

Table 15-2 *QuickFinder Security Configuration Settings*

Configuration Setting	Possible Value	Default Value	Recommended Value for Best Security
QFind.cfg > MsgDetail	"Level of detail in indexing logs:" -1 through 5	3	5
QFind.cfg > AutoTimeOut	"Maximum query duration (seconds)." Any integer.	30 seconds	15
QFind.cfg > CheckRights	"Authorization checking:" Off | Index | ResultItem	Off	ResultItem
QFind.cfg > CheckRightsFilePath	"by Index (Only users that have read access to the following file have access to the index)" Any string value	Points to `/docs/index.html`	A file that is more rights-controlled
QFind.cfg > CheckRightsFilter	"Unauthorized hits filtered by" Engine | Templates	Engine	Engine
QFind.cfg > UserID	"Basic Authentication: User ID" Any string value	""	If possible, crawl public only.
QFind.cfg > Password	"Basic Authentication: Password" Any string value	""	If possible, crawl public only.
QFind.cfg > AuthFields	"Form-based Authentication" fields Any string value	""	If possible, crawl public only.
QFind.cfg > LoginURL	"Alternate Login URLs" Any string value	""	If possible, crawl public only.

Configuration Setting	Possible Value	Default Value	Recommended Value for Best Security
QFind.cfg > HTTPHeaders	“HTTP Headers:” any string value	“”	“”
QFind.cfg > IndexLocation	“Location of index files:” Any directory in the file system Any string value	<qfsearch>/Sites/<*VSS name*>/indexes/<*index name*>	Any directory that is secured.
QFind.cfg > CanBeMirrored	“Index may be copied to other clustered servers:” true | false	FALSE	FALSE
AdminServlet.properties > AdminServlet.RequireSSL	“Require HTTPS when administering QuickFinder Server:” true | false	TRUE	TRUE
AdminServlet.properties > AdminServlet.Authenticate	“Require authorization when administering QuickFinder Server:” true | false	TRUE	TRUE
AdminServlet.properties > AdminServlet.ProductUpdates.Enabled	“Check for product updates:” true | false	TRUE	FALSE
AdminServlet.properties > AdminServlet.ProductUpdates.URL	Not in UI Any string value	http://search.novell.com/qfsearch/UpdateServlet	“”
???????.properties > GeneralServlet.Errors.Enabled QueryLog.Enabled ClusterServices.Log.Enabled	“Log Enabled” true | false	TRUE	TRUE
????????.properties > GeneralServlet.Errors.Destination ClusterServices.Log.Destination	“Log to:” File | Console | Both	Both	Both
???????.properties > GeneralServlet.Errors.LogFile.DeleteOnRestart ClusterServices.Log.DeleteOnRestart	“New log when services load:” true | false	TRUE	FALSE
???????.properties > GeneralServlet.Errors.LogFile.MaxSize ErrorLog.MaxSize ClusterServices.Log.MaxSize	“Maximum log size (bytes):” Any positive integer	30000	30000
GeneralServlet.properties > GeneralServlet.Mail.Enabled	“Enable e-mail services:” true | false	FALSE	TRUE

Configuration Setting	Possible Value	Default Value	Recommended Value for Best Security
GeneralServlet.properties > GeneralServlet.Mail.SMTPHost	“Outgoing SMTP Host name:” Any string value	“”	The name of the local SMTP mail server.
GeneralServlet.properties > GeneralServlet.Mail.SMTPPort	“Outgoing SMTP Port #:” Any positive integer	25	The correct port of the local SMTP server.
GeneralServlet.properties > GeneralServlet.Mail.SMTPUserID	“Outgoing SMTP User ID (optional):” Any string value	“”	The UserID of the local SMTP mail server.
GeneralServlet.properties > GeneralServlet.Mail.SMTPPassword	“Outgoing SMTP Password (optional):” Any string value	“”	The password of the local SMTP mail server.
Cluster.properties > ClusterServices.Send.Enabled	“Will this machine send cluster data:” true | false	FALSE	FALSE
Cluster.properties > ClusterServices.Receive.Enabled	“Will this machine receive cluster data:” true | false	TRUE	FALSE
Cluster.properties > ClusterServices.RequireHTTPS	“Require HTTPS for all cluster communications:” true | false	FALSE	TRUE
Cluster.properties > ClusterServices.Authentication.RequireAuthentication	“Require admin authorization when receiving cluster data:” true | false	TRUE	TRUE
SiteList.properties > SiteList.GlobalSite	“Default location of virtual search servers:” Any string value	<QFSearch>/sites	A protected file system location.
SiteList.properties > Monitor.SiteCache.DynamicUpdates	“Detect manual search server changes:” true | false	TRUE	FALSE
SiteList.properties > Monitor.TemplateCache.DynamicUpdates	“Detect template changes:” true | false	TRUE	FALSE
General.properties > Monitor.SiteCache.PurgeSiteSeconds	Not in UI Defaults to 20 minutes Any non-negative integer	1200	60
General.properties > QueryLog.Mail.Enabled	“E-mail log reports:” true | false	FALSE	FALSE
General.properties > Recipients	“...enter recipients...”	“”	“”
General.properties > Site.Cluster.Send.Enabled	“Let Virtual Search Server send cluster data:” true | false	FALSE	FALSE

Configuration Setting	Possible Value	Default Value	Recommended Value for Best Security
General.properties > Site.Cluster.Send.ClusterNames	“Name of clusters to send to:” Any string value	“”	“”
General.properties > Site.Cluster.Receive.Enabled	“Let Virtual Search Server receive cluster data:” true | false	FALSE	FALSE
General.properties > Response.PotentialHits.Max	“Refuse queries if potential hits exceed:” Any positive integer.	10000	5000
Search.properties > Print.properties > Search.Request.NumHits.Max Search.BestBet.NumHits.Max Print.Request.NumHits.Max	“Maximum number of results per page:” Any positive integer up to “Highest allowed result number:”	200 5 200	100 5 25
Search.properties > Print.properties > Search.Request.LastHitNum.Max Print.Request.LastHitNum.Max	“Highest allowed result number:” Any positive integer	1000 1000	200 200
Search.properties > Print.properties > Search.Templates.Directory Print.Templates.Directory Highlighter.Templates.Directory	“Templates directory:” Any string value	<qfsearch>/ Templates	A rights-controlled directory.
Print.properties > Print.Response.PrintSize.Warning	“Print job size warning (bytes):” Any positive integer	102400	40960
Print.properties > Print.Response.PrintSize.Max	“Maximum print job size (bytes):” Any positive integer	2097152	512500
Security.properties > Security.Authentication.RealmString	“Authentication realm string:” (Only used if LoginType is set to basic, 0) Any string value	“QuickFinder Server”	Best if it matches the Web server’s realm.
Security.properties > Security.LoginType	Not in UI 0 - login type basic 1 - login type form	1	1
Security.properties > Security.CheckRightsByDir	“Check authorization by directory:” true | false	TRUE	FALSE
Security.properties > Security.AutoLogoutTime	“Auto-logout time (minutes):” Any positive integer -1 == never log out	30	5
Security.properties > Security.RequireHTTPS	“Require https:” true | false	TRUE	TRUE

15.3.2 Configuration Settings for Other Products

The following table lists the configuration settings for other products that impact the security of QuickFinder.

Table 15-3 *Configuration Settings for Other Products*

Product Name	Configuration Setting	Default Value	Recommended Value for Best Security
PAM	Rights to the shadow group on Linux	Not made a member of this group.	Not made a member of this group, but requires that eDirectory is installed.

15.4 Other Security Considerations

- Servers should be kept in a physically secure location with access by authorized personnel only.
- The corporate network should be physically secured against eavesdropping or packet sniffing. Any packets associated with the administration of QuickFinder should be the most secured.
- Access to QuickFinder indexes, configuration settings, and logs should be restricted. This includes file system access rights, FTP access, access via Web utilities including QuickFinder Manager, and any other type of access to these files.
- Configuration settings that serve to send QuickFinder data to other servers or e-mail accounts or that protect QuickFinder data should be examined periodically to ensure that they have not been tampered with.
- When synchronizing QuickFinder indexes, configuration settings, or templates to servers outside the corporate firewall, both QuickFinder Authentication and the HTTPS protocol should be employed (see “Modifying Administrator Authentication Settings” on page 52). Because this ultimately sends the entire QuickFinder configuration of a server to another server, great security precautions should be taken.
- When QuickFinder is administered by users outside of the corporate firewall, both QuickFinder Authentication and the HTTPS protocol should be used. A VPN should also be employed.
- If a server is accessible from outside the corporate network, a firewall should be employed to prevent direct access by a would-be intruder.
- Audit logs and query reports should be kept and analyzed periodically.
- Previous versions of QuickFinder stored username and passwords in config files. After updating to QuickFinder 5.0, the username and passwords are still stored in the config files until the config file is updated. At that point, the passwords are moved to CASA. Until they are moved to CASA, users who gain access to the `qfind.cfg` file can use the specified UserIDs and passwords to access those sites themselves.
- A memory walker could discover usernames and passwords. If a user somehow has access to the server and could manipulate the memory, he or she could possibly get the passwords because passwords are not immediately cleared from memory or obfuscated in any way.
- Usernames and Passwords are stored in CASA by using the Tomcat user. Other applications running in Tomcat, or users that know the Tomcat user credentials, could potentially get the usernames and passwords for the remote servers. In order to take advantage of this, some

application would have to be installed as an Apache application and then specifically call CASA with the correct ID to get the passwords.

- Although a QuickFinder Manager administrator might not have access rights to the entire server, QuickFinder Engine's File System repository and indexer might. The QuickFinder administrator can generate an index of the entire server, then see the first 255 bytes of the file (search descriptions).
- Users might try to gain access to the index files. If a very sophisticated user properly decompiles the indexes, he or she might be able to discover the structure of the indexed files and rebuild the content of the files. An administrator should know where the index files are stored and what permissions are granted to those index files. An administrator should also be aware that an index file stored on a file system that is actually a remote mount might expose the index files to the security vulnerabilities of that remote mounted file system. For example, if the remotely mounted file system uses unsecured NFS, a network listener could listen for the packets on the network and reassemble the index or simply do a remote mount as the Apache user with no authentication required. Even users who cannot rebuild the contents of the files can still see the first 255 relevant text bytes by simply performing searches and reading the descriptions. To the best of our knowledge, even if users have access to the index (and the index format spec), they would still have a very difficult time rebuilding the file contents.

Troubleshooting QuickFinder

A

This section provides some troubleshooting topics that can help you resolve search and print performance issues. It offers possible causes for issues you might encounter, as well as suggested actions for resolving them.

If the problem you are working with doesn't appear in this appendix, visit the Novell® Support Web site (http://www.novell.com/support).

- "Characters of descriptions or titles appear as unintelligible characters" on page 195
- "Several titles or descriptions contain the same text" on page 195
- "Some titles are returned as the URL of the document instead" on page 195

Characters of descriptions or titles appear as unintelligible characters

Possible Cause: You've probably indexed documents written in multiple languages and encodings. QuickFinder can index most of the world's languages and encodings. However, QuickFinder™ needs to know the encoding of each document.

Possible Cause: Some of your documents were probably not tagged with an encoding or were incorrectly tagged.

Action: Make sure all of your documents contain the correct Content-Type meta tag. If your international documents do not contain a Content-Type meta tag, either add it or use the *Encoding (If Not in META Tags)* index definition option to specify the default encoding.

For more information, see Chapter 14, "Internationalizing Search Services," on page 177

Several titles or descriptions contain the same text

Possible Cause: If search results include duplicate titles or descriptions, your description fields (description, summary, or abstract) might include boilerplate information.

Action: The more accurate your meta tag description fields are, the better your search results are. Where possible, consider adding descriptions to your document's meta tags.

Possible Cause: It could also be that you have indexed the same document more than once, or several links throughout your Web site might point to the same document but do so using different character cases each time.

Action: To solve the latter problem, try using the *URLs Are Case Sensitive* option to direct QuickFinder to turn off case-sensitive crawling. Also, remove any duplicate backup files you might have and exclude any backup directories from your index definition.

Some titles are returned as the URL of the document instead

Possible Cause: QuickFinder pulls document titles from within each document that it indexes. If your document doesn't have a title, QuickFinder uses the URL or path of the

document instead. If the URL is unavailable, a Title Unavailable message is returned.

Action: Make sure all of the documents you index have specifically defined titles.

Combined Character Sets for Use with QuickFinder

B

The following tables list the character set encoding names and aliases that QuickFinder™ recognizes when indexing, searching, displaying, or printing files. This information is a subset of the character names registered by the Internet Assigned Numbers Authority (IANA).

Whenever possible, the items listed in the first column of each table are the preferred MIME names listed in the Internet Assigned Numbers Authority (IANA) Character Sets registry. If a preferred MIME name is not available, items in the first column represent the primary registered names.

Items in the second column of each table are aliases that are sometimes used to identify that encoding.

Not all aliases exactly represent the parent encoding under which they are listed. In these cases, they overlap significantly enough to be handled identically by the various QuickFinder engines.

TIP: Character encodings appear in the exact case specified in the Internet Assigned Numbers Authority (IANA) Character Sets registry (http://www.iana.org/assignments/character-sets). Some uses of these encodings are case sensitive. However, QuickFinder ignores the case of these encodings.

- Section B.1, “ASCII Character Set,” on page 198
- Section B.2, “Arabic Character Set,” on page 199
- Section B.3, “Chinese (Simplified) Character Set,” on page 199
- Section B.4, “Chinese (Traditional) Character Set,” on page 200
- Section B.5, “Cyrillic Character Set,” on page 201
- Section B.6, “European Character Set,” on page 202
- Section B.7, “Greek Character Set,” on page 205
- Section B.8, “Hebrew Character Set,” on page 206
- Section B.9, “Japanese Character Set,” on page 206
- Section B.10, “Korean Character Set,” on page 207
- Section B.11, “Thai Character Set,” on page 208
- Section B.12, “Turkish Character Set,” on page 209
- Section B.13, “Vietnamese Character Set,” on page 209

B.1 ASCII Character Set

Table B-1 *ASCII Character Set*

Preferred MIME Name or Primary Registered Name	Encoding Names
US-ASCII (MIBenum: 3)	ANSI_X3.4-1968
A MIBenum is a record number corresponding to an entry in IANA's Management Information Base.	ANSI_X3.4-1986
	ASCII
	ascii7
	iso_646-us
	ISO646-US
	ISO_646.irv:1991
	iso-ir-6
	646
	us
	IBM367
	cp367
	csASCII
IBM437 (MIBenum: 2011)	ibm-437
	cp437
	437
	csPC8CodePage437

B.2 Arabic Character Set

Table B-2 *Arabic Character Set*

Preferred MIME Name or Primary Registered Name	Encoding Aliases
ISO-8859-6 (MIBenum: 9)	ISO_8859-6:1987
A MIBenum is a record number corresponding to an entry in IANA's Management Information Base.	ISO_8859-6
	iso8859-6
	iso8859_6
	8859_6
	IBM1089
	ibm-1089
	cp1089
	1089
	iso-ir-127
	ECMA-114
	ASMO-708
	arabic
	csISOLatinArabic
Windows-1256 (MIBenum: 2256)	cp1256
	win1256
	ms1256

B.3 Chinese (Simplified) Character Set

Table B-3 *Chinese (Simplified) Character Set*

Preferred MIME Name or Primary Registered Name	Encoding Aliases
gb2312 (MIBenum: 2025)	csGB2312
gb_2312-80 (MIBenum: 57)	iso-ir-58
A MIBenum is a record number corresponding to an entry in IANA's Management Information Base.	chinese
	csISO58GB231280
	gb2312-80
	gb2312-1980
	gb-2312-80

Preferred MIME Name or Primary Registered Name	Encoding Aliases
gbk	GBK
	windows-936
	ms936
	cp936
	cp-936
euc-cn	EUC_CN
	euccn
	euc-gb

B.4 Chinese (Traditional) Character Set

Table B-4 *Chinese (Traditional) Character Set*

Preferred MIME Name or Primary Registered Name	Encoding Aliases
big5 (MIBenum: 2026)	Big5
A MIBenum is a record number corresponding to an entry in IANA's Management Information Base.	windows-950
	win950
	ms950
	csBig5
IBM950	ibm-950
	cp950
	cp-950
	950

B.5 Cyrillic Character Set

Table B-5 *Cyrillic Character Set*

Preferred MIME Name or Primary Registered Name	Encoding Aliases
ISO-8859-5 (MIBenum: 8)	ISO_8859-5:1988
A MIBenum is a record number corresponding to an entry in IANA's Management Information Base.	ISO_8859-5
	iso8859-5
	iso8859_5
	8859-5
	iso-ir-144
	IBM915
	ibm-915
	cp915
	915
	cyrillic
	csISOLatinCyrillic
KOI8-R (MIBenum: 2084)	koi8_r
	koi8
	cp878
	cp-878
	csKOI8R
Windows-1251 (MIBenum: 2251)	win1251
	cp1251
	ms1251

B.6 European Character Set

Table B-6 *European Character Set*

Preferred MIME Name or Primary Registered Name	Encoding Aliases
Windows-1252 (MIBenum: 2252)	cp1252
A MIBenum is a record number corresponding to an entry in IANA's Management Information Base.	ms1252
	win1252
	ansi
	ansi-1252
ISO-8859-1 (MIBenum: 4)	ISO_8859-1:1987
	ISO_8859-1
	iso8859-1
	iso8859_1
	8859_1
	iso-ir-100
	IBM819
	ibm-819
	CP819
	819
	l1
	latin1
	csISOLatin1

Preferred MIME Name or Primary Registered Name	Encoding Aliases
ISO-8859-2 (MIBenum: 5)	ISO_8859-2:1987
A MIBenum is a record number corresponding to an entry in IANA's Management Information Base.	ISO_8859-2
	iso8859-2
	iso8859_2
	8859_2
	iso-ir-101
	IBM912
	ibm-912
	cp912
	912
	l2
	latin2
	csISOLatin2
ISO-8859-3 (MIBenum: 6)	ISO_8859-3:1988
	ISO_8859-3
	iso8859-3
	iso8859_3
	8859-3
	iso-ir-109
	IBM913
	ibm-913
	cp913
	913
	l3
	latin3
	csISOLatin3

Preferred MIME Name or Primary Registered Name	Encoding Aliases
ISO-8859-4 (MIBenum: 7)	ISO_8859-4:1988
	ISO_8859-4
	iso8859-4
	iso8859_4
	8859-4
	iso-ir-110
	IBM914
	ibm-914
	cp914
	914
	l4
	latin4
	csISOLatin4
Windows-1250 (MIBenum: 2250)	cp1250
	ms1250
	win1250
IBM850 (MIBenum: 2009) (UNICODE)	ibm-850
	cp850
	850
	csPC850Multilingual
IBM852 (MIBenum: 2010)	ibm-852
	cp852
	852
	csPCp852
IBM860 (MIBenum: 2048)	ibm-860
	cp860
	860
	csIBM860
IBM863 (MIBenum: 2050)	ibm-863
	cp863
	863
	csIBM863

Preferred MIME Name or Primary Registered Name	Encoding Aliases
IBM865 (MIBenum: 2052)	ibm-865
	cp865
	865
	csIBM865

B.7 Greek Character Set

Table B-7 *Greek Character Set*

Preferred MIME Name or Primary Registered Name	Encoding Aliases
ISO-8859-7 (MIBenum: 10)	ISO_8859-7:1987
A MIBenum is a record number corresponding to an entry in IANA's Management Information Base.	ISO_8859-7
	iso8859-7
	8859_7
	IBM813
	ibm-813
	cp813
	813
	iso-ir-126
	ELOT_928
	ECMA-118
	greek
	greek8
	csISOLatinGreek
Windows-1253 (MIBenum: 2253)	cp1253
	ms1253
	win1253

B.8 Hebrew Character Set

Table B-8 *Hebrew Character Set*

Preferred MIME Name or Primary Registered Name	Encoding Aliases
ISO-8859-8 (MIBenum: 11)	ISO_8859-8:1988
A MIBenum is a record number corresponding to an entry in IANA's Management Information Base.	ISO_8859-8
	iso8859-8
	8859_8
	ibm916
	ibm-916
	cp916
	916
	iso-ir-138
	hebrew
	csISOLatinHebrew
Windows-1255 (MIBenum: 2255)	win1255
	cp1255
	ms1255

B.9 Japanese Character Set

Table B-9 *Japanese Character Set*

Preferred MIME Name or Primary Registered Name	Encoding Aliases
ISO-2022-JP (MIBenum: 39)	iso2022-jp
A MIBenum is a record number corresponding to an entry in IANA's Management Information Base.	iso-2022-jis
	junet
	jis
	jis_encoding
	csJISEncoding
	csISO2022JP
ISO-2022-JP-2 (MIBenum: 40)	iso-2022-jp2
	csISO2022JP2

Preferred MIME Name or Primary Registered Name	Encoding Aliases
Shift_JIS (MIBenum: 17/2024)	sjis
	shift-jis
	ShiftJis
	x-sjis
	x-shift-jis
	windows-31j
	csWindows31J
	ms932
	cp932
	win932
	windows-932
	MS_Kanji
	csShiftJIS
	pck
	\u30b7\u30d5\u30c8\u7b26\u53f7\u5316\u8868\u73fe
EUC-JP (MIBenum: 18)	Extended_UNIX_Code_Packed_Format_for_Japanese
	eucjp
	x-euc-jp
	euc_jpnew 10/18/99
	x-eucjp
	eucjis
	csEUCPkdFmtJapanese

B.10 Korean Character Set

Table B-10 *Korean Character Set*

Preferred MIME Name or Primary Registered Name	Encoding Aliases
euc-kr (MIBenum: 38)	euc_kr
A MIBenum is a record number corresponding to an entry in IANA's Management Information Base.	euckr
	csEUCKR

Preferred MIME Name or Primary Registered Name	Encoding Aliases
ks_c_5601-1987 (MIBenum: 36)	ks_c_5601-1989
	ksc5601-1987
	ksc5601_1987
	ksc_5601
	ksc5601
	5601
	korean
	csKSC56011987
IBM949	ibm-949
	cp949
	cp-949
	949
Windows-949	win949
	ms949

B.11 Thai Character Set

Table B-11 *Thai Character Set*

Preferred MIME Name or Primary Registered Name	Encoding Aliases
IBM874	ibm-874
	cp874
	874
Windows-874	win874
	ms874

B.12 Turkish Character Set

Table B-12 *Turkish Character Set*

Preferred MIME Name or Primary Registered Name	Encoding Aliases
ISO-8859-9 (MIBenum: 12) A MIBenum is a record number corresponding to an entry in IANA's Management Information Base.	ISO_8859-9:1989 ISO_8859-9 iso8859-9 8859_9 ibm920 ibm-920 cp920 920 iso-ir-148 l5 latin5 csISOLatin5
Windows-1254 (MIBenum: 2254)	win1254 cp1254 ms1254

B.13 Vietnamese Character Set

Table B-13 *Vietnamese Character Set*

Preferred MIME Name or Primary Registered Name	Encoding Aliases
Windows-1258 (MIBenum: 2258) A MIBenum is a record number corresponding to an entry in IANA's Management Information Base.	win1258 ms1258 cp1258 cp-1258

Disabling QuickFinder Server

C

This section contains information on disabling QuickFinder™ Server on Novell® Open Enterprise Server (OES) 2 for Linux.

- Section C.1, “Disabling QuickFinder Server,” on page 211
- Section C.2, “Consequences of Disabling QuickFinder Server,” on page 211

C.1 Disabling QuickFinder Server

1. Delete the symbolic link `/var/opt/novell/tomcat5/webapps/qfsearch`.
2. Restart Tomcat.

   This removes the QuickFinder registration with Tomcat, and QuickFinder will no longer load.

C.2 Consequences of Disabling QuickFinder Server

- Disabling QuickFinder Server does not effect any OES 2 service except those using search.
- Clicking the QuickFinder link in iManager results in an error.