BEA AquaLogic® Interaction

Administrator Guide

Version 6.1 MP1
Document Revised: December 14, 2007

Table of Contents

1. Welcome

2. Overview of Portal Administration

3. Managing Portal Users and Groups

4. Managing Portal Content

5. Automating Administrative Tasks

6. Migrating, Backing Up, and Restoring Portal Objects

A. Configuring Advanced Properties and Logging

B. Using the Counter Monitoring System

C. Localizing Your Portal

D. Deploying Single Sign-On

E. Default Behavior of Search Service

F. Common Questions and Answers

Index

Welcome

This book describes how to perform initial tasks required to prepare your portal for use and how to perform ongoing portal management tasks.

When you first deploy your portal, you can use this book to perform the following set-up tasks:

1. Configure display, navigation, and branding for the default experience definition and any additional experience definitions.
2. Change the default Administrator password and delegate administrator roles.
3. Populate the portal with administrative users and browsing users. Configure groups, users, user profiles, and Access Control Lists (ACLs) to enable managed access.
4. Populate the portal with documents. Configure ACLs to manage access.
5. Set up automated system maintenance, such as user synchronization, search updates, document refresh, and housekeeping jobs.

When you have completed your initial portal deployment, you can use this guide as a reference for tasks you repeat or tasks that do not apply during initial deployment but might apply later when you extend your base portal deployment to include users from new authentication sources, new content types, documents from new content sources, or search among federated portals.

The appendixes in this guide provide reference information on topics that might apply to your portal but do not apply to all portal deployments. These topics describe localization, single sign-on (SSO), and advanced configuration file settings.

How to Use This Book

Audience

This book is written for portal administrators who are responsible for managing portal users and documents, as well as for system administrators who are responsible for maintaining the portal hardware, integrating the portal with back-end systems (for example, single sign-on), and other advanced configuration.

Typographical Conventions

This book uses the following typographical conventions.

Table 1-1 Typographical Conventions

Convention	Typeface	Examples/Notes
• File names • Folder names • Screen elements	**bold**	• Upload **Procedures.doc** to the portal. • The log files are stored in the **logs** folder. • To save your changes, click **Apply Changes**.
• Text you enter	`computer`	Type `Marketing` as the name of your community.
• Variables you enter	`computer with angle brackets (<>)`	Enter the base URL for the Remote Server. For example, `http://<my_computer>/`.
• New terms • Emphasis • Object example names	*italic*	• *Portlets* are Web tools embedded in your portal. • The URI *must* be a unique number. • The example Knowledge Directory displayed in Figure 5 shows the *Human Resources* folder.

BEA Documentation and Resources

The tables in this section list other documentation and resources provided by BEA.

Table 1-2 Documentation

Resource	Description
Installation and Upgrade Guide	This book is written for system administrators responsible for installing or upgrading AquaLogic Interaction. It is available on edocs.bea.com/alui/ali/docs61.
Release Notes	These files are written for AquaLogic Interaction administrators. They include information about new features and known issues in the release. They are available on edocs.bea.com/alui/ali/docs61 and on any physical media provided for delivering the application.
Online Help	The online help is written for all levels of AquaLogic Interaction users. It describes the user interface for AquaLogic Interaction and gives detailed instructions for completing tasks in AquaLogic Interaction. To access online help, click the help icon.
Deployment Guide	This document is written for business analysts and system administrators. It describes how to plan your AquaLogic User Interaction deployment. It is available on edocs.bea.com/alui/deployment/index.html.

Table 1-3 Other BEA Resources

Resource	Description
Developer Guides, Articles, API Documentation, Blogs, Newsgroups, and Sample Code	These resources are provided for developers on the BEA dev2dev site (dev2dev.bea.com). They describe how to build custom applications using AquaLogic User Interaction and how to customize AquaLogic User Interaction products and features.
AquaLogic User Interaction Support Center	The AquaLogic User Interaction Support Center is a comprehensive repository for technical information on AquaLogic User Interaction products. From the Support Center, you can access products and documentation, search knowledge base articles, read the latest news and information, participate in a support community, get training, and find tools to meet most of your AquaLogic User Interaction-related needs. The Support Center encompasses the following communities: **Technical Support Center** Submit and track support incidents and feature requests, search the knowledge base, access documentation, and download service packs and hotfixes. **User Group** Visit the User Group section to collaborate with peers and view upcoming meetings. **Product Center** Download products, read Release Notes, access recent product documentation, and view interoperability information. **Developer Center** Download developer tools and documentation, get help with your development project, and interact with other developers via BEA's dev2dev Newsgroups. **Education Services** Find information about available training courses, purchase training credits, and register for upcoming classes. If you do not see the Support Center when you log in to http://support.plumtree.com, contact ALUIsupport@bea.com for the appropriate access privileges.

Table 1-3 Other BEA Resources

Resource	Description
Technical Support	If you cannot resolve an issue using the above resources, BEA Technical Support is happy to assist. Our staff is available 24 hours a day, 7 days a week to handle all your technical support needs. E-mail: ALUIsupport@bea.com Phone Numbers: U.S.A. +1 866.262.7586 or +1 415.263.1696 Europe +44 1494 559127 Australia/NZ +61 2.9923.4030 Asia Pacific +61 2.9931.7822 Singapore +1 800.1811.202

Overview of Portal Administration

This chapter provides an overview of the portal and the administrative tasks you perform to manage portal users and documents. It includes the following topics:

- Overview of Portal Components
- Overview of Web Services Architecture
- Overview of Portal Security
- Overview of the Portal Browsing User Interface
- Overview of the Portal Administrative User Interface and Tools

Overview of Portal Components

The following table describes the portal components you must install and configure before you can use the procedures provided in this guide to manage portal users and documents.

Table 2-1 Summary of Portal Components

Component	Description
Portal Database	Stores portal objects, such as user and group configurations, document records, and administrative objects. The portal database does not store the documents available through your portal. Source documents are left in their original locations. For information on setting up the portal database, refer to the *Installation and Upgrade Guide for AquaLogic Interaction*.
Administrative Portal	Handles portal setup, configuration, and content. The Administrative Portal enables administrative functions, such as creating and managing portlets and other Web services. For information on installing the Administrative Portal, refer to the *Installation and Upgrade Guide for AquaLogic Interaction*.
Portal	Serves end user portal pages and content. The portal enables end users to access portal content via My Pages, community pages, the Knowledge Directory, and search. The portal also enables some administrative actions, such as setting preferences on portlets or managing communities. For information on installing the portal, refer to the *Installation and Upgrade Guide for AquaLogic Interaction*. For information on advanced portal configuration, see Appendix A, “Configuring Advanced Properties and Logging.”
Image Service	Serves static content used or created by portal components. The Image Service serves images and other static content for use by the AquaLogic User Interaction system. For information on installing the Image Service, refer to the *Installation and Upgrade Guide for AquaLogic Interaction*. Whenever you extend the base portal deployment to include additional components, such as portal servers or integration products, you may have to install additional Image Service files. For information on installing the Image Service files for those components, refer to the documentation included with the component software.

Table 2-1 Summary of Portal Components

Component	Description
Search Service	Returns indexed content stored in the portal. The Search Service returns content that is indexed in the AquaLogic User Interaction system from the portal, Collaboration, and Publisher. Content that is indexed in the AquaLogic User Interaction system includes documents, portlets, communities, and users as well as many other AquaLogic User Interaction objects. For information on installing the Search Service, refer to the *Installation and Upgrade Guide for AquaLogic Interaction*. For information on advanced Search Service configuration, see Appendix A, “Configuring Advanced Properties and Logging.”
Automation Service	Runs jobs and other automated portal tasks. The Automation Service runs jobs that perform tasks such as crawling documents into the Knowledge Directory, synchronizing groups and users with external authentication sources, and maintaining the search collection. For information on installing the Automation Service, refer to the *Installation and Upgrade Guide for AquaLogic Interaction*. For information on configuring Automation Service jobs, see Chapter 5, “Automating Administrative Tasks.”
Document Repository Service	Stores documents uploaded by portal components. The Document Repository Service stores content uploaded into the portal, such as images or documents uploaded into Collaboration or Publisher. For information on installing the Document Repository Service, refer to the *Installation and Upgrade Guide for AquaLogic Interaction*. For information on managing the portal Knowledge Directory that resides on the Document Repository Service, see “Configuring Content Sources” on page 4-10
API Service	Provides access to the SOAP API. For information on installing the API Service, refer to the *Installation and Upgrade Guide for AquaLogic Interaction*.
Web Services	• World Wide Web (WWW) Content Service The WWW Content Service enables the manual upload (or automatic crawling) of document records into the Knowledge Directory from WWW locations. • Content Upload Service The Content Upload Service enables the manual upload (or automatic crawling) of document records from an internal network. It may be particularly useful if some users do not typically have access to the internal network (for example, if you are running an extranet).

Table 2-1 Summary of Portal Components

Component	Description
AquaLogic Interaction Integration Services	AquaLogic Interaction Integration Services enable integration with third-party applications/repositories. AquaLogic Interaction Integration Services include a set of Identity Services and Content Services: • AquaLogic Interaction Identity Services – Active Directory - Enables the authentication and synchronization of users between Microsoft's Active Directory (AD) and the portal. – LDAP - Enables the authentication and synchronization of users between LDAP (Lightweight Directory Protocol) and the portal. • AquaLogic Interaction Content Services Content Services scan third-party systems/applications for new content, categorizing links to third-party content in the organized, searchable structure of the portal's Knowledge Directory. This allows customers to streamline costs by publishing and leveraging third-party content within the portal. The following sources can be scanned with AquaLogic Interaction Content Services: – Windows File Systems – Documentum – Lotus Notes – Microsoft Exchange • IDK and AquaLogic Interaction Service Station Use these to develop and test custom integration services.
Portlets	Portlets provide portal users customized tools and services as well as information. Portlets allow you to integrate applications, tools, and services into your portal, while taking advantage of portal security, caching, and customization. For information on portlets, see "Extending Portal Services with Portlets" on page 4-18. AquaLogic Interaction provides portlet software packages that integrate back-end content sources with the portal. For information on these portlets, visit the AquaLogic User Interaction Support Center.

Table 2-1 Summary of Portal Components

Component	Description
Remote Servers	Remote servers host Web services that enable portal access to back-end, or remote, content sources. Remote servers do not need to be visible from beyond your firewall. The portal can function as a gateway to the content on the remote servers. For information on remote servers included with AquaLogic Interaction Integration Products, visit the AquaLogic User Interaction Support Center.
AquaLogic Interaction Activity Servers	AquaLogic Interaction Collaboration, Publisher, and Studio extend portal functionality to enable collaboration and branding. For information on these products, visit the AquaLogic User Interaction Support Center.

Figure 2-1 provides a high-level summary of the portal components and content sources that portal administrators manage to enable the portal user experience. All components except for the back-end sources are summarized in the preceding table and described in detail in this guide. For information on installing and using the back-end sources, refer to documentation included with the component software.

Figure 2-1 Summary of Portal Components

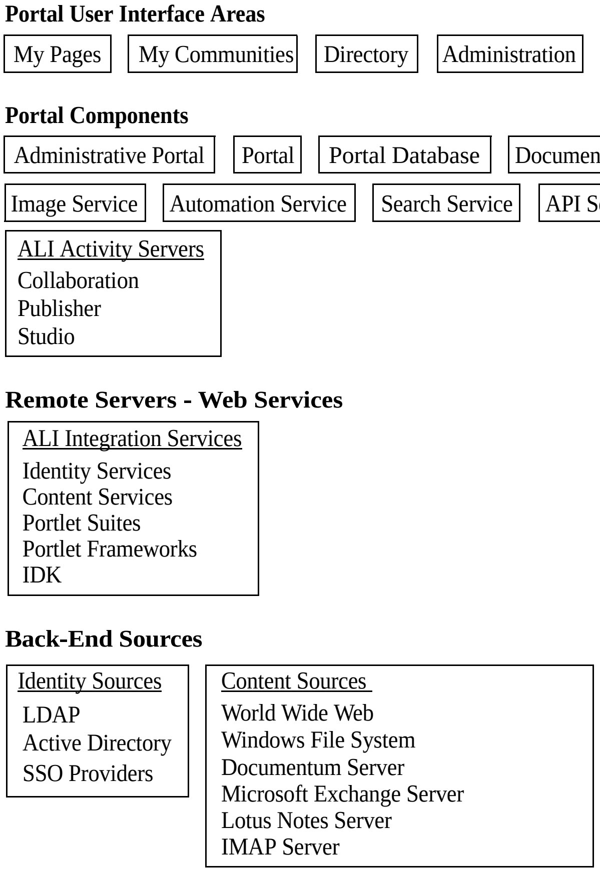

Overview of Web Services Architecture

Many of the objects in the portal utilize *Web services*, which are components that run on a logically separate computer from the one that runs the portal and communicate with the portal via HTTP. We refer to this separate computer as a *remote server*. The Web services architecture allows multiple types of remote services (identity services, content services, portlets) to share a logical remote server, making it easier to manage the computers that make up the portal.

Figure 2-2 shows how multiple portal objects can be set up with the same Web service.

Figure 2-2 Web Service Architecture

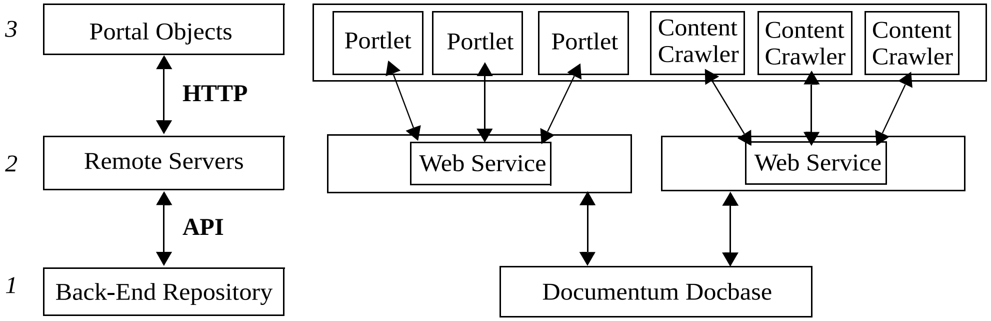

In the example:

1. The back-end repository contains documents stored in a remote location—in this example, a Documentum Docbase.
2. The remote server computer hosts the Web service software—in this example, AquaLogic Interaction software.
3. In the portal, you register the remote server and Web service by importing a migration package (**.pte** file) that defines the base URL and credentials for the remote server and configuration settings for the Web service, such as connection information for the back-end repository. In this case, you register the Web service by configuring connection information for the Documentum Docbase.

   In addition, you create the specific objects that use the Web service by configuring additional administrative preferences and user preferences. In this case, you create multiple portlet and content crawler objects by configuring portlet-specific and content crawler-specific settings.

As you see in Figure 2-2, Web services allow you to share settings (sometimes rather complex settings) among the objects created from those services. In the case of portlets and content crawlers, Web services allow you to create templates that business users can customize to implement specific functionality based on specific preferences.

In addition, Web services enable you to create composite applications that utilize functionality from multiple Web services. For example, you might have several Web services accessing an application that requires user credentials. Rather than creating a separate configuration page for each Web service and requiring users to specify the same information multiple times, you can create a link to these shared settings, allowing users to specify the information only once for all of these Web services.

For more information on configuring administrative and user preferences for the objects that use Web services, refer to the BEA AquaLogic User Interaction Development Center (http://dev2dev.bea.com/aluserinteraction/).

For information on settings specific to a particular kind of Web service described in this guide, refer to the documentation for the particular Web service.

Overview of Portal Security

AquaLogic Interaction provides many ways to secure your portal and its content:

- Activity security in the form of activity rights. See “Delegating Activity Rights” on page 3-4.
- Audit records, which you should periodically review to keep track of actions performed by users. See “Auditing User Accounts and Actions” on page 3-19.
- Automatic user lockout. See “Locking and Unlocking User Accounts” on page 3-21.
- Web application credential management in the form of lockboxes. See “Managing User Credentials for Existing Applications” on page 3-23.
- Object level security in the form of Access Control Lists (ACLs). See “Setting User Access Privileges” on page 3-24.
- Document security imported from source repositories. See “Enabling Document Discovery with Content Crawlers and Content Services” on page 4-40.
- Single sign-on. See Appendix D, “Deploying Single Sign-On.”

In addition to the security available through the portal, you must also secure your hardware and back-end systems (for example, your portal and user databases) to fully protect your portal. You should follow all security guidance provided in your hardware and software documentation.

You must also create strong passwords not only for administrators, but for all portal users and you must advise everyone to keep their passwords safe.

Overview of the Portal Browsing User Interface

This section describes the user interface for portal browsing users.

The following topics provide an overview of basic portal user interface features and options:

- "Portal Areas" on page 2-9
- "Display Options" on page 2-10
- "Navigation Options" on page 2-12
- "Branding" on page 2-17
- "Locale Settings" on page 2-18

You define your user interface by configuring the default *experience definition* and any additional experience definitions. An experience definition is a portal object that contains a collection of settings that determine what users see when they use the portal. Users can be associated with many experience definitions. For information on configuring experience definitions, see "Configuring Experience Definitions" on page 2-18.

Portal Areas

Figure 2-3 shows the main portal areas, which are summarized in this section.

Figure 2-3 Main Portal Areas

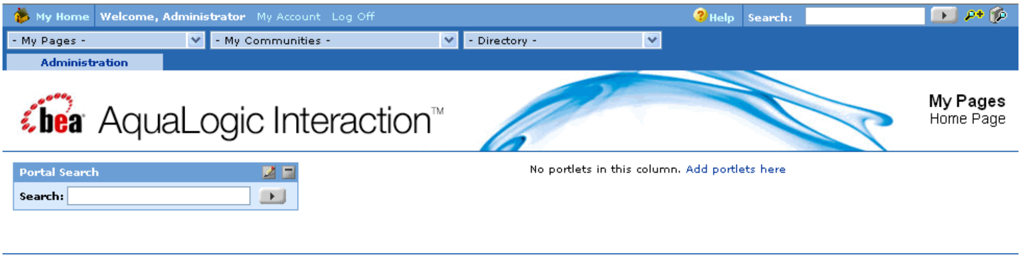

No portlets in this column. Add portlets here

The following table summarizes the purpose of the main portal areas.

Table 2-2 Main Portal Areas

Main Area	Purpose
My Pages	My Pages display a personalized view of the portal. Users can customize their My Pages by changing their My Account preferences and adding portlets. Users can create up to six My Pages to organize their content and services.
My Communities	While My Pages offer users a personalized view of the portal, communities offer a view of the portal that is shared by many users with a common interest. Any user with the Create Community activity right can create a community and populate it with portlets that display content relevant to the community or that allow members of the community to work together.
Directory	Your portal Knowledge Directory displays links to documents. These documents can be external or internal Web pages, Microsoft Office documents, or any files of interest. The Knowledge Directory gives structure for these links with folders and subfolders, much like a file system. This structure makes it easy for people to browse your portal when they are not looking for a particular document, but want to see what is available on a given subject.
Administration	The Administration area provides access to the Administrative Object Directory, which stores portal objects (such as content crawlers, portlets, and users), and access to portal utilities (such as server configuration utilities). Depending on the permissions granted to portal administrators, they see different menu items in the Administration area. Information on the different objects and utilities available through Administration is provided throughout this guide.

Display Options

The following table summarizes portal display options.

Table 2-3 Display Options

Option	Description
Standard Portal	The fully-featured user interface for the AquaLogic Interaction software. Use it to provide the richest user interface experience for internal and external users. This version does not support assistive technologies.
Assistive Technology Portal	Designed for people with disabilities. It supports only portlets that meet requirements for use with assistive technologies. Section 508 of the Rehabilitation Act is a federal statute requiring federal agencies' electronic and information technology to be accessible to people with disabilities, including employees and members of the public. The federal criteria for Web-based technology are based on access guidelines developed by the Web Accessibility Initiative of the World Wide Web Consortium (W3C). Designed to adhere to the federal criteria for Web-based technology, the Assistive Technology Portal allows end users with visual disabilities to access the portal through assistive browsing technologies, such as screen readers, screen magnifiers, voice recognition and Braille devices. The interface is text-based with a linear presentation of information, and with no embedded client-side JavaScript or Java applets.
Low Bandwidth Portal	Accommodates users with slower Internet connections. This version supports all AquaLogic User Interaction portlets, but does not support assistive technologies. Remote users have two options for viewing portal pages—the standard version and the Low Bandwidth version. Users can switch from one version to the other during a portal session and the change occurs immediately. The Low Bandwidth Portal provides better performance for end users accessing the portal remotely when network performance is slow due to low bandwidth or heavy traffic. This version presents a user interface with far fewer graphics and no embedded JavaScript or Java applets.

Navigation Options

The portal includes navigation schemes that allow you to select the menu layout and core navigation structure most appropriate for your bandwidth constraints, browser requirements, design needs, deployment size, and end-user expectations. You can also create your own navigation schemes, using the existing code as a starting point. For information on customizing navigation and other user interface elements, see the BEA AquaLogic User Interaction Development Center (http://dev2dev.bea.com/aluserinteraction/).

The navigation schemes included with the portal can be divided into horizontal and vertical groups, based on the alignment of the navigational elements. In horizontal navigation, links to My Pages, communities, the Knowledge Directory, Administration, and any mandatory links you specify appear at the top of the page in drop-down menus, maximizing the space available for portlets. In vertical navigation, links appear on the left side of the screen.

Any navigation scheme (except the No Navigation scheme) can include mandatory links to Web sites, user profiles of portal experts, documents from the portal Knowledge Directory, and pages in communities. These links display in the navigation scheme under a category (like My Pages, My Communities, or Directory) with the name of your choosing. You might want to use these links to promote new portlets, communities, or important documents.

The following navigation options are specified in experience definitions.

Horizontal Combo Box Drop-Down Navigation

This navigation scheme uses standard HTML controls to place navigational elements in drop-down menus. Because it does not use JavaScript for rendering menus, this option is bandwidth-efficient.

Figure 2-4 Horizontal Combo Box Drop-Down Navigation

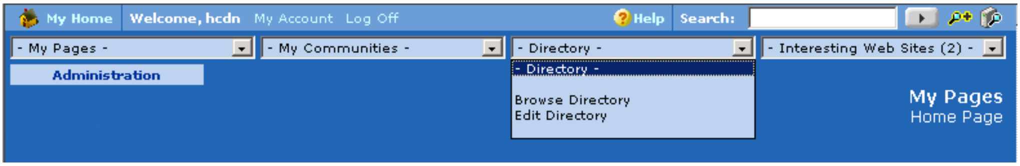

Tabbed Section Left Vertical Navigation

This navigation scheme uses horizontal tabs at the top for the main portal areas, which, when clicked, display links on the left to the options available within that portal area. This scheme is similar to the navigation for sites such as Amazon.com and MSN.

Figure 2-5 Tabbed Section Left Vertical Navigation

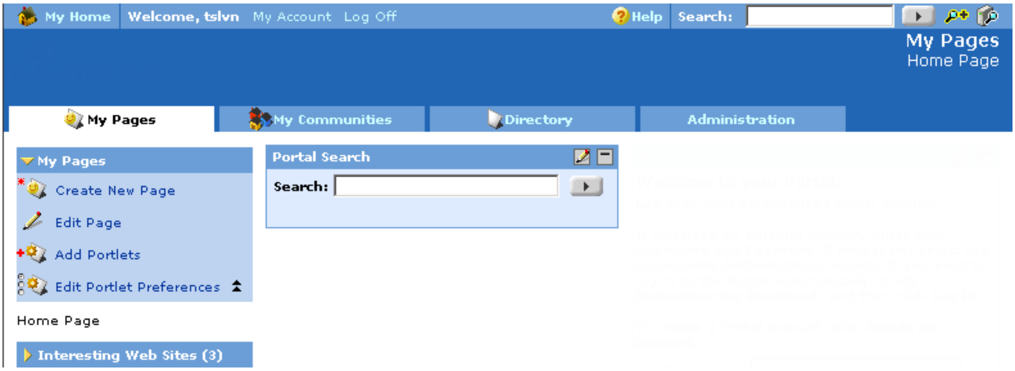

Left Vertical Navigation

This navigation scheme lists all available links unless the user minimizes particular elements. It is very easy to use, because users see all links without additional clicks. Because it does not use JavaScript for rendering menus, this option is bandwidth-efficient. However, if users join a large number of communities, they have to scroll to see some of the links.

Figure 2-6 Left Vertical Navigation

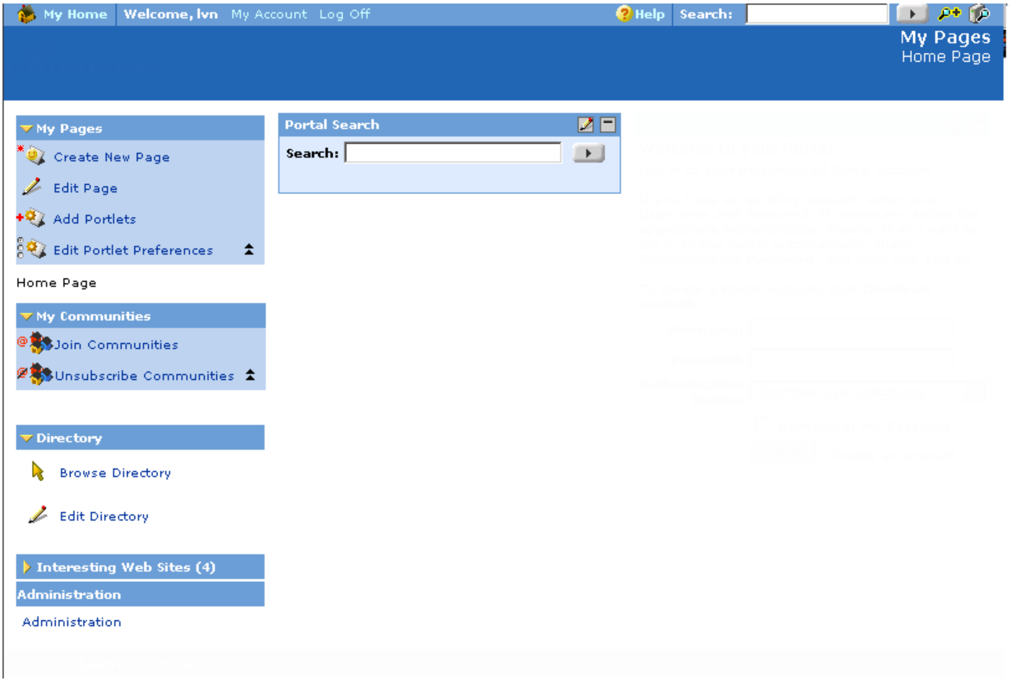

Mandatory Links Only

This navigation scheme displays only the mandatory links (which you specify in the experience definition) using the same menu style used in Horizontal Drop-Down Navigation. Users can see only their home page (the page that displays when they log in) and any areas for which you have created mandatory links. However, they can still access documents through search and might be able to access other areas if those areas are available through portlets. You might use this scheme if you want to severely limit portal access to users. For example, you might want a group of customers to access only a particular community to learn about a new product.

Figure 2-7 Mandatory Links Only

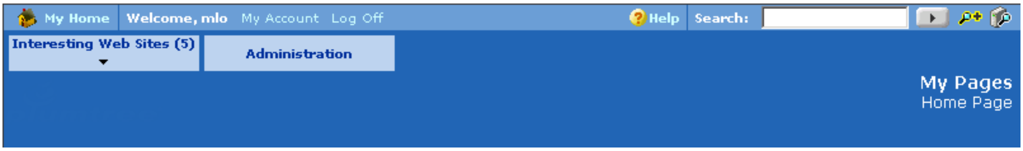

No Navigation

This navigation scheme displays no navigation, but includes the top bar. However, there is a link to Administration if the user has access. As with the Mandatory Links Only navigation scheme, users can access portal content and areas through search and portlets.

Figure 2-8 No Navigation

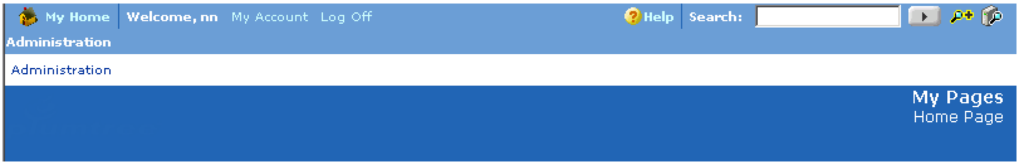

Horizontal Drop-Down Navigation

This navigation scheme uses horizontal tabs and JavaScript-based drop-down menus to access navigation elements. Clicks, not mouse-overs, display the menus. The drop-down menus expand both vertically and horizontally, but cover only the portal's banner to avoid covering the portlets. If a user belongs to more communities than can fit in the allotted space, a vertical scroll bar appears in the drop-down. You can configure the extent of the vertical and horizontal tiling of the drop-down menus.

Figure 2-9 Horizontal Drop-Down Navigation

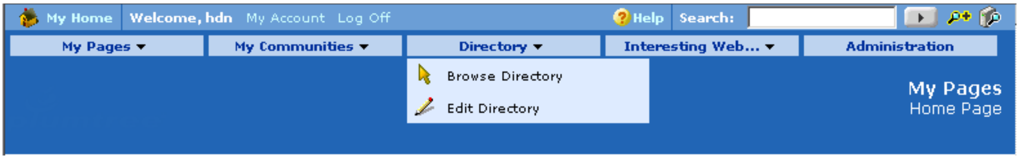

Low Bandwidth and Accessibility Navigation

Low Bandwidth and Accessibility Navigation is used by low bandwidth and accessibility modes of the portal. This navigation is used by those modes no matter which navigation is selected by the experience definition for standard mode.

Portlet-Ready Navigation

Portlet-Ready Navigation disables all navigation areas except the header and footer. The top bar, which includes the search box, is also disabled. This navigation scheme is only used when navigation is controlled by portlets (usually header or footer portlets) using navigation tags. Navigation tags provide developers a faster, easier way to customize navigation than modifying the other available navigation schemes.

For more information on navigation tags, see the BEA AquaLogic User Interaction Development Center (http://dev2dev.bea.com/aluserinteraction/).

Branding

Branding customizes the look of your portal, experience definitions, and communities, through the use of headers and footers. For example, you probably want to add your company logo and tagline to the header and you might want to add contact information or copyrights to the footer.

When you create or edit a branding portlet, you set up the properties, HTML, and default values for the portlet. Experience definition and community administrators can then add these branding portlets to their experience definitions or communities. Community administrators can additionally customize the portlets to fit their needs if they have the proper permissions.

Note: You cannot change the property values for branding portlets in experience definitions. Therefore, branding portlets in experience definitions always display the default property values.

The security set in the Portlet Editor allows users to view the portlet or add it to their community pages. With the proper permissions, users can go into the Community Editor, on the Header and Footer page, click the name of the branding header or footer portlet, and edit the default property values to change portlet content (for example: text, images, or color) without changing the portlet’s overall design.

Creating and Editing Branding Portlets

If you have administrative rights to the Publisher administrative folder, you can create or edit a branding portlet:

1. Click **Administration**.
2. Open an administrative folder.
3. Open the Portlet Editor. To edit an existing branding portlet, click the branding portlet name. To create a new branding portlet:
    a. In the Create Object drop-down list, click **Portlet**.
    b. In the Choose Template or Web Service dialog box, expand the **Publisher | Published Content Portlets | <language>** folders choose **Header** or **Footer** and click **OK**.
4. Under Web Service, to the right of Configure this Portlet, click **Edit**.
5. In the Configure Portlet Wizard, define your portlet as described in the online help.

Note: For more information on branding, refer to the AquaLogic Interaction Publisher documentation.

Locale Settings

Users can choose a locale to determine:

- The language for the portal user interface.
- The format conventions for portal entries.

For example, if you choose British English, the portal language is English and dates appear in the DD/MM/YYYY format; in American English, the portal language is English and dates appear in the MM/DD/YYYY format.

For information on localizing objects, refer to Appendix C, “Localizing Your Portal.”.

Configuring Experience Definitions

Experience definitions provide multiple user experiences within a single portal. Experience definitions do not pre-configure the portal experience with portlets. Multiple guest users, with unique My Pages, can be created and associated with different experience definitions. Each guest user can experience a different experience definition depending on rules you create.

Experience definitions and experience rules allow you to configure the following:

- Portal functions that are visible when users are in that experience definition

  Note: If you disable the Knowledge Directory, users cannot browse document folders, but they can still search for portal documents.

- Navigation scheme
- Branding scheme
- Mandatory links
- Default home page displayed (such as a My Page, a particular community, or a Knowledge Directory folder)
- Login settings specifying which guest user, and associated My Page, is used
- Folders associated with the experience definition
- Which experience definition to display to users when specific conditions are met

Creating Experience Definitions

Note: To create an experience definition, a user must have Admin access to the folder that will contain the experience definition.

To create an experience definition:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Experience Definition**.
4. Configure the experience definition features, navigation, and branding for the experience definition as described in the online help.
5. Click **Finish**.

Directing Users to Experience Definitions

Users can be directed to experience definitions in three ways:

- The users satisfy rules you create in the Experience Rules Manager. These rules may specify the URL used to access the portal, a community the user accesses, a group to which the user belongs, or the user’s IP address.
- The users are stored in a folder associated with an experience definition.
- If no experience definitions are explicitly applied, users see the default experience definition.

Experience rules are applied first, followed by folder association. If users are not matched with any experience rule or folder association, they see the default experience definition, which is created during installation.

To create an experience definition rule:

1. Click **Administration**.
2. In the Select Utility drop-down menu, select **Experience Rules Manager**.
3. Click **Create New Rule**, define the rule as described in the online help, then click **Finish**.
4. Order the rule within the existing rules as described in the online help, then click **Finish**.

To associate an experience definition to users created in a specific folder:

1. Create an experience definition and associate it with an administrative folder as described in the online help.
2. Add users to the folder associated with the experience definition:
   - For manually created users, create them in, or move them to, the folder associated with the experience definition.
   - For authentication sources, in the Authentication Source Editor, on the Main Settings page, under Default Profiles, choose the folder to which you applied the experience definition as the Target Folder for new users.
   - For invitations, in the Invitation Editor, on the Main Settings page, choose the folder to which you applied the experience definition as the Folder for Invitees.

   New users are created in the experience definition folder using one of the methods above, and therefore they see the experience definition interface when they log in.

For information on the procedures for adding users to the portal, see Chapter 3, “Managing Portal Users and Groups.”

Overview of the Portal Administrative User Interface and Tools

This section describes the user interface for portal administrative users and summarizes the tools administrators use to manage the portal. It includes the following topics:

- Administration Section of the Portal
- Administration Utilities in the Portal Installation

Administration Section of the Portal

The Administration section of the portal displays administrative objects and utilities you use to perform portal administration tasks. All procedures in this guide begin with “Click **Administration**” to display the Administration section of the portal.

When you click **Administration**, the portal displays the Administrative Objects Directory. Figure 2-10 shows an administrative folder (*Portal Resources*), the Create Object drop-down list, and the Select-Utility drop-down list (expanded).

Figure 2-10 Administration Section of the Portal

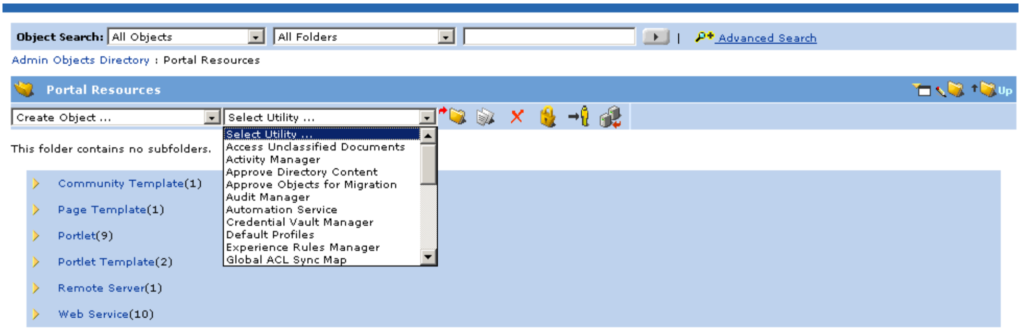

The folders displayed in the Administrative Objects Directory and the utilities available through the Create Object drop-down list and the Select Utility drop-down list depend on activity rights associated with the user name you used to log in to the portal. For information on configuring activity rights for administrator groups and users, see Chapter 3, “Managing Portal Users and Groups.”

The following table describes the objects you can create with the Create Object drop-down list.

Table 2-4 Create Object Drop-Down List

Object	Description
Administrative Folder	Create a folder in the Administrative Objects Directory.
Authentication Source - Remote	Configure a portal authentication source object that references an authentication source on a remote server.
Community	Configure a portal community.
Community Template	Configure a template on which to base a community object.
Content Crawler - Remote	Configure a content crawler instance for targets located on remote hosts.
Content Crawler - WWW	Configure a content crawler instance for targets that are WWW locations.
Content Source - Remote	Configure a portal content source object that references a remote repository.
Content Source - WWW	Configure a portal content source object that references a WWW location.
Content Type	Configure content type identifiers.
Experience Definition	Configure a user experience including such things as branding and navigation
External Operation	Configure a job that calls an external script.
Federated Search - Incoming	Configure an incoming search relationship with a federated portal.
Federated Search - Outgoing	Configure an outgoing search relationship with a federated portal.
Filter	Configure document filters.
Group	Configure group ACL and membership.
Invitation	Configure an invitation to become a portal user or view portal objects.
Job	Configure a job to automate administrative tasks.

Table 2-4 Create Object Drop-Down List

Object	Description
Page	(Only displays when in a community folder) Configure a page to add to a community.
Page Template	Configure a page template on which to base community pages.
Portlet	Configure an instance of a portlet.
Portlet Bundle	Configure a bundle of portlets that must be selected together.
Portlet Template	Configure a template on which to base portlets.
Profile Source - Remote	Configure a portal profile source object that references a profile repository on a remote host.
Property	Configure a portal property.
Remote Server	Configure a remote server host for Web services located on a remote host.
Snapshot Query	Configure a query and publishing location to broadcast availability of portal content.
User	Configure a portal user.
Web Service - Authentication	Configure a Web service portal object that references authentication services on a remote host.
Web Service - Content	Configure a Web service portal object that references content services on a remote host.
Web Service - Intrinsic Portlet	Configure a Web service portal object that references portlet services hosted on the computer hosting the portal.
Web Service - Profile	Configure a Web service portal object that references profile services on a remote host.
Web Service - Remote Portlet	Configure a Web service portal object that references portlet services on a remote host.
Web Service - Search	Configure a Web service portal object that references a search service on a remote host.

The following table describes the administrative utilities provided in the Select Utility drop-down list.

Table 2-5 Select Utility Drop-Down List

Utility	Description	Activity Right or Group Required
Access Unclassified Documents	Access documents imported by a content crawler and placed in the Unclassified Documents folder in the Knowledge Directory.	Access Unclassified Documents
Activity Manager	Create, modify, or delete activities.	Create Activities
Approve Directory Content	Approve directory content submitted to the Knowledge Directory.	Access Utilities
Approve Objects for Migration	Approve migration packages.	Administrators Group Only
Audit Manager	Audit user activity or object activity.	Administrators Group Only
Automation Service	Configure and run jobs.	Administrators Group Only
Credential Vault Manager	Manage lockboxes corresponding to external systems that users can access through the portal.	Administrators Group Only
Default Profiles	Configure default user profiles.	Create User
Experience Rules Manager	Define and prioritize Experience Rules.	Access Experience Rules Manager
Global ACL Sync Map	Configure the global access control list (ACL) synchronization map.	Administrators Group Only
Global Content Type Map	Configure the Global Content Type Map.	Administrators Group Only
Global Document Property Map	Configure the global document property map.	Administrators Group Only
Global Object Property Map	Configure the global object properties map.	Administrators Group Only
Knowledge Directory Preferences	Configure Knowledge Directory preferences.	Administrators Group Only
Localization Manager	Localize the portal.	Administrators Group Only

Table 2-5 Select Utility Drop-Down List

Utility	Description	Activity Right or Group Required
Migration - Export	Create a portal export package.	Administrators Group Only
Migration - Import	Import a portal export package.	Administrators Group Only
Object Migration Status	View the status of portal objects that have been requested for migration.	Access Utilities
Portal Settings	Modify Portal settings.	Administrators Group Only
Release Disabled Logins	Manage user locks.	Administrators Group Only
Release Item Locks	Manage object locks.	Administrators Group Only
Search Cluster Manager	Check status and manage search topology and checkpoints.	Administrators Group Only
Search Results Manager	Manage search results preferences.	Access Search Results Manager
Search Service Manager	Manage Search Service settings.	Administrators Group Only
Smart Sort	Run the Smart Sort utility.	Access Smart Sort
System Health Monitor	View diagnostic information.	Administrators Group Only
Tag Library Manager	Displays the tag libraries loaded on the computer that hosts the portal.	Administrators Group Only
User Profile Manager	Modify the user profiles map.	Access User Profile Manager

Administration Utilities in the Portal Installation

The following table summarizes administration utilities provided in the **<PT_HOME>/ptportal/6.1/bin** directory of the host computer for the portal, and in the portal administrative interface.

Table 2-6 Administration Utilities in the Portal Installation

Utility	Purpose
cryptoutil.sh	The Cryptographic Password utility generates the passwords you might set during installation. To display the man pages for the Cryptographic Password utility, enter the following command: `<PT_HOME>/ptportal/6.1/bin/cryptoutil.sh -h` <PT_HOME> specifies the installation root for the installation of the portal, for example `C:\bea\alui` or `/opt/bea/alui`.
portalenv.sh	The Portal Environment utility sets the portal environment for tools in the **<PT_HOME>/ptportal/6.1/bin** directory. To display the man pages for the Portal Environment utility, enter the following command: `<PT_HOME>/ptportal/6.1/bin/portalenv.sh -h` <PT_HOME> specifies the installation root for the installation of the portal, for example `C:\bea\alui` or `/opt/bea/alui`.
ptmigration.sh or ptmigration.bat	The Migration Wizard manages import packages that enable you to migrate portal objects to new host portals, such as migration from a development environment to a QA environment or production environment, or from a remote server host computer to the portal host computer. The command-line interface (CLI) of the Migration Wizard enables you to import migration packages from the command line. To display the man pages for the Migration Wizard CLI, enter the following command: `<PT_HOME>/ptportal/6.1/bin/ptmigration.sh -h` <PT_HOME> specifies the installation root for the installation of the portal, for example `C:\bea\alui` or `/opt/bea/alui`. For information on object migration, see Chapter 6, “Migrating, Backing Up, and Restoring Portal Objects.”

Table 2-6 Administration Utilities in the Portal Installation

Utility	Purpose
ptspy.sh or ptspy.bat	The graphical user interface (GUI) version of the AquaLogic Interaction Logging Utilities provides an X-windows GUI that enables you to troubleshoot any system errors that might occur when you deploy your portal. To launch AquaLogic Interaction Logging Utilities, enter the following command: `<PT_HOME>/ptlogging/6.0/bin/ptspy.sh` <PT_HOME> specifies the installation root for the installation of the portal, for example `C:\bea\alui` or `/opt/bea/alui`. Alternatively, in Windows, go to **Start | Programs | BEA | ALI Logging Utilities | Logging Spy** For information on modifying the servers to which AquaLogic Interaction Logging Spy listens, see the *Installation and Upgrade Guide for AquaLogic Interaction*.
ptverify.sh	The Installation Verification utility allows you to verify connectivity for installation components and the portal database. To display the man pages for the Installation Verification utility, enter the following command: `<PT_HOME>/ptportal/6.1/bin/portalenv.sh -h` <PT_HOME> specifies the installation root for the installation of the portal, for example `C:\bea\alui` or `/opt/bea/alui`. For details, see the *Installation and Upgrade Guide for AquaLogic Interaction*.
automation service daemon	The Automation Service daemon ensures the Automation Service is running. For details, see the *Installation and Upgrade Guide for AquaLogic Interaction*. For information modifying Automation Service defaults, see “Configuring the Automation Service” on page A-16.
upgrade.sh	The Upgrade Database utility upgrades the portal database from 4.5 versions to the current version. For details, see the *Installation and Upgrade Guide for AquaLogic Interaction*.
Tag Library Manager	This utility allows you to view the tag libraries for this portal machine. To access the Tag Library Manager, in the portal go to: **Administration | Select Utility | Tag Library Manger**

CHAPTER 3

Managing Portal Users and Groups

This chapter describes the portal conventions for user and group management and provides the steps you take to implement managed access to portal objects. It includes the following sections:

About Portal Roles, Groups, Users, and Profiles

In the portal, a *role* is not an object, rather a way of thinking about administrative responsibilities. For example, the Knowledge Directory manager role is not an object you define; it relates to administrative responsibilities for those who manage contents of the Knowledge Directory.

A *group* is an object you configure. A group contains other groups and users, as well as any activity rights assignment for group members. A group can have static membership and membership that changes dynamically based on properties of users' profiles or their memberships in other groups.

A *user* is an object that corresponds to a user account. You configure new users based on profile templates, known as default profiles. *Profiles* are objects you configure. They contain information about users, such as name, job title, and so forth.

Activity rights determine which portal objects a user can create and which portal utilities a user can execute to create or modify portal objects.

Access privileges determine which portal objects a user can browse or edit, which objects appear in search results, and which can be added to My Pages and Community pages.

You implement security for portal access and portal activities in a manner similar to implementing security for other network, domain, and system objects—by managing a hierarchy of the objects that determine access privileges and activity rights.

The default portal installation creates the following group, user, and profile objects, with default access privileges and activity rights.

Table 3-1 Installed Groups, Users, and Profiles

Group, User, or Profile	Default Access Privileges and Activity Rights
Administrator Group	All
Everyone Group	Read and Edit Own Profile access
Administrator	All
Default Profile	n/a
Guest Profile	n/a

As a portal administrator, one of your first tasks is to delegate and share the administrative role by defining the access privileges and activity rights for administrative groups, and then assigning to these groups the users and groups of users that are the basis for your deployment plan.

Caution: By default, you can log in to the administrative portal as Administrator with no password. If the default Administrator password has not yet been changed, you should do so as soon as possible. Make sure that you document the change and inform the appropriate portal administrators.

Before you begin the task of managing portal groups and users, familiarize yourself with your deployment plan. If you plan to leverage group configurations from an existing Active Directory or LDAP authentication source, you might proceed differently from an administrator who plans to manage portal groups using the default portal authentication source. For example, if you plan on using an LDAP authentication source, you might import users, groups, and profile information first and then proceed with configuration of role-based group rights assignments. If you plan to add users to your portal by invitation and manage them through the portal authentication source, you might create the groups first and then add the invited users to these groups.

For detailed information on developing a plan to manage the administrative roles, groups, and users for your enterprise portal, refer to the *Deployment Guide for BEA AquaLogic User Interaction G6*.

The topics in this chapter describe the following basic steps you take to add users and to manage user access privileges and activity rights:

1. Develop a plan to delegate portal administrative roles.
2. Configure role-based groups and associate them with activity rights.
3. Configure default profiles for user account types.
4. Add groups and users to these groups.
5. Configure Access Control Lists (ACLs) for access privileges in the Administrative Object directory.

Delegating Activity Rights

Before you create portal groups, you should become familiar with the definition and scope of the administrative tasks you plan to delegate, and you should develop a plan for assigning the responsibilities and rights for these administrative roles to groups of users. In terms of portal objects, roles are not objects that you configure. Roles are associations between groups and activity rights. In "Configuring Groups" on page 3-9, you configure group objects to enable administrative roles. In "Adding Users" on page 3-12, you add child groups and users to the groups. This section describes the activity rights that are defined by default during installation. Use the information provided in this section to develop your plan to delegate specific activity rights to the administrative groups and users you configure in the sections that follow.

Table 3-2 summarizes the activity rights that are defined by default in the portal. If you encounter cases that require rights not covered by the defaults, you can create custom activity rights. For information on creating custom activity rights, see "Creating Custom Activity Rights" on page 3-8.

Table 3-2 also provides an example map between activity rights and administrative roles. Roles are related to the specific activity rights required to perform a job function. In the example, the role called Content Manager provides the activity rights required to populate the portal with document records crawled from remote content sources: Access Administration, Access Utilities, Create Admin Folders, Create Content Types, Create Content Crawlers, Create Content Sources, and Create Jobs. A separate role called Knowledge Directory Manager provides the activity rights required to create Knowledge Directory structure: Access Smart Sort, Access Unclassified Documents, Access Utilities, Advanced Document Submission, Create Filters, Create Folders, Edit Knowledge Directory, and Self-Selected Experts. Although some users might fill both roles, others might not. By creating two separate roles, you can assign rights for one or both roles to one or many groups.

Table 3-2 Mapping a Relationship Between Rights and Roles

Activity Rights	Example Roles		
	Portal Manager	Content Manager	Knowledge Directory Manager
Access Administration Allows users to see the Administration tab and access the administrative object hierarchy.	x	x	
Access Experience Rules Manager Allows users to access the Experience Rules Manager administrative interface.	x		
Access Search Results Manager Allows users to access the Search Results Manager administrative interface.	x		
Access Smart Sort (Create Taxonomists) Allows users to create new smart-sort.	x		x
Access Unclassified Documents Allows users to see the Unclassified Documents folder in Edit mode of the Knowledge Directory.	x		x
Access User Profile Manager Allows users to access the User Profile Manager.	x		
Access Utilities Allows users to see the Select Utility drop-down list in Administration. Users can also access the Approve Directory Content and Object Migration utilities, but need additional rights to access other utilities.	x	x	x
Advanced Document Submission Allows users to specify advanced options when submitting a document to the Knowledge Directory.	x		x

Table 3-2 Mapping a Relationship Between Rights and Roles

Activity Rights	Example Roles		
	Portal Manager	Content Manager	Knowledge Directory Manager
Create Activities Allows users to create new portal activities. A user must have this right to use the Activity Manager utility.	x		
Create Admin Folders Allows users to create new administrative folders.	x	x	
Create Authentication Sources Allows users to create new authentication sources.	x		
Create Communities Allows users to create new communities and subcommunities.	x		
Create Community Infrastructure Allows users to create new community templates and page templates.	x		
Create Content Crawlers Allows users to create new content crawlers.	x	x	
Create Content Sources Allows users to create new content sources.	x	x	
Create Content Types Allows users to create new content types.	x	x	
Create Experience Definitions Allows users to create new experience definitions.	x		
Create External Operation Allows users to create new external operations.	x		
Create Federated Searches Allows users to create new incoming and outgoing federated searches.	x		

Table 3-2 Mapping a Relationship Between Rights and Roles

Activity Rights	Example Roles		
	Portal Manager	Content Manager	Knowledge Directory Manager
Create Filters Allows users to create new filters.	x		x
Create Folders Allows users to create new Knowledge Directory folders.	x	x	x
Create Groups Allows users to create new groups.	x		
Create Invitations Allows users to create new invitations.	x		
Create Jobs Allows users to create new jobs.	x	x	
Create Portlets Allows users to create new portlets.	x		
Create Profile Sources Allows users to create new profile sources.	x		
Create Properties Allows users to create new properties.	x	x	
Create Snapshot Queries Allows users to create new snapshot queries and corresponding results portlet.	x	x	
Create Users Allows users to create new users and profiles.	x		
Create Web Service Infrastructure Allows users to create new remote servers, Web services, portlet templates, and portlet bundles.	x	x	

Table 3-2 Mapping a Relationship Between Rights and Roles

Activity Rights	Example Roles		
	Portal Manager	Content Manager	Knowledge Directory Manager
Delegate Rights Allows users to delegate activity rights to other users. Users can delegate only activities to which they have rights themselves.	x	x	x
Edit Knowledge Directory Allows users to enter Edit mode in the Knowledge Directory.	x		x
Edit Own Profile Allows users to modify the values of their own user profiles.	x		
Edit Profile Layout Allows users to modify the Profile sections in the User Profile Manager.	x		
Self-Selected Experts Allows users to specify themselves as experts on Knowledge Directory folders.	x		x

Use Table 3-2 to create your own map that delegates administrative rights and responsibilities to roles.

Creating Custom Activity Rights

You can also create custom portal activities. For example, if you have an inventory control system accessed through the portal and only certain users are allowed to edit it, you can create an *Edit Inventories* activity. You can then create inventory-control portlets that verify whether a user has the correct activity right prior to receiving access to the portlet.

To create, modify, or delete custom portal activity rights:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Activity Manager**.
3. Manage the activity rights as described in the online help, and click **Finish**.

Configuring Groups

This section provides procedures for creating the portal groups to which you grant role-specific activity rights. Before you perform the task of creating portal groups, make sure you have created a plan to share portal administration responsibilities, following the guidelines in the *Deployment Guide for BEA AquaLogic User Interaction G6* and the overview of roles, groups, and users in "About Portal Roles, Groups, Users, and Profiles" on page 3-2.

If you plan to import users and groups from an existing authentication source, such as LDAP or Active Directory, you might want to import them first and then follow the procedures in this section to add users to groups. For information on importing users and groups from an authentication source, see "Importing Users and Groups from Authentication Sources" on page 3-12.

To create a group:

1. Click **Administration**.
2. From the Create Objects drop-down list, choose **Administrative Folder** and create a folder to store the groups you will create. For example, you might want to name the folder *Roles*.
3. Open the folder you just created.
4. From the Create Objects drop-down list, choose **Group**.
5. Specify a name that describes the group, such as *Content Managers*.
6. If you have already imported or created the users for this group, add them as described in the online help. Otherwise, you can assign group membership when you import or create the users.
7. Click **Finish**.

Configuring Dynamic Group Membership

You may want to have users automatically added to, or removed from, groups based on properties in their user profiles. For example, you may want to give users access to a community based on their location, title, department, or any other property in their profile. Because some properties can change frequently, you can set up dynamic group membership rules so that when selected properties change, users are automatically added to, or removed from, a group.

To create dynamic group membership rules:

1. Click **Administration**.
2. Open the Group Editor:
   - To edit an existing group, navigate to the group and click the group name.
   - To create a new group, navigate to an existing administrative folder or create a new one in which to store the group. In the Create Object drop-down list, click **Group**.
3. On the left, under Edit Object Settings, click **Dynamic Membership Rules**.
4. Set up rules as described in the online help, and click **Finish**.

If you add or change dynamic membership rules for a group, dynamic members are updated for this group after you click Finish. Otherwise, dynamic memberships are updated for all groups as part of a job (the Dynamic Membership Update Agent). When user profile data changes, the resulting dynamic group membership changes are updated as part of this job. For more information, see “About Jobs” on page 5-1.

Configuring Default User Profiles

Default profiles are templates for new users. Default profiles configure the initial My Account settings, the name and number of My Pages, and the layout of the portlets on those My Pages.

Before you add users to your portal, configure the default profiles you want to apply to the users you will add. For example, if all technical writers should have certain mandatory portlets on their default My Pages, configure a default profile for this purpose and apply the profile to these users when you add them.

Create multiple profiles to apply to the different types of users you anticipate.

To configure default profiles:

1. Click **Administration**.
2. In the Select Utility drop-down list, select **Default Profiles**.
3. In the Create Object drop-down list, select **User**.
4. Enter a name for the default profile, and click **Finish**.

To edit the layout of a default profile:

1. Click **Administration**.
2. In the Select Utility drop-down list, select **Default Profiles**.
3. Select the profile that you want to customize. You can only edit the layout of one profile at a time.
4. Click **Edit Profile Layout**.
5. Specify My Account settings, create My Pages, and change the My Pages’ layouts.
6. Click **Finish**.

Adding Users

This section provides procedures for adding users to the portal. When you add users, you configure group memberships and apply a default profile. This section describes the following options for adding users to the portal:

- "Maintaining Groups, Users, and Profiles with Identity Services" on page 3-12
- "Creating Users" on page 3-15
- "Allowing Users to Create Their Own Accounts" on page 3-16
- "Adding Users with Invitations" on page 3-16

Maintaining Groups, Users, and Profiles with Identity Services

This section provides procedures for adding the users and groups that are already defined in your enterprise in existing authentication sources, such as Active Directory, LDAP, or Windows domain sources. This section includes the following topics:

- "Importing Users and Groups from Authentication Sources" on page 3-12
- "Importing User Profiles from Profile Sources" on page 3-14

For information on installing AquaLogic Interaction Identity Services on a remote server host computer, refer to the product documentation provided with your software.

Importing Users and Groups from Authentication Sources

An authentication source enables you to import users, groups, and group memberships into the portal from an external authentication server. After you have imported the users, the authentication source authenticates portal logins.

If you plan to import users with an AquaLogic Interaction Identity Service, such as AquaLogic Interaction Identity Service - LDAP or AquaLogic Interaction Identity Service - Active Directory, follow the product documentation provided with that software instead of the procedures in this guide.

The following table describes the steps you take to import users, groups, and group memberships from a remote authentication server.

Table 3-3 Importing Users, Groups, and Group Memberships from an External Authentication Server

Basic Step	Procedure
Create a remote server.	1. Click **Administration**. 2. Navigate to an existing administrative folder or create a new one in which to store the portal objects needed for authentication. 3. In the Create Object drop-down list, select **Remote Server**. 4. Configure connection information for the remote server as described in the online help. 5. Click **Finish**.
Create an authentication Web service.	1. In the Create Object drop-down list, select **Web Service - Authentication**. 2. Associate the Web service with the remote server you just created, and configure connection information for the Web service as described in the online help. 3. Click **Finish**.
Configure an authentication source and associated synchronization job.	1. In the Create Object drop-down list, select **Authentication Source - Remote**. 2. In the Select Web Service dialog box, select the Web service you just created. 3. Configure authentication and synchronization preferences, apply default profiles, and associate a synchronization job according to the online help. Before you can run the synchronization job, you must associate the folder that contains the job with an Automation Service. For information on associating folders with an Automation Service, see Chapter 5, "Automating Administrative Tasks." 4. Click **Finish**.
Verify that the correct profiles were applied to users and that portal groups contain only the groups and users you specified when you configured the authentication source.	1. Run the synchronization job. 2. Navigate to the administrative folder that contains the users you imported. 3. Click a user account that you are familiar with and verify the correct group and profile configuration has been applied.

Importing User Profiles from Profile Sources

A profile source enables the portal to use an external source to define user properties that can be searched by portal users, forwarded to portlets to authenticate portlet access, or for other purposes.

If you plan to import profile information with an AquaLogic Interaction Identity Service, such as AquaLogic Interaction Identity Service - LDAP, follow the product documentation provided with that software instead of the procedures in this guide.

The following table describes the steps you take to import user properties from an external source.

Table 3-4 Importing User Properties from an External Source

Basic Step	Procedure
Create a remote server.	1. Click **Administration**. 2. Navigate to an existing administrative folder or create a new one in which to store the portal objects needed for importing user profiles. 3. In the Create Object drop-down list, select **Remote Server**. 4. Configure connection information for the remote server as described in the online help. 5. Click **Finish**.
Create a profile Web service.	1. In the Create Object drop-down list, select **Web Service - Profile**. 2. Associate the Web service with the remote server you just created, and configure connection information for the Web service as described in the online help. 3. Click **Finish**.

Table 3-4 Importing User Properties from an External Source

Basic Step	Procedure
Configure a profile source.	1. In the Create Object drop-down list, select **Profile Source - Remote**. 2. In the Select Web Service dialog box, select the Web service you just created. 3. Map fields in the profile source to portal profile properties, and associate a job to import profile properties as described in the online help. Before you can run the job, you must associate the folder that contains the job with an Automation Service. For information on associating folders with an Automation Service, see Chapter 5, "Automating Administrative Tasks." 4. Click **Finish**.
Verify that the profile properties you imported have been properly applied.	1. Run the synchronization job. 2. Navigate to the administrative folder that contains the users you imported. 3. Click a user account that you are familiar with and verify the profile is configured as expected.

Creating Users

If your enterprise does not use third-party authentication sources, you can use a portal utility to create users. Users you create with the portal utility, users who self-register, and users added by invitation are included in the AquaLogic Interaction Authentication Source.

To create a user:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **User**.
4. Edit the Main Settings page as described in the online help.
5. Click **Finish**.

Allowing Users to Create Their Own Accounts

The portal enables users to create their own accounts by clicking **Create an Account** on the Login page. These users are stored in the portal's **Default Experience Definition** folder and are are included in the AquaLogic Interaction Authentication Source.

Users who self-register are granted access privileges based on the settings for the default profile named **Default Profile**. Based on this security, users can personalize their views of the portal with My Pages, portlets, and community memberships, and can view portal content.

Adding Users with Invitations

Invitations allow you to direct potential users to your portal, making it easy for them to create their own user accounts and letting you customize their initial portal experiences with content that is of particular interest to them.

You should create a single invitation for all potential users who should be added to the same portal groups and should see the same communities, portlets, and My Pages when they first log in to your portal.

To accept the invitation, the user clicks the link included in the e-mail and follows the directions to create a new user to log in to the portal. When the user logs in, the portlets, content, and communities specified in the invitation are displayed to the new user.

Users added by invitation are included in the AquaLogic Interaction Authentication Source.

To create an invitation:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Invitation**.
4. Configure the profile and group information for the invited user as described in the online help.
5. Click **Finish**.

After creating an invitation, you need to send the invitation.

To send an invitation:

1. Click **Administration**.
2. Navigate to the invitation you want to send.
3. Select the invitation you want to send and click **Send Invitation**.
4. Click **Create New Invitation Link**.
5. Specify the maximum number of times you want this link to be accessed and the date you want this link to expire.
6. To generate the link, click **Finish**.
7. To display the invitation link, click its name. Copy and paste the invitation link into an e-mail.
8. To close the Invitation Link dialog box, click **Finish**.
9. To close the Send Invitation Editor, click **Finish**.

Managing User Profiles and User Accounts

This section provides procedures for managing user profiles and user accounts. It includes the following topics:

- "Deleting Users" on page 3-18
- "Managing User Profiles" on page 3-19
- "Auditing User Accounts and Actions" on page 3-19
- "Locking and Unlocking User Accounts" on page 3-21
- "Managing User Credentials for Existing Applications" on page 3-23

Deleting Users

To delete a user:

1. Click **Administration**.
2. Navigate to the user.
3. Select the user you want to delete and click ✗ .

To delete a user whose account is locked:

1. Click **Administration**.
2. In the Select Utilities drop-down list, click **Release Disabled Logins**.
3. Select the user you want to delete and click ✗ .
4. Click **Finish**.

Managing User Profiles

Profile information, such as name and job title, is stored with user objects as properties. You can use the User Profile Manager to specify which properties are sent to portlets when requested.

The values for specific properties are set either by the user on the Edit User Profile screen or by a profile source.

To specify which user properties are sent to portlets:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **User Profile Manager**.
3. Define how user profiles are displayed and which user properties are sent to portlets as described in the online help.
4. Click **Finish**.

Auditing User Accounts and Actions

The portal logs user activities, which allows you to query for actions taken by particular users, actions taken on a particular administrative object, or actions taken within a specified time period.

Note: You should configure activity logging to adequately meet the security auditing needs of your portal deployment and then implement procedures for periodically reviewing the audit records.

Configuring User Activity Auditing

To configure user activity auditing:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Audit Manager**.
3. Complete the configuration according to the online help.
4. Click **Finish**.

Archiving Audit Messages

The Audit Log Management agent moves audit messages from the portal database into a collection of archive files and deletes old archive files based on the settings you configure in the Audit Manager. The Audit Log Management agent runs in the Audit Log Management Job, created upon installation and stored in the **Intrinsic Operations** folder. By default, this job runs daily. For information on configuring the Audit Log Management agent, see "Running Portal Agents" on page 5-4.

Querying Audit Information

To query the database for audit entries:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Audit Manager**.
3. On the left, under Edit Utility Settings, click **Create Audit Query**.
4. Define your query as described in the online help.
5. Click **Finish**.

Deleting Audit Messages and Archives

When you configure user activity auditing, you can specify the frequency with which audit messages are deleted automatically.

To delete messages and archives immediately:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Audit Manager**.
3. Specify the messages and archives to delete as described in the online help.
4. Click **Finish**.

Locking and Unlocking User Accounts

You lock user accounts to disable access to the portal. You can configure automatic locking based on repeated failed login attempts, or you can lock user accounts any time with the User Editor.

Automatically Locking User Accounts

To configure account locking for failed login attempts:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Portal Settings**.
3. On the User Settings Manager page, enable account locking and specify how long failed logins are tracked, the total number of failed logins required before an account will be locked, and the number of minutes for which automatically locked accounts remain locked.

   Your individual security needs will determine what settings to use for automatic account locking. For example, to meet a strength of password function rating of SOF-basic as defined in the Common Criteria for Information Technology Security Evaluation, Version 2.3, August 2005 (found at http://niap.bahialab.com/cc-scheme/cc_docs/), you might set the following values:

   – **Minutes to track failed Logins:** 60 minutes or more

   – **Number of failed Login attempts allowed:** 5 or fewer

   – **Minutes to keep user account locked:** 60 minutes or more
4. Click **Finish**.

Manually Locking User Accounts

To lock a user account:

1. Click **Administration**.
2. Navigate to the user whose account you want to lock and click the user name.
3. Select **Disable Login**.
4. Click **Finish**.

Unlocking User Accounts

The lock on accounts that are locked automatically will eventually expire, but you can remove account locks with the Release Disabled Logins utility or the User Editor.

The following table describes how to unlock user accounts that have been locked in particular ways.

Table 3-5 Unlocking User Accounts

Type of Lock	To remove the lock:
Admin Lock A portal administrator locked the user account.	To remove a lock with the Release Disabled Logins utility: 1. Click **Administration**. 2. In the Select Utilities drop-down list, click **Release Disabled Logins**. 3. Manage the lock as described in the online help. 4. Click **Finish**. To remove a lock with the User Editor: 1. Click **Administration**. 2. Navigate to the user whose account you want to unlock and click the user name. 3. Clear the check box next to **Disable Login**. 4. To immediately release the user account lock, click **Finish**.
Automatic Lock If the user repeatedly types the wrong user name or password when logging into the portal, the portal locks the account. The number of login attempts allowed before the user is locked out is determined in the Portal Settings utility. Locks on accounts that are locked automatically eventually expire.	
Agent Lock A user account might be locked if it is not found in the external authentication server during a synchronization job. This lock might be unexpected if the synchronization job did not find the user because the job failed. Users can remove the lock by specifying the correct credentials the next time they log in.	To remove the locks for all affected users: 1. Click **Administration**. 2. Navigate to the authentication source and click its name. 3. On the left, under Edit Object Settings, click **Fully Synchronized Groups**. 4. Click **Re-Enable Users**. Unlocking these accounts may take a few minutes. 5. Click **Finish**.

Managing User Credentials for Existing Applications

You can enable users to access existing Web applications through the portal. For example, users may need to access an employee benefits system. If they access the benefits system through the portal, they do not have to enter their login credentials separately for that application, and can continue to have the convenience of the portal context, personalization, and navigation.

To manage user credentials, you can create a lockbox for each application the user needs to access through the portal. Then, users enter their credentials for each lockbox in their My Account settings.

For more information on integrating applications, see "Using Portlets to Access Existing Web Applications" on page 4-20.

To create a lockbox:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Credential Vault Manager**.
3. Click **New Lockbox** and create a lockbox as described in the online help.
4. Click **Finish** to close the Credential Vault Manager.

To supply login credentials for lockboxes, users do the following:

1. Click **My Account**.
2. Click **Password Manager**.
3. For each application listed (corresponding to a lockbox), enter the user name and password used to access that application.
4. Click **Finish**.

Setting User Access Privileges

After you have imported users and groups into the portal, you can configure access control lists (ACL) to manage privileges to folders in the Administrative Objects Directory. You can also set access control lists for objects within folders. The default portal installation includes the following folders in the Administrative Objects Directory.

Table 3-6 Folders in the Administrative Objects Directory

Folder	Default ACL
Administrative Resources This folder contains the following objects created at installation: users, groups, the AquaLogic Interaction Authentication Source, the WWW content source, properties, content types, and federated search objects.	Administrators Group - Admin access Everyone Group - Read access
Intrinsic Operations This folder contains external operations and intrinsic jobs, such as Search Update, Document Refresh, and Weekly Housekeeping. The folder is registered with the primary Automation Service.	Administrators Group - Admin access
Portal Resources This folder contains intrinsic portlets and Web services, as well as page, community, and portlet templates.	Administrators Group - Admin access Everyone Group - Read privilege
Default Experience Definition This folder contains the users associated with the default experience definition. Upon installation, one user is associated with the default experience definition—Administrator.	Administrators Group - Admin access Everyone Group - Read access

Users in the Administrators group have full access to all portal objects. Other users can be assigned the following access privileges:

- **Read** allows users or groups to see an object.
- **Select** allows users or groups to add an object to other objects. For example, it allows users to add portlets to their My Pages, add users to groups, or associate remote servers with Web services. Object selection lists display only those objects to which you have Select access. Select privileges include Read privileges.
- **Edit** allows users or groups to modify an object, including moving or copying an object. Edit privileges include Select and Read privileges.
- **Admin** allows users or groups full administrative control of an object, including deleting the object or approving it for migration. Admin privileges include Read, Select, and Edit privileges.

By default, newly created objects inherit the ACL configuration of their parent folder.

The following table describes the minimum access required to perform actions on an object.

Table 3-7 Access Requirements

Action	Minimum Access Required on:		
	Object	Source Folder	Target Folder
View	Read	Read	n/a
Modify	Edit	Read	n/a
Create	n/a	Edit	n/a
Copy	Admin	Read	Edit
Move	Admin	Edit	Edit
Delete	Admin	Edit	n/a

The Everyone group always has Read access to the following types of portal objects:

- Content types. All users must have Read access to all content types because content types are used to submit documents to the Knowledge Directory.
- Filters. All users must have Read access to all filters to ensure that the Folder Editor displays accurate information.
- Invitations. All users must have Read access to invitations because invitations are accessed before users log in and therefore cannot take advantage of portal security.
- Properties. All users must have Read access to all properties to ensure that the Filter Editor displays accurate information. Filters are based on properties.

To set the ACL on folders in the Administrative Objects Directory:

1. Click **Administration**.
2. Select the folder and click the **Edit Subfolder** button.
3. On the left, under Edit Standard Settings, click **Security**.
4. Edit the ACL by adding groups, users, and their privileges according to the online help.
5. Click **Finish**.

By default, newly created objects inherit the ACL configuration of their parent folder. If subfolders require different configuration from their parent, modify the ACL for the subfolders as needed. You can also set the security on objects within a folder by editing the Security page in that object’s editor.

Directing Users to Experience Definitions

Experience definitions define the user experience by controlling the branding, styles, navigation, and features of the portal pages the user sees. You can set up rules to evaluate which experience definition any given user should see. The rules can be based on the URL the user uses to access the portal, the user’s IP address, the user’s group memberships, or when the user navigates to a specific community. For more information, see “Configuring Experience Definitions” on page 2-18.

CHAPTER 4

Managing Portal Content

This chapter explains the design of managed content availability in the portal and provides the steps you take to make content available to users. The chapter includes the following topics:

About Portal Content

The portal is designed to enable users to discover all of the enterprise content related to their employee role by browsing or searching portal areas.

Portal users should be able to assemble a My Page that provides access to all of the information they need. For example, to write user documentation, technical writers need to be able to assemble a My Page that includes portlet- or community-based access to documentation standards and conventions, solution white papers, product data sheets, product demonstrations, design specifications, release milestones, test plans, and bug reports, as well as mail-thread discussions that are relevant to customer support and satisfaction. To perform their role, technical writers do not need access to the personnel records that an HR employee or line-manager might require, or to the company financial data that the controller or executive staff might need, for example. A properly designed enterprise portal, then, would reference all of these enterprise documents so that any employee performing any function can access all of the information they need; but a properly designed enterprise portal would also ensure that only the employee performing the role can discover the information.

To enable such managed access to enterprise content:

- Enable discovery of content through browsing or searching the portal.
- Configure Access Control Lists (ACLs) to manage access to these documents.

This chapter describes the following tasks you complete to enable managed discovery of enterprise content through the portal:

- For all content types you plan to support in your portal, configure document properties that enable document filters used by the Knowledge Directory, content crawlers, the Smart Sort utility, and the Search Service.

  For information on document properties and content types, see "Configuring Content Types and Document Properties" on page 4-3.

- Configure access to content sources that can be selected by users or content crawlers to add document records to the Knowledge Directory and Search Index.

  For information on content sources, see "Configuring Content Sources" on page 4-10

- Enable users to upload document records to the Knowledge Directory.

  For information on the Knowledge Directory, see "Managing the Knowledge Directory" on page 4-13.

- Configure portlets that users can add to their My Pages.

  For information on portlets, see “Extending Portal Services with Portlets” on page 4-18.

- Create communities that users can add to their My Communities list.

  For information on Communities, see “Managing Communities” on page 4-29.

- Configure content crawlers and crawl jobs to create links to back-end content sources, such as WWW locations, file system locations, Documentum Content Servers, Exchange Servers, Lotus Notes Servers, or other IMAP-compliant servers.

  For information on content crawlers, see “Enabling Document Discovery with Content Crawlers and Content Services” on page 4-40.

- Run a Search Update job to index these documents so that they can be discovered with the search.

  For information on search, see “Working with Search” on page 4-48.

Configuring Content Types and Document Properties

This section describes how to configure the content type objects and document properties objects that enable document filters used by the Knowledge Directory, content crawlers, the Smart Sort utility, and the Search Service. Filters and returned search results are based on the associated portal properties, not properties defined in the source document.

When you add documents to the portal, the portal maps source document fields to portal properties according to mappings you specify in the Global Content Type Map, the particular content type definition, the Global Document Property Map, and any content crawler-specific content type mappings.

To enable content type and property mapping:

1. Configure the Global Content Type Map.

   Add and configure additional content types, as needed.

   For details, see Configuring the Global Content Type Map

2. Configure the Global Document Property Map.

   Add and configure additional document properties, as needed.

   For details, see “Configuring the Global Document Property Map” on page 4-4

Configuring the Global Content Type Map

The Global Content Type Map allows you to map source document identifiers (for example, file extensions) to content types. The content type associated with a source document determines how metadata in the source document is mapped to portal properties.

To configure the Global Content Type Map:

1. Click **Administration**.
2. In the Select Utility drop-down list, choose **Global Content Type Map**.
3. Configure identifiers for content types as described in the online help.
4. Click **Finish**.

To create a new content type:

1. Click **Administration**.
2. In the Create Object drop-down list, choose **Content Type**.
3. Select an appropriate Accessor, configure a property map, and specify default behavior for populating portal properties from the source documents as described in the online help.
4. Click **Finish**.

Configuring the Global Document Property Map

The Global Document Property Map provides default mappings for properties common to the documents in your portal. These mappings are applied after the mappings in the content type.

When you import a document into the portal, the portal performs the following actions:

1. The portal determines which content type to use, based on the Global Content Type Map or the content type settings for the content crawler.
2. The portal maps source properties to portal property values based on the mappings in the content type.
3. If the Global Content Type Map includes properties that are not in the source documents, the portal populates the portal property values for these documents based on defaults you configure.
4. Click **Finish**.

To configure the Global Document Property Map:

1. Click **Administration**.
2. In the Select Utility drop-down list, choose **Global Document Property Map**.
3. Create mappings between portal properties and document attributes as described in the online help.
4. Click **Finish**.

To create new properties:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Property**.
4. Configure the property settings as described in the online help.
5. Click **Finish**.

Mapping HTML Page Properties

Generally, you will be able to determine what source document attributes can be mapped to portal properties, but this might not be as clear in HTML documents. Table 4-1 provides suggestions for mapping HTML attributes to portal properties.

The HTML Accessor handles all common character sets used on the Web, including UTF-8.

Table 4-1 Mapping HTML Attributes to Portal Properties

HTML Attribute	Portal Property
<Title> Tag	The HTML<TITLE> tag maps to the portal property *Title*.

<Meta> Tag	You can add property information to an HTML page using the <META> tag, as shown in this example:
	`<HTML>`
	`<HEAD>`
	`<TITLE>Press Release – Company X Promotes Five Vice Presidents and Elects Six New Corporate Officers </TITLE>`
	`<META NAME=”corporate_information_class” CONTENT=”Press Relations”>`
	`<META NAME=”creation_date” CONTENT=”18-Jan-2004”>`
	`<META NAME=”stop_date” CONTENT=”18-Jan-2005”>`
	`<META NAME=”next_check_date” CONTENT=”18-Jan-2005”>`
	`<META NAME=”last_check_date” CONTENT=”18-Jan-2004”>`
	`<META NAME=”web_author_id” CONTENT=”ktstatha”>`
	`<META NAME=”language” CONTENT=”English”>`
	`<META NAME=”country” CONTENT=”USA”>`
	`</HEAD>`
	Using this meta information and setting up the appropriate content type allows content crawlers and filters to be much more effective. For example, you could map the <META> tag creation_date to the portal property Created; this allows you to automatically sort documents into the correct monthly folder, such as Jan 2004.

Table 4-1 Mapping HTML Attributes to Portal Properties (Continued)

HTML Attribute	Portal Property

Headline Tags	The Accessor returns a value for each headline tag (<H1>, <H2>, <H3>, <H4>, <H5>, and <H6>) and each bold tag (<B>). The attribute name returned by the Accessor is the name of the tag followed by an ordinal, one-based index in parentheses, and the value is the contents of the tag. For example, an HTML document contains:
	`<H1>Value 1</H1>`
	`<H3>Value 2</H3>`
	`<H1>Value 3</H1>`
	`<B>Value 4</B>`
	The HTML Accessor returns the following source document attribute-value pairs:
	`<h1>(1)        Value 1`
	`<h3>(1)        Value 2`
	`<h1>(2)        Value 3`
	`<B>(1)         Value 4`
	If on a particular news site, the second <H2> tag contains the name of the article and the third <B> tag contains the name of the author, you could map the portal property *Title* to *<H2>(2)* and the portal property *Author* to *<B>(3).*

Table 4-1 Mapping HTML Attributes to Portal Properties (Continued)

HTML Attribute	Portal Property

HTML Comments	It has become a common practice to store metadata in HTML comments using the following format:
	`<!-- Writer: jm -->`
	`<!-- AP: md -->`
	`<!-- Copy editor: mr -->`
	`<!-- Web editor: ad -->`
	In other words, the format is the HTML comment delimiter followed by the name, a colon, the value, and a close comment delimiter. The HTML Accessor parses data in this format and returns source document attribute-value pairs:
	`Writer jm`
	`AP md`
	`Copy editor mr`
	`Web editor ad`

Table 4-1 Mapping HTML Attributes to Portal Properties (Continued)

HTML Attribute	Portal Property
Parent URL	Documents imported via Web crawl return an attribute named *Parent URL* with the value of the URL of the parent page that contains a link to the document.
Anchors	The HTML Accessor provides special handling for internal anchors (<a name=”target”>) and URLs that reference them (http://server/page#target). You might map anchors to portal attributes in the following ways: • Alternate Sources for the portal Title attribute When the document URL for an HTML document contains a fragment identifier (for example, #target in the example above) and the Accessor finds that anchor in the document, it discards all title and headline tags preceding the anchor and returns, as the suggested document title, the first subsequent headline tag. All subsequent tags are indexed relative to the anchor tag, so mapping a property to <H1>(2) means “use the second <H1> tag after the anchor tag named in the document URL.” • Mapping Anchor Section to Document Description or Summary The HTML Accessor returns an attribute named Anchor Section containing text immediately following the named anchor tag (stripped of markup tags and HTML decoded). Mapping this property to the document description allows the portal to generate a relevant description for each section of a large document. The HTML Accessor generates its own summary by returning the first summary-sized chunk of text in the document stripped of HTML markup tags and correctly HTML decoded. It returns this summary as an attribute named Summary. The Accessor executes the DocumentSummary method, which returns the value of the Anchor Section attribute, if available. If this attribute is not available, its second choice is the value of the Description attribute from the <META NAME=”description”> tag. If this is not available, its third and final choice is the Summary attribute.

Configuring Content Sources

This section describes how to configure content sources that enable portal access to content on WWW locations, file systems, and back-end content servers. This section includes the following topics:

- About Content Sources
- "Configuring Web Content Sources" on page 4-11
- "Configuring Remote Content Sources" on page 4-12

About Content Sources

Content sources provide access to external content repositories, allowing users and content crawlers to add document records and links in the Knowledge Directory. For example, a content source for a secured Web site can be configured to fill out the Web form necessary to gain access to that site.

Register a content source for each secured Web site or back-end repository from which content can be imported into your portal.

Content Source Histories

Content sources keep track of what content has been imported, deleted, or rejected by content crawlers accessing the content source. It keeps a record of imported files so that content crawlers do not create duplicate links. To prevent multiple copies of the same link being imported into your portal, set multiple content crawlers that are accessing the same content source to only import content that has not already been imported.

Content Sources and Security

Because a content source accesses secured documents, you must secure access to the content source itself. Content sources, like everything in the portal, have security settings that allow you to specify exactly which portal users and groups can see the content source. Users that do not have Read access to a content source cannot select it, or even see it, when submitting content or building a content crawler.

Using Content Sources and Security to Control Access

You can create multiple content sources that access the same repository of information. For example, you might have two Web content sources accessing the same Web site. One of these content sources could access the site as an executive user that can see all of the content on the site. The other content source would access the site as a managerial user that can see some secured content, but not everything. You could then grant executive users access to the content source that accesses the Web site as an executive user, and grant managerial users access to the content source that accesses the Web site as a managerial user.

Note: If you crawled the same repository using both of these content sources, you would import duplicate links into your portal. Refer to "Content Source Histories" on page 4-10.

Configuring Web Content Sources

Web content sources allow users to import content from the Web into the portal through Web content crawlers or Web document submission. When you install the portal, the World Wide Web content source is created. This content source provides access to any unsecured Web site.

To create a Web content source:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Content Source - WWW**.
4. When prompted, select the Web service for the **World Wide Web**.
5. Define your Web content source as described in the online help.
6. Click **Finish**.

Configuring Remote Content Sources

Remote content sources allow users to import content from an external content repository into the portal through remote content crawlers or remote document submission.

The following table describes the steps you take to configure a remote content source.

Table 4-2 Steps to Configure a Remote Content Source

Basic Step	Procedure
Create a remote server to use for both user document submission and content crawlers. If you are using an AquaLogic Interaction Content Service, you can import the remote server when you register the AquaLogic Interaction Content Service with the portal. For details, refer to the documentation provided with the software.	1. Click **Administration**. 2. Navigate to an existing administrative folder or create a new one in which to store the portal objects needed for importing content. 3. In the Create Object drop-down list, select **Remote Server**. 4. Configure connection information for the remote server as described in the online help. 5. Click **Finish**.
Create a Web service to use for both user document submission and content crawlers. If you are using an AquaLogic Interaction Content Service, you can import the Web service when you register the AquaLogic Interaction Content Service with the portal. For details, refer to the documentation provided with the software.	1. In the Create Object drop-down list, select **Web Service - content**. 2. Configure connection information for the Web service as described in the online help. 3. Click **Finish**.
Configure a remote content source.	1. In the Create Object drop-down list, click **Content Source - Remote**. 2. Define the remote content source as described in the online help. 3. Click **Finish**.

Managing the Knowledge Directory

This section describes how to set up and manage the portal Knowledge Directory. It includes the following topics:

- About the Knowledge Directory
- "Setting Knowledge Directory Preferences" on page 4-13
- "Creating Folders" on page 4-14
- "Submitting Documents" on page 4-14
- "Controlling Document Placement with Filters" on page 4-14
- "Maintaining Document Links" on page 4-17

About the Knowledge Directory

The Knowledge Directory is a portal area that users can browse to discover documents that have been uploaded by users or imported by content crawlers. This information is organized into subfolders in a manner similar to file storage volumes and shares, but you might want to organize it in a more granular fashion to allow you to delegate administrative responsibility and facilitate managed access with ACLs.

The default portal installation includes a Knowledge Directory root folder with one subfolder named Unclassified Documents. Before you create additional subfolders, define a taxonomy, as described in the *Deployment Guide for BEA AquaLogic User Interaction G6*.

Setting Knowledge Directory Preferences

You can specify how the Knowledge Directory displays documents and folders, including whether to generate the display of contents from a Search Service search or a database query, by setting Knowledge Directory preferences.

To set Knowledge Directory preferences:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Knowledge Directory Preferences**.
3. Specify preferences as described in the online help.
4. Click **Finish**.

Creating Folders

To create a Knowledge Directory folder:

1. Click **Directory | Edit Directory**.
2. Navigate to the folder into which you want to place a new subfolder.
3. Click *📁. to launch the Folder Editor.
4. Provide a name and description and click **OK**.
5. Click the Edit Details icon ✏ and complete the settings as described in the online help.

   If you want to modify the ACL that is inherited from the parent folder by default, click **Security**.

Submitting Documents

To submit (upload) a document:

1. Browse to the folder where you want to place the document.
2. From the Submit Document drop-down list, choose **Simple Submit** or choose a content source.
3. Complete the submission forms as described in the online help.

Controlling Document Placement with Filters

Use filters to control what content goes into which folder when crawling in documents or using Smart Sort. A filter sets conditions to sort documents into associated folders in the Knowledge Directory.

A filter is a combination of a basic fields search and statements. The basic fields search operates on the name, description, and full-text content fields associated with documents. Statements can operate on both the content and the properties of documents. Statements can be grouped together in *groupings*. Groupings are containers for statements or other groupings allowing you to create complex filters. Groupings are analogous to parentheses in mathematical equations.

A single filter can be used by multiple folders. You can also apply multiple filters to one folder.

Creating Filters

To create a filter:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Filter**.
4. Define your filter as described in the online help.
5. Click **Finish**.

Assigning Filters to a Folder

After you create a filter, you assign it to folders. You can assign filters to any Knowledge Directory folder to which you have the appropriate access. If you assign more than one filter to a folder, you must specify whether content must pass all filters, or at least one filter.

To assign a filter to a folder:

1. Click **Directory | Edit Directory**.
2. Navigate to the folder you want to assign the filter to and click ✎ next to that folder. This launches the Folder Editor.
3. In the filters section, click + **Add Filter**.
4. In the Add Filter dialog box, expand the folders, as necessary, and select any filters you want to add to this folder.
5. When you are finished adding filters, click **OK**.
6. In the Folder Editor, click **Finish**.

Any content that passes the filters of the destination folder, but does not pass the filters of the destination folder's subfolders can be placed into a default folder.

To specify a default folder:

1. Click **Directory | Edit Directory**.
2. Navigate to the folder and click the Edit Details icon ✎ next to that folder.
3. In the Default Folder drop-down list, choose where you want to store documents that do not pass the filters of any subfolders. If this folder does not have any subfolders, the default folder will be this folder. If this folder has subfolders, you can choose whether to make this folder or one of the subfolders the default folder.
4. Click **Finish**.

Using Filters to Organize Crawled Content

You can organize crawled content into subcategories by creating a folder in the Knowledge Directory, and using filters on that folder's subfolders. For example, you can create a content crawler that crawls a news Web site and places content into a folder, then use filters on that folder's subfolders to separate the content into Politics, Sports, and Travel.

Note: For information on content crawlers, see "About Content Crawlers" on page 4-40.

To use filters to organize content using this example:

1. Create a folder in the Knowledge Directory and call it *News*.
2. Create three subfolders under News called: *Politics*, *Sports*, and *Travel*.
3. Create a content crawler that crawls the news Web site, and choose the News folder as its destination folder.
4. Create a filter to assign to each of the subfolders. Each filter should limit the content of the associated folder to the desired news category: politics, sports, or travel.
5. Assign each filter to the appropriate subfolder. When content is crawled in from the news site into the News folder, it will automatically be filtered into the appropriate subfolder according to the filters you created and assigned to those subfolders.

Sorting Content into Folders with the Smart Sort Utility

You can use the Smart Sort Utility to redistribute content in your portal from one folder to another, applying filters according to your needs.

To redistribute content with the Smart Sort Utility:

1. Click **Administration**.
2. From the Select Utility drop-down list, choose **Smart Sort**.
3. Configure source and destination details as described in the online help.

Maintaining Document Links

The Document Refresh Agent is an intrinsic job that updates the document links in the Knowledge Directory. The Document Refresh Agent visits every link in your portal. For each link, the Document Refresh Agent first determines if the link requires refreshing based on the setting for the document record that was imported into the Knowledge Directory. If the link requires refreshing, the Document Refresh Agent looks at the source document. If the source document has changed, any changed content is updated in the search index, and, optionally, the portal properties are regenerated from the source document. For example, if someone adds a line to the source document or changes the author, as soon as the link is refreshed, portal users can locate the document by searching for this new line of text or searching for the new author.

The Document Refresh Agent also deletes links with missing source documents and links that have expired.

You should run the Document Refresh Agent as frequently as you expect your links to require updates. The Document Refresh Agent knows if other copies of the agent are running and will distribute the work across these agents. However, the more copies of this agent that are running, the more CPU cycles that are used by the Automation Service, so you should limit the number of agents to fit your CPU resources.

For more information on the Document Refresh Agent job, see “Running Portal Agents” on page 5-4.

To examine the refresh settings for a document:

1. Click **Directory | Edit Directory**.
2. Navigate to the document, select it, and click **Document Settings**. The options on the Document Settings page determine how the document is refreshed.

Extending Portal Services with Portlets

This section describes how to set up and manage the availability of portlets. It includes the following topics:

- "About Portlets" on page 4-18
- "Creating Portlet Web Services, Portlet Templates, Portlets, and Portlet Bundles" on page 4-24
- "Requiring and Recommending Portlets" on page 4-27

About Portlets

Portlets provide customized tools and services, as well as information. The portal comes with many portlets, but you can also create your own, have a Web developer or an AquaLogic User Interaction portlet developer create portlets for you, or download portlets from the AquaLogic User Interaction Support Center.

For information on installing and configuring portlets provided as a software package, refer to the portlet software documentation instead of the procedures in this guide.

For information on developing portlets, see the BEA AquaLogic User Interaction Development Center (http://dev2dev.bea.com/aluserinteraction/).

There are several steps involved in making a portlet available for users to add to My Pages or community pages:

1. Install the portlet software.
2. Create a remote server and portlet Web service to define the functional settings.
3. Optionally, create a portlet template to define display settings on which to base multiple portlets.
4. Create a portlet to define the portlet display settings.
5. Optionally, add portlets to a portlet bundle to allow users to easily add groups of related portlets to My Pages or community pages.

Portlet Characteristics

The following table describes some of the characteristics of portlets you might use in your deployment.

Table 4-3 Some Characteristics of Portlets

Characteristic	Description
Intrinsic or Remote	• Intrinsic portlets are included with the default portal and are installed on the computer that hosts the portal application. • Remote portlets extend the base functionality of the default portal and are hosted on a remote server. When a user displays a My Page or community page that includes a remote portlet, the portal contacts the remote server via HTTP to obtain updated portlet content.
Community or Personal	• Community portlets can be added only to community pages. You can also create portlets through the Community Editor; these portlets can be used only in that community. • Personal portlets can be added to My Pages or community pages.
Narrow or Wide	• Narrow portlets can be added to narrow or wide columns. Columns extend to fit portlet content; therefore, if you choose Narrow for a portlet that produces wide content, your portal might look awkward. • Wide portlets can be added only to wide columns.
Header, Footer, or Content Canvas	• Header portlets can be added to communities, community templates, and experience definitions to change the branding of these objects by replacing the banner at the top of the page (so that it differs from the top banner displayed by the main portal). • Footer portlets can be added to communities, community templates, and experience definitions to change the branding of these objects by replacing the banner at the bottom of the page (so that it differs from the bottom banner displayed by the main portal). • Content canvas portlets can be added below the top banner on community pages that include a content canvas space in the page layout. You cannot add more than one content canvas portlet per page. **Note:** Header, footer, and content canvas portlets and portlet templates are included with the default portal.

Using Portlets for Navigation and Login

AquaLogic Interaction provides tags that can be used in portlets as an easy way for developers to customize navigation and login components (such as name field, login field, and so on). Two portlets are included in your portal to provide examples of using these tags:

- Navigation Tags Header Portlet - Included as an example of using navigation tags. Use this portlet with the Portlet-Ready Navigation scheme (set in an experience definition). The Tag Navigation experience definition is also included in the portal as a convenience when you are using portlets for navigation. This experience definition uses the Portlet-Ready Navigation scheme and has the Navigation Tags Header Portlet set as its header.
- Login Portlet - Included as an example of using login tags.

For more information on portal navigation, see "Navigation Options" on page 2-12.

For more information on using tags, see the BEA AquaLogic User Interaction Development Center (http://dev2dev.bea.com/aluserinteraction/).

Using Portlets to Access Existing Web Applications

You can enable users to access existing Web applications through the portal. For example, users may need to access an employee benefits system. If they access the benefits system through the portal, they do not have to enter their login credentials separately for that application, and can continue to have the convenience of the portal context, personalization, and navigation.

To surface an existing application through the portal:

1. (Recommended) Create a lockbox in the portal for the existing application, and have users supply their login credentials for that lockbox.

   To create a lockbox:

   a. Click **Administration**.

   b. In the Select Utility drop-down list, click **Credential Vault Manager**.

   c. Click **New Lockbox** and create a lockbox as described in the online help.

   d. Click **Finish** to close the Credential Vault Manager.

To supply login credentials for lockboxes, users do the following:

a. Click **My Account**.

b. Click **Password Manager**.

c. For each application listed (corresponding to a lockbox), enter username and password used to access that application.

d. Click **Finish**.

2. Create a remote server in the portal for the existing application:

   a. Click **Administration**.

   b. Navigate to or create the administrative folder for this server.

   c. In the Create Object drop-down list, select **Remote Server**.

   d. Configure connection information for the remote server as described in the online help.

   e. Click **Finish**.

3. Create a remote portlet Web service in the portal to associate with a portlet you will create for the existing application:

   a. Click **Administration**.

   b. Navigate to or create the administrative folder for this Web service.

   c. In the Create Object drop-down list, select **Web Service - Remote Portlet**.

   d. Associate the Web service with the remote server you created in the previous step. Configure the rest of the Web service as described in the online help. You can use the lockbox you created for this application to supply the user credentials for authenticating to this application.

   e. To display the existing application’s content in the entire area between the portal header and footer, choose **Hosted Display Mode** in the HTTP Configuration page of the Web Service Editor. This allows users to see a larger view of the application while preserving portal navigation. Otherwise, the content is displayed within the portlet you will create for this application.

   f. Click **Finish**.

4. Create a portlet based on the above Web service:
   a. Click **Administration**.
   b. Navigate to or create the administrative folder for this portlet.
   c. In the Create Object drop-down list, select **Portlet**.
   d. In the Choose Template or Web Service dialog box, select the Web service you created in the previous step, and click **OK**.
   e. Configure the portlet as described in the online help.
   f. Click **Finish**.
5. Add the portlet to My Pages or communities.

   You can let users add the portlet on their own (**My Pages | Add Portlets** or **My Communities | Add Portlets**), or you can make the portlet mandatory. See "Defining Mandatory Portlets" on page 4-27.

Portlet Content Caching

Caching some portlet content can greatly improve the performance of your portal. When you cache portlet content, the content is saved on the portal for a specified period of time. Each time a user requests this content—by accessing a My Page or community page that includes the cached portlet—the portal delivers the cached content rather than running the portlet code to produce the content.

When you create a portlet, you can specify whether or not the portlet should be cached, and if it is cached, for how long. You should cache any portlet that does not provide user-specific content. For example, you would cache a portlet that produces stock quotes, but not one that displays a user e-mail box.

If you develop portlet code, you can and should define caching parameters.

For more information on portlet caching, refer to the BEA AquaLogic User Interaction Development Center (http://dev2dev.bea.com/aluserinteraction/) or the documentation provided with the portlet software.

Portlet Preferences

You can configure the following types of preferences for portlets.

Table 4-4 Portlet Preferences

Preference Type	Description
Administrative Preferences *E-mail portlet example:* Setting which e-mail server to connect to	These preferences are set by the portlet creator on the Main Settings page of the Portlet Editor. They affect everyone's view of the portlet. Users with administrative rights can edit these preferences from **My Pages | Edit Portlet Preferences**, or by clicking the edit icon in a portlet's titlebar.
Personal Preferences *E-mail portlet example:* Setting how many e-mails are displayed in the portlet	These preferences are set by the user from **My Page | Edit Portlet Preferences** or **My Communities | Edit Portlet Preferences**. These preferences affect that user's view of the portlet.
Community Preferences *E-mail portlet example:* Setting a specific public e-mail folder to display, and a shared login/password for that folder	These preferences are set by the community administrator on the **Portlet Preferences** page of the Community Editor. This page can include community preferences for portlets specific to that community or for other portlets. Community preferences affect everyone's view of portlets in that community. When in a community, community administrators can edit these preferences from **My Communities | Edit Portlet Preferences**, or by clicking the edit icon in a portlet's titlebar.
Portlet Template Preferences *Example:* Which portlet Web service to use	These preferences are set by the portlet template creator on the **Main Settings** page of the Portlet Template Editor. They affect the portlet template itself and all portlets created from that template. If you change these preferences after portlets have been created from this template, the change will affect only new portlets. Portlets created from this template before the change was made will not be affected.

Creating Portlet Web Services, Portlet Templates, Portlets, and Portlet Bundles

Creating Portlet Web Services

Portlet Web services allow you to specify functional settings for your portlets in a centralized location, leaving the display settings to be set in each associated portlet.

Intrinsic portlets are installed on the portal.

To create a Web service for an intrinsic portlet:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Web Service - Intrinsic Portlet**.
4. Define the portlet Web service as described in the online help.
5. Click **Finish**.

Remote portlets extend the base functionality of the default portal and are hosted on a remote server.

To create a Web service for a remote portlet:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Web Service - Remote Portlet**.
4. Define the portlet Web service as described in the online help.
5. Click **Finish**.

Creating Portlet Templates

Portlet templates allow you to create multiple instances of a portlet, each sharing much of the basic configuration but displaying slightly different information. For example, you might want to create a Regional Sales portlet template, from which you could create different portlets for each region to which your company sells. You might even want to include all the Regional Sales portlets on one page for an executive overview.

After you have created a portlet from a portlet template, there is no further relationship between the two objects. If you make changes to the portlet template, these changes are not reflected in the portlets already created with the template.

To create a portlet template:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Portlet Template**.
4. Define your portlet template as described in the online help.
5. Click **Finish**.

Creating Portlets

To create a portlet (intrinsic or remote):

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Portlet**.
4. Define your portlet as described in the online help.
5. Click **Finish**.

Creating Portlet Bundles

Portlet bundles are groups of related portlets, packaged together for easy inclusion on My Pages or community pages. When users add portlets to their My Pages or community pages, they can add all the portlets in a bundle or select individual portlets from a bundle. You might want to create portlet bundles for portlets that have related functions or for all the portlets that a particular group of users might find useful. This makes it easier for users to find portlets related to their specific needs without having to browse through all the portlets in your portal.

To create a portlet bundle:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Portlet Bundle**.
4. Add portlets to the bundle.
5. Click **Finish**.

Requiring and Recommending Portlets

This section describes how to require or recommend portlets to groups or users. It includes the following topics:

- "Defining Mandatory Portlets" on page 4-27
- "Recommending Portlets" on page 4-27

Defining Mandatory Portlets

You can force users or groups to include a portlet on their default My Page by making it mandatory for those users or groups. Mandatory portlets display above user-selected portlets. Users cannot remove mandatory portlets from their My Pages.

Because mandatory portlets are added to My Pages, the following portlet types cannot be mandatory: Header, Footer, Content Canvas, and community-only portlets.

To make a portlet mandatory for a particular group:

1. Click **Administration**.
2. Navigate to the portlet you want to make mandatory and click its name.
3. Click the **Security** page.
4. Define mandatory settings for users and groups as described in the online help.
5. Click **Finish**.

Recommending Portlets

You can *recommend* portlets to encourage users to add them to their My Pages. Users can recommend any portlet that can be added to a My Page and to which they have access.

Because recommended portlets are added to My Pages, the following portlet types cannot be recommended: Header, Footer, Content Canvas, and community-only portlets.

To recommend a portlet:

1. From the Add Portlets page or from within the Portlet Editor, click ✉. This displays text, including a URL, that you can paste into an e-mail and send to users.
2. E-mail the link or add it to a community links portlet.

Adding Multiple Portlets to Multiple Groups' My Pages

You can add one or more portlets to one or more groups' My Pages as a bulk operation.

To add multiple portlets to multiple groups:

1. Click **Administration**.
2. Navigate to an administrative folder containing portlets, or search for portlets.
3. Select one or more portlets, and click →👤 .
4. Select the portlets to add to selected groups' My Pages as described in the online help.
5. Click **Finish**. Users will be able to remove the portlets you push to them in this way from their My Pages.

Managing Communities

This section describes how to set up portal communities and how to enable content managers to create and manage additional communities. It includes the following topics:

About Communities

A community is similar to a My Page in that it displays portlets. However, communities provide content and services to a *group* rather than to just an individual user.

You might create communities based on departments in your company. For example, the Marketing department might have a community containing press information, leads volumes, a trade show calendar, and so on. The Engineering department could have a separate community containing project milestones, regulatory compliance requirements, and technical specifications.

You might create communities based on projects your company is working on. For example, a member of the Professional Services department working with a customer to deploy a system could create a community where that group could collaborate on deployment issues. You would probably delete this type of community when the project ends.

Each community is based on a *community template*, which consists of one or more *page templates*, which can include portlets. Each page template you add to a community (either through the community template or through the community itself) appears as a link at the top of the community.

Individual community pages have their own security settings, so you can use pages, as well as subcommunities, to control access to different areas of the community.

The first page you add becomes the community Home Page—the default page that displays to users when they visit your community.

Communities can also include the following features:

- Subcommunities, which allow you to set different security on areas of your community
- Headers and footers, which visually differentiate the community from the rest of the portal
- Portlets and groups created specifically for the community
- A community Knowledge Directory, which can include a list of community members, links to related Web sites, portal documents, and other community pages

Creating Page Templates and Community Templates

Creating Page Templates

Page templates include portlets and layout settings that are used as the basis to create pages in communities. A single page template can be used by many different communities, allowing you to keep similar types of pages looking analogous. For example, you might want each department to create a community in which the first page lists the general duties of the group, the department members, and the current projects owned by the department.

Each page template specifies a particular page layout. The page layout determines where particular types of portlets can be displayed on the page. For example, if you want to include a Content Canvas portlet on a page, you must choose a page layout that allows you to do so.

There are three possible parts to a page layout, which are combined in different ways in the available page layouts:

- wide column, to which wide or narrow portlets can be added; each page layout includes a wide column.
- narrow column, to which only narrow portlets can be added; some page layouts include a narrow column on the right, left, or both sides of the wide column.
- content canvas area, to which a Content Canvas portlet can be added; some page layouts include a content canvas area, which sits at the top of the page layout and can straddle some or all of the columns.

The following page layouts are available (the dark gray sections are content canvas areas).

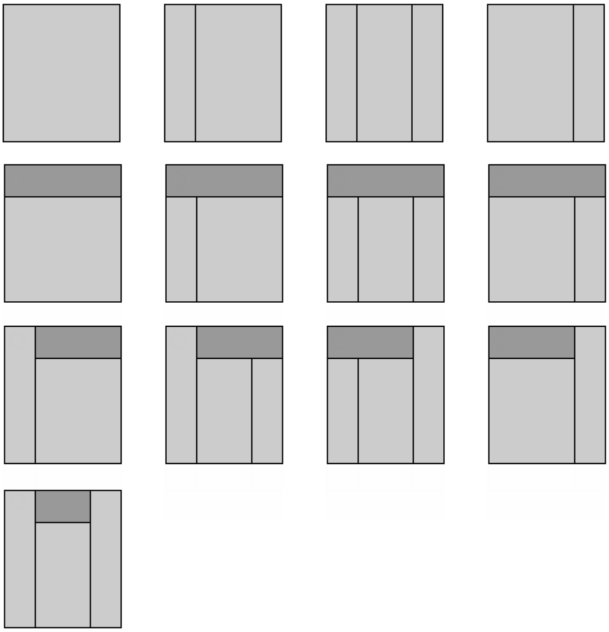

To create a page template:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Page Template**.
4. Define your page template as described in the online help.
5. Click **Finish**.

When you create community pages based on a page template, you have the option to have the pages you are creating inherit any future changes to the template. For example, if you choose to have a community page inherit the template, when you add a portlet on the template, the portlet is added to the associated community pages.

Creating Community Templates

When you create a community, it is based on a community template. Community templates allow you to define the minimum requirements for communities, including page templates and, optionally, a header or footer for the community page. Community creators can add new content and services, but cannot remove the content, services, or design provided by the community template. A single community template can be used by many different communities, allowing you to keep similar types of communities looking similar. For example, you might want all communities based on departments to look similar and contain similar content, while you might want communities based on projects to look different.

You can add Header and Footer portlets to a community in one of two ways:

- You can add them to the community template, forcing each community created from the template to display the same basic header and footer.
- If no header or footer is specified in the community template, you can add them directly to the community.

If you use branding portlets (the Header, Footer, and Content Canvas portlets provided with your portal), community administrators can edit portlet settings such as the text, icon, and color of the header or footer. This allows communities to have similar, but distinct headers and footers.

To create a community template:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Community Template**.
4. Define your community template as described in the online help.
5. Click **Finish**.

If you create a community based on a community template, you can choose to have the community you are creating inherit any future changes to the template. If you choose to inherit changes, any change applied to the community template affects the community. For example, if a page template is removed from a community template, the page created from this template will be removed from your community as well.

Creating Communities

You must have Edit privilege to the community and Create Communities activity right to create a community or a subcommunity.

To create a community:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Community**.
4. Define your community as described in the online help.
5. Click **Finish**.

Creating Subcommunities

Subcommunities (along with Pages) allow you to create separately-secured subsections in a community, so it can have a more restrictive security than the main community. For example, you might have a Marketing Community that includes an Advertising Subcommunity. This subcommunity might have distinct owners or might be accessible to only a subset of the Marketing Community.

A subcommunity is just a community folder stored in another community folder. Therefore, the subcommunity inherits the security and design of the parent community, but you can then change these settings to suit the needs of the subcommunity. You can also change the relationships of communities and subcommunities just by rearranging the folder structure.

Note: If you choose to display a community Knowledge Directory in the subcommunity, it is separate from the community Knowledge Directory in the parent community.

User community access determines subcommunity access:

- All users that have access to the subcommunity can view or join subcommunities through the Join Communities link on the My Communities section of the portal.
- Users that have access to the whole community can also access subcommunities through the subcommunities tab in the parent community.

You must have Create Communities activity right to create a subcommunity.

To create a subcommunity in a new community:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Community**.
4. Define the pages for your community as described in the online help.
5. Click the subcommunities page.
6. Click **Create Subcommunity**.
7. Define the Subcommunity as described in the online help.
8. Click **Finish** in the Subcommunity Editor.
9. Click **Finish** in the Community Editor.

Note: Subcommunities can be nested up to 10 levels deep.

Caution: The Related Communities tab displays peer communities—the communities that are stored in the same administrative folder as your community. For this reason, consider carefully where to store communities and the administrative folder structure necessary to make related communities useful.

Community Pages

Community pages appear as links in a community. You can create a community page in a community folder or in a community editor. Like communities, pages are based on templates from which you can choose whether or not to inherit future changes. Like other portal objects, community pages can be copied (to another community folder), localized, migrated, and can have unique security settings.

To create a community page:

1. Click **Administration**.
2. Go to a community folder.
3. From the **Create Object** drop-down list, select **Page**.
4. Choose whether or not to **Inherit the Template**.

   If you inherit the page template, you cannot delete portlets associated with the page template, but you can add portlets to the page created from the template. If you do not inherit the page template, you can delete portlets associated with the template, add new portlets, and change the page layout.
5. Define the page as described in the online help.
6. Click **Finish**.

Creating Community Groups

A group is a set of portal users to whom you grant specific access privileges. You can create community groups without affecting portal groups. You create community groups so that you can easily assign responsibilities to community members. For example, you might have a group that is responsible for maintaining schedules in the community. If you later want to make your community group available outside of the community, you can move the group from the community folder to another administrative folder.

You must have the Create Groups activity right to create a community group.

To create a community group:

1. Click **My Communities** and select the community you want to edit.
2. Open the Community Editor by clicking ✎ **Edit this Community** on the right.
3. On the left, under Edit Community Settings, click **This Community's Groups**.
4. Create the Community Group as described in the online help.
5. To save the group, click **Finish** and complete the **Save Object** dialog box.

Creating Community Portlets

You can create and manage portlets in the community. You need access to portlet Web services or portlet templates and must have Create Portlets activity right to create portlets.

Portlets created in the Community Editor are only available within the community. If you later want to make portlets available outside of the community, you can move the portlet from the community folder to a higher level administrative folder.

Note: Removing community portlets from the community deletes them from the portal.

To create portlets available only to this community:

1. Click **My Communities** and select the community you want to edit.
2. Open the Community Editor by clicking ✎ **Edit this Community** on the right.
3. On the left, under Edit Community Settings, click **This Community's Portlets**.
4. Create Community portlets as described in the online help.
5. After creating your portlet, click **Finish**.

To display these portlets to community users, you must add these portlets to the appropriate community page.

Managing Community Users and Groups

Community membership controls the community selection in the My Community section of the portal. It also controls the mandatory tabs in the community navigation. You can control who can join, edit, and administer the community.

Users must have Select rights to join the community.

To change the access rights of each member of the community:

1. Click **Administration**.
2. Open the administration folder that contains the community.
3. Click the community.
4. Click ✎.
5. Under Edit Standard Settings in the Community Editor, click **Security**.
6. Configure the ACL as described in the online help.
7. Click **Finish**.

Requiring Communities for Groups

You can make a community mandatory for the members of one or more groups. Users cannot remove themselves from mandatory communities. You can also display tabs for mandatory communities in the banner at the top of the portal, alongside the My Pages and My Communities tabs.

To make a community mandatory for a particular group:

1. Click **Administration**.
2. Open the administration folder that contains the community.
3. Click the community.
4. Click ✎.
5. Under Edit Standard Settings in the Community Editor, click **Security**.
6. Define mandatory settings for users and groups.
7. Click **Finish**.

Recommending Communities

You can *recommend* communities to encourage users to join them. Users can recommend any community to which they have access.

To recommend a community:

1. From the Join Communities page or from within the Community Editor, click to display text, including a URL, that you can paste into an e-mail and send to users.
2. E-mail the link or add it to a community links portlet.

Subscribing Multiple Groups to Multiple Communities

You can subscribe one or more groups to one or more communities as a bulk operation.

To add multiple communities to multiple groups:

1. Click **Administration**.
2. Navigate to an administrative folder containing communities, or search for communities.
3. Select one or more communities, and click .
4. Select the communities to add to selected groups as described in the online help.
5. Click **Finish**. Users will be able to unsubscribe from the communities you push to them in this way.

Managing the Community Knowledge Directory

The community Knowledge Directory is an optional part of a community that allows you to provide access to additional community-specific content through a folder hierarchy. There are two folders that are always present in a community Knowledge Directory:

- Members, which displays a list of all the users that belong to the community (either individually or as part of a group for which the community is mandatory).
- Subcommunities, which displays a list of this community's subcommunities. If there are no subcommunities, this folder still displays.

You can also create your own folders and fill them with links to Web sites, user profiles of community experts, documents from the portal Knowledge Directory, and pages in other communities. Users can browse these links from the community Knowledge Directory, or you can display the links in a Community Links portlet.

Note: You might want to create a Community Links portlet that includes links to important secondary community pages and then invite users to add the portlet to their My Pages. This provides direct access to those community pages; users do not have to navigate to the community home page and then click the community page they want.

To create community Knowledge Directory folders and or to create a community Links portlet:

1. Click **My Communities** and then the community you want to edit.
2. Click **Community Members and Knowledge Directory** to the right of the community page links.
3. Click **Edit** (not Edit This Community).
4. Click .
5. Type a name and description for the folder and click **OK**.
6. Open the folder by clicking its name.
7. Add links to the folder:
   - To add links to Web sites, click **Add Links** and complete the settings as described in the online help.
   - To add links to specific user profiles, click **Add Experts** and complete the settings as described in the online help.
   - To add links to documents in the portal Knowledge Directory, click **Add Documents** and complete the settings as described in the online help.
   - To add links to pages in another community, click **Add Pages** and complete the settings as described in the online help.
8. If you want to display the links in this folder in a portlet in your community, click **Content Snapshot** and specify the community page to which you want to add the portlet.

Enabling Document Discovery with Content Crawlers and Content Services

This section describes how to crawl WWW locations, file system locations, and back-end content and mail servers to make documents in these repositories available through portal links. This section includes the following topics:

- "About Content Crawlers" on page 4-40
- "Developing Content Services to Target Specific Content" on page 4-42
- "Configuring a Content Crawler" on page 4-43
- "Testing Content Crawlers" on page 4-45

For a summary of AquaLogic Interaction content crawlers, as well as guidelines on best practices for deploying content crawlers, see the *Deployment Guide for BEA AquaLogic User Interaction G6*.

For information on installing and configuring AquaLogic Interaction remote content crawlers, follow the product documentation included with your software instead of the documentation in this guide.

About Content Crawlers

Content crawlers import, from back-end content sources, document records that contain descriptive information, such as content type and properties, document ACL (read access only), and links to these documents into Knowledge Directory subfolders according to property-based filters, as shown in the following figure.

Figure 4-1 Content Crawler Flow Chart

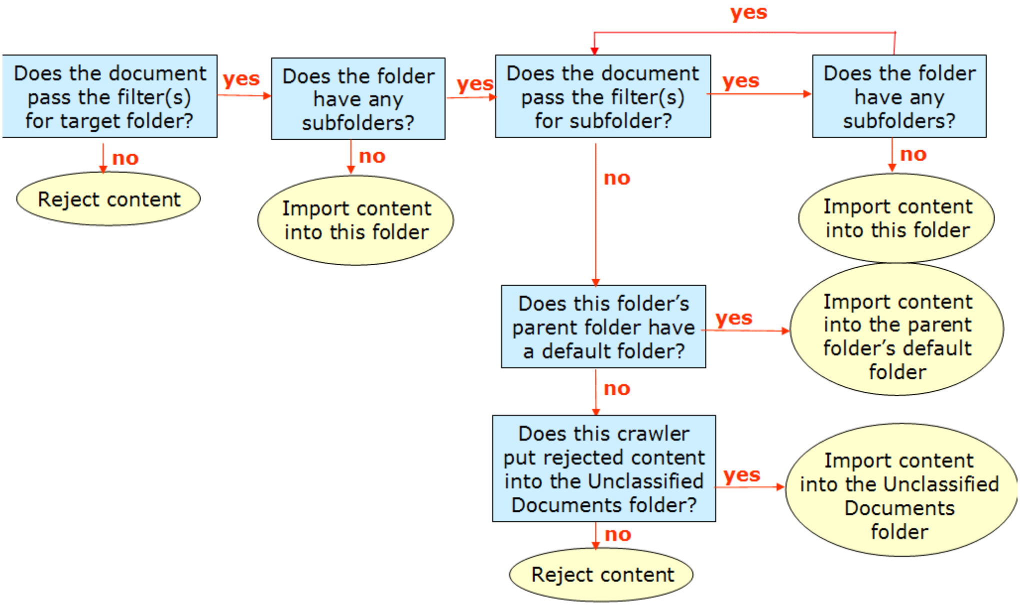
Does the document pass the filter(s) for target folder?
yes
no
Reject content
Does the folder have any subfolders?
yes
no
Import content into this folder
Does the document pass the filter(s) for subfolder?
yes
yes
no
Does the folder have any subfolders?
no
Import content into this folder
Does this folder's parent folder have a default folder?
yes
Import content into the parent folder's default folder
no
Does this crawler put rejected content into the Unclassified Documents folder?
yes
Import content into the Unclassified Documents folder
no
Reject content

There might be cases where imported content does not pass the filters on any folder, even the destination folder. In these cases you can either choose to not import the rejected content, or to place the rejected content into the Unclassified Documents folder. If you place the rejected content into the Unclassified Documents folder, you can view this content in the Knowledge Directory edit mode. You can later move these document records into the Knowledge Directory.

The following table summarizes the metadata AquaLogic Interaction content crawlers can import.

Table 4-5 Types of Metadata that can be Imported by AquaLogic Interaction Content Crawlers

Content Crawler	Import Links to Documents	Import Document Security	Import Folder Security
Web Content Crawler	Yes	No	No
Remote Windows Content Crawler	Yes	Yes (Windows)	Yes (Windows)
Remote Exchange Content Crawler (Windows)	Yes	No	No
Remote Lotus Notes Content Crawler (Windows)	Yes	Yes	No
Remote Documentum Content Crawler	Yes	Yes	Yes

Content crawlers also index the full document text, and this index is used by the Search Service to make documents available through the Search tool.

Developing Content Services to Target Specific Content

To facilitate maintenance, we recommend you implement several instances of each content crawler type, configured for limited, specific purposes.

For file system content crawlers, you might want to implement a content crawler that mirrors an entire file system folder hierarchy by specifying a top-level starting point and its subfolders. Although the content in your folder structure is available on your network, replicating this structure in the portal offers several advantages:

- Users are able to search and access the content over the Web.
- Interested users can receive regular updates on new content with snapshot queries.
- You can use default profiles to direct new users to important folders.

However, you might find it easier to maintain controlled access, document updates, or document expiration by creating several content crawlers that target specific folders.

If you plan to crawl WWW locations, familiarize yourself with the pages you want to import. Often, you can find one or two pages that contain links to everything of interest. For example, most companies offer a list of links to their latest press releases, and most Web magazines offer a list of links to their latest articles. When you configure your content crawler for this source, you can target these pages and exclude others to improve the efficiency of your crawl jobs.

If you know that certain content will no longer be relevant after a date—for example, if the content is related to a fiscal year, a project complete date, or the like—you might want to create a content crawler specifically for the date-dependent content. When the content is no longer relevant, you can run a job that removes all content created by the specific content crawler.

For remote content crawlers, you might want to limit the target for mail content crawlers to specific user names; you might want to limit the target for document content crawlers to specific content types.

For additional considerations and best practices, see the *Deployment Guide for BEA AquaLogic User Interaction G6*.

Configuring a Content Crawler

Content services allow you to specify general settings for your remote content repository, leaving the target and security settings to be set in the associated remote content crawler. This allows you to crawl multiple locations in the same content repository without having to repeatedly specify all the settings.

If you plan to use an AquaLogic Interaction content Web service (AquaLogic Interaction Content Services) to crawl document repositories, follow the product documentation provided with that software instead of the procedures in this guide. AquaLogic Interaction remote content crawlers include a migration package that enables you to import pre-configured remote server and Web service objects.

The following table describes the steps you take to configure a target-specific content service.

Table 4-6 Configuring a Target-Specific Content Service

Basic Step	Procedure
Create a remote server.	1. Click **Administration**. 2. Navigate to or create the administrative folder for content service objects. 3. In the Create Object drop-down list, select **Remote Server**. 4. Configure connection information for the remote server as described in the online help. 5. Click **Finish**.
Create a content service.	1. Click **Administration**. 2. Navigate to the administrative folder for the group of content service objects you are configuring. 3. In the Create Object drop-down list, select **Web Service - Content**. 4. Configure connection information for the Web service as described in the online help. 5. Click **Finish**.
Configure a content source.	1. Click **Administration**. 2. Navigate to the administrative folder for the group of content service objects you are configuring. 3. In the Create Object drop-down list, click **Content Source - WWW** for a Web site or **Content Source - Remote** for a back-end content repository. 4. Configure connection information for the Web service as described in the online help. 5. Click **Finish**.
Configure a document property map.	If you need to define a new content type and properties for your content, follow the procedures in "Configuring Content Types and Document Properties" on page 4-3.

Table 4-6 Configuring a Target-Specific Content Service

Basic Step	Procedure
Configure a content crawler and crawl job. When you configure a content crawler, you specify: • location of source documents • content types that determine how source document properties are mapped to portal properties • document security settings • document sorting, refresh, and purge settings	1. Click **Administration**. 2. Open an administrative folder. 3. In the Create Object drop-down list, click **Content Crawler - WWW** for a Web site or **Content Crawler - Remote** for a back-end content repository. 4. Define your Web content crawler as described in the online help. 5. On the Set Job page, add this operation to a Job and schedule the Job to run. 6. Click **Finish**. **Note:** For information on how to organize crawled content in folders, see "Using Filters to Organize Crawled Content" on page 4-16.
To import security, the domain and group information for the source being crawled must be mapped to an authentication source prefix in the global ACL sync map. If you run a content crawler and find that some or all of the security has not been imported, map the domain in the global ACL sync map and run the content crawler again.	1. Click **Administration**. 2. In the Select Utility drop-down list, choose **Global ACL Sync Map**. 3. Add the domain prefix and group mappings for the content source as described in the online help. 4. Click **Finish**.

Testing Content Crawlers

Before you have a content crawler import content into the public folders of your portal, test it by running a job that crawls document records into a temporary folder.

When you create the test folder, remove the Everyone group, and any other public groups, from the Security page on the folder to ensure that users cannot access the test content.

The following table provides a summary test plan for your content crawlers.

Table 4-7 A Test Plan for Content Crawlers

Test Objective	Steps
Make sure the content crawler creates the correct links.	Examine the target folder and ensure the content crawler has generated records and links for desired content and has not created unwanted records and links. If you iterate this testing step after modifying the content crawler configuration, make sure you delete the contents of the test folder and clear the deletion history for the content crawler as described in "Clearing the Deletion History" on page 4-47.
Make sure the content crawler creates correct metadata.	Make sure that all documents are given the right content types, and that these content types correctly map properties to source document attributes. Go to the Knowledge Directory, and look at the properties and content types of a few of the documents this content crawler imported to see if they are the properties and content types you expected. To view the properties and content type for a document: 1. Click **Directory** and navigate to the folder that contains the document whose properties and content type you want to view. 2. Click **Properties** under the document to display the information about the document. The properties are displayed in a table along with their values. The content type is displayed at the bottom of the page. If you iterate this testing step after modifying the content crawler configuration, make sure you configure the content crawler to refresh these links. For information on refreshing links, see "Keeping Document Records Up-to-Date" on page 4-47.
Test properties, filters, and search.	To test that document properties have been configured to enable filters and search, browse to the test folder, and perform a search using the same expression used by the filter you are testing. Either cut and paste the text from the filter into the portal search box or use the Advanced Search tool to enter expressions involving properties. Select the **Search Only in this Folder** option. The links that are returned by your search are for the documents that will pass your filter.

Maintaining Content Imported by Content Crawlers

This section describes how to maintain document records imported by content crawlers. It includes the following topics:

- "Keeping Document Records Up-to-Date" on page 4-47
- "Clearing the Deletion History" on page 4-47
- "Removing Document Records" on page 4-48

Keeping Document Records Up-to-Date

The Document Refresh Agent is an intrinsic job that updates the records in the Knowledge Directory. The Document Refresh Agent examines every link in your portal. For each link, the document refresh agent first determines if the link requires refreshing based on the document record setting set when the file was uploaded or by the content crawler that created the link.

To administer refresh settings for the content crawler:

1. Click **Administration**.
2. Navigate to the content crawler whose document records you want to refresh.
3. On the Document Settings page and Advanced Settings page, configure document refresh attributes.
4. Click **Finish**.

Clearing the Deletion History

Content crawlers keep a history of actions performed on crawled document records, including the deletion history. If you delete records, the content crawler remembers that the content was imported and deleted and it will not attempt to re-import this content. If you later decide to import records for that content, you must clear the deletion history.

To clear the deletion history:

1. Click **Administration**.
2. Navigate to the content crawler and click its name.
3. On the Advanced Settings page, click **Clear Deletion History**.
4. Click **Finish**.

Removing Document Records

By carefully targeting your content crawler to generate content on only one topic, you allow for the easy removal of a topic that becomes irrelevant, without disturbing unrelated content.

To remove all the content ever imported by a particular content crawler:

1. Click **Administration**.
2. Navigate to the content crawler whose imported links you want to remove and click its name. This launches the Content Crawler Editor.
3. On the Document Settings page, set the generated content to be deleted immediately, and select **Apply these settings to existing documents created by this content crawler**.
4. Click **Finish**.

   The next time the Document Refresh Agent runs, it will delete all of the records created by this content crawler.

Working with Search

This section describes how to implement search for documents that reside in the Knowledge Directory, in communities, or in the collection of crawled links. It includes the following topics:

- "Customizing Search Service Behavior" on page 4-49
- "Managing Grid Search" on page 4-58
- "Creating Snapshot Queries" on page 4-69
- "Implementing Federated Searches" on page 4-70

Customizing Search Service Behavior

This section describes how to customize portal search. It includes the following topics:

- "Configuring Best Bets and Top Best Bets" on page 4-49
- "Modifying the Properties Searched and the Relevance Weight for Properties" on page 4-51
- "Enabling Spell Correction" on page 4-52
- "Implementing a Search Thesaurus" on page 4-53
- "Customizing Categorization of Search Results" on page 4-57

For information on default behavior for search syntax and results ranking, see Appendix E, "Default Behavior of Search Service."

Configuring Best Bets and Top Best Bets

You configure best bets with the Search Results Manager. Best bets associate specific search phrases you specify with a set of search results, in rank order. In addition, users can go directly to the highest ranking result, the top best bet, instead of seeing the normal search results.

When end-users enter a banner search query that matches a best bet search phrase, the best bet results appear as the first results in the relevance-ranked result list. The phrase "Best Bet" appears next to each best bet result to inform the user that the result has been judged especially relevant to his or her query.

Best bets apply only to the portal banner search box and search portlet. Best bets are not used by other portal search interfaces, such as advanced search and object selection search.

Note: Best bets are case-insensitive.

To create a best bet:

1. Click **Administration**.
2. From the Select Utility drop-down list, select **Search Results Manager**.
3. Launch the Best Bet Editor by clicking **New Best Bet**.
4. Complete the best bet settings as described in the online help.
5. Click **Finish** to save your best bet settings.
6. Click **Finish** in the Search Results Manager.

You can create hundreds of best bets, each mapping to a maximum of 20 results.

Since best bets are handled by the Search Service and are not managed portal objects, best bets do not migrate from development to production environments; you must re-create them in the production environment.

Working With Top Best Bets

The highest ranking best bet result for a given search term is the top best bet. If best bets are set for a term, instead of seeing search results, users can go directly to the top best bet result (an object such as a community or document) by doing one of the following:

- Type “>” as a prefix to the search term in the banner search box or search portlet, and click ▶. For example:

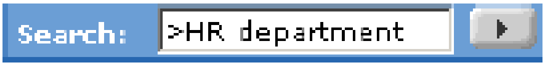

Search: >HR department ▶

- Insert “tbb=*SearchTerm*” in the portal URL (may include spaces). For example:

  `http://portal.company.com/portal/server.pt?tbb=HR department`

  Note: If your search term contains spaces, they will be converted to `%20`.

- If the ◈ button is visible, enter the search term and click ◈ instead of ▶.

  Note: For information on how to enable this button using the Search tag, see the *BEA AquaLogic User Interaction Development Center* at http://dev2dev.bea.com/aluserinteraction/.

If there are no best bets set for the term the user entered, the search results for the term are displayed instead.

If an object is a top best bet for any search terms, those terms are listed on the Properties and Names page of the object’s editor.

Modifying the Properties Searched and the Relevance Weight for Properties

When a user enters a query into a search box in the portal, the portal searches the properties specified on the Banner Fields page of the Search Results Manager. The default banner field properties are Name, Description, and Full-Text Content. However, you can also add other properties, such as Keyword, Department, or Author, to further refine the search results.

Another way of controlling the search results is by modifying the relevance weight for banner field properties. Overweighting a property increases its relevancy ranking; and underweighting it decreases it. For example, you can manipulate the search to first return documents whose *content* matches the search string (by overweighting the Full-Text Content property) followed by documents whose *name* matches the search string (by underweighting the Name property). When users type `widgets`, documents with `widgets` in the content appear first in a relevance-ranked search result; they are followed by documents or files with `widgets` in their names.

Banner field settings apply to the banner search box, advanced search, object selection search, or any other portal search interfaces.

To configure the weights of existing banner fields:

1. Click **Administration**.
2. From the Select Utility drop-down list, select **Search Results Manager**.
3. Under Edit Utility Settings, click **Banner Fields**.
4. Complete the Banner Field settings as described in the online help.
5. Click **Finish**.

To add new banner fields:

1. Click **Administration**.
2. From the Select Utility drop-down list, select **Search Results Manager**.
3. Under Edit Utility Settings, click **Banner Fields**.
4. Click **Add Field**.
5. From the drop-down list that appears, select the banner field that you want to add.
6. Complete the Banner Field settings as described in the online help.
7. Click **Finish**.

Since banner fields and relevance weights are a Search Service setting and not managed portal objects, the settings do not migrate from development to production environments; you must re-create them in the production environment.

Enabling Spell Correction

Automatic spell correction is applied to the individual terms in a basic search when the terms are not recognized by the Search Service. Spell correction is not applied to quoted phrases.

For example, if a user queries for `portel server` but the term `portel` is unknown to the Search Service, items matching the terms `portal` and `server` would be returned instead. The same applies to Internet style mode and query operators mode. So, for instance, a search for `portel <NEAR> server` would return documents containing the terms `portal` and `server` in close proximity, but only if there are no matches for `portel` and `server` in close proximity.

Automatic spell correction is enabled by default. You can disable it from the Search Results Manager in the administrative portal user interface.

To disable the automatic spell correction:

1. Click **Administration**.
2. From the Select Utility drop-down list, select **Search Results Manager**.
3. Under Edit Utility Settings, click **Thesaurus and Spell Correction**.
4. Clear the **Apply Spell Correction** check box.
5. Click **Finish**.

Implementing a Search Thesaurus

The Search Service allows you to create a thesaurus (or synonym list), load it into the server, and enable thesaurus expansion for all user queries. Thesaurus expansion allows a term or phrase in a user's search to be replaced with a set of custom related terms before the actual search is performed. This feature improves search quality by handling unique, obscure, or industry-specific terminology.

For example, with conventional keyword matching, a search for the term `gadgets` might not return documents that discuss `portlets` or `Web services`. But, by creating a thesaurus entry for `gadgets`, it is possible to avoid giving users zero search results because of differences in word usage. The entries allow related terms or phrases to be weighted for different contributions to the relevance ranking of search results. For example, `gadgets` is not really a synonym for `Web services`, so a document that actually contains `gadgets` should rank higher than one that contains `Web services`.

The entries are lower-case, comma-delimited lists of the form:

```
gadgets,portlets,web services[0.5]
```

In this example, the number [0.5] corresponds to a non-default weighting for the phrase `web services`.

Note: Thesaurus entries must be lower-case.

Thesaurus entries can be created to link closely related terms or phrases, specialized terminology, obsolete terminology, abbreviations and acronyms, or common misspellings. The expansion works by simply replacing the first term in an entry with an OR query consisting of all the terms or phrases in the entry. The weights are then taken into consideration when matching search results are ranked.

The thesaurus expansion feature is best used for focused, industry- or domain-specific examples. It is not intended to cover general semantic relationships between words or across languages, as with a conventional paper thesaurus. Although the Search Service thesaurus expansion can definitely improve search quality, adding entries for very general or standard terms can actually degrade search quality if it leads to too many search result matches.

Enabling the Thesaurus

To enable the thesaurus:

1. Click **Administration**.
2. From the Select Utility drop-down list, select **Search Results Manager**.
3. Under Edit Utility Settings, click **Thesaurus and Spell Correction**.
4. Select **Use the Thesaurus**.
5. Click **Finish**.

After you enable this feature, you must create the synonym list in the database, described next.

Setting Up the Synonym List for the Thesaurus

To set up the search thesaurus:

1. Create a comma-delimited, UTF-8 file containing the desired thesaurus entries.

   Note: Thesaurus entries must be in lower-case.

   The thesaurus is a comma-delimited file, also known as a CDF. Each line in the file represents a single thesaurus entry. The first comma-delimited element on a line is the name of the thesaurus entry. The remaining elements on that line are the search tokens that should be treated as synonyms for the thesaurus entry. Each synonym can be assigned a weight that determines the amount each match contributes to the overall query score. For example, a file that contains the following two lines defines thesaurus entries for couch and dog:

   ```
   couch,sofa[0.9],divan[0.5],davenport[0.4]
   dog,canine,doggy[0.85],pup[0.7],mutt[0.3]
   ```

   Searches for `couch` generate results with text matching terms `couch`, `sofa`, `divan`, and `davenport`. Searches for `dog` generate results that have text matching terms `dog`, `canine`, `doggy`, `pup`, and `mutt`. In the example shown, the term `dog` has the same contribution to the relevance score of a matching item as the term `canine`. This is equivalent to a default synonym weighting of 1.0. In contrast, the presence of the term `pup` contributes less to the relevance score than the presence of the term `dog`, by a factor of 0.7 (70%).

   The example thesaurus entries constitute a complete comma-delimited file. No other information is needed at the beginning or the end of the file.

Entries can also contain spaces. For example, a file that contains the following text creates a thesaurus entry for New York City:

```
new york city,big apple[0.9],gotham[0.5]
```

Searches for the phrase `“new york city”` will return results that also include results containing `“big apple”` and `“gotham.”` Thesaurus expansion for phrase entries only occurs for searches on the complete phrase, not the individual words that constitute the phrase. Similarly, the synonym entries are treated as phrases and not as individual terms. So while a search for `“new york city”` returns items containing `“big apple”` and `“gotham,”` a search for `new` (or for `york`, or for `city`, or for `“new york”`) will not. Conversely, an item that contains `big` or `apple` but not the phrase `“big apple”` will not be returned by a search for `“new york city.”`

Comma-delimited files support all UTF8-encoded characters; they are not limited to ASCII. However, punctuation should not be included. For example, if you want to make `ne’er-do-well` a synonym of `wastrel`, replace the punctuation with whitespace:

```
wastrel,ne er do well[0.7]
```

This matches documents that contain `ne’er-do-well`, `ne er do well` or some combination of these punctuations and spaces (such as `ne’er do well`). If you want your synonym to match documents that contain `neer-do-well`, which does not separate the initial `ne` and `er` with an apostrophe, you must include a separate synonym for that, such as:

```
wastrel,ne er do well[0.7],neer do well[0.7]
```

Finally, comment lines can be specified by beginning the line with a “#”:

```
# furniture entries

couch,sofa[0.9],divan[0.5],davenport[0.4]

#chair,stool[5.0]

# animal entries

dog,canine[0.9],doggy[0.85],pup[0.7],mutt[0.3]
```

In this example, the Search Service parses two thesaurus entries: couch and dog. There will be no entry for chair.

These examples are of entries that contain only ASCII characters. This utility supports non-ASCII characters as well, as long as they are UTF8-encoded.

Note: Some editors, especially when encoding UTF-8, insert a byte order mark at the beginning of the file. Files with byte order marks are not supported, so remove the byte order mark before running the customize utility.

A CDF thesaurus file can have at most 50,000 distinct entries (lines). Each entry can have at most 50 comma-delimited elements (including the name of the entry). If either of these limits are exceeded, the customize utility will exit with an appropriate error message.

2. Stop the Search Service.

   The comma-delimited file is converted to a binary format in the next step. The conversion removes and replaces certain files used by the Search Service, and this removal and replacement cannot be done while the Search Service is running.

3. At a command prompt, run the customize utility.

   The customize utility can be found in the **bin\native** directory of the Search Service installation, for example, C:\bea\alui\ptsearchserver\6.1\bin\native\customize.exe. The utility must be run from a command prompt, taking command-line arguments for the thesaurus CDF file and the path to the Search Service installation:

   ```
   customize -r <thesaurus file> <SEARCH_HOME>
   ```

   where SEARCH_HOME is the root directory of the Search Service installation, for example, C:\bea\alui\ptsearchserver\6.1. This is not an environment variable that needs to be set; the directory merely needs to be specified directly on the command line. For example, if your thesaurus file is located in **\temp**, you enter:

   ```
   customize -r \temp\thesaurus.cdf C:\bea\alui\ptsearchserver\6.1
   ```

   When you run the customize utility, the files in SEARCH_HOME\common are removed and replaced by files of the same name, though their contents now represent the mappings created by the customize utility. The customize utility has a command-line mode for reverting to the set of mappings files that shipped with the Search Service (and hence removing any thesaurus customizations). This mode uses the `-default` flag in place of `-r <thesaurusfile>`, but otherwise is identical to the invocations shown above:

   ```
   customize -default C:\bea\alui\ptsearchserver\6.1
   ```

4. Restart the Search Service.

   The files produced by the customize utility are loaded when the Search Service starts.

Customizing Categorization of Search Results

Users can use the Sort By drop-down list on the search results page to sort results by object type or by folder location in the Knowledge Directory or Administrative Object Directory. You can customize this drop-down list to include additional categories relevant for your users. If you use a property in your portal documents named Region, for example, you can customize the Sort By drop-down list to include Sort By Region: New England, Midwest, and so forth.

The first issue to consider when assessing whether categorizing search results by a particular property is a good idea is whether the property will be defined for a substantial percentage of all search results. For instance, if 90% of search results do not have the property defined, then when categorizing by that property, most everything will fall under "All Others", and the categorization will not be very useful. For that reason, as a rule of thumb it is not generally recommended to add a custom categorization option for a property which is undefined for more than half of all documents and administrative objects.

The other issue to consider is whether the values for the property will make reasonable category titles. In order for categorization to work well for a property, each value should be a single word or a short noun phrase, for example, New England, Midwest, Product Management, Food and Drug Administration, and so forth. The values should not be full sentences or long lists of keywords, for example, "This content crawler crawls the New York Times finance section". The entire contents of the property value for each item will be considered as a single unit for the purposes of categorization, so it will look odd if a full sentence is returned as a category title.

Setting Up Property Data for Categorization

The first step in the process of adding a new categorization option is to ensure that documents and objects include the property you want to use to sort by category. For information on setting up maps from source document attributes to portal properties, see "Configuring Content Types and Document Properties" on page 4-3. Ensure that the property that defines the category for sorting has the following configuration:

- Supported for use with documents
- Visible in the user interface
- Searchable
- Mandatory (since the search results categorization will only be valuable if there are many items with defined values for the property, and will be of maximum value if everything has a value for the property)
- Named appropriately

Enabling Results Sorting

To enable results sorting by property, add the following settings within the <Search> section of **portalconfig.xml**:

```
<CategoryName_1 value="CategoryName"/>
<CategoryField_1 value="PTObjectID"/>
```

CategoryName is the name you want to appear in the Sort By drop-down list, for example, Region.

ObjectID is the integer that identifies the property object. To find the object ID, right-click the link to the property object and then choose **Properties**. This will yield a link that looks something like this:

```
http://portal.company.com/portal/server.pt?open=36&objID=200&parentname=Ob
jMgr&parentid=5&mode=1&in_hi_userid=1&cached=true
```

The `objId` argument is the one containing the integer you want. In this link, the object ID is `200`, so complete the CategoryField entry as follows:

```
<CategoryField_1 value="PT200"/>
```

You can add multiple custom categorization options by adding analogous tags named CategoryName_2, CategoryField_2, CategoryName_3, CategoryField_3, and so forth. In **portalconfig.xml**, the Category tags must be numbered consecutively without skipping. For example, if there is a <CategoryName_3> tag, there must be tags for Category 1 and 2.

For more information about the portalconfig.xml file, see Appendix A, "Configuring Advanced Properties and Logging."

Managing Grid Search

Grid search consists of shared files (for example, C:\cluster) and search nodes. When you start up the Search Service, it looks at the **cluster.nodes** file in the shared files location to determine the host, port, and partition of each node in the cluster. It monitors and communicates the availability of the search nodes and distributes queries appropriately.

The Search Service also automatically repairs and reconciles search nodes that are out of sync with the cluster. At startup, nodes will check their local TID against the current cluster checkpoint and index queues. If the current node is out-of-date with respect to the rest of the cluster, it must recover to a sufficiently current transaction level (at or past the lowest cluster node TID) before servicing requests for the cluster. Depending upon how far behind the local TID is, this operation may require retrieval of the last-known-good checkpoint data in addition to replaying queued index requests.

Although the Search Service performs many actions automatically to keep your cluster running properly, there are some maintenance and management tasks you perform manually to ensure quality search in your portal. This section includes the following topics:

- “Updating the Search Collection” on page 4-59
- “Repairing Your Search Index” on page 4-60
- “Managing Checkpoints and Restoring Your Search Collection” on page 4-61
- “Managing Search Cluster Topology” on page 4-61
- “Monitoring Search Activity with Logs” on page 4-62
- “Using the Command Line Admin Utility” on page 4-62

Updating the Search Collection

As users create, delete, and change objects in the portal, the search index gets updated. In some cases, the portal updates the search index immediately; in other cases, the search is not updated until the next time the Search Update Agent runs. The following table describes the cases in which the search index is updated immediately (I) or updated by the Search Update Agent (SU).

Table 4-8 How the Search Index is Updated

Object	Create	Delete	Move	Change Name or Description	Change Other Properties
Document	I	SU	SU	I	I
Directory Folder	I	SU	SU	I	SU
Administrative Folder	I	I	I	I	I
Administrative Object	I	I	I	I	I

Note: If the Knowledge Directory preferences are set to use the search index to display browse mode, changes will not display until the Search Update Agent runs. The Knowledge Directory edit mode and the Administrative Object Directory display objects according to the database, and therefore show changes immediately.

The Search Update job is located in the Intrinsic Operations administrative folder. It performs the following actions on the search index:

- Updates the index
- Releases expired locks on users and objects
- Repairs the search index according to the Search Service Manager repair settings

The default frequency of the Search Update job is one hour, which is suitable for most portal deployments; but, if your search index is very large, the Search Update Agent might not be able to finish in one hour. For information on modifying Search Update job settings, see "Running Portal Agents" on page 5-4.

Repairing Your Search Index

Your search index might get out of sync with your database if, during the course of a crawl, the Search Service became unavailable or a network failure prevented an indexing operation from completing. Another possibility is that a Search Service with empty indices was swapped into an existing portal with pre-existing documents and folders.

The Search Service Manager lets you specify when and how often the Search Update Agent repairs your search index. Instead of just synchronizing only particular objects, the repair also synchronizes all objects in the database with the search index. Searchable objects in the database are compared with IDs in the search index. If an object ID in the database is not in the search index, the Search Update Agent attempts to re-index the object; if an ID in the search index is not in the database, the Search Update Agent removes the object from the search index.

Run the Search Update Agent for purposes of background maintenance or complete repopulation of the search index.

To configure Search Repair:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Search Service Manager**.
3. Under the Search Repair Settings, specify when your search index should repair itself and the interval of the next repair sessions.
4. Click **Finish**.

Managing Checkpoints and Restoring Your Search Collection

A checkpoint is a snapshot of your search cluster that is stored in the **cluster** folder (for example, C:\bea\alui\cluster), a shared repository available to all nodes in the cluster. When initializing a new cluster node, or recovering from a catastrophic node failure, the last known good checkpoint will provide the initial index data for the node's partition and any transaction data added since the checkpoint was written will be replayed to bring the node up to date with the rest of the cluster.

You manage checkpoints on the Checkpoint Manager page of the Search Cluster Manager. You can perform the following actions with the Checkpoint Manager:

- Manually create an individual checkpoint or schedule checkpoints to be automatically created on a periodic basis.
- Restore your search collection from a checkpoint.

Note: For instructions on using the Search Cluster Manager, refer to online help.

Since checkpoint data is of significant size, limit the number of checkpoints maintained by the system. Specify how many checkpoints to keep on the Settings page of the Search Cluster Manager. Refer to online help for details.

Managing Search Cluster Topology

Your search cluster is made up of one or more partitions, each of which is made up of one or more nodes. As your search collection becomes larger, the collection can be partitioned into smaller pieces to facilitate more efficient access to the data. As the Search Service becomes more heavily utilized, replicas of the existing partitions, in the form of additional nodes, can be used to distribute the load. Additional nodes also provide fault-tolerance; if a node becomes unavailable, queries are automatically issued against the remaining nodes.

Note: If a partition becomes unavailable, the cluster will continue to provide results; however, the results will be incomplete (and thus indicated in the query response).

You manage the partitions and nodes in your search cluster on the Topology Manager page of the Search Cluster Manager. You can perform the following actions with the Topology Manager:

- Add or delete a node.
- Repartition the cluster (add or delete partitions).

  Caution: Adding a partition to the cluster requires redistributing of potentially hundreds of thousands of documents.

- Assign a node to a different partition.

Monitoring Search Activity with Logs

Search logs are kept for the search cluster as well as for each node in the search cluster. The cluster logs are stored in the **\cluster\log** folder, for example, C:\bea\alui\cluster\log\cluster.log. The cluster logs include cluster-wide state changes (such as cluster initialization, node failures, and node recoveries), errors, and warnings.

The node logs are stored in the node's **logs** folder, for example, C:\bea\alui\ptsearchserver\6.1\node1\logs. There are two kinds of node logs: event logs and trace logs. Event logs capture major node-local state changes, errors, warnings, and events. Trace logs capture more detailed tracing and debugging information.

There are several ways to view the logs:

- You can open the log file in a text reader.
- You can view search logging through PTSpy.
- You can set up another OpenLog listener to receive logging information.

A new cluster log is created with each new checkpoint. The log that stores all activity since the last checkpoint is called **cluster.log**. When a new checkpoint is created, the cluster.log file is saved with the name <checkpoint>.log, for example, 0_1_5116.log.

Using the Command Line Admin Utility

The Command Line Admin Utility allows you to perform the same functions you can perform in the Search Cluster Manager as well as the following additional functions:

- Changing the run level of the cluster
- Purging and resetting the search collection

The Command Line Admin Utility is located in **bin\native** folder in the Search Service installation folder, for example, C:\bea\alui\ptsearchserver\6.1\bin\native\cadmin.exe. Invoking the command with no arguments displays a summary of the available options:

```
% $RFHOME/bin/cadmin
Usage: cadmin <command> [command-args-and-options] [--cluster-home
<CLUSTER_HOME>]
```

Requesting Cluster Status

The status command displays the status of the cluster. By default, the status command displays a terse, one-line summary of the current state of the cluster:

```
% cadmin status --cluster-home=/shared/search

2005-04-22 13:54:13 checkpoint_xxx 0/1/198 0/1/230 impaired
```

If you add the verbose flag, the status command displays the full set of information, including the status of every node in the cluster:

```
% cadmin status --verbose --cluster-home=/shared/search

2005-04-22 13:54:13 /shared/search checkpoint_xxx

cluster-state: impaired

cluster-tid: 0/1/198 0/1/230

partition-states: complete impaired

node p0n0: 0 192.168.1.1 15244 0/1/198 0/1/460 run

node p0n1: 0 192.168.1.2 15244 0/1/198 0/1/460 run

node p1n0: 1 192.168.1.3 15244 0/1/198 0/1/230 run

node p1n1: 1 192.168.1.4 15244 0/1/100 0/1/120 offline
```

You can also use the status command to repeatedly emit status requests at a specified interval:

```
% cadmin status --period=10 --count=5

2005-04-22 13:54:13 checkpoint_xxx 0/1/198 0/1/230 impaired

2005-04-22 13:54:23 checkpoint_xxx 0/1/198 0/1/230 impaired

2005-04-22 13:54:33 checkpoint_xxx 0/1/198 0/1/230 impaired

2005-04-22 13:54:43 checkpoint_xxx 0/1/198 0/1/230 impaired

2005-04-22 13:54:53 checkpoint_xxx 0/1/400 0/1/428 complete
```

Requesting Specific Cluster Status Information

You can request information about specific nodes within the cluster. This displays the same type of information that is displayed as part of the verbose cluster status request:

```
% cadmin nodestatus p0n0 p1n0

node p0n0: 0 192.168.1.1 15244 0/1/198 0/1/460 run

node p1n0: 1 192.168.1.3 15244 0/1/198 0/1/230 run
```

As with cluster status, you can request periodic status output:

```
% cadmin nodestatus p0n0 p1n0 --period=10

2005-04-22 13:54:13 p0n0 0 192.168.1.1 15244 0/1/198 0/1/460 run

2005-04-22 13:54:13 p1n0 0 192.168.1.1 15244 0/1/198 0/1/460 run

2005-04-22 13:54:23 p0n0 0 192.168.1.1 15244 0/1/198 0/1/460 run

2005-04-22 13:54:23 p1n0 0 192.168.1.1 15244 0/1/198 0/1/460 run
```

Changing the Run Level of the Cluster

You can modify the run level of the cluster, or of individual nodes within the cluster. For example, you might want to place nodes in standby mode prior to changing cluster topology or shutting them down. Transitioning from standby to any of the operational modes (recover, readonly, stall, run) will validate the node's state against the cluster state and will trigger a checkpoint restore if one is warranted.

Transitions to readonly or offline modes are also potentially useful: readonly mode halts incorporation of new index data on a node; offline mode will cause the search server to exit.

To set run level of p0n0 and p1n0 to standby:

```
% cadmin runlevel standby p0n0 p1n0
```

To set run level of the entire cluster to run (affects only non-offline nodes):

```
% cadmin runlevel run
```

Purging and Resetting the Search Collection

You can purge the contents of the search collection. You might want to purge the cluster in staging or development systems, or if you want to clean out the search collection without re-installing all the nodes. Purging the search collection may also be useful in a dire situation where the contents of the cluster are corrupted beyond repair and good checkpoints are not available for recovery.

By default, the checkpoints and index queue are left in place. This allows you to rebuild the local index on a node whose archive appears to be corrupted.

To purge the search collection, but keep checkpoints:

```
% cadmin purge
```

Caution: As a safeguard against performing this operation by accident, all cluster nodes must be in standby mode and you must confirm the action before the purge command is sent out.

The purge command causes a node to generate empty archive collections (document, spell, and mappings) and perform a soft-restart to load them into memory. Before reloading, the admin utility updates the checkpoint files in the shared repository to prevent the nodes from automatically reloading from an existing checkpoint.

To purge the search collection and delete existing checkpoints:

```
% cadmin purge --remove-checkpoints
```

Initiating a Cluster Checkpoint

You can request a cluster checkpoint at any time (in addition to any periodic checkpoints initiated by the cluster):

```
% cadmin checkpoint
```

Since creating a checkpoint is a time-consuming process, the admin utility displays its progress:

```
Checkpoint using nodes: p0n0 p1n1 p2n0

Node p0n0 copying data

Node p1n1 copying data

Node p2n0 copying data

0%..10%..20%..30%..40%..50%..60%..70%..80%..90%..100%

Checkpoint complete in \\cluster_home\checkpoint_xxx
```

If the cluster has insufficient active nodes to perform the checkpoint, the admin utility displays appropriate feedback:

```
Node p0n0 is offline

Node p0n1 is offline

Unable to checkpoint at this time: partition 0 is unavailable
```

Any error messages encountered during the checkpoint process also display:

```
Checkpoint using nodes: p0n0 p1n1 p2n0

Node p0n0 copying data

Node p1n1 copying data

Node p2n0 copying data

0%..10%..20%..

Node p1n1 is offline

Checkpoint aborted
```

Reloading from a Checkpoint

You can request a checkpoint restore at any time.:

```
% cadmin restore
```

Since restoring from a checkpoint is a time-consuming process, the admin utility displays its progress:

```
Restoring cluster from \\cluster_home\checkpoint_xxx

Node p0n0 retrieving data

Node p0n1 retrieving data

0%..10%..20%..30%..40%..50%..60%..70%..80%..90%..100%

Node p0n0 restarted

Node p0n1 restarted

Restoration complete
```

Changing Cluster Topology

You use the same command to add or remove nodes from the search cluster as you do to repartition the cluster:

```
% cadmin topology new.nodes
```

The difference is how you change the **cluster.nodes** file:

- To add new nodes to the search cluster (for failover capacity), install a new node, and edit the **cluster.nodes** file to include the node as a peer on an existing partition.

  Issue a “soft reset” to the cluster through the command line utility, which causes all nodes to re-examine the cluster topology file and thus recognize the new node. When the new node receives a soft reset, it recognizes that it needs to catch up to the rest of the cluster and begins the automated index recovery process from the last checkpoint.

- To repartition the cluster, edit the number of partitions in the **cluster.nodes** file. You will be asked to confirm the action and the admin utility will confirm that a checkpoint exists before performing the repartitioning operation.

Since changing cluster topology can be a time-consuming process, the admin utility displays its progress. Here’s an example of what the output might be when you add and remove nodes:

```
Current topology:

<contents of current cluster.nodes file>

New topology:

<contents of new.nodes file>

Nodes to add: p0n2, p1n2, p2n2

Nodes to remove: p0n0, p1n0, p2n0

Is this correct (y/n)? y

Applying changes…

p0n2 has joined

p2n0 has left

...

Changes applied successfully
```

Here’s an example of what the output might be when you repartition the cluster:

```
Current topology:
<contents of current cluster.nodes file>
New topology:
<contents of new.nodes file>
Nodes to add: p3n0, p3n1
Is this correct (y/n)? y
CAUTION: the requested changes require repartitioning the search collection
The most recent checkpoint is checkpoint_xxx from 2004-04-22 16:00:00
Is this correct (y/n)? y
Repartitioning from 3 partitions into 4
0%
5%
<progress messages>
100%
Repartitioning successful
Applying changes…
p0n2 has joined
p2n0 has left
...
Changes applied successfully
```

If the repartition fails, the search collection leaves the cluster in its original state, if at all possible, and provides information about the failure. The cluster.nodes file is rolled back to the previous state after making sure that the last-known good checkpoint refers to an un-repartitioned checkpoint directory.

Aborting a Checkpoint or Reconfiguration Operation

You can abort a long-running checkpoint or cluster reconfiguration operation by exiting from the command line utility with Control-C. The cluster will be restored to its state prior to attempting the checkpoint or topology reconfiguration.

In the case of a checkpoint operation, the utility sends a “checkpoint abort” command to the checkpoint coordinator to cleanly abort the checkpoint create/restore operation.

In the case of a cluster reconfiguration, the utility restores the original **cluster.nodes** file and initiates a soft restart of the affected cluster nodes to restore the cluster to its previous configuration.

Creating Snapshot Queries

A snapshot query allows you to display the results of a query in a portlet or e-mail the results to users. You can select which repositories to search (including Publisher and Collaboration), and limit your search by language, object type, folder, property, and text conditions.

To create a snapshot query:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, select **Snapshot Query**.
4. Complete the query and results format according to the online help.
5. Click **Finish**.
6. When prompted, specify a folder to which to save the snapshot query.

   The editor prompts you to send an invitation to view the query. Follow the editor instructions to do so.

Implementing Federated Searches

This section describes *federated searches*, which allow your users to search external repositories for content or allow users of other portals to search your portal for content. This section includes the following topics:

- "About Federated Searches" on page 4-70
- "Configuring Federated Searches" on page 4-71
- "Searching Non-Portal Repositories" on page 4-72

About Federated Searches

Federated searches connect separate AquaLogic Interaction portals with one another and with external repositories. Federated searches empower dispersed organizations to deploy multiple portals and link them together, thereby combining local control over content with global scope. Federated searches provide end-users a single interface and unified result set for searches over multiple AquaLogic Interaction portals, as well as parallel querying of external Internet and intranet-based search engines.

When you install the portal, the Public Access Incoming Federated Search is created. This allows other AquaLogic Interaction portals to search this portal as the Guest user.

To allow other search relationships, you must create new incoming or outgoing federated searches. Whether your portal is requesting or serving content, you and the other administrators involved need to agree upon the following issues prior to establishing federated searches:

- Which portals will serve content?
- Which portals will request content?
- Do the portals share a common external database of users?

  If both portals share a common external database of users, such as an LDAP server or NT domain, you must grant the shared users access to the appropriate content on the serving portal. This provides the greatest degree of content security without requiring any additional administrative work.

  If the portals involved do not share a database of user information, you must create one or more portal users in the serving portal that can be impersonated by users of the requesting portal.

- If you do not use an authentication source, what portal ID and password will be used to identify the portals?

  For every request issued, the requesting portal sends an ID and password to identify itself to the serving portal. You must enter the same ID and password in both the requesting portal outgoing federated search and the serving portal incoming federated search.

  Incoming federated searches can be configured to allow unauthenticated users to search the portal as a guest.

- What content from the serving portal will be available to the requesting portal?

Configuring Federated Searches

AquaLogic Interaction portals can use federated search to search other AquaLogic Interaction portals. To enable this, you must configure a trust relationship between the searching (outgoing) and searched (incoming) portals. To establish the trust relationship, the two participating portals must agree upon a name and password combination that will be used to ensure that requests are coming from a trusted source. This information is recorded as the portal identification name and password.

There are outgoing and incoming federated searches:

- An *outgoing federated search* allows users of your portal to search other AquaLogic Interaction portals or other external repositories. To enable a federated search between AquaLogic Interaction portals, a relationship must be established between the outgoing and incoming federated portals. To enable a federated search between a AquaLogic Interaction portal and another portal or an external repository, a search Web service must be created.

- An *incoming federated search* allows other AquaLogic Interaction portals to search your portal.

Configuring Outgoing Federated Searches

To create a search Web service:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Web Service - Search**.
4. Define the search Web service as described in the online help.
5. Click **Finish**.

To create an outgoing federated search:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Federated Search - Outgoing**.
4. When prompted, select the search Web service.
5. Define your outgoing federated search as described in the online help.
6. Click **Finish**.

Configuring Incoming Federated Searches

To create an incoming federated search:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Federated Search - Incoming**.
4. Define your incoming federated search as described in the online help.
5. Click **Finish**.

Searching Non-Portal Repositories

If there is a non-portal repository that you want to search, BEA or another vendor might have written a search Web service to access it. If not, BEA provides an Enterprise Web Development Kit that allows you to easily write your own Search Web services in Java or .NET. For details, visit the BEA AquaLogic User Interaction Development Center (http://dev2dev.bea.com/aluserinteraction/).

To create an outgoing federated search that accesses a non-portal repository.

Table 4-9 Creating an Outgoing Federated Search (Non-Portal Repository)

Task	Steps
Install the search provider.	For details, refer to documentation from the search provider.
Configure a remote server.	1. Click **Administration**. 2. Navigate to or create the administrative folder for this content source. 3. In the Create Object drop-down list, select **Remote Server**. 4. Configure connection information for the remote server as described in the online help. 5. Click **Finish**.
Configure a search Web service.	1. Click **Administration**. 2. Open an administrative folder. 3. In the Create Object drop-down list, click **Web Service - Search**. 4. Define the search Web service as described in the online help. 5. Click **Finish**.
Create an outgoing federated search.	1. Click **Administration**. 2. Open an administrative folder. 3. In the Create Object drop-down list, click **Federated Search - Outgoing**. 4. When prompted, select the search Web service. 5. Define your outgoing federated search as described in the online help. 6. Click **Finish**.

Automating Administrative Tasks

This chapter provides the steps you take to set up the Automation Service and schedule jobs that perform routine portal administration tasks. It includes the following sections:

About Jobs

Jobs allow you to schedule portal management operations. A job is a collection of related operations. Each operation is one task, such as a crawl for documents, an import of users, or one of the system maintenance tasks.

You must run jobs to perform the following actions:

- Import or synchronize users and groups through an authentication source
- Import or refresh documents through a content crawler
- Perform external operations
- Run and store content for some portlets
- Import user information through a profile source
- Move or copy content through a smart sort (the portal creates and runs the job automatically)

Registering Automation Services

The Automation Service must be registered with the portal before you can run jobs. The primary Automation Service is registered when you install the Automation Service and execute the related database scripts described in the *Installation and Upgrade Guide for AquaLogic Interaction*.

To register additional Automation Services:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Automation Service**.
3. Click **Add Automation Service**.
4. Complete the configuration according to the online help.
5. Click **Finish**. Note that you cannot run jobs on this Automation Service until you assign job folders to it.

Setting Up Job Folders

Jobs can run only if the folder in which they are stored is assigned to an Automation Service. All of the jobs in a folder are run by *one or more* Automation Services. If multiple Automation Services are associated with a single folder, the BEA ALI Automation Service assigns jobs according to the resources available on each Automation Service.

To manage folder assignment for an Automation Service:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Automation Service**.
3. Click the name of the Automation Service to which you want to assign folders.
4. Add or remove folders as necessary.
5. Click **Finish**.

Starting the BEA ALI Automation Service

The Automation Service runs as a Windows service. Ensure the BEA ALI Automation Service is configured to start automatically when you boot your system. For information on configuring the BEA ALI Automation Service to start automatically, see the *Installation and Upgrade Guide for AquaLogic Interaction*.

Running Portal Agents

The portal comes with four jobs that are created upon installation. Two other jobs are created when you perform bulk subscriptions or create dynamic group membership rules. All these jobs are stored, by default, in the Jobs folder under the Intrinsic Operations folder in Administration:

- The **Audit Log Management** Agent archives old audit messages into files and deletes old audit files.

  The **Audit Log Management** Agent also archives and deletes audit files according to the schedule set in the Audit Manager utility.

- The **Bulk Subscriptions** Agent subscribes users in bulk to the communities and portlets you specify in the Bulk Add editor.

- The **Document Refresh** Agent performs background maintenance on your Knowledge Directory, such as refreshing document links and properties, and deleting expired documents.

- The **Dynamic Membership** Update Agent updates dynamic portal group memberships.

- The **Search Update** Agent makes sure the search collection is synchronized with the database. You can run multiple instances of this job at the same time.

  The Search Update Agent also repairs the search index according to the frequency set in the Search Service Manager utility.

- The **Weekly Housekeeping** Agent performs weekly housekeeping on your system, such as deleting expired invitation codes and old job logs and removing community members who no longer have access to a community.

To run these agents, click the job, and schedule it.

Creating and Running Jobs

When you create portal objects that require related jobs, the Create Object editor includes a page to configure and schedule the related job. If you want to create additional jobs independently of the Create Object editors, follow the instructions in this section.

To create and run a job:

1. Click **Administration**.
2. Open an administrative folder.
3. In the Create Object drop-down list, click **Job**.
4. Define your job as described in the online help.
5. Click **Finish**.

Before you can run jobs, you must:

1. Confirm that the BEA ALI Automation Service is running on the Automation Service machine. If it is not running, start it now, as described in "Starting the BEA ALI Automation Service" on page 5-3.
2. Register the Automation Service with the portal, as described in "Registering Automation Services" on page 5-2.
3. Assign administrative folders to the registered Automation Services, as described in "Setting Up Job Folders" on page 5-3.

Configuring External Operations to Run as Jobs

An external operation allows you to run command-line actions through the portal and schedule these actions through portal jobs. For example, you might want to create scripts that query documents, ping portals, e-mail snapshot query results to users, or any custom script on a specified schedule.

To create an external operation:

1. Click **Administration**.
2. Open an administrative object folder.
3. In the Create Object drop-down list, click **External Operation**.
4. Define your external operation and related job as described in the online help.
5. Click **Finish**.

The portal includes two example external operations in the **Admin Objects - Intrinsic Operations** folder:

- **Saved Search Mailer:** This is a sample external operation that e-mails the results of snapshot queries to users. For details on customizing this script, see the comments in the ***PT_HOME*/scripts/SavedSearchMailer.sh** file invoked by the Saved Search Mailer operation.
- **Weekly Search Log Report:** This is an external operation that runs a summary report of the Search Service logs for the previous week, or you can configure it to report on the period you want. For details on customizing this script, see the comments in the ***PT_HOME*/scripts/WeeklySearchLogReport.sh** file invoked by the Weekly Search Log Report operation.

Reviewing Jobs Status and Job Logs

The Job History page of the Automation Service utility provides information about in-process and completed jobs as well as an option to abort jobs.

To view status and logs for in-process and completed jobs:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Automation Service**.
3. On the left, under Edit Utility Settings, click **Job History**.
4. View the history of jobs that have run and the logs for individual jobs or abort in-process jobs.
5. Click **Finish**.

Migrating, Backing Up, and Restoring Portal Objects

This chapter provides the steps you take to migrate (export and import), back up, and restore portal objects. It includes the following sections:

About Object Migration

Object migration lets you copy resources from one portal to another. You might want to do this for several reasons. You might have multiple portals to handle a global deployment or you might want to create multiple portals to separate development, testing, and production.

You can copy resources from one portal to another by creating migration packages, which can be used to:

- Export objects created in a development portal and import them to your production portal when they have been properly tested.
- Import portal objects in order to install new features on your portal. For example, you might want to install a portlet suite and register those portlets in your portal.

Table 6-1 summarizes the features of object migration.

Table 6-1 Migration Features

Migration Feature	Description
Portal objects that can be included in the package	All objects
Collaboration and Publisher information	Can migrate Collaboration or Publisher information
Requests and approval	Users with at least Edit access to objects can request migration, but only members of the portal administrators group can approve objects for migration. An administrator selects approved objects to add to a migration package, and can also add object to the package without making a migration request. Users with the Access Utilities activity right can check the status of their migration requests.
Creating a migration package	Only users with access to the Access Utilities portal activity can create a migration package. An administrator can add objects that do not have migration requests to a migration package (bypassing the request and approval process).
Object dependencies	Dependencies always maintained. Dependent objects can be included in a migration package, but do not need to be.
Unique universal identifiers (UUIDs) and their effect on subsequent importing migration packages	By default UUIDs are maintained, so that subsequent migrations overwrite previously migrated objects. However, if you do not want to overwrite previously migrated objects, you have the option of creating a new instance of the same object, with a new UUID.

Using Migration Packages

The Migration - Export utility in the portal lets you create migration packages. To import objects from a migration package, you use the Migration - Import utility.

There are several things you can do to make migration as easy and effective as possible:

- Keep source and target portals as similar as possible to reduce the mapping required.
- Create migration packages as soon as possible after approving objects. The object settings in a migration package are those present at package creation, not those present at object approval. Creating migration packages soon after approval minimizes the chance that object settings have changed since approval.
- You can selectively import objects in a migration package, so if you want to import content crawlers and communities separately, you can import a package twice, and select different objects each time you import it.

Creating a Migration Package in the Portal

You can create a migration package that includes portal resources as well as Publisher and Collaboration information.

To create a migration package in the portal:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Migration - Export**.
3. Add portal resources, and if applicable, Publisher and Collaboration information, as described in the online help.
4. Click **Finish**. A status message is displayed as the migration package is being created. When the migration package is created, you can download it to your desktop.

   Note: If you are also migrating Collaboration or Publisher objects, those will be written to a .zip file on the machine where Collaboration or Publisher is installed. You must move this file from this location to the target location.

Creating a Migration Package Using a Command Line Tool

To create a migration package that includes portal resources:

1. Log in to the host computer for the portal as the user who owns the portal installation.
2. Use the command **ptmigration.bat** (for Windows) or **./ptmigration.sh** (for Unix) with the following parameters:

   ```
   ./ptmigration.sh [username] [password] -export [migration package
   name] [log file name] <-exportdependencies>
   ```

   Where the parameters are as follows:

Parameter	Description
migration package name	Required. The name and path of the migration package to be created
log file name	Required. The name and path of the log file to be created. The path to the log file must be different from that of the migration package.
-exportdependencies	Optional. Use this parameter to export any additional objects upon which the objects you are exporting depend.

Note: You cannot export Collaboration or Publisher objects using the command line tool. To export those objects, use the Migration - Export utility in the portal's administrative UI. See "Creating a Migration Package in the Portal" on page 6-3.

3. Press **Enter**. All objects approved for migration are exported into the migration package. The migration utility updates the migration status in the source portal.

Importing Objects in the Portal

To import objects, including Publisher and Collaboration information, from a migration package:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Migration - Import**.
3. Upload the migration package, and select the objects you want to import, as described in the online help.
4. Resolve any unresolved dependencies as described in the online help.
5. Click **Finish**.

Importing Objects Using a Command Line Tool

To import portal resources from a migration package:

1. Copy the migration package to the target portal host computer.
2. Log in to the host computer for the portal as the user who owns the portal installation.
3. Use the command **ptmigration.bat** (for Windows) or **./ptmigration.sh** (for Unix) with the following parameters:

   ```
   ptmigration.bat [username] [password] -import [migration package
   name] [log file name] <-noacl> <-overwriteremoteservers>
   <-createnewobjectinstances>
   ```

   Where the parameters are as follows:

Parameter	Description
migration package name	Required. The name and path of the migration package to be created
log file name	Required. The name and path of the log file to be created. The path to the log file must be different from that of the migration package.
-noacl	Optional. Use this parameter if you do not want to import the Access Control Lists (security data) associated with the objects you are importing.
-overwriteremoteservers	Optional. Specifies that existing remote server objects should be overwritten by remote server objects in the migration package. The default is that existing remote servers are not overwritten.
-createnewobjectinstances	Optional. Use this parameter if you want to create new object instances instead of overwriting objects that may already exist on the importing portal.

4. Press **Enter**. All the objects in the migration package are imported. The imported objects are located in the same folders on the target portal as on the source portal. Objects with missing dependencies will be skipped and not imported. Check the migration log to see which ones were skipped.

Backing Up and Restoring the Portal

You can back up your system without taking it offline.

To back up your portal:

1. Back up your database according to your database vendor documentation and best practices.
2. Back up your search collection to another location or tape backup.

To restore your portal to a previously saved state:

1. Stop the Web service on all machines hosting the portal application.
2. Stop the BEA ALI Automation Service on all Automation Services.
3. Stop the BEA ALI Search service.
4. If you need to rebuild your portal database, use your database software to restore from a previously saved database.
5. Replace your search collection with backups as close as possible to the time of the database backup you are using.

Your database backup might not exactly match your search collection backup, so the restored database and search collection will be out of sync. To correct this, rebuild the search collection from scratch.

To rebuild the search collection:

1. Reinstall the Search Service and select **Overwrite the existing search index**. For details, refer to the *Installation and Upgrade Guide for AquaLogic Interaction*.
2. Configure the Search Service to run an update:
   a. Click **Administration**.
   b. From the Select Utility drop-down list, choose **Search Service Manager**.
   c. Under Search Repair Settings, change **Next Repair Date** to a time in the past.
   d. Click **Administration**.
   e. Click the **Intrinsic Operations** folder.
   f. Click a Search Update job and schedule it to run immediately.

Configuring Advanced Properties and Logging

This appendix describes how to modify the default configuration for portal components, and configure AquaLogic Interaction Logging Spy. It includes the following sections:

- "Configuring Portal Components" on page A-1
- "Fine-Tuning the Search Service Configuration" on page A-17
- "Using the Config Tool to Modify Database and Portal Logging Settings" on page A-23
- "Configuring AquaLogic Interaction Logging Spy" on page A-24

Configuring Portal Components

The default settings enable your portal to function fully. This section provides the following topics that describe how to customize the installation default portal settings:

- "Using the Portal Utility" on page A-2
- "Modifying Portal Configuration Files" on page A-2
- "Configuring the Automation Service" on page A-16

Using the Portal Utility

If host names or IP addresses change after your initial deployment, you can use the Portal Settings utility to modify portal URLs.

To modify portal URLs:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Portal Settings**.
3. Modify portal URLs as needed.
4. Click **Finish**.

Modifying Portal Configuration Files

If your configuration requirements change after initial deployment, you can modify the configuration in files located in the ***PT_HOME*/settings** directory. PT_HOME is the AquaLogic User Interaction installation directory, for example, C:\bea\alui or /opt/bea/alui.

The following topics in this section describe the configuration files:

- "portalconfig.xml" on page A-3
- "serverconfig.xml" on page A-15

portalconfig.xml

This section describes the following elements of the **portalconfig.xml** file:

- "MainURLs" on page A-4
- "SystemProperties" on page A-4
- "URLMapping" on page A-5
- "CachedSettings" on page A-6
- "Authentication" on page A-7
- "Security" on page A-8
- "International" on page A-9
- "Documents" on page A-10
- "SystemProperties" on page A-4
- "Content Crawlers" on page A-10
- "Style" on page A-10
- "Communities" on page A-10
- "Administration" on page A-11
- "Authentication" on page A-11
- "Invitations" on page A-14

MainURLs

- **WebHome:** The home directory for the portal JAR or DLL files.
- **ImageServerBaseURL:** The base URL of the Image Service.
- **ImageServerSecureBaseURL:** The base URL of the Image Service running HTTPS.
- **ImageServerConnectionURL:** The base URL that is used when the portal connects to the Image Service to retrieve JSRegistry information. In many configurations, this URL is the same as the Image Service Base URL.
- **ImageServerConnectionURLTimeout:** The timeout for ImageServerConnectiontionURL in seconds. A value of -1 means do not check during startup.
- **AdminSiteBaseURL:** The base URL of the Administrative Portal.
- **TempHome:** The temporary file directory to be used by the portal. This should not be Web accessible.

SystemProperties

- **ServerName:** The name of the proxy server running the portal, that is, the mapped server.
- **MachineName:** The server machine name, that is, the name of the portal behind the proxy server.
- **PerformanceComments:** The setting for enabling or disabling the performance comments in the HTML source. 1= the comments are enabled, 0= the comments are disabled.
- **DoctypeSpecification:** 1= none, 2= HTML 3.2, 3= HTML 4.0 Transitional, 4= HTML 4.0 Frameset, 5= HTML 4.0 Strict.
- **VirtualDirectoryPath:** The virtual path to the portal, usually `/portal/`.
- **HTTPEntryPoint:** Portal main Servlet mapping name. For Java (portalconfig.xml), this has to be the same mapping name as the HTTPServlet in web.xml. For .Net, the application mapping name has to be the same mapping name as the httpHandlers in web.config.
- **HTTPPort:** The port number of the portal running HTTP.
- **HTTPSecurePort:** The port number of the portal running HTTPS.
- **SSOVirtualDirectoryPath:** Single sign-on (SSO) virtual directory path.

- **SSOServletName:** SSO Servlet mapping name. For Java (portalconfig.xml), this has to be the same mapping name as the SSOLoginPage in web.xml. For .Net, this has to be the same name in SSOLogin.aspx file in the **Default Web Site/portal/sso** virtual directory.

URLMapping

The URL mapping determines the portal base URL according to the requested URL from the request object. The portal base URL is the base URL for every single link and redirection. For example, if your portal base URL were: `http://portal.company.com/portal/server.pt`, then the link to the default My Page would be: `http://portal.company.com/portal/server.pt?space=MyPage`.

You can add as many entries as you want to the mapping. Mapping URLs should start with http:// or https:// and end with the HTTP entry point name (unless it's the default value, “*”).

For more information on URL Mapping, visit the BEA AquaLogic User Interaction Development Center (http://dev2dev.bea.com/aluserinteraction/).

- **ApplicationURL0:** In SecurityModes 2 and 3, ApplicationURL should be set to the same value as SecureApplicationURL. In mode 0, ApplicationURL0 might be equal to “*”. In this case, the Application Base URL will be the URL from the Request object.

- **SecureApplicationURL0:** In SecurityMode 0, SecureApplicationURL is not used. In modes 2 and 3, SecureApplicationURL0 might be equal to “*”. In this case, the Application Base URL will be the URL from the request object.

CachedSettings

Add entries for personal settings that should be cached on the http session of each user. Personal settings that are not included in this list will be retrieved from the portal every time they are requested. Settings that are on this list are obtained from the portal on login and are cached for the duration of the user's http session.

Note: AccessStyle, Locale, and TimeZone should always be cached by the server.

Users can customize these settings by clicking **My Account** in the portal interface. For instructions and a detailed description of any page in the portal, click ? **Help**.

The following settings are cached by default:

- **Greeting:** Stores the personalized greeting that displays on the portal banner when the user logs into the portal.
- **AccessStyle:** Stores the user display option.
- **Locale:** Stores the preferred language of the portal interface. Portlet names and content display in the selected language only if the language is supported by the portlets. It also stores the format for portal entries (including search requests). For example, if the user chose British English, the portal displays and expects dates in the DD/MM/YYYY format (whereas in American English, the portal displays and expects dates in the MM/DD/YYYY format).
- **TimeZone:** Stores the time zone of the user.
- **PortletTimeout:** Stores the maximum time to wait for a portlet to load.
- **CollapsedGadgets:** Stores which portlets were minimized by the user.
- **COMCollapsedGadgets:** Stores which portlets in the community page were minimized by the user.
- **MyPageRefreshRate**: Stores the refresh rate of user My Pages.

Authentication

- **AllowGuestAccess:** Allow the guest user to access the portal. If guest access is not allowed, the portal will always prompt for login information.
- **GuestPassword:** This is the password for the guest user. Changing this password requires that it be changed here and in the portal database itself.
- **GuestRedirectToLogin:** If the guest user does not specify a space in the URL query string, this setting determines the initial page the user sees.

  Users can navigate to different portal pages by including “space=xxxx” strings in the URL query string. For example, if the user were to type: `http://MYSERVER/portal/server.pt?space=MyPage`, the user will be directed to the My Page (the access privileges of that page will be in effect).

  However, if the user did not include a “space=xxxx” string in the query string (that is, the user only typed: `http://MYSERVER/portal/server.pt)`, the portal directs them to a default page, depending on the GuestRedirectToLogin setting and the experience definition settings, as shown below:

GuestRedirectToLogin Setting	Description
0	The portal will redirect the guest user to the home page defined in the current experience definition (usually a My Page or community page).
1	The portal will redirect the guest user to the login page as defined in the current experience definition.

- **RedirectOnLogout:** After logging out the user is redirected to a default page as shown below:

RedirectOnLogout Setting	Description
0	The portal will redirect the guest user to the home page defined in the current experience definition (usually a My Page or community page).
1	The portal will redirect the guest user to the login page as defined in the current experience definition.

Security

SecurityMode: This setting determines which pages will use SSL encryption. You must install a digital certificate and enable SSL on your Web server before changing the default value of 0.

Note: Changing the security mode affects the URL Mapping. For more information, see Appendix A, "Configuring Advanced Properties and Logging."

Table A-1 Security Modes

Security Mode	Description
0	The portal does not check the security of incoming requests. In mode 0, ApplicationURL0 and SecureApplicationURL0 may be equal to "*". In this case, the Application Base URL will be the base URL from the Request object.
1	Selected pages that involve sensitive information such as passwords use SSL, while other pages are sent unencrypted for better performance. Only pages of Activity Spaces listed in SecureActivitySpaces.xml (which is located in the same folder as portalconfig.xml) are sent through HTTPS. The portal verifies that links and redirections to Secure Activity Spaces uses HTTPS. If a secure Activity Space were requested through a non-secure URL, the portal would redirect the same request to HTTPS. If XPRequest.GetRequestURL() is equal to URLFromRequest0, ApplicationURL0 and SecureApplicationURL0 might both be the Application Base URL, depending on the security of the Activity Space. You must install a digital certificate and enable SSL on your Web server.

Table A-1 Security Modes

Security Mode	Description
2	Every page uses SSL.The portal verifies that every single incoming request uses HTTPS. If it does not, the portal will redirect this request to HTTPS. This setting is best for very secure applications where performance is not a major concern. If the URL from the Request object is equal to URLFromRequest0, SecureApplicationURL0 will be the Application Base URL. URLFromRequest0 has to be equal to "*". This is the default entry. It will be used if no mapping entry matched the URL from the Request object. You must install a digital certificate and enable SSL on your Web server.
3	Select this mode if you are using an SSL Accelerator. Because the portal is behind an SSL Accelerator, the security of the incoming requests is not verified. The portal trusts every request from the SSL Accelerator. All the links and redirections are in HTTPS. If URL from the Request object is equal to URLFromRequest0, SecureApplicationURL0 will be the Application Base URL. URLFromRequest0 has to be equal to "*". This is the default entry. It will be used if no mapping entry matched the URL from the Request object. You must install a digital certificate and enable SSL on your Web server.

International

- **Locale:** The default locale for new users. If you set locale to *UseBrowser*, the portal derives the locale from the browser language settings. The user locale setting, stored when creating a user from a user profile (or configured in **My Account** | **Edit Locale** in the portal) overrides the locale setting in the portalconfig.xml file.
- **TimeZone:** The default time zone for the user. The user time zone settings (configured in **My Account** | **Edit Locale** in the portal) overrides this setting. The portal uses the settings in this file only if the user has not configured his or her default time zone.
- **MandatoryObjectLanguage:** This setting allows the administrator to set the language for all new objects. If it is blank, the user creating the object can choose the language for the object. If it is not blank, the value will be used as the language for all new and edited objects. The value should be a locale string (that is, it should match the name of a folder under the msgs directory).

Documents

- **NewDocumentTime:** The number of days that a document or folder displays “new” after its name.
- **DocumentLastUpdated:** The number of days that a document displays “updated” after its name.
- **OpenNewWindow:** The default setting for the open in new window preference. 0= same window, 1= new window. Users can override this setting by changing their personal preferences.

Content Crawlers

- **MaxWebCrawlRadius value:** The setting for the maximum number of links away from the target page that you want to crawl. For example, if you select 1, the content crawler attempts to import every page directly linked to the target page; if you select 2, the content crawler attempts to import every page directly linked to the target page, and every page directly linked to those linked pages. This setting corresponds to the **Crawl Radius** drop-down list in the Web Content Crawler Editor. The default maximum is 4, which means the drop-down list allows a crawl radius of 1 to 4.

Style

- **StyleSheetName:** The name for the portal's default stylesheet.

Communities

- **DefaultCommunityID**: Configure this setting only if your navigation scheme is Tabbed Section Left Vertical Navigation or if you use a custom navigation scheme that uses the IPluggableNavModelRO.GetDefaultCommunity() method. In these cases, the setting specifies the ID of the default community to display when a user clicks the **Community** tab.
- **CommKnowledgeDirLinksPerPage:** The number of links to display on one screen in the community Knowledge Directory.

Administration

- **IsAdminSite:** Determines whether the computer is an administrative portal or a browsing-only portal.

Table A-2 Portal Modes

Portal Mode	Description
0	Sets the computer into a browsing-only portal.
1	Sets the computer into a browsing and administrative portal.

- **AdminObjectsPerPage:** Determines the number of administrative objects of a single type (for example, content crawlers or content sources) allowed to display on one screen. This controls the number of items displayed on the administrative interface. The default is 10.

Authentication

- **AllowDefaultLoginPageAuthSource:** Controls the use of the default authentication source for portals (that do not use single sign-on) on the login page and Login Portlet. It also lets you configure the authentication source drop-down list.

Table A-3 Authentication Source Modes

Authentication Source Mode	Purpose	Appearance	Required Actions
0	The portal does not use the default authentication source.	The drop-down list has no special ordering.	Default mode.
1	The portal uses the default authentication source.	The drop-down list is hidden, but it displays a link that brings up the authentication source drop-down list. This lets users select a non-standard authentication source.	You must turn off the caching on the Login Portlet or disable the Login Portlet. You must set the DefaultAuthSourcePrefix tag to the prefix of the authentication source that is the default for all users.

Table A-3 Authentication Source Modes

Authentication Source Mode	Purpose	Appearance	Required Actions
2	The portal uses the default authentication source.	The drop-down list is not hidden, and the default authentication source is pre-selected.	You must turn off the caching on the Login Portlet or disable the Login Portlet. You must set the DefaultAuthSourcePrefix tag to have the prefix of the authentication source that is the default for all users.
3	The portal uses the default authentication source.	The drop-down list is permanently turned off.	

- **DefaultAuthSourcePrefix:** Sets the default authentication source prefix that will be prepended to the login name when users log into your system, unless they select another authentication source from the drop-down list on the login page. In the case of SSO, this is the authentication source category for all of your SSO users.

  You can use AuthSourcePrefix tags to order the items in the authentication source drop-down list. Entries in the list should look like the following:

  ```
  <AuthSourcePrefix[i] value="Auth Source Prefix">
  </AuthSourcePrefix[i]>
  ```

  where [i] is replaced with the items' order in the drop-down list (starting with 1).

  To include the AquaLogic Interaction Authentication Source in the list, simply make an entry with “AquaLogic Interaction Authentication Source” as the value. The AquaLogic Interaction Authentication Source is for users who are created in the portal, manually, through invitations, or through the Create an Account page. For example, to include the AquaLogic Interaction Authentication Source as the third item in the drop-down list, use the following tag:

  ```
  <AuthSourcePrefix3 value=“AquaLogic Interaction Authentication
  Source”>
  </AuthSourcePrefix3>
  ```

  Authentication source prefixes in the ordered list are displayed first in the drop-down list and are followed by any authentication sources not included in the ordered list.

- **AllowAutoConnect:** Setting for saving passwords in cookies.

Table A-4 Auto Connect Modes

Auto Connect Mode	Description
0	Turns off the option of saving passwords in a cookie.
1	Users will see a "Remember my password" check box on the login page of your portal. Passwords are saved as cookies for users that select this check box, which lets users who navigate to your portal be logged in automatically.

- **CaptureBasicAuthenticationForPortlets:** Determines whether or not to capture basic authentication information (login and password) and store it in the session (to send to portlets). The basic authentication information cannot be captured when users select "Remember my password" to login via a cookie. 1= , 0= .

Table A-5 Session Store Modes

Session Store Mode	Description
0	The authentication information will *not* be stored in the session.
1	The authentication information will be stored in the session.

- **RememberPassword:** This setting allows you to set how long the portal remembers your login password. The value should be formatted in minutes. The default is one week.
- **SSOVendor**. Sets the single sign-on configuration. For information on SSO, see Appendix D, "Deploying Single Sign-On."

Invitations

IsInvitationURLSecure: Sets the security of the invitation URL.

Use a secure URL only if you disable http.

Table A-6 Invitation Security Modes

Invitation Security Mode	Description
0	Unsecure. The URL uses http://. You can use an unsecure invitation URL with any security mode, so long as you do not disable http or have http URLs redirect to https.
1	Secure. The URL uses https://. If the SecurityMode setting in your portalconfig.xml file does not allow http, you must select this mode.

serverconfig.xml

You can edit **serverconfig.xml** to modify the following connection and path information for portal components.

Table A-7 Elements of serverconfig.xml

Element	Description
Logging	• <ServerName> specifies the name of the portal host computer on which logs are collected. • <LocalMachineOnly>, when set to `true`, specifies log collection for only the local host, not the network of portal hosts.
Database	• <Common> contains elements to configure connection information for the portal database, such as database type, database name, host name, database user and password, as well as login timeout and pool connection settings.
SysManagement	• <JMXView> indicates whether or not the JMX-based management is enabled. • <JMXRemoteService> configures the port on which the RMI registry listens.
HTTP	• <DefaultProxy> configures http proxy settings for the portal. • <HttpContentCache> configures the cache size.
Content Crawlers	Contains default settings for content crawler transactions.
Gateway	Contains default settings for gateway transactions.
Search	Contains default settings for search transactions.

Configuring the Automation Service

The default settings enable your Automation Service to function fully. If your configuration requirements change after initial deployment, you can modify the *PT_HOME*/settings/**serverconfig.xml** file to specify a new configuration for the Automation Service. The AutomationServers element contains the following settings.

Table A-8 Elements of Automationserver.xml

Element	Description
Nodes	Configures the port number for the Automation Service.
MaxConcurrentJobs	Configures the number of jobs that can run at the same time.
VirtualMachineArgs	Configures tuning values for the JVMs of individual jobs run by the Automation Service.

Fine-Tuning the Search Service Configuration

The installer sets most Search Service configuration parameters to useful defaults. In addition to the default configuration file, the ***PT_HOME*/ptsearchserver/6.0/config** directory includes template configuration files for Search Service deployments. The templates include settings appropriate for a number of operating systems and RAM configurations. RAM determines the recommended maximum number of documents in the search collection, and this collection size determines many of the settings in the template configuration files. Examine the contents of these files, choose the one appropriate for your deployment, and rename the template **ignite.ini** (the active configuration file).

Note: If the Search Service component resides on the same host computer as other portal components, consider using a template tuned for a smaller amount of memory to prevent system paging that adversely affects Search Service performance.

In some cases you might be able to further improve performance by modifying some of the values in the **ignite.ini** file. This section includes the following topics that describe the parameters in **ignite.ini**:

- “Default Search Service Parameters” on page A-17
- “Optional Search Service Parameters” on page A-19

Default Search Service Parameters

The following parameters appear in **ignite.ini** by default:

- **RFINDEX:** Directory used to store Search Service index files. By default, the installer puts these files in the ***PT_HOME*/ptsearchserver/6.0/index** subdirectory. The directory should have sufficient space for the collection you are indexing.

  You should not change this parameter unless you move your index files or are instructed to do so by customer support (ALUIsupport@bea.com).

- **RFPORT:** Port that the Search Service uses for communication with other processes (mainly the portal). The installer prompts for this port number during installation. This value displays in the Search Service Manager, on the Host Settings page. If you change this value in **ignite.ini**, you must also change the value in the Search Service Manager or the portal will malfunction.

- **RF_MAPPING_TOKEN_CACHE_SIZE:** Specifies the size of the cache of mapping tokens. These tokens represent thesaurus and Best Bets entries read from the mappings collection. The default value is 5000. This parameter is chosen in the configuration file templates to reflect the expected number of tokens associates with the maximum supported collection size. The value of this parameter does not have a large effect on Search Service performance. Each cache element is 120 bytes in size, so the default mapping cache will occupy 600 kilobytes of memory.
- **RF_LOG_VERBOSITY:** Numeric parameter that determines how much information is logged in the Search Service logs. Values range from 0 to 5. The default is 3 (high verbosity). You generally do not need to change this parameter. We recommend you not set this below 3. If RF_LOG_VERBOSITY is set below 3, the reports generated by the Search Log Analysis external operation (and viewable on the Search Results Manager) will not contain all of the information needed to support log analysis.
- **RF_DOCUMENT_TOKEN_CACHE_SIZE:** Numeric parameter that specifies the size of the cache of document tokens. These tokens are words from the actual indexed content in the Search Service. The default value is 250000. This parameter has a significant effect on Search Service indexing and query performance, with larger values providing better performance. This parameter is chosen in the configuration file templates to reflect the expected number of tokens associates with the maximum supported collection size. Each cache element is 120 bytes in size, so the default document token cache occupies 29 megabytes of memory.
- **RF_SPELL_TOKEN_CACHE_SIZE:** Numeric parameter that specifies the size of the cache of spelling tokens. These tokens are word fragments from the spelling data derived from the indexed content. The default value is 250000. This value does not need to be larger than the number of tokens in the spell collection, and it does not need to exceed the value in the configuration file templates provided in the **config** directory. This parameter has a significant effect on indexing performance, spell checking, and wild card queries. If these operations seem particularly slow, you can increase the value specified by this parameter. In practice, values larger than 1000000 provide diminishing return while consuming significant amounts of memory. Each cache element is 120 bytes in size, so the default spell cache occupies 29 megabytes of memory.
- **RF_INDEX_CACHE_BYTES:** Numeric parameter specifying the size of the index cache in bytes. The default value is 78643200 (75 megabytes). The value of this parameter has a significant effect on Search Service query performance. The index and docset (see RF_DOCSET_CACHE_BYTES) caches should be made as large as possible while leaving sufficient memory available for the Search Service's other needs.

- **RF_DOCSET_CACHE_BYTES:** Numeric parameter specifying the size of the document cache in bytes. The default value is 2614400 (25 megabytes). The value of this parameter has a significant effect on Search Service query performance. The index and docset (see RF_INDEX_CACHE_BYTES) caches should be made as large as possible while leaving sufficient memory available for the Search Service's other needs.

Optional Search Service Parameters

Optionally, if advised by customer support, you can add the following values to the **ignite.ini** file:

- **RFLOG:** Directory where the Search Service writes its logs. The default is the **<SearchServerInstall>/logs** subdirectory. Edit this value only if you change this directory; the new directory must exist and must be writable by the Search Service.
- **RF_HIGH_PRIORITY:** If this parameter is set to any value, Search Service attempts to increase its process priority over that of other processes. This is not normally necessary, but might be useful on a computer where other processes compete for resources with the Search Service.
- **RF_MAX_WILDCARD_EXPANSIONS:** When a user enters a query that uses a wildcard (for example, “plum*”), this parameter determines the maximum number of terms into which the wildcard is expanded (for example, plum, plums, plumber). The default is 100 terms. This limit keeps overly general queries (“*ing”) from expanding into a huge number of terms and consuming too much time and memory. In large installations, you might need to increase this value.
- **RF_MAX_QUERY_MSECS:** Maximum time, in milliseconds, for user queries. The default is 10000 (ten seconds). After processing the query for this much time, the Search Service returns results it has found so far. You might want to lower the value of this parameter if end users complain that ten seconds is too long to wait for query results.
- **RF_MAX_TOTAL_RESULTS:** Maximum number of results returned by a query. The default is 10000. You do not generally need to change this parameter because the portal displays fewer results than the Search Service maximum.
- **RF_MAX_NUM_STATIC_ARCHIVES:** Maximum number of static archives per collection created in the index directory. The default is 50; this means that there will be up to 50 archive.NNN.docset files (where NNN is a number), archive.NNN.index files, and so on. There can also be up to 50 spell.NNN.docset files, spell.NNN.index files, and so forth. You do not generally need to change this parameter; the only reason might be an operating system (such as Solaris 2.6) that does not allow the Search Service to use enough file descriptors to open all the files at once. Lowering this number causes archive merges to use more memory and disk space.

- **RF_QUERY_THREADS:** Number of threads to dedicate to query processing. The default is 8. You might need to increase this parameter if your Search Service is under heavy load. The value should represent the expected number of simultaneous queries. If this value is too low, incoming queries will wait on a queue for the next free query thread and users will experience longer query times (possibly several seconds).

- **RF_QUERY_QUEUE_SIZE:** When all Search Service query threads (see RF_QUERY_THREADS) are busy, incoming query requests are placed on a queue to wait for the next available query thread. This parameter determines the length of that queue. The default value is 20 and usually does not need to be changed. Should the query queue ever become full, additional query requests are rejected and a message is written to the Search Service logs. If this happens, you can increase RF_QUERY_QUEUE_SIZE.

- **RF_INDEX_THREADS:** Number of threads to dedicate to indexing requests. The default is 2. You might need to increase this parameter if the indexing performance of your Search Service is too low. However, devoting additional system resources to indexing will reduce query performance. Ideally, the value of RF_INDEX_THREADS should not exceed the number of CPUs on the system.

- **RF_INDEX_QUEUE_SIZE:** When all Search Service index threads (see RF_INDEX_THREADS) are busy, incoming index requests are placed on a queue to wait for the next available index thread. This parameter determines the length of that queue. The default value is 20. To estimate a good value, add the number of threads in all content crawlers that might be running simultaneously (you can request up to four threads when setting up a content crawler). To be conservative, make your estimates high. If this parameter is too low, content crawlers can fill the index queue and the Search Service rejects additional index requests until the queue has room for more requests.

  Note: It is better to schedule nonoverlapping crawls than to set a high value for RF_INDEX_QUEUE_SIZE; consider changing the crawl schedule before modifying this parameter.

- **RF_HANDSHAKE_THREADS:** Number of threads to dedicate to servicing incoming socket connections. The default is 5. This value should never need to be changed.

- **RF_HANDSHAKE_QUEUE_SIZE:** Socket connections from Search Service clients are placed on this queue to await acknowledgement by one of the handshake threads (see RF_HANDSHAKE_THREADS). This parameter determines the length of that queue. The default value is 20. Once successfully acknowledged, the connections are assigned to either the query or index queue. Under exceptionally high loads, this queue might fill up and the Search Service will reject new connections until the queue has room for more requests. Should this happen, you can increase the value of this parameter.

- **RF_TOKEN_LEXICON_REBUILD_LIMIT:** Maximum lexicon size, measured in number of tokens, to rebuild automatically. If the Search Service detects that the lexicon was closed improperly, the lexicon is rebuilt as part of the startup process. This can be time consuming. The default value is 400000, ensuring that the rebuild requires no more than a few minutes. Larger lexicons needing repair must be rebuilt with the standalone *examinearchive* utility. You might change the value or set it to zero to allow automatic rebuild of arbitrarily large lexicons.

  In Windows Systems: Setting this value too large can result in error dialogs being posted by the Windows Service Control Manager when the Search Service is run as a Windows service and a lexicon rebuild is performed. These error dialogs indicate that the service is failing to start in a timely manner. They can be disregarded.

- **RF_USE_DATA_FILE_CACHE:** Numeric parameter indicating whether the Search Service should use caches when accessing index and document data. A value of zero disables the caches and causes the Search Service to use memory-mapping. A nonzero value activates the caches. The default value is 1. We strongly recommend you do not change this value.

  This parameter serves as the master on/off switch for RF_INDEX_CACHE_BYTES, RF_DOCSET_CACHE_BYTES, RF_INDEX_CACHE_MAX_PAGES_PER_BLOCK, and RF_DOCSET_CACHE_MAX_PAGES_PER_BLOCK. When RF_USE_DATA_FILE_CACHE is zero, these other parameters have no effect.

  Disabling the caches for small search collections (less than 1 gigabyte of data) might provide a slight improvement in performance depending upon the amount of available physical memory on the Search Service host. In memory-mapped mode, the Search Service fails if the index and document data, plus the Search Service's internal data structures should exceed 2 or 3 gigabytes (depending upon the operating system configuration).

- **RF_REQUIRED_DISK_SPACE:** Amount of disk space (in KB) required to start a dynamic index merge. When merging dynamically indexed data into the search collection, the Search Service verifies that this amount of free space is available on the volume containing the search collection. If the space is not available, the merge process aborts, the Search Service enters read-only mode, and further dynamic indexing requests are rejected. The default value for this parameter is 40000 and should not need to be changed.

When you modify cache settings, keep the following important values and relationships in mind:

- A RF_DOCSET_CACHE_BYTES:RF_INDEX_CACHE_BYTES ratio of 1:3 has been empirically determined to provide near-optimal cache performance for typical search collections.
- Token cache entries occupy 108 (32-bit platforms) or 120 (64-bit platforms) bytes.
- Reasonable values for RF_MAPPING_TOKEN_CACHE_SIZE are 500 to 10000.
- For performance reasons, document offsets and index offsets data is accessed via memory-mapping, regardless of the setting of RF_USE_DATA_FILE_CACHE in the **ignite.ini** file. This means that the memory footprint of a running Search Service depends on the size of the search collection, and this consideration has been calculated in the settings for the configuration file templates in the **config** directory. The amount of memory needed for these mappings can be calculated approximately as (Size of *.docsetOffsets files in bytes) + 0.006 * (Size of *.index and *.key.index files in bytes)
- Leave sufficient address space (and, ideally, physical RAM) available for the number of query and index threads. Allow 10 MB per query thread (see RF_QUERY_THREADS) and 50 MB per index thread (see RF_INDEX_THREADS).

Using the Config Tool to Modify Database and Portal Logging Settings

The Config Tool provides a graphical user interface to simplify making changes to the database settings, and to test the database connection. In addition, you can use this tool to enable logging messages from the portal server to be viewed remotely.

The changes you make using the Config Tool affect settings in the serverconfig.xml file.

To use the Config Tool:

1. From a command prompt, enter **ptconfig.exe** (Windows) or **ptconfig.sh** (Unix).
2. In the Config Tool window:
   - To change the database login and password settings, click the **Database Settings** tab, and enter the desired settings, along with the database host, port number, and database type.
   - To test the database connection, on the Database Settings tab, click **Test DB Connection**. A status message is displayed.
   - To enable remote viewing of portal logging messages, click the Logging Settings tab. Enter the server name, and under Local Machine Only, click **False**.

   Note: This only enables remote logging for the portal server, not the automation or the API servers. To enable remote logging for those servers, modify the serverconfig.xml file:

   - To enable remote logging for the automation server, in the **automation server logging** section, set **logging:local-only** to **false**.
   - To enable remote logging for the API server, in the **wsserver logging** section, set **logging:local-only** to **false**.
3. Click **OK**.
4. Restart the portal for the configuration changes to take effect.

Configuring AquaLogic Interaction Logging Spy

Portal is installed with AquaLogic Interaction Logging Utilities, and by default is set to send log messages to those utilities. You can configure the portal to disable logging or to enable remote logging (to have the portal send log messages to a separate machine on a network). For information on how to configure portal logging, see the *Installation and Upgrade Guide for AquaLogic Interaction Logging Utilities*.

AquaLogic Interaction Logging Spy (formerly PTSpy) is the primary log message receiver in the AquaLogic Interaction Logging Utilities. AquaLogic Interaction Logging Spy provides a graphical user interface for displaying log messages as they stream in from the portal and other log message senders (such as AquaLogic Interaction Collaboration or AquaLogic Interaction Publisher).

Configuring AquaLogic Interaction Logging Spy to Display Portal Messages

To configure AquaLogic Interaction Logging Spy to display portal messages:

1. Launch AquaLogic Interaction Logging Spy by navigating to **Start | All Programs | BEA | ALI Logging Utilities | Logging Spy**. For more information on using AquaLogic Interaction Logging Spy, see the Online Help provided with AquaLogic Interaction Logging Spy.

2. Open the **Filter Settings** window by selecting **View | Set Filters**.

3. To add a logging sender, right-click anywhere in the Filter Settings window. The context menu appears.

4. Select **Add Message Sender**. The Add Message Sender dialog box appears.

5. Type a server name or select it from the list of names and click **OK**.

   Server names exist in the <*logging:server-name*> nodes of the *Logging* sections in the portalconfig.xml file.

   When you add a server as a message sender, it appears in a tree structure in the Filter Settings window. Click the plus sign to expand the server and see a list of its message-sending components.

6. In the **Filter Settings** window, expand each component under a server to see the selected logging levels for that component.

The checkbox next to each component has three states:

- Gray with a check mark - some, but not all, logging levels are selected
- Clear with a check mark - all of the logging levels are selected
- Clear - none of the logging levels are selected

You can toggle through these states by clicking the checkbox next to the component.

7. You can perform the following additional actions in the Filter Settings window:
   - To remove a message-sending server and its components, right-click on the server name, and select **Remove Message Sender**.
   - To enable a selected logging level for all components of a server, right-click on the server name, and select Enable <*LoggingLevel*>, for example, **Enable Performance**.
   - To enable or disable logging levels for a single component, expand the component, and select or clear the checkbox next to the logging level.
   - To clear all logging levels for all components of a server, right-click on the server name, and select **Clear All Filters**. Then click **OK** when asked to confirm. This prevents those components from sending logging messages to this instance of Logging Spy.
   - To reset logging levels for all components of a server to the original four levels, right-click on the server name, and select **Reset Filters**. Then click **OK** when asked to confirm.
8. Click **OK** when finished.

Using the Counter Monitoring System

This chapter describes how to use the Counter Monitoring System to view real time statistical data on your portal, reported by various performance counters. This chapter includes the following sections:

About Counter Monitoring

The Counter Monitoring System collects information from various performance counters for portal applications and exposes them for diagnosis and review. This system can be used to examine counters from any AquaLogic User Interaction application that resides on a remote host, provided the two machines are on a network in which they can reach each other via UDP.

With the Counter Monitoring System you can:

- Set up counter logging files in your desired format to view counter information.
- Use the Counter Monitoring console to request specific counter data in real time.
- If you use a Windows system, use the Windows Perfmon utility to view portal counter data.

The following table lists the key counters provided with the portal. Each category of performance has one or more instances. Each instance in a category can be monitored with the counters for that category.

Note: You can get a complete list of available counters using the **info** command in the Counter Monitoring console. See "Running the Counter Monitoring Console" on page B-6.

Table B-1 Key Performance Counters

Category	Instance(s)	Counter(s)
Cache Many UI objects and pages have their own individual cache systems. Cache counters track each individual cache.	CommunityInfoCache - The cache for a PT Community GuestLoginInfoCache - The cache for a Guest Login HTTP_CACHE - The cache for HTTP requests, for remote portlets or Web services PreferenceCache - The cache for any preference page SubportalInfoCache - The cache for any experience definition	**Size** - The number of items currently in the cache **MaxSize** - The maximum number of items in the cache before it gets flushed **NumSearches** - Increments every time the cache is accessed **NumHits** - Increments every time a cache is accessed and cached contents are found **NumInserts** - Increments every time a cache is accessed and no cached contents are found
Opendb_SQLstats Database statistics for OpenDB	**SQLSelectStats** - SQL queries that are "SELECT" statements	**NumOperations** - The number of SQL operations that occurred
OpenHTTPLowLevelNetwork Counter Basic HTTP information, including usage, connections, transactions	**Total** - There is one instance per remote host. Total aggregates all of the statistics.	**BytesReceived** - Number of bytes received from the remote host **BytesSent** - Number of bytes sent to the remote host **OpenConnections** - The number of open connections to remote hosts

Table B-1 Key Performance Counters (Continued)

Category	Instance(s)	Counter(s)
OpenHTTPHttpLevelstatistics HTTP requests statistics	**Total** - There is one instance per remote host. Total aggregates all of the statistics.	**RequestsActive** - The number of HTTP requests that are active **RequestsProcessed** - The number of HTTP requests that have been processed
portalpages Statistics related to portal pages	NA - Single instance	**CommunityPages** - How many times a community page was hit **LoginsFailure** - How many times a user login attempt failed **LoginsSuccessful** - How many times users logged in **MyPages** - How many times a My Page was hit **TotalHits** - How many times any portal page was hit **TotalOpensessions** - How many open sessions there are currently

Setting up Counter Log Files

You can specify the location, size, and format of the counter log file(s), how often the counter values are polled, as well as filter which counters are written to the file(s).

To set up the counter log file(s):

- In the PORTAL_HOME\bin directory, enter the following command:

  opencounterslogger.bat (or opencounterslogger.sh on Unix) *ServerName ContextName* **-d** *LogDirectory* **-l** *LogOutputStyle* **-s** *MaxLogSize* **-r** *PollingInterval* **-f** *FilterExpression*

  Where these parameters are:

  ServerName - The name of the server where the portal you want to monitor is installed, for example: ptserver2. The value of this name is set in the **opencounters:application-name** element in the
  serverconfig.xml file (in *PT_HOME*\settings\common). This value is case-sensitive.

  ContextName - The name of the AquaLogic User Interaction application you want to monitor, for example: *portal*. The value of this name is set in the **context** element in the serverconfig.xml file. This value is case-sensitive.

  -d *LogDirectory* - (Optional) The local directory in which to create counter log files. The directory must exist prior to executing the opencounterslogger.bat command.

  -l *LogOutputStyle* - (Optional) The log output style can be any combination of the following:

  – **c** - Log counter values in a .csv file, with counters sorted by counter name. CSV signifies a comma-delimited file, which can be read by applications like Excel, for convenient graphing of the values.

  – **t** - Log counter values in a .csv file, with counters sorted by time stamp.

  – **f** - Stream counter information to the file, with each line consisting of a counter name/value pair

  – **s** - Stream counter information to the screen, with each line consisting of a counter name/value pair.

  -s *MaxLogSize* - (Optional) An integer that specifies the maximum size of any log file in kilobytes. Once this log file size is reached, the log is rolled over to a new file.

  -r *PollingInterval* - (Optional) The rate, in seconds, at which counter values should be logged to file.

Note: decreasing the polling interval (i.e., increasing the polling rate) can affect overall portal performance. Retaining the out-of-the-box setting should only affect performance by a maximum of 2%.

-f *FilterExpression* - (Optional) An expression that filters which counters are logged. Expressions are case-insensitive. Use Table B-1 to look up Category, Instance, and Counter names, then write the expression in the following format:

*CategoryNameContains***:***InstanceNameContains***:***CounterNameContains*

Each condition in the above expression is optional, and is used to limit the counter values returned.

For example:

-f Cache::Num will match all counters with a category name that contains *Cache* and a counter name that contains *Num*. The instance name will not be filtered.

Note: You can see a list of all available counters using the **info** command in the Counter Monitoring console, see "Running the Counter Monitoring Console" on page B-6.

Example of the opencounterslogger Command:

opencounterslogger.bat *PtServer2admin collab* **-d** *C:\logdir* **-l** *t* **-s** *1000* **-r** *10* **-f** *open***:***sql***:**

This command does the following:

- Listens to the *collab* application on *PtServer2admin* (in this example, AquaLogic Interaction Collaboration is installed)
- Logs all counters to *C:\logdir*
- Generates a log file sorted by *timestamp*
- Has a maximum log size of one megabyte, or *1000 kilobytes*
- Logs values every *10 seconds*
- Limits output to log counters with categories names that contain *open* and with instance names that contain *sql* (for example, category: Opendb_SQLstats, instance: SQLSelectStats)

When you enter this command, you should see the following:

```
Starting counter logger...

Log files will be written to directory: C:\logdir

Logging rate (seconds): 10

Log file rollover size (kilobytes): 1000

Logging from host: PtServer2.collab
```

At this point, if the logger is able to successfully connect to the server, you will see the following:

```
Logging counters...
```

You can now check the logging directory for log files. Log files will not appear until there is at least one counter created on the counter server.

Running the Counter Monitoring Console

In addition to monitoring counter log files, you can view specified counter values through a console.

To view counter values through a console:

1. In the PORTAL_HOME\bin directory, enter the following command:

   opencountersconsole.bat (opencountersconsole.sh on Unix) *ServerName ContextName*

   Where these parameters are:

   ServerName - The name of the server where the portal you want to monitor is installed, for example: ptserver2. The value of this name is set in the **opencounters:application-name** element in the
   serverconfig.xml file (in *PT_HOME*\settings\common). This value is case-sensitive.

   ContextName - The name of the AquaLogic User Interaction application you want to monitor, for example: *portal*. The value of this name is set in the **context** element in the serverconfig.xml file. This value is case-sensitive.

   Once you enter the command, you see the counter console startup messages. If the connection to the specified server succeeds, you should see the following (with the server and application names you are monitoring displayed before the command prompt):

   ```
   ... Connection success!

   [ServerName.ContextName]>
   ```

Example:

opencountersconsole.bat *PtServer2 collab*

This command opens a console to monitor the *collab* application on *PtServer2*. In this example, AquaLogic Interaction Collaboration is installed.

2. From the console command prompt, use any of the available commands (shortcut keys are in parentheses):

   - **help** (h) - Displays the help menu
   - **last** (l) - Performs the last command that was entered
   - **num** (n) - Gets the total number of counters that are available in the host server.
   - **info** (i) - Returns all counter names with their current values and additional description information. Optionally, you can enter a filter expression after the command to limit the information returned. A filter expression has the following format:

     *CategoryNameContains***:***InstanceNameContains***:***CounterNameContains*

     Each condition in the above expression is optional, and matches a substring of the category, instance, or counter names. See Table B-1 for some key category, instance, and counter names.

     Example: **info cache:pref:num**

   - **values** (v) - Returns all counter names with their current values. Optionally, you can enter a filter expression after the command to limit the information returned. See the above example.

     Example: **values cache:pref:num**

   - **filterset** (fs) - Sets the current filter for the category name, instance name, and counter name. This filter limits the information returned when using the **filterget** command. If **filterset** is not specified, the **filterget** command matches all available counters. This command uses a filter expression as described under the **info** command above.

     Example: **filterset cache:pref:num**

     Using this example command, followed by the command **filterget**, gives the same result as using the command **values cache:pref:num** by itself.

     The filterset you specify remains in effect until you set a new one. To reset the filterset to the default value, enter **filterset** by itself. The default filterset matches all counters.

– **filtercategory** (fc) - Sets the current category name filter. This affects the counters that are returned by the **filtergetvalues** command. The category names returned contain the substring you enter.

  Example: **filtercategory Open**

  The category filter you specify remains in effect until you set a new one.

– **filterinstance** (fi) - Sets the current instance name filter. This affects the counters that are returned by the **filtergetvalues** command. The instance names returned contain the substring you enter.

  Example: **filterinstance Total**

  The instance filter you specify remains in effect until you set a new one.

– **filtercounter** (fr) - Sets the current counter name filter. This affects the counters that are returned by the **filtergetvalues** command. The counter names returned contain the substring you enter.

  Example: **filtercounter Bytes**

  The counter filter you specify remains in effect until you set a new one.

– **filtergetvalues** (fg) - Returns counter names with their current values. This command only retrieves counters that match the filters you have set using any of these commands: **filterset**,
filtercategory, **filterinstance**, **filtercounter**. If no filters have been set, the command matches all counters.

– **verbose** (vb) - Toggles verbosity level on and off (default is OFF). When verbosity is set to On, the **values** command returns counter values with additional counter information (such as a counter description). When verbosity is set to Off, the **values** command returns counter values without additional counter information.

– **connect** (c) - Connects to another host server. This disconnects from the current host server if one is already connected.

  Example: **connect PtServer4 portal**

– **disconnect** (d) - Disconnects from the current host server, if connected.

– **exit** (e) - Exits the console application

Using Windows Perfmon to View Counter Data

The Counter Monitoring System integrates with the Windows Perfmon application. Once you start the portal, the Perfmon adaptor will add AquaLogic User Interaction counters to the list of possible counters to monitor. You can then start Perfmon (or any other monitoring application that works with Windows Performance Counters) and see AquaLogic User Interaction counters in the list of available counters.

To start Windows Perfmon:

1. Go to **Start | Run**.
2. Enter **perfmon.exe** and click **OK**.

   Note: In Perfmon, the category name is prefixed by the context name. The context name is set in the **context** element in the serverconfig.xml file (in *PT_HOME*\settings\common).

Disabling the Perfmon Adaptor

The Perfmon adaptor adds a few percentage points of overhead to overall system performance. If you want to disable the Perfmon adaptor, in *PT_HOME*\settings**serverconfig.xml**, set **opencounters:perfmon-enabled** to **false**.

Localizing Your Portal

This chapter describes how you can internationalize and localize your portal, and includes the following sections:

The Purpose of Localization

Each object can have an alternate name and description for each language that is supported in the portal. When you create an object (for example, a portlet) in the portal, you supply a primary name and description for the object, which displays to users, regardless of their locale choice. However, to accommodate users whose portal interface displays in a different language, you can also supply localized names and descriptions.

For example, if you have an object called "Engineering," then you might have the French translation, "Ingénierie," as the alternate name for that object. So when users log in using the French user interface, they see the object names and descriptions associated with French.

Localizing Objects

Although you can supply the localized terms on an object-by-object basis through the object editor, you might find it more efficient to edit all objects at the same time through the Localization Manager.

The Localization Manager allows you to download an .xml file that includes the names and descriptions for all objects that support localized names. Configure each object to allow localization before downloading the .xml file. You can then edit the list and upload it back into the portal.

Enabling Object Localization

To enable the object to be localized:

1. Click **Administration**.
2. Navigate to an administrative folder.
3. Either edit an existing object (like a content crawler, a portlet, or a job) by clicking an object, or create a new object by selecting an object from the **Create Object** drop-down list. This launches an object editor.
4. Complete the configuration properties, including information about locales, according to the online help.

Localizing All Objects

To localize objects using the Localization Manager:

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Localization Manager**.
3. In the Localization Manager, click **Download**.
4. Save the .xml file to your computer, as prompted.
5. Use a text editor to edit the entries. For example, to edit the French term, you might make the following change:

   ```
   <target language=”fr”>Tous</target>
   ```

6. Save your changes.

7. In the Localization Manager, click **Browse** to browse to your edited file, and then click **Upload**. A message displays at the bottom of the editor as to whether the upload was successful.

8. Click **Finish**.

For instructions and a detailed description of any page in the portal, click ? **Help**.

Localization Manager XML

Each localized object has an entry in the following format:

```
<segment classid=”2” itemid=”51” stringid=”0”>
  <source language=”en”>Everyone</source>
  <target language=”de” />
  <target language=”en” />
  <target language=”es” />
  <target language=”fr” />
  <target language=”it” />
  <target language=”ja” />
  <target language=”ko” />
  <target language=”pt” />
  <target language=”zh” />
  </segment>
```

The first line displays information about the object entry:

- The classid represents the object type (in this example, a group).
- The itemid represents the object ID in the portal database.
- The stringid is “0” for the object name and “1” for the object description.

The second line displays the primary language term (in this example, the primary language is English and the term is Everyone).

The remaining lines display the available languages:

Table C-1 Available Languages

Value	Language
de	German
en	English
es	Spanish
fr	French
it	Italian
ja	Japanese
ko	Korean
nl	Dutch
pt	Portuguese
zh	Simplified Chinese
zh-tw	Traditional Chinese

Localizing Each Object

To specify a name and description, language, localized names and descriptions, and properties of this object in the Properties and Names page of the object editor (for instructions on navigating to this editor, see "Enabling Object Localization" on page C-2), do the following:

1. Next to **Primary Language**, you see one of the following:
   - If the MandatoryObjectLanguage in your portalconfig.xml file has been set, you see the mandatory language. You must supply a name (and, optionally, a description) for this language.
   - If the mandatory language has not been set, you see a drop-down list. Select the language for this object's default name and description.

   The primary language sets the default language, so users whose locale settings are outside of the supported languages will see the name and description in the primary language you have selected.
2. Select the **Supports Localized Names** check box. A new Localized Name and Description section displays, allowing you to add names and descriptions for additional languages.
3. Under the Localized Name and Description section, click **New Localized Name**.
4. In the **Name and Description** dialog box, type the localized name for this object.

5. In the **Language** drop-down list, choose the language for which you are adding a name and description.
6. In the **Description** text box, type the localized description for this object.
7. When you are done, click **Finish**.

   Note: If this object has already been localized into all supported languages, you cannot add more localized names and descriptions. (The new Localized Name button is not displayed.)

Internationalizing the Search Service

Unicode characters are used to store and retrieve text, and the system has access to linguistic rules for multiple languages during full-text indexing. This makes it possible to have documents of different languages within the same search collection, with significantly improved results.

The portal provides support for 61 languages. Of these, the following languages include support for word stemming and compound decomposition. This additional information is used to enhance results of the full-text index.

- Chinese (Simplified)
- Chinese (Traditional)
- Czech
- Danish
- Dutch
- English
- Finnish
- French
- German
- Greek
- Hungarian
- Italian
- Japanese
- Korean
- Norwegian (Bokmal)
- Norwegian (Nynorsk)
- Polish
- Portuguese
- Russian
- Spanish
- Swedish
- Turkish

The following languages are supported at a reduced level.

- Afrikaans
- Albanian
- Arabic
- Basque
- Belarusian
- Bengali
- Bulgarian
- Catalan
- Cornish
- Croatian
- Esperanto
- Estonian
- Faeroese
- Gallegan
- Hebrew
- Hindi
- Icelandic
- Indonesian
- Irish
- Kalaallisut
- Konkani
- Latvian
- Lithuanian
- Macedonian
- Maltese
- Manx
- Marathi
- Persian
- Romanian
- Serbian
- Serbian-Croatian
- Slovak
- Slovenian
- Swahili
- Tamil
- Telugu
- Thai
- Ukranian
- Vietnamese

Displaying International Text in the Portal

All content in the search collection is stored as Unicode characters. The user interface handles text using the UTF-8 encoding, so search results are always displayed correctly, assuming that the appropriate fonts are available to the Web browser.

Searching International Documents

For information on searching international documents, see “Search Language” on page E-10.

Crawling International Document Repositories

Web and file content crawlers are associated with a specific language. All documents processed by a content crawler are indexed using the linguistic rules appropriate for the specified language. For optimal results, create a separate content crawler to handle documents of different languages. For most European languages, mixing languages within a single crawl will not render the content unsearchable; however, word stemming and decomposition information stored in the documents will be missing for languages other than the content crawler's designated language. Avoid indexing Asian language documents with a content crawler configured for a European language, as special tokenization rules are required for processing the Asian languages.

Submitting International Documents to the Knowledge Directory

When you use the Submit Document utility to add documents to the Knowledge Directory, you specify the document language by choosing from a pop-up list of the supported languages. As with content crawlers, this language should be set to the actual language of the document for optimal results. Correct specification of language is particularly crucial for proper indexing of Asian language content.

Deploying Single Sign-On

This appendix describes how to deploy Single Sign-On (SSO) capabilities in the portal environment. This appendix includes the following sections:

- "About SSO" on page D-1
- "Configuring an SS0 Authentication Provider for Use with the Portal" on page D-2
- "Configuring an Authentication Source" on page D-28
- "Configuring the Portal for SSO" on page D-29

About SSO

SSO is an authentication system that permits users to access multiple servers in a domain through a single point of entry. When SSO is deployed in the portal, user sessions are authenticated transparently by an SSO service against authentication sources commonly deployed in the enterprise, such as LDAP or Active Directory.

To deploy SSO in your portal, configure the following resources so that they can share user and domain authentication information:

1. SSO Authentication Server. See "Configuring an SS0 Authentication Provider for Use with the Portal" on page D-2.
2. Authentication Source. See "Configuring an Authentication Source" on page D-28.
3. Portal. See "Configuring the Portal for SSO" on page D-29.

Configuring an SSO Authentication Provider for Use with the Portal

This section describes how to configure authentication servers to protect the portal. The following topics provide configuration details for supported and unsupported servers:

- “Configuring the Windows Integrated Authentication Service” on page D-3
- “Configuring Netegrity SiteMinder 4.6” on page D-4
- “Configuring Netegrity SiteMinder 5.5” on page D-9
- “Configuring a Netegrity SiteMinder 5.5 SP3 and 6.0.2.5” on page D-14
- “Configuring an Oblix Authentication Provider” on page D-20
- “Integrating With Other Authentication Providers” on page D-27

Caution: You configure the SSO authentication server to protect the portal. You do not need to configure the SSO server to protect the Image Service. If you do, communication between the portal and Image Service can result in errors. The Image Service contains only static public content that ships with every portal installation. No data specific to users or to your organization is ever stored on the Image Service.

Configuring the Windows Integrated Authentication Service

Follow these steps to configure the Windows Integrated Authentication (WIA) service (formerly known as Windows NT LAN Manager (NTLM)) for your SSO deployment:

1. Open the IIS Internet Service Manager.
2. In the left pane, expand the Web server folder and then its portal virtual directory subfolder; right-click the **sso** folder and choose **Properties**.
3. Click the **Directory Security** tab.
4. In the Anonymous access and authentication control group, click the **Edit** button.
5. In the Authenticated Access group, select the **Integrated Windows Authentications** box. Leave the remaining boxes unselected.
6. Click **OK** to accept changes to Directory Security settings.
7. Click **OK** to accept changes to Properties.

   Note: In the portal environment where SSO via WIA is deployed, Internet Explorer users are not prompted for a user name and password when they attempt to log into the portal Web site if the following provisions have been made by the client:

   – The user must be logged into a Windows NT domain as a user that has rights to access the portal.

   – An HTTP proxy must not reside between the client computer and the portal Web site.

   – Internet Explorer must be configured to recognize the portal Web site as a local intranet site. If the site is not in the local intranet zone by default, add it from the browser:

   a. Choose **Tools | Security | Local Intranet**.

   b. Click **Sites**.

   c. Click **Advanced**.

   d. Add the address for the portal Web site.

   Note: Netscape Navigator versions prior to 7.1 are not supported. For Netscape Navigator 7.1 (and later), users are prompted for user name and password when they attempt to log in to the portal Web site.

   Note: In the portal environment where SSO via WIA is deployed, the portal does not pass user passwords as Basic Authentication Headers to remote servers.

Configuring Netegrity SiteMinder 4.6

Follow these basic steps to configure Netegrity SiteMinder 4.6 for use with the portal:

1. Configure SiteMinder Policy Server. For detailed instructions, see "Configuring Netegrity SiteMinder Policy Server" on page D-4.
2. Configure SiteMinder Web Agent. For detailed instruction, see "Configuring Netegrity SiteMinder Web Agent (4.6)" on page D-7.

Configuring Netegrity SiteMinder Policy Server

To configure SiteMinder Policy Server for your SSO deployment:

1. Install the server as described in Netegrity documentation.
2. Open the SiteMinder administrative tool and log in as a user that can create objects.
3. Create the following objects in the order they are presented.

Table D-1 Procedures for Creating Objects in Netegrity SiteMinder Policy Server 4.6

Object	To create the object:
Agents	1. In the left pane, right-click Agents and choose **Create Agent**. 2. In the Name box, type the name of the portal. 3. In the address box, type the IP address of the portal or use SiteMinder controls to perform DNS lookup. 4. In the shared secret box, type a string that matches one to be set on the Web Agent host. 5. Click **Apply** and then **OK**.
User Directory	1. In the left pane, right-click User Directories and choose **Create User Directory**. 2. In the Name box, type a descriptive name for the object, for example `Iplanet`. 3. In the NameSpace box, choose the appropriate namespace, for example: – **WinNT**. If you choose WinNT, specify the Windows NT domain name in the Windows Domain text box. – **LDAP.** If you choose LDAP, specify the IP address and port number for server that hosts the LDAP user directory and use the LDAP Search and LDAP User DN Lookup group controls to configure search and lookup according to the conventions and examples described in the context-sensitive online help for the Netegrity administration tool. 4. Click **Apply**. 5. To display user groups that have been imported into the Policy Server, click **View Contents**. 6. To verify that users have been imported, click **Search** and query LDAP for specific users. 7. Click **Apply** and then **OK**.
Policy Domain	1. In the left pane, right-click Policy Domain and choose **Create Policy Domain**. 2. In the name box, type a descriptive name for the domain, for example `Portal`. 3. In the add directory box, specify the User Directory created above. 4. Click **Apply** and then **OK**.

Table D-1 Procedures for Creating Objects in Netegrity SiteMinder Policy Server 4.6

Object	To create the object:
Realm	1. In the left pane, click the **Domains** tab. 2. Right-click the domain created above and choose **Create Realm**. 3. In the name box, type a descriptive name for the realm, for example `SSO`. 4. In the Resource group, from the Agent drop-down list box, select the agent you created above. 5. In the Resource Filter box, type `/portal/sso`, which is the directory the portal uses to authenticate against SSO services. 6. In the Authentication Scheme box, choose **Basic Authentication**. 7. Click **Apply** and then **OK**.
Rule for the Realm	1. In the left pane, expand the policy domain tree so it displays named realms; right-click the realm created above and choose **Create Rule under Realm**. 2. In the name box, type a descriptive name for the rule, for example `Allow Access`. 3. In the Realm and Resource group, choose the realm created above; in the resource box, specify `/*`. 4. In the Allow/Deny and Enable/Disable group, enable the rule and set it to **Allow Access** when the rule fires. 5. In the Action box, click **Web Agent Actions** and then **GET**, **POST**, and **PUT**. 6. Click **Apply** and then **OK**.
Policy for the Realm	1. Under the domain created above, right-click Policies and choose **Create Policy**. 2. In the Name box, type a descriptive name for the Policy, for example `Normal Case`. 3. Click the **Users** tab and use the controls to add users or groups for whom this policy applies. 4. Click the **Rules** tab and use **Add/Remove Rules** button to add the rule you created above. 5. Click **Apply** and then **OK**.

Configuring Netegrity SiteMinder Web Agent (4.6)

To configure SiteMinder Web Agent for your deployment:

1. Install the Web Agent setup program on the same host as the portal.
2. When setup is complete, you are prompted to complete the Web Agent Configuration Wizard. If you choose to run the wizard at a different time, click **Start | Programs | SiteMinder | Web Agent Configuration Wizard** to open the wizard.

   When the wizard prompts for components to configure, choose **IIS**, click **Configure,** and complete the following configuration information.

Table D-2 Netegrity SiteMinder Web Agent Configuration Wizard Pages

Prompt	Value
Policy Server	Type the name of the Policy Server you set up in the previous section.
Agent name	Type the name of the agent you created in the previous section.
Cookie domain	Type the fully qualified domain name for which you want the authentication cookie to be forwarded. For example, if you specify `.company.com`, the cookie enables access to all domains that end in **company.com**. If you specify `.sub.company.com`, the cookie enables access only to domains that end in **sub.company.com**.
IIS Proxy User Name and Password	Type a user name and password to run the SiteMinder ISAPI filter on IIS. This user must have administrator rights on the IIS host.
Shared Secret	Type a string that exactly matches the name of the Agent object you created on the Policy Server.

3. To open the IIS Web Agent Management Console, click **Start | Programs | SiteMinder | IIS Web Agent Management Console**.

   Modify the configuration as described in the following table.

Table D-3 Web Agent Management Console Configuration

Console Control	To modify settings:
Settings tab	Click the **Settings** tab and specify the following settings: • Enable Web Agent • Enforce Policies
Single Sign On tab	Click the **Single Sign On** tab and specify the following settings: • Select **Require Cookies** and clear the other two boxes. • Set the Cookie Domain to the fully qualified domain name for which you want the cookie to be forwarded. For example, if you specify `.company.com`, the cookie enables access to all domains that end in **company.com**. If you specify `.sub.company.com`, the cookie enables access only to domains that end in **sub.company.com**.

4. Restart IIS to apply the modified settings.

Configuring Netegrity SiteMinder 5.5

Follow these basic steps to configure Netegrity SiteMinder Policy Server 5.5 for use with the portal:

1. Configure the Netegrity SiteMinder Policy Server.

   For detailed instructions, see "Configuring Netegrity SiteMinder Policy Server" on page D-9.

2. Configure Netegrity SiteMinder Web Agent.

Configuring Netegrity SiteMinder Policy Server

To configure SiteMinder Policy Server for your deployment:

1. Install the server as described in Netegrity documentation.
2. Open the SiteMinder administrative tool and log in as a user that can create objects.
3. Create the following objects in the order they are presented.

Table D-4 Procedures for Creating Objects with Netegrity SiteMinder Policy Server 5.5

Object	To create the object:
Host Conf Objects	1. In the left pane, click **Host Conf Object**. 2. In the right pane, right-click the Default Host Configuration Object and choose **Duplicate Configuration Object**. 3. In the Name box, type a host name, such as `policyserver`. 4. In the Configuration Values group, double-click the **policyserver** object and use the controls to set the IP address for the policy server address and the three ports, typically 44441, 44442, 44443. For example, type `10.1.140.124,44441,44442,44443`. 5. Click **Apply** and then **OK**.
Agents	1. In the left pane, right-click Agents and choose **Create Agent**. 2. In the Name box, type the name of the portal. 3. Click **Apply** and then **OK**.

Table D-4 Procedures for Creating Objects with Netegrity SiteMinder Policy Server 5.5

Object	To create the object:
AgentConf Objects	1. In the left pane, click **Agent Conf Objects**. 2. In the right pane, right-click the type of server that approximates the default settings (for example, IISDefaultSettings) and choose **Duplicate Configuration Object**. 3. In the Name box, type a descriptive name for the object, typically the host name followed by the configuration object name, for example `PortalServerIISDefaultSettings`. 4. In the Configuration Values group, double-click **DefaultAgentName**; uncomment the parameter (remove the leading # from) and specify as its value the name of the agent created in the previous step 5. Click **Apply** and then **OK**.
User Directory	1. In the left pane, right-click User Directories and choose **Create User Directory**. 2. In the Name box, type a descriptive name for the object, for example `Iplanet`. 3. In the NameSpace box, choose the appropriate namespace, for example: • **WinNT**. If you choose WinNT, specify the Windows NT domain name in the Windows Domain text box. • **LDAP.** If you choose LDAP, specify the IP address and port number for the server that hosts the LDAP user directory and use the LDAP Search and LDAP User DN Lookup group controls to configure search and lookup according to the conventions and examples described in the context-sensitive online help for the Netegrity administration tool. 4. Click **Apply**. 5. To display user groups that have been imported into the Policy Server, click **View Contents**. 6. To verify that users have been imported, click **Search** and query LDAP for specific users. 7. Click **Apply** and then **OK**.
Policy Domain	1. In the left pane, right-click Policy Domain and choose **Create Policy Domain**. 2. In the name box, type a descriptive name for the domain, for example `Portal`. 3. In the add directory box, specify the User Directory created above. 4. Click **Apply** and then **OK**.

Table D-4 Procedures for Creating Objects with Netegrity SiteMinder Policy Server 5.5

Object	To create the object:
Realm	1. In the left pane, click the **Domains** tab. 2. Right-click the domain created above and choose **Create Realm**. 3. In the name box, type a descriptive name for the realm, for example `SSO`. 4. In the Resource group, from the Agent drop-down list box, select the agent you created above. 5. In the Resource Filter box, type `/portal/sso`, which is the directory the portal uses to authenticate against SSO services. 6. In the Authentication Scheme box, choose **Basic Authentication**. 7. Click **Apply** and then **OK**.
Rule for the Realm	1. In the left pane, expand the policy domain tree so it displays named realms; right-click the realm created above and choose **Create Rule under Realm**. 2. In the name box, type a descriptive name for the rule, for example `Allow Access`. 3. In the Realm and Resource group, choose the realm created above; in the resource box, specify `/*`. 4. In the Allow/Deny and Enable/Disable group, enable the rule and set it to **Allow Access** when the rule fires. 5. In the Action box, click **Web Agent Actions** and then **GET**, **POST**, and **PUT**. 6. Click **Apply** and then **OK**.
Policy for the Realm	1. Under the domain created above, right-click Policies and choose **Create Policy**. 2. In the Name box, type a descriptive name for the Policy, for example `Normal Case`. 3. Click the Users tab and use the controls to add users or groups for whom this policy applies. 4. Click the **Rules** tab and use **Add/Remove Rules** button to add the rule you created above. 5. Click **Apply** and then **OK**.

Configuring Netegrity SiteMinder Web Agent (5.5)

To configure SiteMinder Web Agent for your deployment:

1. Install the Web Agent setup program on the same host as the portal.
2. When setup is complete, you are prompted to complete the Web Agent Configuration Wizard. If you choose to run the wizard at a different time, click **Start | Programs | SiteMinder | Web Agent Configuration Wizard** to open the wizard.

   When the wizard prompts for components to configure, choose **IIS**, click **Configure,** and complete the following configuration information.

Table D-5 Netegrity SiteMinder Web Agent Configuration Wizard Pages

Prompt	Value
Policy Server	Type the name of the Policy Server you set up in the previous section.
Agent name	Type the name of the agent you created in the previous section.
Cookie domain	Type the fully qualified domain name for which you want the authentication cookie to be forwarded. For example, if you specify `.company.com`, the cookie enables access to all domains that end in **company.com**. If you specify `.sub.company.com`, the cookie enables access only to domains that end in **sub.company.com**.
IIS Proxy User Name and Password	Type a user name and password to run the SiteMinder ISAPI filter on IIS. This user must have administrator rights on the IIS host.
Shared Secret	Type a string that exactly matches the name of the Agent object you created on the Policy Server.

3. To open the IIS Web Agent Management Console, click **Start | Programs | SiteMinder | IIS Web Agent Management Console**.

   Modify the configuration as described in the following table.

Table D-6 Web Agent Management Console Configuration

Console Control	To modify settings:
Settings tab	Click the **Settings** tab and specify the following settings: • Enable Web Agent • Enforce Policies
Single Sign On tab	Click the **Single Sign On** tab and specify the following settings: • Select **Require Cookies** and clear the other two boxes. • Set the Cookie Domain to the fully qualified domain name for which you want the cookie to be forwarded. For example, if you specify `.company.com`, the cookie enables access to all domains that end in **company.com**. If you specify `.sub.company.com`, the cookie enables access only to domains that end in **sub.company.com**.

4. Restart IIS to apply the modified settings.

Configuring a Netegrity SiteMinder 5.5 SP3 and 6.0.2.5

These instructions apply to SSO integration with Netegrity SiteMinder 5.5 SP3, and 6.0.2.5.

Follow these basic steps to configure SiteMinder Policy Server for use with the portal:

1. Install and configure the SiteMinder Policy Server. For details, see "Configuring Netegrity SiteMinder Policy Server" on page D-14.
2. Install and configure a corresponding SiteMinder Web Agent. For details, see "Configuring Netegrity SiteMinder Web Agent (5.5 SP3 or 6.0) on Linux/Unix" on page D-17.

Configuring Netegrity SiteMinder Policy Server

To configure SiteMinder Policy Server for your deployment:

1. On a remote host computer, install SiteMinder Policy Server 5.5 and SP3, or 6.0.2.5, as described in Netegrity documentation.
2. Open the SiteMinder administrative tool and log in as a user that can create objects.
3. Create the following objects in the order they are presented.

Table D-7 Procedures for Creating Objects with Netegrity SiteMinder Policy Server

Object	To create the object:
Host Conf Objects	1. In the left pane, click **Host Conf Object**. 2. In the right pane, right-click the Default Host Configuration Object and choose **Duplicate Configuration Object**. 3. In the Name box, type a host name, such as `policyserver`. 4. In the Configuration Values group, double-click the **policyserver** object and use the controls to set the IP address for the policy server address and the three ports, typically 44441, 44442, 44443. For example, type: `10.1.140.124,44441,44442,44443` 5. Click **Apply** and then **OK**.
Agents	1. In the left pane, right-click Agents and choose **Create Agent**. 2. In the **Name** box, type the name of host computer for the portal. 3. Click **Apply** and then **OK**.

Table D-7 Procedures for Creating Objects with Netegrity SiteMinder Policy Server

Object	To create the object:
AgentConf Objects	1. In the left pane, click **Agent Conf Objects**. 2. In the right pane, right-click the type of server that approximates the default settings (for example, ApacheDefaultSettings) and choose **Duplicate Configuration Object**. 3. In the Name box, type a descriptive name for the object, typically the host name followed by the configuration object name, for example `PortalServerApacheDefaultSettings`. 4. In the Configuration Values group, double-click **DefaultAgentName**; uncomment the parameter (remove the leading # from) and specify as its value the name of the agent created in the previous step. 5. In the Configuration Values group, double-click **BadURLChars** and modify its value to the following: `//,./,/.,/*,*.,~,\` 6. In the Configuration Values group, configure **LogAppend**, **LogConsole**, **LogFileName**, **LogLevel**, and **Logfile** to your preferences. For details, refer to the online help. 7. Click **Apply** and then **OK**.
User Directory	1. In the left pane, right-click User Directories and choose **Create User Directory.** 2. In the **Name** box, type a descriptive name for the object, for example `Iplanet`. 3. In the **NameSpace** box, choose the appropriate namespace, for example, if you choose LDAP, specify the IP address and port number for the server that hosts the LDAP user directory and use the LDAP Search and LDAP User DN Lookup group controls to configure search and lookup according to the conventions and examples described in the context-sensitive online help for the Netegrity administration tool. 4. Click **Apply**. 5. To display user groups that have been imported into the Policy Server, click **View Contents**. 6. To verify that users have been imported, click **Search** and query LDAP for specific users. 7. Click **Apply** and then **OK**.
Policy Domain	1. In the left pane, right-click Policy Domain and choose **Create Policy Domain**. 2. In the name box, type a descriptive name for the domain, for example `Portal`. 3. In the add directory box, specify the User Directory created above. 4. Click **Apply** and then **OK**.

Table D-7 Procedures for Creating Objects with Netegrity SiteMinder Policy Server

Object	To create the object:
Realm	1. In the left pane, click the **Domains** tab. 2. Right-click the domain created above and choose **Create Realm**. 3. In the name box, type a descriptive name for the realm, for example `SSO`. 4. In the Resource group, from the Agent drop-down list box, select the agent you created above. 5. In the Resource Filter box, type `/portal/SSOServlet`, which is the service the portal uses to authenticate against SSO services. 6. In the Authentication Scheme box, choose **Basic Authentication**. 7. Click **Apply** and then **OK**.
Rule for the Realm	1. In the left pane, expand the policy domain tree so it displays named realms; right-click the realm created above and choose **Create Rule under Realm.** 2. In the name box, type a descriptive name for the rule, for example `Allow Access`. 3. In the Realm and Resource group, choose the realm created above; in the resource box, specify `*`. 4. In the Allow/Deny and Enable/Disable group, enable the rule and set it to **Allow Access** when the rule fires. 5. In the Action box, click **Web Agent Actions** and then **GET**, **POST**, and **PUT**. 6. Click **Apply** and then **OK**.
Policy for the Realm	1. Under the domain created above, right-click Policies and choose **Create Policy**. 2. In the Name box, type a descriptive name for the Policy, for example `Normal Case`. 3. Click the Users tab and use the controls to add users or groups for whom this policy applies. 4. Click the **Rules** tab and use **Add/Remove Rules** button to add the rule you created above. 5. Click **Apply** and then **OK**.

Configuring Netegrity SiteMinder Web Agent (5.5 SP3 or 6.0) on Linux/Unix

To set up SiteMinder Web Agent:

1. If you have not done so already, on the host computer for the portal, install Apache HTTPd. For details, see the *Installation and Upgrade Guide for AquaLogic Interaction*.

2. On the portal host computer, install:

   Either these components:

   – SiteMinder Web Agent 5.5

   – SiteMinder Web Agent 5.5 Quarterly Maintenance Release (QMR) 6

   – SiteMinder Web Agent 5.5 QMR 6 CR007

   Or these components:

   – SiteMinder Web Agent 6.0

   – SiteMinder Web Agent 6.0 Quarterly Maintenance Release (QMR) 2

   – SiteMinder Web Agent 6.0 QMR 2 CR005

   For details on installing Netegrity SiteMinder components, refer to the Netegrity Customer Care Web site and Netegrity SiteMinder documentation.

   For an example of installing the SiteMinder Web Agent, Web Agent QMR, and hotfix, see Knowledge Base article DA_236222, “Netegrity Siteminder 5.5 Web Agent Installation Tips.”

3. Invoke the SiteMinder configuration utility, for example, from the command-line, enter:

   ```
   ./nete-wa-config
   ```

   When prompted, enter your preferences but be sure to specify the settings you configured in the previous section for the following objects:

   – Policy Server IP Address

   – Host Configuration Object name

   When prompted, enter the location for the Apache Web server root directory, for example, **/opt/httpd/**.

4. Source the Netegrity environment script. From the command-line, enter:

   ```
   source /opt/netegrity/siteminder/webagent/nete_wa_env.sh
   ```

5. Modify the Apache Web server **httpd.conf** configuration file to enable the SiteMinder Web Agent. The lines in the following excerpt show an **httpd.conf** file that enables the SiteMinder Web Agent:

```
...
LoadModule sm_module /opt/netegrity/siteminder/webagent/lib/mod2_sm.so


...
# SSO Configuration
SmInitFile /opt/httpd/conf/WebAgent.conf
Alias /siteminderagent/pwcgi/ “/opt/netegrity/siteminder/webagent/pw”
<Directory “/opt/netegrity/siteminder/webagent/pw”>
        Options Indexes MultiViews ExecCGI
        AllowOverride None
        Order allow,deny
        Allow from all
</Directory>

Alias /siteminderagent/pw/  “/opt/netegrity/siteminder/webagent/pw”
<Directory “/opt/netegrity/siteminder/webagent/pw”>
        Options Indexes MultiViews ExecCGI
        AllowOverride None
        Order allow,deny
        Allow from all
</Directory>
```

```
Alias /siteminderagent/  “/opt/netegrity/siteminder/webagent/samples/
<Directory “/opt/netegrity/siteminder/webagent/samples/”>
        Options Indexes MultiViews
        AllowOverride None
        Order allow,deny
        Allow from all
</Directory>

##SITEMINDER .exe ##
AddHandler cgi-script .exe

##SITEMINDER .fcc ##
AddHandler smformsauth-handler .fcc

##SITEMINDER .scc ##
AddHandler smadvancedauth-handler .scc

##SITEMINDER .ccc ##
AddHandler smcookieprovider-handler .ccc

...
```

The lines that configure SiteMinder Web Agent must be *before* lines that include a Web application server configuration file, such as **bea.conf**.

6. Modify the following settings in **/opt/httpd/conf/WebAgent.conf**:
   - Ensure: `enablewebagent=”YES”`
   - Add: `ServerPath=”/opt/httpd/conf/httpd.conf”`
7. Restart the Apache Web server.

Configuring Netegrity SiteMinder Web Agent on (5.5 SP3 or 6.0) Windows

To configure SiteMinder Web Agent for your deployment:

1. Install the Web Agent setup program on the same host as the portal. See "Configuring Netegrity SiteMinder Web Agent (5.5 SP3 or 6.0) on Linux/Unix" on page D-17 for the supported Netegrity versions.
2. When setup is complete, you are prompted to complete the Web Agent Configuration Wizard. If you choose to run the wizard at a different time, click **Start | Programs | SiteMinder | Web Agent Configuration Wizard** to open the wizard.
3. When prompted, specify the settings you configured in the previous section for the following objects:
   - Policy Server IP Address
   - Host Configuration Object name
4. In the **<siteminder_webagent_install_location>\IIS\bin\WebAgent.conf** file, set the **EnableWebAgent** parameter to `yes`.

Configuring an Oblix Authentication Provider

AquaLogic Interaction can integrate Oblix 6.5, 7.0, and 7.0.4 for SSO.

Follow these basic steps to configure the Oblix components for use with the portal:

1. Install and configure the Oblix suite 6.5, 7.0, or 7.0.4. For details, see "Configuring Oblix Access Server for a Portal Running on Tomcat" on page D-20 or "Configuring Oblix Access Server for a Portal Running on IIS" on page D-24.
2. Install and configure a corresponding Oblix Webgate for Apache. For details, see "Configuring Oblix WebGate for Apache" on page D-26.

Configuring Oblix Access Server for a Portal Running on Tomcat

Follow these basic steps to configure Oblix Access Server for use with a portal running on Tomcat:

1. Install Oblix suite 6.5, 7.0, or 7.0.4. Each Oblix suite includes the following components: COREid, WebPass, Access Server, and Access Manager.

   For information on installing Oblix products, refer to Oblix product documentation.
2. Open NetPoint Access Manager, typically **http://oblix_access_server:port/access/oblix**.

3. Create the following objects in the order they are presented.

Table D-8 Procedures for Creating Objects with Oblix NetPoint Policy Server

Object	To create the object:
Policy Domain	1. In the left pane, click **Create Policy Domain**. 2. Type a descriptive name for the policy domain and click **Save**. 3. Click **Modify**. 4. Enable the policy domain and click **Save**.
HTTP Resource	1. Click the **Resources** tab and then click **Add**. 2. In the Resource drop-down list, choose **HTTP**. 3. In the URL-prefix drop-down list, choose the backslash(/) and type **portal**. 4. Click **Save**.
Policy	1. Click the **Policies** tab and then click **Add**. 2. In the **Name** box, type a descriptive name for the policy, for example `allow access`. 3. In the Resource Type drop-down list, choose **HTTP**. 4. In the Resource Operations group, click **GET** and **POST**. 5. In the Resource group, select **all**. 6. In the URL Pattern field, enter **{SSOServlet[;!]*,SSOServlet/.../*,SSOServlet}** URLS that contain the string “SSOServlet” and other variations of “SSOServlet” will be forced to authenticate. Without a URL pattern, Oblix issue an authentication prompt for all requests, however, this would prevent portal Guest functionality. 7. Leave the other fields blank and click **Save**.

Table D-8 Procedures for Creating Objects with Oblix NetPoint Policy Server

Object	To create the object:
Authorization Rules	The Oblix server parses these rules to send the user name attribute to the portal, enabling the user to log in as a known user with user-defined roles, privileges, views, and so forth. 1. Click the **Authorization Rules** tab, and then click **Add**. 2. In the name box, type a descriptive name for the rule, for example `allow everyone rule`. 3. Enable the rule and click **Save**. 4. Click the **Allow Access** sub-tab and then click **Add** to display controls for adding users or groups to allow access to this resource. 5. Add appropriate users and groups. Any user/group that needs access to the portal should be included here. Change the “Role” drop down to “Any one” and click **Save**. **Note:** The name of the user displayed here does not necessarily match the name you enter during Oblix login.

Table D-8 Procedures for Creating Objects with Oblix NetPoint Policy Server

Object	To create the object:
Default Rules	1. Click **Default Rules** tab, then click **Add**. 2. On the **Authentication Rule** sub-tab, enter a Name and Description. 3. In the Authentication Scheme drop-down menu, select **NetPoint None Authentication**. This authentication scheme is created automatically by the Oblix installation. 4. Click **Save**. 5. Click the **Authorization Expression** sub-tab, and click **Add**. 6. On the **Expression** sub-sub-tab, add the `Allow Everyone Rule` you created previously (or select the rule by the name you originally gave it). You must click **Add** so that the name appears in the Authorization Expression box. 7. Click the **Actions** sub-sub-tab, then click **Add**. 8. Under Authorization Success, fill in the last line of fields: – For **Type**, enter a descriptive name such as `headerVar`. – For **Name**, enter **UID**. – For **Return Attribute**, enter the value of the UID header sent to portal. Typically, this value is **UID** if using LDAP or **samaccountname** if using Active Directory. 9. Click **Save**.
Authentication Rule	1. Click the **Policies** tab. 2. Click the name of the policy created above to select it. 3. Click the **Authentication Rule** sub-tab and add a method appropriate to your configuration, for example, Basic over LDAP. **Note:** Basic over LDAP is a rule created during Oblix installation. If it is not available, Oblix was not installed properly. 4. Click **Save**.

Configuring Oblix Access Server for a Portal Running on IIS

Follow these basic steps to configure Oblix Access Server for use with a portal running on IIS:

1. Install Oblix suite 6.5, 7.0, or 7.0.4. Each Oblix suite includes the following components: COREid, WebPass, Access Server, and Access Manager.

   For information on installing Oblix products, refer to Oblix product documentation.

2. Open NetPoint Access Manager, typically **http://oblix_access_server:port/access/oblix**.

3. Create the following objects in the order they are presented.

Table D-9 Procedures for Creating Objects with Oblix NetPoint Policy Server

Object	To create the object:
Policy Domain	1. In the left pane, click **Create Policy Domain**. 2. Type a descriptive name for the policy domain and click **Save**. 3. Click **Modify**. 4. Enable the policy domain and click **Save**.
HTTP Resource	1. Click the **Resources** tab and then click **Add**. 2. In the Resource drop-down list, choose **HTTP**. 3. In the URL-prefix drop-down list, choose the backslash(/) and type the path to the SSO virtual directory in the adjacent text box, for example, **portal/sso** for typical AquaLogic Interaction deployments. **Note:** Do not enter the full path to the server. 4. Click **Save**.
Policy	1. Click the **Policies** tab and then click **Add**. 2. In the **Name** box, type a descriptive name for the policy, for example `allow access`. 3. In the Resource Type drop-down list, choose **HTTP**. 4. In the Resource Operations group, click **GET** and **POST**. 5. In the Resource group, select the resource you created above (in this example, **/portal/sso**). 6. Leave the other fields blank and click **Save**.

Table D-9 Procedures for Creating Objects with Oblix NetPoint Policy Server

Object	To create the object:
Authorization Rules	The Oblix server parses these rules to send the user name attribute to the portal, enabling the user to log in as a known user with user-defined roles, privileges, views, and so forth. 1. Click the **Authorization Rules** tab, and then click **Add**. 2. In the name box, type a descriptive name for the rule, for example, `forward user name`. 3. Enable the rule and click **Save**. 4. Click **Actions**, then click **Add**. 5. Under Authorization Success, fill in the last line of fields: – For **Type**, enter a descriptive name such as `headerVar`. – For **Name**, enter **UID**. – For **Return Attribute**, enter the name of the attribute used by the authentication source to map to the user name in the user directory. For example, IPlanet LDAP uses the **uid** attribute by default. Other LDAP repositories, and Active Directory, use **cn** or **samaccountname** by default. **Note:** Do not configure an action to return a value. 6. Click **Save**. 7. Click the **Allow Access** sub-tab and then click **Add** to display controls for adding users or groups to allow access to this resource. 8. Add appropriate users and groups. Any user/group that needs access to the portal should be included here. 9. Click **Save**.
Authentication Rule	1. Click the **Policies** tab. 2. Click the name of the policy created above to select it. 3. Click the **Authentication Rule** sub-tab and add a method appropriate to your configuration, for example, Basic over LDAP. **Note:** Basic over LDAP is a rule created during Oblix installation. If it is not available, Oblix was not installed properly. 4. Click **Save**. 5. Click the **Authorization Expression** sub-tab, then click **Add**. 6. In the Select the Authorization Rule box, choose the rule you created above. 7. Click **Save**.

Configuring Oblix WebGate for Apache

Use the version of Oblix WebGate that is compatible with your Oblix suite. For example, if you use Oblix NetPoint 6.5, configure Oblix WebGate 6.5; if you use Oblix COREid 7.0, use Oblix WebGate 7.0.

To set up Oblix WebGate for Apache:

1. On the host computer for the portal, install the version of Apache required by Oblix WebGate:
   - For WebGate 6.5, install Apache 1.3.
   - For WebGate 7.0 or 7.0.4, install Apache 1.3 or Apache 2.0.

   Note: The version of Apache provided by BEA and described in the *Installation and Upgrade Guide for AquaLogic Interaction* cannot be used with the Oblix WebGate. You must download the required Apache version from the Apache Web site.

2. On the host computer for the portal, install Oblix WebGate for Apache. For details, refer to Oblix documentation.

3. On the host computer for the portal, on the Web application server to which the portal application is deployed, modify the Web application server setting to turn off URL rewrites.

   For information about modifying the Web application server to turn off URL rewrites, refer to the Web application server documentation or Knowledge Base article DA_239501, “Configuring Web Application Servers to not Rewrite URLs.”

Configuring Oblix WebGate to Work with Remote Servers

To enable SSO token delegation to a remote tier, for all remote portlet servers that have WebGate installed, turn off IP validation in the WebGate configuration file:

1. Open the `webgatestatic.lst` file in the ../netpoint/webcomponent/access/oblix/apps/webgate/ directory.

2. At the beginning of the file, set **IPValidation** to **false**. The beginning of the file should look something like this:

   ```
   BEGIN:vCompoundList
   DenyOnNotProtected:false
   CachePragmaHeader:no-cache
   CacheControlHeader:no-cache
   IPValidation:false
   ```

3. Save the `webgatestatic.lst` file, and restart the remote server. You do not need to restart the portal.

Integrating With Other Authentication Providers

AquaLogic Interaction does not provide out of the box integrations with other authentication vendors. However you can integrate with other vendors using the following SSO configuration options:

- **BasicSSO**. We recommend that you implement SSO for unsupported authentication servers using the BasicSSO method whenever possible.

  Follow these basic steps for the BasicSSO method:

  a. Install and configure your authentication server as described in your vendor’s documentation.

  b. Follow standard instructions for identifying your authentication source in “Configuring an Authentication Source” on page D-28

  c. Follow special configuration requirements for BasicSSO described in “Modifying the Portal Configuration for BasicSSO” on page D-29.

- **CustomSSO**. We recommend that you implement SSO for unsupported authentication servers using the BasicSSO method whenever possible. Use CustomSSO only if BasicSSO is not sufficient to implement the SSO solution needed for your deployment.

  Follow these basic steps for the CustomSSO method:

  a. Install and configure your authentication server as described in your vendor’s documentation.

  b. Follow standard instructions for identifying your authentication source in “Configuring an Authentication Source” on page D-28

  c. Follow special configuration requirements for CustomSSO described in “Modifying the Portal Configuration for CustomSSO Service” on page D-43.

  For information on developing CustomSSO integration code, see “Developing Custom SSO Objects to Integrate Third-Party SSO Servers with the Portal: DA_217750,” which is available from the Knowledge Base of the Support Center.

Configuring an Authentication Source

Follow the steps described in the following table to set up an authentication source for SSO in the portal.

Table D-10 Procedures to Set Up an Authentication Source for SSO in the Portal

Step	Procedure
Configure the authentication Web service (AWS) and select the authentication source to use for SSO.	For detailed procedures, see "Importing Users and Groups from Authentication Sources" on page 3-12. **Note:** When you configure the remote server and Web service, specify the configuration settings for the SSO partner you set up in "Configuring an SS0 Authentication Provider for Use with the Portal" on page D-2. **Note:** When you configure the authentication source: • On the Main Settings page, specify a string for Category. Make a note of this string, In most cases, this string is the value you must configure for the **DefaultAuthSourcePrefix** setting in the **portalconfig.xml** file. • On the Synchronization page, in the This Authentication Source supports section, select **Synchronization**. • On the Synchronization page, from the Authentication Partners drop-down list, select **SSO Authentication Source**. • Set other synchronization preferences as you like, but do not set up authorization for users.

Configuring the Portal for SSO

This section describes how to modify the portal configuration to enable SSO in the following cases:

- "Modifying the Portal Configuration for BasicSSO" on page D-29
- "Modifying the Portal Configuration for Integration with Netegrity Authentication Servers" on page D-38
- "Modifying the Portal Configuration for Integration with Oblix Authentication Servers" on page D-40
- "Modifying the Portal Configuration for CustomSSO Service" on page D-43

Note: You must configure the appropriate SSO settings described in this section on each portal server for which you want to deploy SSO.

Modifying the Portal Configuration for BasicSSO

AquaLogic Interaction provides a built-in BasicSSO service that enables integration with any authentication server. To configure the BasicSSO service, you configure **portalconfig.xml** so the portal can derive authentication information from the remote authentication source or LDAP configuration source you configured in "Configuring an Authentication Source" on page D-28.

Configuring portalconfig.xml

Configure settings in the **portalconfig.xml** file, as described in the following table and subsequent example.

Table D-11 SSO Settings in portalconfig.xml

Setting	Values

SSOVendor	`<setting name="SSOVendor">`
	`   <value xsi:type="xsd:integer">50</value>`
	`</setting>`

DefaultAuthSourcePrefix	This setting can be omitted if the value of the PrefixHeader setting matches the Authentication Source Category string you configured for your remote or LDAP authentication source.
	Otherwise, set the value of this setting to the Authentication Source Category string.
	For example, if your Authentication Source Category string is **HQ**, set DefaultAuthSourcePrefix to `HQ`, as shown in the following example:
	`<setting name="DefaultAuthSourcePrefix">`
	`   <value xsi:type="xsd:string">HQ</value>`
	`</setting>`
	For details on configuring an authentication source, see "Importing Users and Groups from Authentication Sources" on page 3-12

CookiePath	Set this value to `/`. Specify a different setting only if your SSO authentication server requires a different convention.
	Example:
	`<setting name="CookiePath">`
	`   <value xsi:type="xsd:string">/</value>`
	`</setting>`

Table D-11 SSO Settings in portalconfig.xml

Setting	Values

CookieDomain	Set this value to the fully qualified domain name for which you want the cookie to be forwarded. For example, if you specify `.company.com`, the cookie enables access to all domains that end in **company.com**. If you specify `.sub.company.com`, the cookie enables access only to domains that end in **sub.company.com**.
	The string must start with a period (.) and include a minimum of two periods.
	Example:
	`<setting name="CookieDomain">`
	`   <value xsi:type="xsd:string">.plumtree.com</value>`
	`</setting>`

SSOCookieIsSecure	Set this value to 0 or 1.
	`0` (the default) specifies the connection to the remote server does not require SSL for the cookie to be forwarded.
	`1` specifies SSL is required.
	Example:
	`<setting name="SSOCookieIsSecure">`
	`   <value xsi:type="xsd:integer">0</value>`
	`</setting>`

The following example configuration enables BasicSSO:

```
<setting name="SSOVendor">

       <value xsi:type="xsd:integer">50</value>

</setting>

<setting name="DefaultAuthSourcePrefix">

       <value xsi:type="xsd:string"/>

</setting>

<setting name="CookiePath">

       <value xsi:type="xsd:string">/</value>

</setting>

<setting name="CookieDomain">
```

```
        <value xsi:type="xsd:string">.it.company.com</value>
</setting>
<setting name="SSOCookieIsSecure">
        <value xsi:type="xsd:integer">0</value>
</setting>
```

Next, modify the settings of the **<portal:SSOVendor>** component in the **portalconfig.xml** file, as described in the following table and subsequent example.

Table D-12 <portal:SSOVendor> Component

Setting	Values
NameHeader	Set this value to the name of the user name header your authentication server sends to the portal. The value must be a legal header name. If you want the user name extracted from the Base64-decoded Authentication Header, specify `Authorization`. <NameHeader> requires a valid value if <UseRemoteUser> is not specified or is set to `0`.
PrefixHeader	Set this value to the name of the header containing an authentication source prefix if one is required by remote portlets to authenticate login. If you want the prefix extracted from the Base64-decoded Authentication Header, specify `Authorization`. <PrefixHeader> can be set to an empty string but must be present in **portalconfig.xml**.
PasswordHeader	Set this value to the name of the header containing a password if one is required by remote portlets to authenticate login. If you want the password extracted from the Base64-decoded Authentication Header, specify `Authorization`. <PasswordHeader> can be set to empty string but must be present in **portalconfig.xml**.

Table D-12 <portal:SSOVendor> Component

Setting	Values

Cookie	Set this value to the name of a header containing a cookie if one is required by remote portlets to authenticate login.
	You can configure 0, 1, or many entries using this format:
	`<setting name="Cookie">`
	`<value xsi:type="xsd:string">ssocookie1;ssocookie2</value>`
	`</setting>`
	You configure cookie attributes in the **portalconfig.xml** file. For information, see “Configuring portalconfig.xml” on page D-30.

SecureHeader	Set this value to the name of a header that should not be forwarded to remote portlets. The value you specify is understood as a prefix: headers that start with this value are not forwarded. You can configure 0, 1, or many entries.

UseRemoteUser	To extract the name of the authenticated user from a server variable instead of a user name header, set
	`<setting name="UseRemoteUser">`
	`<value xsi:type="xsd:integer">1</value>`
	`</setting>`.
	For Java implementations, the server variable is REMOTE_USER.
	In .NET, the variable is AUTH_USER.
	By default, this value is set to `0` (false).
	If <NameHeader> is not specified, the value of <UseRemoteUser> must be set to `1`.

The following example configuration summarizes settings for BasicSSO:

```
<component name="portal:SSOVendor"
type="http://www.plumtree.com/config/component/types/portal/ssovendor">

<setting name="NameHeader">

       <value xsi:type="xsd:string">pt_user</value>

</setting>
```

```
<setting name="PrefixHeader">
       <value xsi:type="xsd:string">pt_domain</value>
</setting>
<setting name="PasswordHeader">
       <value xsi:type="xsd:string">pt_password</value>
</setting>
<setting name="Cookie">
       <value xsi:type="xsd:string">PTSSOCookie</value>
</setting>
<setting name="SecureHeader">
       <value xsi:type="xsd:string">authorization</value>
</setting>
<setting name="LogoutURL">
       <value xsi:type="xsd:string"/>
</setting>
<clients>
       <client name="portal"/>
</clients>
</component>
```

To extract the name of the authenticated user from a server variable instead of a user name header:

```
<component name="portal:SSOVendor"
type="http://www.plumtree.com/config/component/types/portal/ssovendor">
<setting name="UseRemoteUser">
       <value xsi:type="xsd:integer">1</value>
</setting>
</component>
```

Configuring Integration with WIA

AquaLogic Interaction provides built-in integration with WIA. Instead of configuring BasicSSO service, follow the procedures in this section to configure SSO integration with WIA.

If you specify the Windows NT domain name as the name for the Authentication Source Category when you set up your authentication source in "Configuring an Authentication Source" on page D-28, you need to configure only settings of **portalconfig.xml**.

Configuring portalconfig.xml

Configure settings in the **portalconfig.xml** file, as described in the following table and subsequent example.

Table D-13 SSO Settings in portalconfig.xml

Setting	Values

SSOVendor	`<setting name="SSOVendor">`
	`   <value xsi:type="xsd:integer">5</value>`
	`</setting>`

DefaultAuthSourcePrefix	This setting can be omitted if you set the Authentication Source Category value to the appropriate Windows NT domain name when you configure your remote authentication source.
	For example, if your Windows domain name is USA and your Authorization Source Category string is USA, this setting can be empty.
	If you set the Authentication Source Category to a value other than the Windows NT domain name, specify that string.
	For example, if your Authentication Source Category string is **HQ**, set DefaultAuthSourcePrefix to `HQ`, as shown in the following example:
	`<setting name="DefaultAuthSourcePrefix">`
	`   <value xsi:type="xsd:string">HQ</value>`
	`</setting>`
	Additionally, set:
	`<setting name="UseDomain">`
	`   <value xsi:type="xsd:integer">0</value>`
	`</setting>`
	For details on configuring an authentication source, see "Importing Users and Groups from Authentication Sources" on page 3-12.

The following example configuration enables integration with WIA:

```
<setting name="SSOVendor">
       <value xsi:type="xsd:integer">5</value>
</setting>
```

```
<setting name="DefaultAuthSourcePrefix">
   <value xsi:type="xsd:string"/>
</setting>
```

Next (if applicable) modify the settings of the <**portal:SSOVendor**> component in the **portalconfig.xml** file, as described in the following table and subsequent example.

Table D-14 <portal:SSOVendor> Settings

Setting	Values

UseDomain	If you set the Authentication Source Category to the Windows NT domain name, you do not need to configure this setting.
	If you set the Authentication Source Category to a value other than the Windows NT domain name, set
	`<setting name="UseDomain">`
	`   <value xsi:type="xsd:integer">0</value>`
	`</setting>.`
	Additionally, configure the <**DefaultAuthSourcePrefix**> setting in **portalconfig.xml**, as described in "Configuring portalconfig.xml" on page D-36.

The following example configuration summarizes settings that can be configured for integration with WIA:

```
<component name="portal:SSOVendor"
type="http://www.plumtree.com/config/component/types/portal/ssovendor">
<setting name="UseDomain">
       value xsi:type="xsd:integer">0</value>
</setting>
</component>
```

Modifying the Portal Configuration for Integration with Netegrity Authentication Servers

AquaLogic Interaction provides built-in integration with Netegrity authentication servers. Instead of configuring BasicSSO service, follow the procedures in this section to configure SSO integration with a Netegrity authentication server.

Configure settings in the **portalconfig.xml** file, as described in the following table and subsequent example.

Table D-15 SSO Settings in portalconfig.xml

Setting	Values

SSOVendor	`<setting name="SSOVendor">`
	`   <value xsi:type="xsd:integer">2</value>`
	`</setting>`

DefaultAuthSourcePrefix	Set this value to a string that matches the value you entered for Authentication Source Category when you configured your authentication source.
	For example, if your Authentication Source Category string is **HQ**, set DefaultAuthSourcePrefix to `HQ`, as shown in the following example:
	`<setting name="DefaultAuthSourcePrefix">`
	`   <value xsi:type="xsd:string">HQ</value>`
	`</setting>`
	For details on configuring an authentication source, see "Importing Users and Groups from Authentication Sources" on page 3-12.

CookiePath	Set this value to `/`. Specify a different setting only if your SSO authentication server requires a different convention.
	Example:
	`<setting name="CookiePath">`
	`   <value xsi:type="xsd:string">/</value>`
	`</setting>`

Table D-15 SSO Settings in portalconfig.xml

Setting	Values

CookieDomain	Set this value to the fully qualified domain name for which you want the cookie to be forwarded. For example, if you specify `.company.com`, the cookie enables access to all domains that end in **company.com**. If you specify `.sub.company.com`, the cookie enables access only to domains that end in **sub.company.com**.
	The string must start with a period (.) and include a minimum of two periods.
	Example:
	`<setting name="CookieDomain">`
	`   <value xsi:type="xsd:string">.company.com</value>`
	`</setting>>`

SSOCookieIsSecure	Set this value to 0 or 1.
	`0` (the default) specifies the connection to the remote server does not require SSL for the cookie to be forwarded.
	`1` specifies SSL is required.
	Example:
	`<setting name="SSOCookieIsSecure">`
	`   <value xsi:type="xsd:integer">0</value>`
	`</setting>`

The following example enables integration with a Netegrity authentication server:

```
<setting name="SSOVendor">

       <value xsi:type="xsd:integer">2</value>

</setting>

<setting name="DefaultAuthSourcePrefix">

       <value xsi:type="xsd:string">HQ</value>

</setting>

<setting name="CookiePath">

       <value xsi:type="xsd:string">/</value>

</setting>

<setting name="CookieDomain">
```

```
        <value xsi:type="xsd:string">.company.com</value>
</setting>
<setting name="SSOCookieIsSecure">
        <value xsi:type="xsd:integer">0</value>
</setting>
```

Modifying the Portal Configuration for Integration with Oblix Authentication Servers

AquaLogic Interaction provides built-in integration with Oblix authentication servers. Instead of configuring BasicSSO service, follow the procedures in this section to configure SSO integration with an Oblix authentication server.

> **Note:** By default, the portal expects the Oblix server to forward the user name header named **uid**. If you configure your Oblix server to forward a user name header with a different name, you must configure your SSO implementation as BasicSSO service. For information about BasicSSO service, see "Modifying the Portal Configuration for BasicSSO" on page D-29.

Configure settings in the **portalconfig.xml** file, as described in the following table and subsequent example.

Table D-16 SSO Settings in portalconfig.xml

Setting	Values

SSOVendor	`<setting name="SSOVendor">`
	`   <value xsi:type="xsd:integer">3</value>`
	`</setting>`

DefaultAuthSourcePrefix	Set this value to a string that matches the value you entered for Authentication Source Category when you configured your authentication source.
	For example, if your Authentication Source Category string is **HQ**, set DefaultAuthSourcePrefix to `HQ`, as shown in the following example:
	`<setting name="DefaultAuthSourcePrefix">`
	`   <value xsi:type="xsd:string">HQ</value>`
	`</setting>`
	For details on configuring an authentication source, see “Importing Users and Groups from Authentication Sources” on page 3-12.

CookiePath	Set this value to `/`. Specify a different setting only if your SSO authentication server requires a different convention.
	Example:
	`<setting name="CookiePath"`
	`   <value xsi:type="xsd:string"/>/</value>`
	`</setting>`

Table D-16 SSO Settings in portalconfig.xml

Setting	Values

CookieDomain	Set this value to the fully qualified domain name for which you want the cookie to be forwarded. For example, if you specify `.company.com`, the cookie enables access to all domains that end in **company.com**. If you specify `.sub.company.com`, the cookie enables access only to domains that end in **sub.company.com**.
	The string must start with a period (.) and include a minimum of two periods.
	Example:
	`<setting name="CookieDomain">`
	`   <value xsi:type="xsd:string">.company.com</value>`
	`</setting>`

SSOCookieIsSecure	Set this value to 0 or 1.
	`0` (the default) specifies the connection to the remote server does not require SSL for the cookie to be forwarded.
	`1` specifies SSL is required.
	Example:
	`<setting name="SSOCookieIsSecure">`
	`   <value xsi:type="xsd:integer">0</value>`
	`</setting>`

The following example enables integration with an Oblix authentication server:

```
<setting name="SSOVendor">

       <value xsi:type="xsd:integer">3</value>

</setting>

<setting name="DefaultAuthSourcePrefix">

       <value xsi:type="xsd:string">HQ</value>

</setting>

<setting name="CookiePath">

       <value xsi:type="xsd:string">/</value>

</setting>

<setting name="CookieDomain">
```

```
        <value xsi:type="xsd:string">.company.com</value>
</setting>
<setting name="SSOCookieIsSecure">
        <value xsi:type="xsd:integer">0</value>
</setting>
```

Modifying the Portal Configuration for CustomSSO Service

AquaLogic Interaction supports custom integration with unsupported authentication servers for which you have developed integration code. For information on developing integration code, see "Developing Custom SSO Objects to Integrate Third-Party SSO Servers with the Portal: DA_217750," which is available from the Knowledge Base in the AquaLogic User Interaction Support Center.

Instead of configuring BasicSSO service, follow the procedures in this section to configure integration with your CustomSSO service.

Configure settings in the **portalconfig.xml** file, as described in the following table and subsequent exampe.l

Table D-17 SSO Settings in portalconfig.xml

Setting	Values

SSOVendor	`<setting name="SSOVendor">`
	`<value xsi:type="xsd:integer">100</value>` (or any value 100 or above)
	`</setting>`

CustomSSOClass	Set this value to the fully qualified class name for the object you developed to integrate an SSO authentication server with the portal.
	Example:
	`<setting name="CustomSSOClass">`
	`<value xsi:type="xsd:string">com.company.portaluiinfrastructure.sso.integrations.SSOTest</value>`
	`</setting>`

Table D-17 SSO Settings in portalconfig.xml

Setting	Values

CustomSSOAssembly	(.NET only)
	Set this value to the name of the assembly that contains the .NET class you specified with the CustomSSOClass setting.
	Example:
	`<setting name="CustomSSOAssembly">`
	`<value xsi:type="xsd:string"/>portaluiinfrastructure</value>`
	`</setting>`

DefaultAuthSourcePrefix	Set this value to a string that matches the value you entered for Authentication Source Category when you configured your authentication source.
	For example, if your Authentication Source Category string is **HQ**, set DefaultAuthSourcePrefix to `HQ`, as shown in the following example:
	`<setting name="DefaultAuthSourcePrefix">`
	`<value xsi:type="xsd:string">HQ</value>`
	`</setting>`
	For details on configuring an authentication source, see "Importing Users and Groups from Authentication Sources" on page 3-12.

CookiePath	Set this value to `/`. Specify a different setting only if your SSO authentication server requires a different convention.
	Example:
	`<setting name="CookiePath"`
	`<value xsi:type="xsd:string"/>/</value>`
	`</setting>`

Table D-17 SSO Settings in portalconfig.xml

Setting	Values

CookieDomain	Set this value to the fully qualified domain name for which you want the cookie to be forwarded. For example, if you specify `.company.com`, the cookie enables access to all domains that end in **company.com**. If you specify `.sub.company.com`, the cookie enables access only to domains that end in **sub.company.com**.
	The string must start with a period (.) and include a minimum of two periods.
	Example:
	`<setting name="CookieDomain"`
	`   <value xsi:type="xsd:string"/>.company.com</value>`
	`</setting>`

SSOCookieIsSecure	Set this value to 0 or 1.
	`0` (the default) specifies the connection to the remote server does not require SSL for the cookie to be forwarded.
	`1` specifies SSL is required.
	Example:
	`<setting name="SSOCookieIsSecure">`
	`   <value xsi:type="xsd:integer">0</value>`
	`</setting>`

The following example uses CustomSSO to enable integration with an unsupported authentication server on a .NET platform:

```
<setting name="SSOVendor">

       <value xsi:type="xsd:integer">100</value>

</setting>

<setting name="CustomSSOClass">

     <value
     xsi:type="xsd:string">com.company.portaluiinfrastructure.sso.integrations
     .SSOTest</value>

</setting>

<setting name="CustomSSOAssembly">

       <value xsi:type="xsd:string">portaluiinfrastructure</value>
```

```
/setting>
<setting name="DefaultAuthSourcePrefix">
       <value xsi:type="xsd:string">HQ</value>
</setting>
<setting name="CookiePath">
       <value xsi:type="xsd:string">/</value>
</setting>
<setting name="CookieDomain">
       <value xsi:type="xsd:string">.company.com</value>
</setting>
<setting name="SSOCookieIsSecure">
       <value xsi:type="xsd:integer">0</value>
</setting>
```

The following example uses CustomSSO to enable integration with an unsupported authentication server on a Java platform:

```
<setting name="SSOVendor">
       <value xsi:type="xsd:integer">100</value>
</setting>
<setting name="CustomSSOClass">
      <value
      xsi:type="xsd:string">com.company.portaluiinfrastructure.sso.integrations
      .SSOTest</value>
</setting>
<setting name="DefaultAuthSourcePrefix">
       <value xsi:type="xsd:string">HQ</value>
</setting>
<setting name="CookiePath">
```

```
       <value xsi:type="xsd:string">/</value>
</setting>
<setting name="CookieDomain">
       <value xsi:type="xsd:string">.company.com</value>
</setting>
<setting name="SSOCookieIsSecure">
       <value xsi:type="xsd:integer">0</value>
</setting>
```

Default Behavior of Search Service

This appendix describes the default behavior of the portal searches. This information is available to users through online help. This appendix includes the following sections:

Types of Search

The portal provides basic and advanced search tools for typical and advanced users, respectively. The fundamental search syntax and behavior are the same in basic and advanced search, but basic search adds automatic broadening, ranking features, and syntax correction. The following table specifies the search type implemented in the search tools available through different areas of the portal.

Table E-1 Portal Areas

Portal Area	Search Type	Description
Banner search	Basic	Searches the following portal objects: banner fields, the Knowledge Directory, portlets, communities, users, Collaboration items, and Publisher items.
Advanced search	Advanced	Allows composition of complex queries on specific document or object properties. Allows searches on date fields as well as text fields. Allows restriction to specific object type. Advanced search also enables searching of all (or any combination of) indexable portal objects, including many which are not searched in banner search, such as content crawlers, jobs, and Web services.
Federated Search	n/a	Federated search allows you to query multiple search Web services and receive collated results. Portal search can be included as one of the search services. The portal search option from this page behaves similarly to basic search, except only documents in the Knowledge Directory are searched. Spell correction, Best Bets, and other customizations made with the Search Results Manager do not apply.
Object selection	Basic	Search functionality enables end users to search for portlets when adding portlets to pages or search for communities when joining communities.
Administrative object search	Basic	Administrators can search the Administrative Objects Directory, optionally filtering by folder and object type. Search for specific kinds of portal objects is also integrated into the creation of various kinds of administrative objects. For instance, when creating a remote content crawler, the administrator is presented with the option of searching the available content source objects.

Table E-1 Portal Areas

Portal Area	Search Type	Description
Filters	Advanced	Allows you to create an advanced search query that documents must match to be allowed into a particular folder in the Knowledge Directory.
Snapshot Query	Advanced	A search query that allows you to specify conditions for searching portal objects and, optionally, display the results in a Content Snapshot Portlet and/or e-mail the results to users. You can limit your search by language, object type, folder, property, and text conditions.

Search Syntax

This section describes the expected behavior of supported search syntax. It includes the following topics:

Operator Modes

The Search Service parses queries to determine which of the following operator modes to use for the query:

- Bag of Words mode: If the query does not include any search operators (+/-, AND, OR, NEAR, etc.), the Search Service parses the query in Bag of Words mode. Each word in the query must be present in all of the search results; the Boolean AND operator is implicit.

- Query Operators mode: If the query includes query operators, the Search Service parses the query in Query Operators mode.

  Query operators *AND*, *OR*, *NOT*, and *NEAR* are spotted without any special marking (for example, `cat AND dog`), but all other operators must be surrounded by angle brackets (for example, <WORD>) to be recognized as having special meaning.

  A query that contains three or more terms and an operator is parsed as if the terms on each side of the operator were quoted phrases. For example:

  ```
  Search Service and Notification
  ```

  This query is parsed as:

  ```
  “Search Service” AND Notification
  ```

  Search operators are localized for the following European languages: English, Danish, Dutch, Finnish, French, German, Italian, Norwegian (Bokmal), Norwegian (Nynorsk), Portuguese, and Spanish. If you put angle brackets around the operators, the English versions are also recognized. For example, in the Spanish locale, the following queries are equivalent: `perro Y gato, perro <AND> gato`, and `perro gato`. However, `perro AND gato` is not equivalent in the Spanish locale, because *AND* is not surrounded by angle brackets.

  Anything enclosed in angle brackets but not recognized as one of the supported operators is ignored.

- Internet Style mode: If the query includes operators common to internet search engines such as AltaVista and Google, the Search Service parses the search in Internet Style mode. All terms preceded by a plus (+) are required. All terms preceded by a minus (-) are excluded. If at least one term is preceded by a +, then any “plain” terms not preceded by a + or - are used to boost ranking of results, but are not required. For example, consider the following query:

  ```
  +dog -cat bird
  ```

This query returns documents that contain *dog* but do not contain *cat*, and ranks documents with both *dog* and *bird* highest. Compare this to a similar query:

```
bird -cat
```

This query returns documents that contain *bird* but do not contain *cat*. Absent any + terms, the plain term *bird* is treated as a required term.

The following table summarizes the behavior of operators.

Table E-2 Search Operators

Operator	Meaning	Alternate
<AND>	Boolean operator that connects terms that must both match the items returned.	AND, '&'(ampersand)
<OR>	Boolean operator that connects terms in which either can match the items returned.	OR, ACCRUE, ANY, '|'(vertical bar), ','(comma)
<NOT>	Items must not match the term.	NOT, AND NOT
<NEAR>	Terms must occur within *N* words of each other, regardless of order.	NEAR, <NEAR/25>
<ORDER>	First term must precede the second term.	
<WORD>	Turns off stemming or alternate case, requiring exact spelling.	
<PHRASE>	Terms must appear as sequential terms in a phrase.	Surround terms in " (double quotes)
<SENTENCE>	Same as <NEAR/10>.	
<PARAGRAPH>	Same as <NEAR/50>.	
+(plus)	Term must appear in the items returned.	
-(minus)	Term must not appear in the items returned.	

Table E-2 Search Operators

Operator	Meaning	Alternate
*(asterisk)	The wildcard specifies that the result must match 0 or more characters at the beginning or end of a word.	

There are certain circumstances in which a user can unintentionally invoke a more advanced search mode by inadvertently using operators. Examples include the following queries:

Table E-3 Example Queries

Query	Equivalent to...
`The young and the restless`	`“the young” <AND> “the restless”`
`File not found`	to `file <AND> <NOT> found`
`Error -217439239`	to `Error <AND> <NOT> 217439239`

In each of these examples, enclosing the query in double quotes yields the desired effect.

Precedence and Parentheses

The Internet Style mode operators '+' and '-' take precedence over the other search operators. For example, `+big dog <order> cat` matches all documents that contain the term `big`, boosting the ranking of any documents that contain any of the three terms `dog`, or `cat`.

Within query operators mode, the operators have the following precedence classes, from greatest to least:

- NEAR, ORDER, PHRASE, SENTENCE, PARAGRAPH
- NOT
- AND
- OR

Parentheses can be used to override operator precedence. The following two queries are equivalent (the parentheses do not effect the semantics of the search).

- `a and b near c or d`
- `(a and (b near c)) or d`

This search matches documents that meet one of two conditions:

- The document contains the term `d`
- The document contains the terms `a`, `b`, and `c`, with `b` and `c` in close proximity.

On the other hand, the parentheses in the following query override the default operator precedence:

```
a and b near (c or d)
```

This search matches documents containing the terms `a` and `b` and either `c` or `d`, where `b` is in close proximity to `c` or `d`.

Punctuation

Punctuation is treated specially. The following rules describe the interpretation of punctuation characters.

- Quotation marks are always interpreted as operators signifying a quoted phrase. It is therefore impossible to search for a quotation mark (there is no escape character, such as a backslash, which would remove the special significance of the quotation marks).
- All other punctuation loses any special operator significance inside of quotation marks. (The same holds for all operators, such as AND.)
- Outside of quotation marks, punctuation either has significance as an operator, or it is ignored. The following punctuation has special operator significance outside of quotation marks:
  - Left and right angle brackets(<>) enclose operators, as in *<NEAR>*
  - Comma (,) is treated as *OR*
  - Ampersand (&) is treated as *AND*
  - Vertical bar (|) is treated as *OR*
  - Plus (+) and minus (-) are interpreted as Internet Style syntax
  - Asterisk (*) is interpreted as a wildcard character

- Punctuation is always split apart from adjoining alpha-numeric characters. For example, an advanced search for `bag-of-words` matches documents containing the three tokens `bag`, `of`, and `words`.
- Underscore is treated as punctuation. This means you must enclose a term containing an underscore in quotes to get an exact match (for example, `“HOST_NAME”` matches `HOST_NAME`, but without the quotes, it also matches `HOST NAME`).

  Symmetrical punctuation tokenization takes place on text stored in the index, so the explosion of a query term such as `bag-of-words` does not prevent the search from matching a document containing the phrase `bag-of-words`.

Case Sensitivity

All searches are case-insensitive, except when the <WORD> operator is used.

Table E-4 Case Sensitivity Examples

Query	Matches
`BEA`	Items containing `BEA`, `bea`, or any other case variant.
`“Search Service”`	Items containing the phrase `Search Service` or any other case variant.
`<WORD> BEA`	Items containing `BEA`, but not `bea` or `Bea`.

Stemming

Word stemming is applied to all individual terms in the search query, except within quoted phrases, or when the <WORD> operator is used. The stemming of query terms means that a query term will match documents containing morphological variants of that term. For example, a search for `dogs AND go` would match a document containing the terms `dog` and `went`. (This example applies to English; stemming employs language-specific information and depends on the user's locale and the language used to index the document.)

Wildcards

The wildcard operator is used to search for prefixes, suffixes, and substrings of indexed terms. Wildcards cannot be used within quoted phrases.

Table E-5 Wildcard Examples

Search Type	Query	Matches
prefix	`cat*`	Finds all documents containing terms that start with `cat`, such as `caterpillar`.
suffix	`*cat`	Finds all documents with terms that end in `cat`, such as `tomcat`.
substring	`*cat*`	Finds all documents with terms that contain `cat`, such as `tomcats`. Mid-string wildcard expressions must contain at least three characters (for example, `*abc*` is legal but `*bc*` is not).

Terms generated by wildcard expansion are not stemmed.

Wildcard expansion is performed internally by replacing each pattern with a limited list of terms that match the pattern before actually executing the query. Very broad wildcard expressions might therefore return a partial list of results.

Quoted Phrases

A quoted phrase in the user search query matches only documents that contain the given sequence of terms. For instance, a search for “`big dog`” will *not* match a document that contains the terms `big` and `dog` if it does not contain the phrase `big dog`. Stemming is not applied to terms within a quoted phrase. Also, wildcards cannot be used within quoted phrases.

Thesaurus Expansion

If thesaurus expansion is enabled, then thesaurus expansion is applied to the individual terms in a basic search. Thesaurus expansion is applied in all three search modes (Internet Style, Query Operators, and Bag of Words). Thesaurus expansion is not applied to quoted phrases. If a term is expanded by a thesaurus entry, then it is not eligible for automatic spelling correction.

Unlike automatic spell correction, which is applied only as a fallback when the non-corrected terms do not match any documents, if thesaurus expansion is enabled, then it is always applied to all individual search terms.

Search Language

Documents and portal objects are indexed with a language setting that determines how word breaking and stemming are applied. When a user issues a search query, word breaking and stemming are applied according to the user account locale settings. Search results are best when the language used for the search matches the language of the documents being searched. However, searches are normally applied to documents in all languages. Cross-language searches do not benefit from localized stemming and word breaking, but can still return useful results.

The advanced search page offers the ability to restrict searches to a particular language.

- The user account search preferences give the option of returning only documents that were indexed using the language of the locale.
- Portal objects can have localized names and descriptions. Basic searches are performed against the default object names and descriptions *and* the names and descriptions of the locale.

When searching portal content via the Search box in the portal banner, the text of the query is processed using the language setting of the user interface. If the portal interface is German, the query is tokenized and stemmed using German language rules, providing optimal search results for documents indexed using German linguistic rules.

If the search collection contains documents in other languages, you can still retrieve them with a query using the appropriate text (assuming the user interface permits entry of the necessary characters). Typing English words into the search box of a portal using a German interface applies German linguistic rules to the query text. Because English stemming is not used, the query is not able to match alternate English word forms; however, English language documents containing the entered words are retrieved.

Although you can enter Asian language text into a European language search box (if a compatible character encoding is used), you should limit the text to a single word or manually separate words with white space to be able to match Asian content in the search collection.

The Advanced Search page provides additional functionality for searching in a multi-language document collection. A pop-up list allows the user to select the language to use for query processing. Linguistic rules for tokenizing and stemming the selected language are used when processing the query text. Among other things, this means that Asian text can be entered without unnecessary white space.

The query operators recognized by Simple Search and Advanced Search are sensitive to the language setting. For example, the AND operator can be specified as “UND” when the query is processed as German. Localized operators are available for the following languages: English, Danish, Dutch, Finnish, French, German, Italian, Norwegian (Bokmal), Norwegian (Nynorsk), Portuguese, and Spanish. All other languages use English operators.

Examples

The descriptions of searches below do not include any of the query expansion or ranking techniques that are employed in basic search. Except where otherwise noted, all matches are case-insensitive.

Table E-6 Query Examples

Query	Expected Behavior
Dog	Searches for documents containing any stem variant of `Dog`.
<WORD> Dog	Searches for documents containing `Dog` as specified exactly with no stemming or lowercasing. This is the only case-sensitive form of search.
Big <PHRASE> Dog	Searches for documents containing the exact phrase `big dog` without stemming.
“Big Dog”	Same as `Big <PHRASE> Dog`.
cat AND dog	Searches for documents containing stem variants of `cat` and `dog`. Equivalent to `cat <AND> dog`.
cat <ALL> dog	Same as `cat AND dog`.
cat OR dog	Searches for documents containing stem variants of `cat` or `dog`.
cat, dog	Same as `cat OR dog`.
cat <ANY> dog	Same as `cat OR dog`.
cat <ACCRUE> dog	Same as `cat OR dog`.
cat NOT dog	Searches for documents containing stem variants of *cat* but not containing stem variants of `dog`.
cat AND NOT dog	Same as `cat NOT dog`.

Table E-6 Query Examples

Query	Expected Behavior
cat NEAR dog	Finds stem variants of `cat` occurring near `dog` (default is within 25 words).
cat NEAR/15 dog	Finds stem variants of `cat` within 15 words of `dog`.
cat <ORDER><NEAR/15> dog	Finds stem variants of *cat* within 15 words before `dog`. Can also use more convenient syntax `cat <ORDER NEAR/15> dog`.
cat <ORDER> dog	Finds stem variants of `cat` anywhere before `dog`.
cat <SENTENCE> dog	Finds stem variants of `cat` within 10 words of `dog`.
cat <PARAGRAPH> dog	Finds stem variants of `cat` within 50 words of `dog`.
cat <XYZ> dog	Finds stem variants of `cat` and `dog`. The unsupported operator XYZ is ignored.
cat*	Finds all documents containing terms that start with `cat`, such as `caterpillar`.
*cat	Finds all documents with terms that end in `cat` such as `tomcat`.
cat	Finds all documents with terms that contain `cat` such as `tomcats`. Mid-string wildcard expressions must contain at least three characters (for example, `*abc*` is legal but `*bc*` is not).
dog *	Finds documents containing stem variants of `dog`. The singleton wildcard is treated as stray punctuation.
dog cat bird	Finds documents containing stem variants of all three terms, `dog`, `cat`, and `bird`. (Bag of Words mode)
big dog AND bird	Finds documents containing the phrase `big dog`, and stem variants of the term `bird`. (Query Operators mode with implicit phrase construction)
dog cat +bird	Finds documents containing stem variants of `bird`. The rank is boosted for documents containing stem variants of `dog` or `cat`. The words `dog` and `cat` are not joined into a phrase in Internet Style mode.
+dog -cat bird	Finds documents that contain stem variants of `dog` but do not contain stem variants of `cat`, and ranks documents with both `dog` and `bird` highest.

Table E-6 Query Examples

Query	Expected Behavior
bird -cat	Finds documents that contain stem variants of `bird` but do not contain stem variants of `cat`.
bag-of-words	Searches for documents containing stem variants of the three terms: `bag`, `of`, and `words`. Punctuation marks are treated as spaces when quotation marks are not present.
“Mr. Jones”	Searches for the phrase `mr. jones`. Punctuation marks are considered part of the search string if they are included within quoted phrases.

Results Ranking

Search results are ranked according to relevance, by default. The following topics in this section describe the factors that determine relevance:

- “Term Frequency” on page E-13
- “Metadata (field) Weighting” on page E-13
- “Phrases and Proximity” on page E-14

Term Frequency

The number of times a query term (or its stemmed and case variant forms) appears in a searchable item has a large influence on the relevance ranking of the item. All other things being equal, items which contain more instances of a query term will rank higher than items containing fewer instances. This is known as *term frequency* based ranking.

Metadata (field) Weighting

Basic searches are performed across several document fields, and some fields are weighted higher than other fields, so that, for instance, a match on an object name ranks higher than a match on an object description. By default, the fields searched are name, description, and full-text content. For information on modifying or adding field weights, see “Modifying the Properties Searched and the Relevance Weight for Properties” on page 4-51.

Phrases and Proximity

In basic search, Bag of Words mode employs special relevancy ranking features which emphasize phrase and proximity matches with the search phrase, even though the user did not employ quotes or proximity operators.

The search phrase terms are used to generate three queries:

- All words joined together as a single phrase
- Stem variants of all words and all quoted phrases *<ORDER><NEAR>* each other
- Stem variants of all words and all quoted phrases joined together with AND.

The three queries combined with the OR operator into a single query, and the relevance ranking are designed to ensure that the results from group #1 always rank above group #2, which rank above group #3.

For example, if you enter `“san francisco” hotels`, the following queries would be generated:

- `“san francisco hotels”`
- `“san francisco” <ORDER><NEAR> hotels`
- `“san francisco” AND hotels`

The search results pages for banner and advanced search allow you to sort the search results by last-modified date, folder, or object type.

Basic Search Behavior

Basic search adds some special features in order to increase the chances that a search will return relevant results. As noted in the previous section, term proximity can boost the relevancy ranking in basic search. Automatic spelling correction is also applied only in basic search.

In basic search, if a user search query causes syntax errors in Internet Style mode or query operators mode, it is automatically retried in Bag of Words mode to be as forgiving as possible of user error. For example, if you enter `dog and`, this query would cause a syntax error in Query Operators mode, because it is missing the right-hand operand to `and`. The query would then be passed to Bag of Words mode, which would attach no special operator significance to `and`, and would therefore retrieve documents containing `dog` and `and`.

Advanced Search Behavior

Advanced search behavior is intended to support complex, precise queries. Therefore it generally does not employ the automatic broadening features of basic search, such as broad cross-field searching or automatic spell correction. Stemming, however, *is* applied in advanced search.

The Text Search portion of advanced search will search across name, description and full text content. Additional property criteria are applied only to the fields specifically selected in each criterion.

User queries that cause syntax errors in Internet Style mode or Query Operators mode will display an error message in the user interface; the search will not fall back to Bag of Words mode.

Common Questions and Answers

This appendix includes the following sections:

- "SSO FAQ" on page F-1
- "Other FAQ" on page F-6

SSO FAQ

This section provides answers to common deployment issues in question and response format.

Question: Why does SSO not work for a particular user?

Answer: Examine the following settings or events to diagnose the cause of this problem:

- In **portalconfig.xml**, the user name prefix must match the value for the Authentication Source Category set in the Authentication Source portal object. Ensure these strings are identical.
- Use AquaLogic Interaction Logging Spy to see if the SSO authentication server is passing the user name to the portal. If you see an error message in red type that indicates "SSO integration returned a null user name. Exiting SSOLoginPage," then there is something wrong with the configuration. Make sure you have configured the authentication server correctly to forward the user name after authentication to the portal. See "Configuring AquaLogic Interaction Logging Spy" on page A-24 for more information.

Question: Why isn't the SSO cookie forwarded to remote servers or portlets?

Answer: Examine the following settings or events to diagnose the cause of this problem:

- In **portalconfig.xml**, ensure the value of the <CookieDomain> element begins with a period.
- In **portalconfig.xml**, ensure the value of the <CookiePath> element is the standard value, `<CookiePath value="/"/>`, or otherwise is a reasonable value.
- In the authentication server, ensure the value of the cookie object enables the cookie to be forwarded.
- Examine the configurations for the authentication server and the portal to ensure fully qualified domain names are specified for all servers.
- If you are unable to diagnose the problem with these methods, use a TCP tracing tool to see the value returned by the SSO provider. The path and domain must match the values for <CookiePath> and <CookieDomain> in **portalconfig.xml**.

Question: Does the portal with SSO support guest user access?

Answer: Guests can access the portal while SSO is enabled. Guest access is controlled by the *AllowGuestAccess* setting in the *Authentication* section of **portalconfig.xml**. When guest access is disabled, users can browse the portal without logging in. When users click **Login** in the portal banner or when they attempt to visit a page for which the guest user does not have access, the portal redirects them to the SSO login page, and they are prompted by the SSO product for their login credentials.

If users already have an SSO cookie from another application, they still browse the portal as the guest user until they click **Login**. At which point, they are logged in without entering their user name and password.

Guest access can be enabled or disabled independently from SSO. If guest access and SSO are both disabled, users have to log in before accessing any part of the portal.

Question: How can a user change login credentials, for example, to become Administrator?

Answer: If users need to log in as Administrator or other portal users from within their SSO session, they can click **Log Off** in the portal banner. This logs them out of the portal and takes them to the portal login page, as if SSO were disabled. From this page they can log in as a non-SSO user or they can browse the portal as guest. When they want to log back in as an SSO user, they can click **Login** in the portal banner and they are automatically logged in to the portal.

Question: Why can’t I access the portal through SSOLogin.aspx or the SSOServlet?

Answer: The first time you access the portal after you deploy SSO, you must access the portal from the main portal URL: **http://<servername>/portal/Server.pt**.

If you try to access the portal through **/portal/sso/SSOLogin.aspx** (.NET) or **/portal/SSOServlet** (Java), your request fails and the following error appears in AquaLogic Interaction Logging Spy trace logs: “The SSO Login Page was unable to retrieve the request URL from the session. Will use a relative redirect to return to the main page.”

Question: If I configure my SSO authentication server to protect the Image Service virtual directory, users encounter JavaScript errors and portal menus fail to load. Why?

Answer: The portal and other AquaLogic User Interaction products, such as Collaboration and Publisher, periodically send HTTP requests to the Image Service to check the version of the JavaScript components stored on the Image Service. These requests are not associated with a particular user's session and do not send an SSO cookie or other credentials. If the Image Service is protected by your SSO solution, the request from the portal is blocked from checking the JavaScript versions. As a result, the portal is unable to load the proper JavaScript files and end users encounter JavaScript errors and possibly other errant behavior. To resolve this problem, do not configure your SSO authentication server to protect the Image Service, but only the portal. You do not need to protect the Image Service as it contains only static public content that ships with every portal installation. No data specific to users or to your organization is ever stored on the Image Service.

Question: How can I debug my SSO deployment?

Answer: The portal provides built-in trace statements that are useful for debugging SSO integration. For example, when a user attempts to log in using SSO, the contents of all headers are traced. To enable this tracing, turn on all tracing for the **Portal UI - Infrastructure** component. See “Configuring AquaLogic Interaction Logging Spy” on page A-24 for more information.

Question: How do I configure reverse proxy with my SSO deployment?

Answer: AquaLogic Interaction has verified deployments with reverse proxy using Apache HTTP server, Oblix Netpoint Access Server (versions 6.1.1 or 6.5) with an Apache WebGate, and a Java-based portal.

Follow these basic steps to complete this configuration:

1. Install Oblix NetPoint Access Server, including NetPoint Access Manager, NetPoint COREid, and Oblix Apache WebGate. WebGate must be installed on the same server as the Apace HTTP server. For detailed instructions, refer to Oblix documentation.
2. Use Oblix Access Manager to create the portal protection policy. For detailed instructions, refer to Oblix documentation.
3. Configure Oblix NetPoint Access Server. For detailed instructions, see “Configuring an Oblix Authentication Provider” on page D-20.
4. Configure the Apache HTTP server for reverse proxy. For detailed instructions, see the procedures that follow these steps.
5. Configure the portal for SSO. For detailed instructions, see “Configuring the Portal for SSO” on page D-29.
6. Configure the portal application server for reverse proxy. For detailed instructions, see the procedures following these steps.
7. Restart services to apply configuration modifications.

To configure the Apache HTTP server for reverse proxy:

1. Install the version of the Apache HTTP server recommended by the Oblix Installation Guide. For Netpoint 6.5, Oblix recommends the latest version of the Apache, v1.3 line. The configuration described in this example has been tested with version v1.3.29.
2. Turn on the proxy module inside of the Apache configuration). To do so, edit **<apache_install_dir>/conf/httpd.conf** to uncomment the lines titled **LoadModule proxy_module modules/mod_proxy.so** and **AddModule mod_proxy.c**. (to uncomment a line, remove the pound symbol (#) at the beginning of the line).
3. Configure Apache to act as a reverse proxy for your portal. To do so, add lines similar to the following example at the end of **httpd.conf**:
4. ProxyRequests Off
5. ProxyPass /portal http://your_portal_server.domain.com:7001/portal

6. ProxyPassReverse /portal http://your_portal_server.domain.com:7001/portal
7. This example configuration redirects requests from the Apache Web server (http://proxy_server.domain.com:80/portal/xyz) to the portal application server (http://your_portal_server.domain.com:7001/portal/xyz).
8. You must specify the fully qualified domain name here and for all other times you type in the server names.
9. For more information on Apache reverse proxy, see http://httpd.apache.org/docs/mod/mod_proxy.html.
10. Start or reboot the Apache HTTP server.

To configure the Java application server for reverse proxy:

1. Open ***PT_HOME*/ptportal/6.1/settings/config/portalconfig.xml** for editing.
2. Configure the <URLMapping> element so that it is similar to the following example:

   ```
   <URLFromRequest0 value="*"/>

   <ApplicationURL0
   value="http://proxy_server.domain.com/portal/server.pt"/>

   <SecureApplicationURL0 value="*"/>
   ```

   Replace `proxy_server.domain.com` with the fully qualified domain name for the Apache HTTP server.

3. Configure the <SSOVirtualDirectoryPath> element so that it is similar to the following example:

   ```
   <SSOVirtualDirectoryPath
   value="http://proxy_server.domain.com/portal/"/>
   ```

   Replace `proxy_server.domain.com` with the fully qualified domain name for the Apache HTTP server.

4. Reboot the application server.

Other FAQ

Question: I get a timeout error when I edit a folder in the Knowledge Directory, change the Security page, and try to apply the security changes to all child objects of the folder.

Answer: This problem may occur if you have many levels of nested subfolders with a large number of child objects (other folders or documents). To work around this problem if it occurs:

1. To find out which folders are updated with the security changes, select all first level subfolders then select the security icon.
2. Scroll through the list to see at which folder the security changes stopped being applied.
3. Work with the remaining folders and apply the needed security changes: open each first level folder separately, set the security, save, and select **Yes** to apply the security changes to all the child objects for that folder.
4. Repeat this process for all remaining first level subfolders that did not get the security applied to them successfully due to the error.

Index