BEA AquaLogic® Interaction

Administrator Guide

Version 6.5
Revised: March 27, 2008

Contents

3. About User Interface Customization

4. Managing Portal Users and Groups

5. Managing Portal Content

6. Managing Administrative Objects and Portal Utilities

7. About Extending Portal Services with Portlets

8. About Providing Content and Services to a Group of Users through Communities

9. Managing Search

10. Automating Administrative Tasks

11. Migrating Portal Objects

A. Configuring Advanced Properties in the Portal Configuration Files

B. Using the Counter Monitoring System

C. Localizing Your Portal

D. Deploying Single Sign-On

E. Default Behavior of Search Service

CHAPTER 1

Introduction to the Administrator Guide for AquaLogic Interaction

The *Administrator Guide for AquaLogic Interaction* describes how to perform initial tasks required to prepare your portal for use and how to perform ongoing portal management tasks.

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.

BEA Documentation and Resources

The following documentation and resources are available from BEA.

Table 1: Documentation

Resource	Description
Installation Guide for AquaLogic Interaction 6.5 on Windows	This guide describes the prerequisites (such as required software) and procedures for installing AquaLogic Interaction 6.5 on Windows machines. It is available on *edocs.bea.com/en/alui/ali/docs65*.
Installation Guide for AquaLogic Interaction 6.5 on Unix and Linux	This guide describes the prerequisites (such as required software) and procedures for installing AquaLogic 6.5 on Unix and Linux machines. It is available on *edocs.bea.com/en/alui/ali/docs65*.
Upgrade Guide for AquaLogic Interaction 6.1 to 6.5 on Windows	This guide describes the prerequisites (such as required software) and procedures for upgrading AquaLogic Interaction from version 6.1 to version 6.5 on Windows machines. It is available on *edocs.bea.com/en/alui/ali/docs65*.
Upgrade Guide for AquaLogic Interaction 6.1 to 6.5 on Unix and Linux	This guide describes the prerequisites (such as required software) and procedures for upgrading AquaLogic Interaction from version 6.1 to version 6.5 on Unix and Linux machines. It is available on *edocs.bea.com/en/alui/ali/docs65*.
Upgrade Guide for AquaLogic Interaction 6.0 to 6.5 on Windows	This guide describes the prerequisites (such as required software) and procedures for upgrading AquaLogic Interaction from version 6.0 to version 6.5 on Windows machines. It is available on *edocs.bea.com/en/alui/ali/docs65*.
Upgrade Guide for AquaLogic Interaction 6.0 to 6.5 on Unix and Linux	This guide describes the prerequisites (such as required software) and procedures for upgrading AquaLogic Interaction from version 6.0 to version 6.5 on Unix and Linux machines. It is available on *edocs.bea.com/en/alui/ali/docs65*.
Release Notes	The release notes provide information about new features, issues addressed, and known issues in the release. They are available on *edocs.bea.com/en/alui/ali/docs65* and on any physical media provided for delivering the application.

Resource	Description
Online Help	The online help is written for all levels of AquaLogic Interaction users. It describes the user interface and gives detailed instructions for completing tasks in AquaLogic Interaction. To access online help, click the help icon.
Deployment Guide	This guide 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 2: Other 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 (ALUI) and AquaLogic Business Process Management (ALBPM) Support Center	The ALUI and ALBPM Support Center is a comprehensive repository for technical information on ALUI and ALBPM 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 ALUI and ALBPM-related needs. The Support Center encompasses the following communities: **Technical Support** Submit online service requests, check the status of your requests, search the knowledge base, access documentation, and download maintenance packs and hotfixes. **User Group** Participate in user groups; view webinars, presentations, the CustomerConnection newsletter, and the Upcoming Events calendar. **Product Center**

Resource	Description
	Download product updates, maintenance packs, and patches; view the Product Interoperability matrix (supported third-party products and interoperability between products). **Developer Center** Download developer tools, view code samples, access technical articles, and participate in discussions. **Education Services** Review the available education options, then choose courses by role and delivery method (Live Studio, Public Classroom Training, Remote Classroom, Private Training, or Self-Paced eLearning). **Profile Center** Manage your implementation details, local user accounts, subscriptions, and more. If you do not see the Support Center when you log in to *one.bea.com/support*, contact *ALUISupport@bea.com* or *ALBPMSupport@bea.com*for the appropriate access privileges.
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* or *ALBPMSupport@bea.com* Phone Numbers: USA, Canada +1 866.262.7586 or +1 415.263.1696 EMEA +44 1494 559127 Asia Pacific +61 2.9931.7822 Australia/NZ +61 2.9923.4030 Singapore +1 800.1811.202

Completing Portal Initial Set-Up Tasks

When you first deploy your portal, you need to perform several set-up tasks before your portal is ready for your users.

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

After you have completed your initial portal deployment, you can 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. You might optionally configure localization, single sign-on (SSO), and advanced configuration file settings.

Overview of AquaLogic Interaction

This chapter provides an overview of the portal and the administrative tasks you perform to manage portal users and documents.

Portal Installer Components

The following table describes the components available in the portal installer. These portal components provide the functionality described in the *Administrator Guide for AquaLogic Interaction* . For information on installing these components, refer to the *Installation Guide for AquaLogic Interaction* .

Component	Description
Administrative Portal	The administrative portal handles portal setup, configuration, and content. It enables administrative functions, such as creating and managing portlets and other web services.
Portal	The portal serves end user portal pages and content. It 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 advanced portal configuration, see *AquaLogic Interaction Configuration Manager* on page 319.

Component	Description
Portal Database (scripts)	The scripts used to configure the database are included in the portal installer. The 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.
Automation Service	The Automation Service runs jobs and other automated portal tasks. You run jobs to 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 configuring Automation Service jobs, see *Automating Administrative Tasks* on page 295.
API Service	The API Service provides access to the SOAP API.
Image Service	The Image Service serves static content used or created by portal components. It serves images and other static content for use by the AquaLogic User Interaction system. 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.
Search Service	The Search Service returns content that is indexed in the AquaLogic User Interaction system from the portal, Collaboration, and Publisher. The indexed content includes documents, portlets, communities, and users as well as many other AquaLogic User Interaction objects. For information on advanced Search Service configuration, see *AquaLogic Interaction Configuration Manager* on page 319.
Document Repository Service	The Document Repository Service stores content uploaded into the portal, Collaboration, or Publisher.
Content Upload Service	The Content Upload Service lets you add files to the portal's Knowledge Directory by uploading them to the Document Repository Service, rather than leaving them in their original locations. This is useful if users need

Component	Description
	to access documents located in an internal network from outside your network.
ALUI Directory Service	The ALUI Directory Service enables AquaLogic Interaction to act as an LDAP server, exposing the user, group, and profile data in the portal database through an LDAP interface. This enables other ALUI products (and other third-party applications) to authenticate users against the portal database.
Remote Portlet Service	The Remote Portlet Service includes the following components: • RSS Reader Portlet The RSS Reader Portlet enables users to specify an RSS or ATOM feed to display on a My Page or community page. • Activity Service The Activity Service includes the User Status portlet, which lets users post their current status; the User Activities portlet, which displays a user's status history and any other recent activities that are submitted by other applications; and a REST-based API for submitting activities into a user's activity stream. **Note:** If you use the REST-based API to submit other activities into the activity stream, those activities will also be displayed in the User Activities portlet.
Notification Service	The Notification Service enables the portal to send e-mail notifications to users upon specified events. There are no portal events that trigger notifications, but other AquaLogic User Interaction events do trigger notifications. For example, AquaLogic Interaction Collaboration can be configured to send notifications to users when documents are uploaded.

Additional AquaLogic User Interaction Components

The following table describes components that provide additional functionality for the portal. For more information on these components or to download the components, visit the AquaLogic User

Interaction (ALUI) and AquaLogic Business Process Management (ALBPM) Support Center (one.bea.com/support).

Component	Description
Integration Services	Integration Services enable integration with third-party applications/repositories. • Identity Services Identity Services let you import users, groups, and user profile information from third-party user repositories into the portal. Identity Services also enable the portal to authenticate users through the third-party user repositories. Identity Services are available for Microsoft's Active Directory (AD) and LDAP (Lightweight Directory Protocol). • Content Services Content Services scan third-party systems/applications for new content, categorizing links to this content in the organized, searchable structure of the portal's Knowledge Directory. Users can then access this content through the portal user interface. Content Services are available for Windows File Systems, Documentum, Lotus Notes, and Microsoft Exchange. • Portlet Suites Portlet Suites are collections of portlets that provide users access to commonly used functions in third-party applications within the portal user interface. Portlet Suites are available for IMAP, Lotus Notes, and Microsoft Exchange. • Portlet Frameworks Portlet Frameworks let users create custom portlets that access information from third-party applications. The portlets can then be added to My Pages or community pages to provide users access to the third-party information. A Portlet Framework is available for Microsoft Excel.

Component	Description
Activity Services	Activity Services extend portal functionality to enable analysis, collaboration, publishing, and simple portlet creation. • Analytics Analytics delivers comprehensive reporting on activity and content usage within portals and composite applications, allowing you to know and meet user information needs. • Collaboration Collaboration helps people to work together via the web, supporting tasks, projects, communities, calendars, discussions, and document sharing with version control. • Publisher Publisher allows publication & management of web content for portals and web applications, with forms-based publishing, branding, templates, workflow, approvals, and content expiration. • Sharepoint Console Sharepoint Console imports, indexes, and returns Microsoft Windows Sharepoint Services resources via AquaLogic Interaction Search. • Studio Studio lets portal managers create portlets, such as telephone lists, work order processes, calendars and surveys, without any coding.
Enterprise Social Computing Products	Enterprise Social Computing Products provide tools that enable users to freely contribute and actively work together. • AquaLogic Ensemble Ensemble is an enterprise system that lets developers create reusable widgets for mashup applications and allows IT to easily manage a diverse set of web resources. • AquaLogic Pages Pages is a simple and powerful system that lets end-user participants create web pages and applications to dramatically improve their productivity for everyday business activities.

Component	Description
	• AquaLogic Pathways Pathways provides social search and expertise discovery by combining social bookmarking and tagging with usage-based enterprise search to help users organize, discover and share information and expertise through social networks.
Developer Tools	The following developer tools help you rapidly build applications through AquaLogic Interaction: • AquaLogic Interaction Development Kit (IDK) The IDK enables Java and .NET developers to rapidly build, deliver, and enhance user-centric composite applications through AquaLogic Interaction. It provides interfaces for Integration Web Services—authentication, profile, crawler, and search—that integrate enterprise systems into AquaLogic Interaction. It provides SOAP-based remote APIs to expose portal, search, and Collaboration features. In addition, the IDK has a portlet API to assist in pagelet, proxied application, and portlet development. • .NET Portlet Toolkit The .NET Portlet Toolkit is used to speed the development of ASP.NET portlets for use natively with AquaLogic Interaction. This product includes the .NET Portlet API and the .NET Web Control Consumer. The .NET Web Control Consumer enables you to create portlets using Microsoft .NET Web Controls. • JSR-168 Container The JSR-168 Container lets you deploy portlets in AquaLogic Interaction that conform to the JSR-168 portlet standard. The JSR-168 Container is included in the AquaLogic Interaction release package. • Logging Utilities The Logging Utilities consist of a set of three tools to receive, display, and store logging messages sent from AquaLogic User Interaction products. • Service Station

Component	Description
	Service Station enables developers of Integration Web Services to easily test and validate custom authentication, crawler, profile, and search web services. Service Station includes two methods for defining tests: an easy to use GUI, which enables users to create, modify, execute, and monitor tests without any programming; and a JUnit based framework which can be run and monitored from a Java IDE (or via Java-based build tools such as Ant). • WSRP Consumer and Producer If you want to deploy portlets in AquaLogic Interaction that conform to the WSRP portlet standard, you need the WSRP Consumer and the WSRP Producer. The WSRP Consumer is available as a stand-alone product and is included in the AquaLogic Interaction release package. The WSRP Producer is available as part of the .NET Application Accelerator. • UI Customization Installer The User Interface Customization Installer (UICI) automates most of the steps required to set up a development environment for the portal. Minimal Eclipse and Ant configuration is still necessary. The UICI is for advanced applications, most UI customizations can be accomplished through adaptive functionality available with AquaLogic Interaction.

Overview of the Browsing User Interface and Features

Navigating the Portal

The portal includes some basic functionality in the portal banner and menus to access the different areas of the portal. The areas you see depend on the portal configuration, whether you are logged in, and your portal access.

- To determine whether you are logged in correctly, look at the portal greeting (at the top of the portal banner).

By default your greeting is `Welcome, user name` where *user name* is the name of the user by which you are logged in. To change your greeting, click **My Account**, then, on the My Account page, click **Display Options**.

- To access the Administrative Objects Directory, where you can create and manage portal objects and access portal utilities, click **Administration**.

  Note: You see **Administration** only if you have the Access Administration activity right.

- To edit your user profile, personalize your display options, set your locale settings, set your search preferences, change your portal password, and manage user names and passwords for external accounts, as well as view your user profile, click **My Account**.
- To display help for the page you are viewing, click **Help**.
- To log in to the portal or log off the portal, click **Log In** or **Log Off**.
- To search for documents or objects in your portal, type your search string in the box in the portal banner and click **Search**.
- To go directly to the result that your portal administrator has set as the top best bet for a term, without first seeing all the search results, type the top best bet operator (>) followed by your search string and click **Search**, or type your search string and click .

  Note: If no top best bet has been set for the term, the regular search results appear.

  Note: The top best bet button ( ) is available only if enabled by your portal developer.

- To search for documents or objects using metadata properties and location, click **Advanced Search**.
- To search other content, portals, and web search engines, click **Federated Search**.

  Note: The **Federated Search** button is available only if enabled by your portal developer.

- To view your user profile, in the **My Profile** menu, click **View User Profile**.
- To view one of your personalized pages, in the **My Pages** menu, click the page you want to view.

  My Pages are your personalized view of the portal. You choose the applications, tools, and services (in the form of portlets) that you want to display on each My Page. For example, you might create a My Page that includes a search tool for all the employees in your company and a portlet that displays the most recent news about your company.

- To view and manage the communities to which you belong, select an option from the **My Communities** menu:

- To view a community, click the community name.
- To join a community, click **Join Communities**, select the communities you want to join, and click **Finish**.
- To unsubscribe from a community, click **Unsubscribe Communities**, select the communities from which you want to unsubscribe, and click **Finish**.

Communities are sites within a portal designed for a specific audience or task, such as collaborative projects. You might have 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 might have a separate community containing project milestones, regulatory compliance requirements, and technical specifications.

- To browse documents in the portal, click **Directory**.

  The Knowledge Directory is similar to a file system tree in that documents are organized in folders and subfolders. A folder can contain documents uploaded by users or imported by content crawlers, as well as links to people, portlets, and communities. If your administrator has given you permission, you might also be allowed to add documents to the Knowledge Directory, or submit yourself as an expert on a particular topic.

Portal Banner Features

There are several features available at the top of your portal that provide access to some basic portal functions (such as help and search).

Feature	Description
Greeting	Lets you know that you are logged in as the correct user. By default your greeting is `Welcome, user name` where *user name* is the name of the user by which you are logged in.
 **Administration**	Provides access to the Administrative Objects Directory, where you can create and manage portal objects and access portal utilities. **Note:**  **Administration** appears only if you have the Access Administration activity right.

Feature	Description
My Account	Lets you edit your user profile, display options, locale settings, and search preferences, as well as view your user profile and change your password. **Note:** You must have the Edit Own Profile activity right to be able to edit your user profile.
Help	Opens the help associated with the displayed page.
Log Off	Logs you out of your portal.
Search box and button	Let you search for documents, document folders, communities, community pages, portlets, and users in your portal.
(top best bet icon)	Takes you directly to the top best bet result for the term you enter in the search box. If no top best bet has been set for the term, you will see the regular search results. **Note:** appears next to the search box only if enabled by portal developers.
Advanced Search	Lets you perform an advanced search by searching the portal for text or specific document properties.
Federated Search	Lets you perform a federated search (if your portal administrator has set up any federated search resources) to search other content, portals, and web search engines. **Note:** **Federated Search** appears only if enabled by portal developers.

Portal Menus

There are several menus available in the portal that provide access to information in your portal (such as communities and documents). Your portal administrator or portal developer controls which menus appear, including custom menus.

Menu	Description
My Profile	Lets you view your user profile. User profiles provide information about users, such as address and position.

Menu	Description
My Pages	Provides access to your My Pages. My Pages are your personalized view of the portal. You choose the applications, tools, and services (in the form of portlets) that you want to display on each My Page. For example, you might create a My Page that includes a search tool for all the employees in your company and a portlet that displays the most recent news about your company.
My Communities	Lets you view and manage the communities to which you belong. Communities are sites within a portal designed for a specific audience or task, such as collaborative projects. You might have 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 might have a separate community containing project milestones, regulatory compliance requirements, and technical specifications.
Directory	Provides access to the Knowledge Directory. The Knowledge Directory is similar to a file system tree in that documents are organized in folders and subfolders. A folder can contain documents uploaded by users or imported by content crawlers, as well as links to people, portlets, and communities. If your administrator has given you permission, you might also be allowed to add documents to the Knowledge Directory, or submit yourself as an expert on a particular topic.

Portal Interface Types

Users can change their portal display to accommodate assistive technologies or slow internet connections.

Interface Type	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

Interface Type	Description
	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.

Editing Locale Settings

Users can change their portal display to accommodate their time zone and locale.

The locale determines:

- The language displayed in the portal interface (portlet names and content display in the language you choose only if those portlets support your chosen language).
- The format for portal entries (including search requests). For example, if you choose 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.

Note: Only the portal interface and localized objects display in the language you choose. Your personal greeting does not change if you change your locale.

1. In the portal banner, click **My Account ➤ Edit Locale Settings** .

2. In the **Your time zone** drop-down list, choose your time zone.
3. In the **Your locale** drop-down list, choose the language in which you want your portal to display content.

About My Pages

My Pages are your personalized view of the portal. You choose the applications, tools, and services (in the form of portlets) that you want to display on each My Page. For example, you might create a My Page that includes a search tool for all the employees in your company and a portlet that displays the most recent news about your company.

Personalizing Your View of the Portal with My Pages

My Pages are your personalized view of the portal. You choose the applications, tools, and services (in the form of portlets) that you want to display on each My Page. For example, you might create a My Page that includes a search tool for all the employees in your company and a portlet that displays the most recent news about your company.

- To view a My Page, in the **My Pages** menu, click the page you want to view.
- To create a My Page, click **Create Page**.

  You can have as many as six My Pages.
- To rename the page or add portlets to the page, click **Edit Page**.
  - If you are editing a My Page that uses an adaptive page layout, the Flyout Editor opens. See *Editing a Page in the Flyout Editor* on page 258.
  - If you are editing a My Page that uses a legacy user interface, the Page Editor opens. See *Editing a Page in the Page Editor* on page 38.
- To change the column layout for the page:

  a) Click **Edit Page**.
  - If you are editing a My Page that uses an adaptive page layout, the Flyout Editor opens. Continue with Step 2.
  - If you are editing a My Page that uses a legacy user interface, the Page Editor opens. Skip Step 2.

  b) Click **Go to Advanced Editor** to open the Page Editor.

c) Click **Select Page Layout**, then, in the Select Page Layout dialog box, select the layout you want and click **Finish**.

Note: Narrow portlets can display in either narrow or wide columns, but wide portlets can display only in wide columns.

- To delete the page, click **Delete This Page**.
- Manage portlets using the buttons in the portlet title bar:

  Note: Some portlets do not include title bars or do not include all the buttons described here.

  - To refresh the content in a portlet, click .
  - To view the help for a portlet, click .
  - To edit preferences for a portlet, click .
  - To collapse a portlet so that only the title bar appears on the page, click .

    To expand the portlet, click .
  - To remove a portlet from the page, click .

Personalizing Your View of the Portal with My Pages

My Pages are your personalized view of the portal. You choose the applications, tools, and services (in the form of portlets) that you want to display on each My Page. For example, you might create a My Page that includes a search tool for all the employees in your company and a portlet that displays the most recent news about your company.

- To view a My Page, in the **My Pages** menu, click the page you want to view.
- To create a My Page, click **Create Page**.

  You can have as many as six My Pages.
- To rename the page or add portlets to the page, click **Edit Page**.
  - If you are editing a My Page that uses an adaptive page layout, the Flyout Editor opens. See *Editing a Page in the Flyout Editor* on page 258.
  - If you are editing a My Page that uses a legacy user interface, the Page Editor opens. See *Editing a Page in the Page Editor* on page 38.
- To change the column layout for the page:

  a) Click **Edit Page**.

- If you are editing a My Page that uses an adaptive page layout, the Flyout Editor opens. Continue with Step 2.
- If you are editing a My Page that uses a legacy user interface, the Page Editor opens. Skip Step 2.

b) Click **Go to Advanced Editor** to open the Page Editor.

c) Click **Select Page Layout**, then, in the Select Page Layout dialog box, select the layout you want and click **Finish**.

Note: Narrow portlets can display in either narrow or wide columns, but wide portlets can display only in wide columns.

- To delete the page, click **Delete This Page**.
- Manage portlets using the buttons in the portlet title bar:

  Note: Some portlets do not include title bars or do not include all the buttons described here.

  - To refresh the content in a portlet, click .
  - To view the help for a portlet, click .
  - To edit preferences for a portlet, click .
  - To collapse a portlet so that only the title bar appears on the page, click .

    To expand the portlet, click .
  - To remove a portlet from the page, click .

Creating a My Page with One Click

You can create a new My Page with one click.

1. Display one of your My Pages.
2. Click **Create Page**.
   The new page is created.
3. To add portlets to the page, click **Edit Page**.

Viewing and Managing Communities and Community Membership

Communities are sites within a portal designed for a specific audience or task, such as collaborative projects. You might have 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 might have a separate community containing project milestones, regulatory compliance requirements, and technical specifications.

- To view the communities to which you belong, open the **My Communities** menu.
- To display a community, in the **My Communities** menu, click the community name.
- To join a community, in the **My Communities** menu, click **Join Communities**.

  Tip: You can also join the community you are viewing by clicking **Join this community**.

- To unsubscribe from a community, in the My Communities menu, click **Unsubscribe Communities**.

  Tip: You can also unsubscribe from the community you are viewing by clicking **Unsubscribe from this community**.

- To display a page in the community, in the community title bar, click the name of the page.
- To view the Community Knowledge Directory, in the community title bar, click **Community Knowledge Directory**.

  This link is available only if the community administrator has enabled the Community Knowledge Directory. The Community Knowledge Directory displays the members of the community, any subcommunities of the community, and any other folders and contents the community administrator added.

- To view the subcommunities to which you belong, open the **Subcommunities** menu.

  This menu appears only if this community contains any subcommunities. Subcommunities are separately secured subsections of a community. For example, you might have a Marketing community that contains the Advertising Subcommunity. The Advertising Subcommunity could have distinct owners; or only a subset of the Marketing community might be entitled to see the Advertising Subcommunity.

- To view a subcommunity, in the **Subcommunities** menu, click the subcommunity name.
- To create a new, blank page in the community you are viewing, click **Create Page**.
- To edit the community page you are viewing (rename the page, add, remove, or rearrange portlets), click **Edit Page**.

- If you are viewing an adaptive page layout, the Flyout Editor opens. See *Editing a Page in the Flyout Editor* on page 258.
- If you are viewing a legacy user interface, the Page Editor opens. See *Editing a Page in the Page Editor* on page 38.

Editing a Page in the Flyout Editor

You can rename a page, add portlets, recommend portlets, and reposition portlets while viewing the page.

Click **Edit Page**. The Flyout Editor appears, enabling you to perform the following actions:

Note: The Flyout Editor appears only if you are viewing an adaptive page layout (not a legacy page layout).

- To rename the page, in the **Change Page Name** box, type the new name.
- To add a portlet to the page, under the portlet name, click **Add to Page**.
  A placeholder for the portlet is added to the page below the Flyout Editor.
- To remove a portlet, under the portlet name, click **Remove**, or click in the portlet's title bar.
- To see what a portlet looks like, under the portlet name, click **Preview**.

  From the Preview Portlet page you can perform the following actions:

  - To add the portlet to your page and close the preview, click **Add this portlet**.
  - To view a description of the portlet, click **View Description**. When you are finished, click **Close**.
  - To return to the list of portlets without adding the portlet to your page, click **Close**.

- To view a list of the portlets in a portlet bundle, under the bundle name, click **Open**.
- To add all the portlets from a bundle, under the bundle name, click **Add**.
  A placeholder for each portlet is added to the page below the Flyout Editor.

  Note: You can remove a portlet added as part of a bundle by clicking in the portlet's title bar.

- To recommend a portlet to other users:
  a) Under the portlet name, click **Invite**.
  b) In the invitation dialog box, copy the text, and click **Close**.
  c) In your e-mail application, paste the text into an e-mail message and send it.

When other portal users click the URL in your e-mail, they are taken to the portlet preview and given the option to add the portlet to one of their My Pages. Users that do not have permission to see the portlet receive an error message.

- To search for portlets and portlet bundles, in the **Search for Portlets** box, type the text you want to search for and click **Search**.

  For searching tips, see *Using Text Search Rules* on page 364.

  To remove your search criteria, click **Search** again.
- To change the sort order of portlets, in the **Sort By** drop-down list, select an option: Item Name Ascending, Item Name Descending, Date Modified Ascending, Date Modified Descending.
- To page through the list of portlets, click **<< Previous**, **Next >>**, or a particular page number.
- To browse through administrative folders, click **Browse All Folders**.

  Note: This link displays folders that might not contain portlets or portlet bundles.

  To view the portlets and portlet bundles in a folder, click the folder name.
- To reposition a portlet, in the area under the Edit Page section, drag the portlet to the desired position (by clicking the portlet, holding down the mouse button, and moving the mouse), then release the mouse button.

  Each column on the page is represented by a gray box. To change the column structure, click **Go to Advanced Editor**, then click **Select Page Layout**.
- To close the Flyout Editor, click **Close**, or, in the Edit Page section, click or **Close Editor**.

Editing a Page in the Page Editor

You can rename a page; add, delete, or reposition portlets; and select a page layout in the Page Editor.

1. On the page you want to edit, click **Edit Page**.
   - If you are editing a page that uses an adaptive page layout, the Flyout Editor opens. Continue with Step 2.
   - If you are editing a page that uses a legacy page layout, the Page Editor opens. Skip Step 2.
2. If the Flyout Editor opens, click **Go to Advanced Editor** to open the Page Editor.

- To change the name of your page, type a new name in the **Page Name** box.

- To add portlets to your page, click  **Add Portlets**.
- To change the format for the columns on your page, click  **Select Page Layout**.
- To see what a portlet looks like, click **Preview**.
- To remove a portlet from the page, click **Remove**.
- To reposition a portlet on your page, drag the portlet to the desired position (by clicking the portlet, holding down the mouse button, and moving the mouse), then release the mouse button.

  Each column on the page is represented by a gray box.

Managing a Portlet Through the Portlet's Title Bar

You can refresh portlet content, edit portlet preferences, hide or show a portlet, and delete the portlet all from the portlet's title bar.

- To refresh a portlet, click  .

  If this portlet does not include refresh functionality, you will not see this button.
- To edit your personal preferences for the portlet, click  .

  If this portlet does not have personal preferences, you will not see this button.
- To minimize a portlet so that only the title bar appears on the page, click  .
- To maximize a portlet so that the entire portlet displays, click  .
- To remove a portlet from the page, click  .

  If you are viewing a community page, you must have at least Edit access to the community to see this button.

Browsing Documents in the Portal Knowledge Directory

The Knowledge Directory is similar to a file system tree in that documents are organized in folders and subfolders. A folder can contain documents uploaded by users or imported by content crawlers, as well as links to people, portlets, and communities. If your administrator has given you permission, you might also be allowed to add documents to the Knowledge Directory, or submit yourself as an expert on a particular topic.

When you open the Directory, you see the folders and subfolders to which you have at least Read access.

- To edit the directory (add or edit folders), click **Edit Directory**.

  Note: You must have the Edit Knowledge Directory activity right to see this button. You must also have at least Edit access to a folder or document to be able to edit the folder or document.

- To open a folder or subfolder, click its name.

  Note: If the folder includes a description, it appears as a tooltip. To view the description, place your mouse over the folder name.

  After you have opened a Directory folder, you see the additional features described next.
- To open a document, click its name.
- To view the properties of a document, click the **Properties** link under the document description.
- To tell other portal users about a document:
  a) Under the document description, click **Send Document Link**.
  b) In the **Document Link** dialog box, copy the text, then click **Close**.
  c) In your e-mail application, paste the text into an e-mail message and send it.

  When other portal users click the URL in your e-mail, the document opens. If a user does not have permission to see the document, an error message is displayed.

- To submit a document to the portal, click **Submit Documents**.

  Note: You must have at least Edit access to the folder and at least Select access to the content source that provides access to the document to be able to submit a document.

- To view another page of items, at the bottom of the list of documents, click a page number or click **Next >>**.
- To change the sort order of documents between ascending and descending, in the **Sort** by drop-down list, select the desired option: **Document Name Ascending** or **Document Name Descending**.
- To change the number of documents that are displayed per page, in the **Items per page** drop-down list, select the desired number.

  By default, 20 items are shown per page.
- To filter the documents by document type (for example, MS Word documents or PDF documents), in the **Show only item type** drop-down list, select the desired document type.
- To open a subfolder, under **Subfolders**, click the subfolder name.

Note: Beneath the banner, you see the hierarchy for the folder you are viewing (sometimes referred to as a breadcrumb trail). To move quickly to one of these folders, click the folder's name.

- To create a subfolder in this folder, under **Subfolders**, click **Create Folder**. In the **Create Document Folder** dialog box, type a name and description for the folder, and click **OK**.
- To view a related community, under **Related Communities**, click the community name.

  Note: If you have at least Select access to the community, you can join the community.
- To open a related folder, under **Related Folders**, click the folder name.
- To preview a related portlet, under **Related Portlets**, click the portlet name.

  Note: If you have at least Select access to the portlet, from the portlet preview page, you can add the portlet to one of your My Pages.
- To view the user profile for a related expert, under **Related Experts**, click the user's name.

  Note: If you have the Self-Selected Experts activity right, and are not already listed as an expert, click **Add Me** to add yourself as an expert on the folder's topic.
- To view the user profile for a related content manager, under **Related Content Managers**, click the user's name.

Directing Users to Communities, Knowledge Directory Folders, or Documents with Friendly URLs

You can direct users to a community, Knowledge Directory folder, or document with a simple URL.

- To direct users to an object, create a link in the following format:
  http://portal.company.com/portal/server.pt/object_token/object_name/object_id
  - Replace *http://portal.company.com/portal/server.pt* with the URL to your portal.
  - Replace *object_token* with the token for the type of object to which you are linking.

    The default values are: `mypage`, `community`, `user`, `directory`, and `document`, but your portal administrator might have customized them. For example, "community" could instead be "site."

- Replace *object_name* with the name of the object. Replace any spaces in My Page, community, user, or Knowledge Directory folder names with underscores (_); replace any spaces in document names with plus signs (+).
- Replace *object_id* with the ID of the object. This element is optional.

Note:

- Users must have at least Read access to the object to which you are directing them.
- If an object cannot be found, the user will receive an error message.
- If more than one object has the name specified in the link and an ID is not specified, the portal displays a list of objects with the same name and the user can select which one to view.

Searching for Documents or Objects

You can search for documents or objects using the portal banner, using the Portal Search portlet, using advanced search, using federated search, or using the Administrative Objects Directory. Each method of search uses the same text search rules.

Note: Only those documents or objects to which you have at least Read access appear in your results.

- To search for documents or objects through the portal banner, type your search string in the box in the portal banner, and click **Search**.
- To go directly to the result that your portal administrator has set as the top best bet for a term, without first seeing all the search results, type the top best bet operator (>) followed by your search string and click **Search**, or type your search string and click  .

  Note: If no top best bet has been set for the term, the regular search results appear.

  Note: The top best bet button () is available only if enabled by your portal developer.
- To search for documents or objects using a saved search, use the Portal Search portlet.
- To search for documents or objects using metadata properties and location, click  **Advanced Search**.
- To search other content, portals, and web search engines, click  **Federated Search**.

  Note: The  **Federated Search** button is available only if enabled by your portal developer.

- To search for objects in the Administrative Objects Directory, click **Administration**, and specify your criteria in the **Object Search** boxes.

Searching for Objects or Documents Using Advanced Search

You can perform an advanced search, using metadata properties and location, to find objects or documents.

In the portal banner or in the Administrative Objects Directory, click **Advanced Search**.

- To search for text in the name or description of an object or document, type the text you want to search for in the **Search** for text box.

  You can use the text search rules.

- To search for property values, click **Add Criteria**, and specify the property criteria in the boxes that appear:

  a) In the first drop-down list (property), select the searchable property for which you want to filter the values.

  b) In the second drop-down list (operator), select the operator to apply to this condition.

  This list will vary depending on the property selected:

  - For any text property you can search for a value that contains your search string (Contains), or you can search for properties that are blank (Contains No Value).

    Note: To exclude results with particular values, select **Contains**, then type "not" followed by words that you want to exclude from your search. For example, if you want to search for documents about retirement benefits, excluding pension plans, then type "retirement benefits" in the Search for text box, select Contains from the drop-down list, and type "not pension plan" in the text box.

  - For any date property you can search for a value that comes after, comes before, is, or is not the date and time you choose or for a value that occurs in the last number of minutes, hours, days, or weeks that you specify.
  - For any number property you can search for a value that is greater than, is less than, is, is not, is greater than or equal to, or is less than or equal to the number you enter in the text box.

  c) In the **Value** text box, enter the value the property must have, or not have, depending on which operator you selected.

  Note: If you are searching for a text property, you can use the text search rules.

- To remove a property condition, select the condition and click  (next to  **Add Criteria**).
- Specify how you want your search criteria handled:
  - To meet all the conditions you define, select **All Criteria**.

    Selecting All Criteria is equivalent to using AND.
  - If you want your search results to meet at least one of the conditions you define, select **Any Criterion**.

    Selecting Any Criterion is equivalent to using OR.
- To restrict your search to specific Knowledge Directory folders, click  **Add Document Folder**. In the Select folder for search dialog box, select the folders you want to search and click **OK**.
- To restrict your search to specific Administrative Objects Directory folders, click  **Add Administration Folder**. In the Select folder for search dialog box, select the folders you want to search and click **OK**.
- To remove folders from your list, select the folders and click  (next to  **Add Administration Folder**).

  To select or clear all folder boxes, select or clear the box next to **Folder Names**.
- Specify whether you want to include subfolders.

  By default, the portal searches subfolders. To exclude subfolders from your search, clear the box next to **Include subfolders**.
- To specify the number of results to display on a page, in the **Results per page** drop-down list, choose a value.
- To restrict your search to a specific language, in the **This language only** drop-down list, choose a language.
- To limit your search to specific object types, in the **Result Types** list, select the object types you want to search.

  To select or clear all object type boxes, select or clear the box next to **Object Type**.
- To set all search conditions back to the defaults, click **Clear**.
- To perform your search, click **Search**.

Complex Property Search Example

You can use multiple property criteria to define complex property searches. For example, if you want to find documents published after a certain date by a specific branch of a company, you could set the property criteria to the following values:

- First Criterion:
  - Property = Object Created
  - Operator = Comes After
  - Value = December 30, 2003

  Your search results would be limited to objects created after December 30, 2003.
- Second Criterion:
  - Property = Company
  - Operator = Contains
  - Value = Company A

  Your search results would be limited to objects where the company property contains Company A.
- Third Criterion:
  - Property = Address
  - Operator = Contains
  - Value = San Francisco

  Your search results would be limited to objects that contain San Francisco in the address.
- You would also want to limit your Result Types to Documents, so that only documents were returned in your results.

Searching Other Content, Portals, or Web Search Engines with Federated Search

You can use federated search to search for content in web search engines (for example, Google or AltaVista), other portals, or other collections of information (for example, a Lotus Notes collection or a set of customer service incident reports).

Your portal administrator must have configured an outgoing federated search and your portal developer must have enabled the **Federated Search** button in the portal banner.

1. In the portal banner, click **Federated Search**.
2. In the **Search for text** box, type the text you want to search for.
3. Select the number of results you want returned per location.

   By default, 5 results per location are returned. To change this number, select a new number in the **Number of results per location** drop-down list.
4. Select the locations you want to search.
   - To select particular search locations, select the locations.
   - To select all search locations, select the box next to **Search Location**.
   - To save your search location settings, select **Always search these locations for me**.
5. Click **Search**.

Viewing Search Results from a Banner Search

When you perform a banner search, each item returned includes an icon to signify what type of document or object it is (for example, signifies a web page and signifies a user), the item name, the item description, when the item was last modified, and a link to view additional item properties.

- To view a result, click its name.

If you click a...	You see...
Knowledge Directory folder	The contents of the folder
Document	The document
Community	The home page for the community **Note:** You can join the community while viewing it.
Community page	The community page **Note:** You can join the community while viewing it.
Portlet	A preview of the portlet **Note:** You can add the portlet to a My Page while previewing it.
User	The user's profile

- To view the properties of an item, click the **Properties** link under the item description.
- If your search returns more than one page of results, click a page number or click **Next >>** to view additional results.
- To change the sort order of your results, in the **Sort by** drop-down list, select the desired option:
  - **Relevance** sorts your results according to how closely they match your search query.

    Note: Best bets are only shown in search results when sorting by relevance.
  - **Name** sorts your results alphabetically by name.
  - **Last Modified Date** displays your results in the order in which they were most recently edited.
- To change the number of results that are displayed per page, in the **Items per page** drop-down list, select the desired number.
- To filter your results by type (for example, documents, communities, portlets), in the **Show only item type** drop-down list, select the desired item type.
- To edit the results to which you have at least Edit access, click  **Edit**.

  Note: You must have the proper permissions to see the  **Edit** button. You must have at least Edit access to some of the results. For documents or document folders, you must have the Edit Knowledge Directory activity right. For communities, community pages, portlets, or users, you must have the Access Administration activity right. If you have only the Edit Knowledge Directory activity right, you must filter your results to display only **Documents in the Directory** or **Document Folders**. If you have only the Access Administration activity right, you must filter your results to display only **Communities**, **Community Pages**, **Portlets**, or **Users**.

Saving a Search

You can save a search and access it later through the Portal Search portlet.

Note: Sort order is not saved when you save your search query.

1. Run a search.
2. On the search results page, click  **Save this Search**.
3. In the dialog box, type a name for this search and click **Save**.
4. To add the Portal Search portlet to your My Page, click the link in the dialog box (**Click here to add the portlet**).

You can run and manage your saved searches through the Portal Search portlet. You can also manage your saved searches through the Search Preferences page.

5. Click **Close Window**.

Linking to a Best Bet Search

You can create a link to go directly to a top best bet. This can be useful if you want to direct users to an object or document related to a particular issue, but the object or document changes frequently.

For example, you might want to direct customers to your current privacy statement, but you need to keep copies of older privacy statements in your portal for internal reference. You could create a top best bet that points to the current privacy statement and add a link to that top best bet on your customer account page. When your privacy statement is updated, you can change the top best bet without having to change any links you made to the privacy statement.

1. Create the top best bet as described in *Creating Best Bets* on page 263.
2. Create the link to the top best by appending `tbb=SearchTerm` to the end of your portal URL.
   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`.

Overview of the Administrative User Interface, Features, and Tools

Portal Objects

The following table describes the portal objects you can create through the **Create Object** drop-down list in the Administrative Objects Directory.

Object	Description
Administrative Folder	Administrative folders provide a hierarchical structure that make it easy to organize portal objects and manage security.

Object	Description
Authentication Source - Remote	Authentication sources enable you to import users, groups, and group memberships that are already defined in your enterprise in existing user repositories, such as Active Directory or LDAP servers. After users are imported, you can authenticate them with the credentials from those user repositories.
Community	Communities are sites within a portal designed for a specific audience or task, such as collaborative projects.
Community Template	Community templates define the basic structure for the resulting communities, such as which page templates to include and, optionally, a header or footer for the community.
Content Crawler - Remote	Remote content crawlers enable you to import content from external content repositories such as a Windows NT file system, Documentum, Microsoft Exchange, or Lotus Notes.
Content Crawler - WWW	Web content crawlers enable you to import content from web sites.
Content Source - Remote	Remote content sources provide access to external content repositories, such as a Windows NT file system, Documentum, Microsoft Exchange, or Lotus Notes.
Content Source - WWW	Web content sources provide access to web sites.
Content Type	Content types specify several options — the source content format (such as Microsoft Office, web page, or Lotus Notes document), whether the text of the content should be indexed for searching, and how to populate values for document properties.
Experience Definition	Experience definitions provide multiple user experiences within a single portal. An experience definition defines certain elements of a user experience, such as adaptive page layout settings, branding style, and navigation.
External Operation	An external operation enables you to run shell scripts (for example, .sh or .bat files) through the portal and schedule these actions through portal jobs. For example, you might want to create a script that queries documents, pings portal servers, e-mails snapshot query results to users,

Object	Description
	or runs some other custom job, then create an external operation that points to the script, and use a job to run the script on a specified schedule.
Federated Search - Incoming	An incoming federated search allows other AquaLogic Interaction portals to search your portal.
Federated Search - Outgoing	An outgoing federated search enables users of your portal to search other AquaLogic Interaction portals or other external repositories.
Filter	Filters control what content goes into which folder when crawling in documents or using Smart Sort to filter content into new folders. A filter sets conditions that document links must pass in order to be sorted into associated folders in the Knowledge Directory.
Group	Groups are sets of users, sets of other groups, or both. Groups enable you to more easily control security because you assign each group different activity rights and access privileges.
Invitation	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.
Job	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.
Page (Only displays when in a community folder)	Community pages let you categorize information for your community audience.
Page Template	Page templates define the basic structure for the resulting community pages, such as the column layout and which portlets to include.
Portlet	Portlets provide portal users customized tools and services as well as information. Portlets let you to integrate applications, tools, and services into your portal, while taking advantage of portal security, caching, and customization.
Portlet Bundle	Portlet bundles are groups of related portlets, packaged together for easy inclusion on My Pages or community pages.

Object	Description
Portlet Template	Portlet templates allow you to create multiple instances of a portlet, each displaying slightly different information.
Profile Source - Remote	Profile sources allow you to import user information (such as name, address, or phone number) that is already defined in your enterprise in existing user repositories, such as Active Directory or LDAP servers. The imported user information can be used to populate user profiles or can be passed to content crawlers, remote portlets, or federated searches as user information.
Property	Properties provide information about, as well as a way to search for, documents and objects in your portal. For example, you might want to create an Author property so users can find all the documents or objects created by a particular user.
Remote Server	Remote servers group together web services that are installed on the same computer and require the same type of authentication. With a remote server, you enter the base URL and authentication settings just once for multiple web services, and, if you need to move the web services, you just need to change the remote server settings.
Snapshot Query	Snapshot portlets enable you to display the results of a search 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..
User	Portal users enable you to authenticate the people who access your portal and assign appropriate security for the documents and objects in your portal. Users can be imported from external user repositories, created through the portal, created through invitations, self-registered, or just guests (unauthenticated users).
Web Service - Authentication	Authentication web services enable you to specify general settings for your external user repository, leaving the more detailed settings (like domain specification) to be set in the associated remote authentication sources. This allows you to create different authentication sources to import each domain without having to repeatedly specify all the settings.
Web Service - Content	Content web services enable you to specify general settings for your external user repository, leaving the target and security settings to be set

Object	Description
	in the associated remote content source and remote content crawler. This allows you to crawl multiple locations of the same content repository without having to repeatedly specify all the settings.
Web Service - Intrinsic Portlet	Portlet web services allow you to specify *functional* settings for your portlets, leaving the *display* settings to be set in each associated portlet. An intrinsic portlet web service references one or more sets of code that are located on the portal computer.
Web Service - Profile	Profile web services enable you to specify general settings for your external user repository, leaving the more detailed settings (like domain specification) to be set in the associated remote profile sources. This allows you to create different profile sources to import information each domain without having to repeatedly specify all the settings.
Web Service - Remote Portlet	Portlet web services allow you to specify *functional* settings for your portlets, leaving the *display* settings to be set in each associated portlet. A remote portlet web service references services hosted by a separate remote server.
Web Service - Search	Search web services allow you to specify general settings for your remote search repository, leaving the security settings to be set in the associated outgoing federated searches. This allows you to segregate access to your search repository through multiple outgoing federated searches.

Portal Objects Created Upon Installation

The default portal installation includes several portal objects that are created upon installation.

Object	Description
Administrative Resources (folder)	This folder contains the following objects created at installation: users, groups, the AquaLogic Interaction Authentication Source, the World Wide Web content source, properties, content types, and federated search objects.

Object	Description
Intrinsic Operations (folder)	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.
Portal Resources (folder)	This folder contains intrinsic portlets and web services, as well as page, community, and portlet templates.
Default Experience Definition (folder)	This folder contains the users associated with the default experience definition. Upon installation, one user is associated with the default experience definition—Administrator.
Audit Log Management (job)	This job archives old audit messages into files and deletes old audit files.
Bulk Subscriptions (job)	This job subscribes users to communities and portlets when you use bulk add.
Document Refresh (job)	This job performs background maintenance on your search index such as refreshing document links and properties and deleting expired documents.
Dynamic Membership Update Agent (job)	This job updates dynamic group memberships as defined on the Dynamic Membership Rules page of the Group Editor.
Search Update (job)	This job makes sure the search collection is synchronized with the database. You can run multiple instances of this job.
Weekly Housekeeping (job)	This job performs weekly housekeeping on your system, such as deleting expired invitation codes and deleting uploaded files for which links have been deleted.
Navigation Tags Header Portlet (portlet)	This portlet is provided as an example of a custom header that includes navigation tags; you can customize it and use it in communities or experience definitions. This portlet is stored in the Portal Resources folder.
Classic Footer Portlet (portlet)	This portlet is provided as an example of a custom footer that you can customize and use in communities or experience definitions.
Classic Header Portlet (portlet)	This portlet is provided as an example of a custom header that you can customize and use in communities or experience definitions.
Layout Footer Portlet (portlet)	This portlet is provided as an example of a custom footer that uses adaptive tags; you can customize it and use it in communities or experience definitions.

Object	Description
Layout Footer Portlet (portlet)	This portlet is provided as an example of a custom header that uses adaptive tags; you can customize it and use it in communities or experience definitions.
Portal Login (portlet)	This portlet allows users to log in to the portal. You probably want to add this to all your guest users' home pages so that users can log in from the default page displayed when they navigate to your portal.
Tag Login Portlet (portlet)	This portlet is provided as an example of a custom login portlet that uses adaptive tags; you can customize it and add it to your guest users' home pages so that users can log in from the default page displayed when they navigate to your portal. This portlet is stored in the Portal Resources folder. For information on adaptive tags, see the *Adaptive Page Layouts* section of the *AquaLogic User Interaction Development Guide*.
Folder Expertise (portlet)	This portlet displays the folders for which the user is an expert. Portal administrators can add users to a folder as an expert through the Related Resources page of the Folder Editor, or, if users have the Self-Selected Experts activity right, they can add themselves as experts when they are browsing folders in the Knowledge Directory. This portlet is stored in the Portal Resources folder and is displayed on the user profile page by default.
General Information (portlet)	This portlet displays user profile information such as name and address, but it is configurable by the portal administrator to display any information. If your portal displays a legacy layout (rather than adaptive layouts), this portlet is displayed on the user profile page by default. This portlet is stored in the Portal Resources folder.
Managed Communities (portlet)	This portlet displays the communities to which the user has Edit or Admin access. If your portal displays a legacy layout (rather than adaptive layouts), this portlet is displayed on the user profile page by default. This portlet is stored in the Portal Resources folder.
Job Histories Intrinsic Portlet (portlet)	This portlet displays the same job history information that is displayed on the Job History page of the Automation Service Manager. This portlet is stored in the Portal Resources folder.
Portal Search (portlet)	This portlet lets users search your portal and access their saved searches. Users might want to add this to their home page for easy access to their saved searches. This portlet is stored in the Portal Resources folder.

Object	Description
RSS Reader Portlet (portlet)	This portlet lets users specify an RSS or ATOM feed to display on a My Page. This portlet is stored in the Portal Resources/RSS Reader folder, but is available only if the portal administrator installed the Remote Portlet Service and imported the RSS Reader migration package.
RSS Community Reader Portlet (portlet)	This portlet lets community managers specify an RSS or ATOM feed to display on a community page. This portlet is stored in the Portal Resources/RSS Reader folder, but is available only if the portal administrator installed the Remote Portlet Service and imported the RSS Reader migration package.
User Status (portlet)	This portlet lets users post their current status. This portlet is stored in the Activity Service folder, but is available only if the portal administrator installed the Remote Portlet Service and imported the Activity Service migration package.
User Activities (portlet)	This portlet displays a user's status history and any other recent activities that are submitted by other applications. This portlet is stored in the Activity Service folder, but is available only if the portal administrator installed the Remote Portlet Service and imported the Activity Service migration package. To view another user's activities, open the user's profile and look at the User Activities portlet displayed in the profile. To subscribe to e-mail notification or an RSS feed of the user's activity, click the appropriate button at the bottom of the user's User Activities portlet.
Community Links Portlet Template (portlet template)	This template is used by the portal to create portlets that display the links saved in a Community Knowledge Directory folder. This portlet template is stored in the Portal Resources folder.
Content Snapshots (portlet template)	This template is used by the portal to create portlets that display the results of a Snapshot Query. This portlet template is stored in the Portal Resources folder.

Portal Utilities

The following table describes the portal utilities accessible through the **Select Utility** drop-down list in the Administrative Objects Directory.

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

Utility	Description
Portal Settings	Modify Portal settings.
Release Disabled Logins	Manage user locks.
Release Item Locks	Manage object locks.
Search Cluster Manager	Check status and manage search topology and checkpoints.
Search Results Manager	Manage search results preferences.
Search Service Manager	Manage Search Service settings.
Smart Sort	Run the Smart Sort utility.
System Health Monitor	View diagnostic information.
Tag Library Manager	Displays the tag libraries loaded on the computer that hosts the portal.
User Profile Manager	Modify the user profiles map.
(Custom Utility)	Portal administrators or portal developers can create custom utilities that display in the **Select Utility** drop-down list.

Administration Utilities in the Portal Installation

AquaLogic Interaction includes several command-line administration utilities in the portal installation directory and a tag library utility in the portal administrative interface.

Note: The command-line utilities are located on the computer that hosts the portal, in *Install_Dir*/ptportal/6.5/bin. *Install_Dir* is the portal installation directory, for example C:\bea\alui or /opt/bea/alui.

Utility (.sh or .bat file)	Purpose
automationserverd	The Automation Service daemon ensures the Automation Service is running. For information on the Automation Service daemon, see the *Installation Guide for AquaLogic Interaction*. For information on modifying Automation Service defaults, see Configuring the Automation Service.
cryptoutil	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: `Install_Dir/ptportal/6.5/bin/cryptoutil.sh -h`
diagnostic	The Diagnostic utility allows you to verify connectivity for installation components and the portal database. To display the man pages for theDiagnostic utility, enter the following command: `Install_Dir/ptportal/6.5/bin/diagnostic.sh -h` For details, see the *Installation Guide for AquaLogic Interaction*.
portalenv	The Portal Environment utility sets the portal environment for tools in `Install_Dir/ptportal/6.5/bin`. To display the man pages for the Portal Environment utility, enter the following command: `Install_Dir/ptportal/6.5/bin/portalenv.sh -h`
ptmigration	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:

Utility (.sh or .bat file)	Purpose
	Install_Dir/ptportal/6.5/bin/ptmigration.sh -h For information on object migration, see *Migrating Portal Objects* on page 305.
Tag Library Manager	This Tag Library Manager allows you to view the tag libraries installed on the computer that hosts the portal. To access the Tag Library Manager, in the portal, click **Administration**, then, in the **Select Utility** menu, click **Tag Library Manger**.

Overview of Portal Security

AquaLogic Interaction provides many features that work together to secure your portal and its content.

- Activity security in the form of activity rights. See *About Activity Rights* on page 61.
- Audit records, which you should periodically review to keep track of actions performed by users. See *Auditing User Accounts and Actions* on page 135.
- Automatic user lockout. See *Automatically Locking User Accounts* on page 98
- Web application credential management in the form of lockboxes. See *Managing User Credentials for External Applications Through the Credential Vault Manager* on page 239.
- Object level security in the form of Access Control Lists (ACLs). See *Setting Security on an Object* on page 221.
- Document security imported from source repositories. See Importing Document Security from External Repositories.
- Single sign-on. See Deploying Single Sign-On.

Important: 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.

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.

About Access Privileges

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.

Access to each object and document in the portal is controlled through the following access privileges:

Access Privilege	Description
Read	Allows users or groups to see the object.
Select	Allows users or groups to add the 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.
Edit	Allows users or groups to modify the object.
Admin	Allows users or groups full administrative control of the object, including deleting the object or approving it for migration.

Note:

- The Everyone group (all users) has mandatory Read access to authentication sources, content types, filters, invitations, and properties.
- If a user is a member of more than one group included in the list, or if they are included as an individual user and as part of a group, that user gets the highest privilege available to the user for the object. For example, if a user is part of the Everyone group (which has Read access) and the Administrators group (which has Admin access), that user gets the higher privilege to the object: Admin.
- Access privileges are based on the security of the folder in which the object is stored. Changes to the security of a folder apply to all the objects within that folder. For example, if a document in the folder is shared with another folder (such as when a document is copied from one folder to another), the security of the document is changed in both locations.

About Activity Rights

Activity rights determine which portal objects a user can create and which portal utilities a user can execute to create or modify portal objects. For example, you can specify that users can create communities, create folders, create content types, and create portlets. Activity rights are global and cumulative. If a user is a member of multiple groups, each with different rights, that user inherits all the activity rights of all the parent groups. That user can exercise all of those rights in any area of the portal to which that user has the appropriate access. Groups can also inherit activity rights.

In addition to the default 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.

Activity Rights Required to Create Portal Objects

To create a portal object, you must have at least Edit access to the parent folder (the folder that will store the object), the Access Administration activity right, and the required activity right listed in the table.

Object	Required Activity Right
Administrative Folder	Create Admin Folders
Authentication Source - Remote	Create Authentication Sources
Community	Create Communities
Community Template	Create Community Infrastructure
Content Crawler - Remote	Create Content Crawlers
Content Crawler - WWW	Create Content Crawlers
Content Source - Remote	Create Content Sources
Content Source - WWW	Create Content Sources
Content Type	Create Content Types
Experience Definition	Create Experience Definitions
External Operation	Create External Operations

Object	Required Activity Right
Federated Search - Incoming	Create Federated Searches
Federated Search - Outgoing	Create Federated Searches
Filter	Create Filters
Group	Create Groups
Invitation	Create Invitations
Job	Create Jobs
Page (Only displays when in a community folder)	No activity right needed; just need at least Edit access to community
Page Template	Create Community Infrastructure
Portlet	Create Portlets
Portlet Bundle	Create Web Service Infrastructure
Portlet Template	Create Web Service Infrastructure
Profile Source - Remote	Create Profile Sources
Property	Create Properties
Remote Server	Create Web Service Infrastructure
Snapshot Query	Create Snapshot Queries
User	Create Users
Web Service - Authentication	Create Web Service Infrastructure
Web Service - Content	Create Web Service Infrastructure
Web Service - Intrinsic Portlet	Create Web Service Infrastructure
Web Service - Profile	Create Web Service Infrastructure
Web Service - Remote Portlet	Create Web Service Infrastructure
Web Service - Search	Create Web Service Infrastructure

Activity Rights and Group Membership Required to Access Portal Utilities

To access a utility, you must have the Access Administration activity right, the Access Utilities activity right, and the required activity right or group membership listed in the table.

Utility	Required Activity Right (AR) or Group Membership (GM)
Access Unclassified Documents	Access Unclassified Documents (AR)
Activity Manager	Create Activities (AR)
Approve Directory Content	Access Utilities (AR)
Approve Objects for Migration	Administrators Group (GM)
Audit Manager	Administrators Group (GM)
Automation Service	Administrators Group (GM)
Credential Vault Manager	Administrators Group (GM)
Default Profiles	Create User (AR)
Experience Rules Manager	Access Experience Rules Manager (AR)
Global ACL Sync Map	Administrators Group (GM)
Global Content Type Map	Administrators Group (GM)
Global Document Property Map	Administrators Group (GM)
Global Object Property Map	Administrators Group (GM)
Knowledge Directory Preferences	Administrators Group (GM)
Localization Manager	Administrators Group (GM)
Migration - Export	Administrators Group (GM)

Utility	Required Activity Right (AR) or Group Membership (GM)
Migration - Import	Administrators Group (GM)
Object Migration Status	Access Utilities (AR)
Portal Settings	Administrators Group (GM)
Release Disabled Logins	Administrators Group (GM)
Release Item Locks	Administrators Group (GM)
Search Cluster Manager	Administrators Group (GM)
Search Results Manager	Access Search Results Manager (AR)
Search Service Manager	Administrators Group (GM)
Smart Sort	Access Smart Sort (AR)
System Health Monitor	Administrators Group (GM)
Tag Library Manager	Administrators Group (GM)
User Profile Manager	Access User Profile Manager (AR)
(Custom Utility)	Read access to the custom utility's web service

Overview of Web Service Architecture

Many of the objects in the portal use 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 service architecture allows multiple types of remote services (authentication sources, content crawlers, outgoing federated searches, portlets, and profile sources) to share a logical remote server, making it easier to manage the computers that make up the portal.

Web services also allow you to share settings (sometimes rather complex settings) with the objects created from those services. For example, administrative users creating portlet web services need a greater understanding of the structure of the portlet, because they need to specify whether the portlet has preferences or whether it sends user information; whereas users creating portlets from that web service might only need to set configuration settings appropriate for a non-technical user.

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.

Objects that use web services follow this general structure:

- The remote server contains the base URL and credentials.
- The web service defines configuration settings for the associated object: remote authentication source, remote content source (used to create remote content crawlers), outgoing federated search, portlet, and remote profile source.
- The associated object defines any remaining configuration settings.

About Remote Servers

Remote servers group together web services that are installed on the same computer and require the same type of authentication. With a remote server, you enter the base URL and authentication settings just once for multiple web services, and, if you need to move the web services, you just need to change the remote server settings.

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.

You can use a remote server to create the following administrative objects:

- Search Web Service
- Profile Web Service
- Authentication Web Service
- Remote Portlet Web Service
- Content Web Service

Specifying the Location and Authentication Settings for a Remote Server

You can specify the location and authentication needed to access a remote server on the Main Settings page of the Remote Server Editor.

1. If the Remote Server Editor is not already open, open it now and display the **Main Settings** page.
2. In the **Base URL** text box, type the URL to the parent folder of the web services installed on this server.

   This can be the root of the Web Server (for example, http://server/) or a specific application or virtual directory (for example, http://server/app/). Because the URL specifies a folder rather than a specific resource, it should always end with a forward slash.

   The portal must be able to resolve the server name. Therefore, you might want to use Fully-Qualified Domain Names (FQDNs) such as "http://server.companyname.com" rather than just "http://server". In some cases, when the portal is in a demilitarized zone (DMZ), you might need to use an IP address like "10.1.2.140".

   You need a remote server for each port (for example, http://server:8082/ requires a different remote server than http://server:7071/). You also need a separate remote server if some services use SSL (for example, https://server/).

   Note: If you are sending any type of basic authentication information (specified in step 2), and you are not using a secured network, such as a separate subnet or a Virtual Private Network (VPN) connection, we strongly recommend that the Base URL use SSL (the URL must begin with https://). Basic authentication uses Base 64 encoding, which can be easily decoded back to clear text.

3. Under **Base Authentication Type**, specify what authentication information, if any, you want this remote server to pass to its associated web services.

   - To use no authentication information, choose **None**.
   - To use credentials from a user's login, choose **User's Basic Authentication Information**.

     Confirm that the portal configuration file has been edited so that the portal stores the user name and password in memory for as long as the user is logged in to the portal (as described in the Installation Guide for AquaLogic Interaction, available on edocs.bea.com). This option is not supported for configurations in which users log in without typing a password (for example, single sign-on or Remember My Password) because the password is not available to the portal.

- To specify a user name and password, choose **Administrator's Basic Authentication Information** and type the user name and password in the associated text boxes.

  This information is encrypted, stored in the portal database, and sent with all requests to this remote server.

4. To send credentials in portlet headers, using RSA public key/private key encryption, in the **Public Encryption Key** box, enter the public key for RSA encryption.

   You must also set up a lockbox in the Credential Vault Manager, associate the lockbox with the remote portlet web service (on the **Authentication Settings** page), and use the IDK to provide the private key for RSA encryption (see the AquaLogic User Interaction Development Center for information).

About User Interface Customization

AquaLogic Interaction provides several features that work together to control your portal user interface.

- *About Customizing the User Interface with Adaptive Layouts* on page 70

  Adaptive layouts let you quickly change the look and feel of areas in the portal user interface using adaptive tags in standard XHTML. Adaptive layouts are displayed in the portal through remote portlet web services. Adaptive page layouts are applied at the experience definition level and affect the entire experience definition. Adaptive portlet layouts are applied at the My Page and community page level and affect only that page.

- *About Controlling the User Interface with Experience Definitions and Experience Rules* on page 76

  Experience definitions provide multiple user experiences within a single portal. An experience definition defines certain elements of a user experience, such as adaptive page layout settings, branding style, and navigation.

- *About Branding with Header and Footer Portlets* on page 88

  Branding portlets customize the look of your portal 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.

- *About Navigation Options* on page 89

  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.

- *Portal Interface Types* on page 31

  Users can change their portal display to accommodate assistive technologies or slow internet connections.

- *Editing Locale Settings* on page 32

  Users can change their portal display to accommodate their time zone and locale.

About Customizing the User Interface with Adaptive Layouts

Adaptive layouts let you quickly change the look and feel of areas in the portal user interface using adaptive tags in standard XHTML. Adaptive layouts are displayed in the portal through remote portlet web services. Adaptive page layouts are applied at the experience definition level and affect the entire experience definition. Adaptive portlet layouts are applied at the My Page and community page level and affect only that page.

You can create different layouts for each area in the portal:

Layout Type	Description
Base Page (adaptive page layout)	Controls the layout of everything surrounding the content area, such as the header, footer, banner, and navigation. **Note:** The base page layout applies to all areas of the portal except for profile pages.
Profile Page (adaptive page layout)	Controls the layout of everything surrounding the content area in user profile pages, such as the header, footer, banner, and navigation.
Knowledge Directory (adaptive page layout)	Controls the layout for the content area in the Knowledge Directory.

Layout Type	Description
Search Results (adaptive page layout)	Controls the layout for the content area in search results.
Portlet Selection (adaptive page layout)	Controls the layout of the flyout editor used to select portlets on My Pages and community pages.
Portlet Layout (adaptive portlet layout)	Controls the column layout for the content area (where portlets are placed) and the look of portlets (the borders and portlet toolbar) in My Pages, profile pages, and community pages. **Note:** As long as adaptive portlet layouts are enabled in the portal configuration file, adaptive portlet layouts can be used in any user interface, whether the interface uses adaptive page layouts or a legacy user interface.

The default adaptive page and portlet layout files are stored on the computer that hosts the Image Service in *Install_Dir*\ptimages\imageserver\plumtree\portal\private\pagelayouts\, where *Install_Dir* is the directory in which you installed the Image Service (for example, C:\bea\alui\ or /opt/bea/alui/). For information on creating adaptive layouts, see the *Adaptive Page Layouts* section of the *AquaLogic User Interaction Development Guide*.

The remote portlet web services that display adaptive layouts must be stored in the **Page Layouts** administrative folder, in the subfolder corresponding to the type of layout. For example, all the layouts stored in the **Base Page Layouts** subfolder are available in the **Base Page Layouts** drop-down list in the Experience Definition Editor. All the layouts stored in the **Portlet Layouts** subfolder are available in the drop-down list in the Select Page Layout dialog box when editing a My Page or community page.

When you create an experience definition that uses adaptive page layouts, you select an appropriate page layout for each area of the portal. When users create a My Page or community page, they select whether to display the adaptive portlet layout or a legacy user interface.

Note: If, for any reason, the page layouts cannot be loaded, the user interface will revert to the legacy user interface.

Customizing the User Interface with Adaptive Layouts

Adaptive layouts let you quickly change the look and feel of areas in the portal user interface using adaptive tags in standard XHTML. Adaptive layouts are displayed in the portal through remote portlet web services. Adaptive page layouts are applied at the experience definition level and affect the entire experience definition. Adaptive portlet layouts are applied at the My Page and community page level and affect only that page.

1. Make sure that adaptive layouts are enabled in the portal configuration file.

   You can enable both adaptive page layouts and adaptive portlet layouts or just one or the other.
2. Create the adaptive layouts as described in the *Adaptive Page Layouts* section of the *AquaLogic User Interaction Development Guide*.
3. Create remote portlet web services that point to the adaptive layouts, as described in *Creating a Remote Portlet Web Service for an Adaptive Layout* on page 72.
4. If you want to use adaptive page layouts, create at least one experience definition that uses them, as described in *Creating an Experience Definition to Display Adaptive Page Layouts* on page 73.

If you enabled adaptive portlet layouts, users can select them for their My Pages, community managers can select them for their community pages, and portal administrators can select them for the profile pages.

Creating a Remote Portlet Web Service for an Adaptive Layout

Adaptive layouts let you quickly change the look and feel of areas in the portal user interface using adaptive tags in standard XHTML. Adaptive layouts are displayed in the portal through remote portlet web services. Adaptive page layouts are applied at the experience definition level and affect the entire experience definition. Adaptive portlet layouts are applied at the My Page and community page level and affect only that page.

The tasks described below assume that you have already created the adaptive layouts as described in the *Adaptive Page Layouts* section of the *AquaLogic User Interaction Development Guide*.

Before you create a remote portlet web service, if necessary, create the remote server the web service will point to. If your adaptive layouts are stored on the computer that hosts the Image Service, you can use the **ImageServer Remote Server**.

To create a remote portlet web service you must have the following rights and privileges:

- Access Administration activity right
- Create Web Service Infrastructure activity right

- At least Edit access to the parent folder (the folder that will store the web service)
- At least Select access to the remote server the web service will point to

1. Click **Administration**.
2. Open the **Page Layouts** folder.

   Remote portlet web services for adaptive layouts must be stored in this folder to be available in the Experience Definition Editor and the Select Page Layout dialog box for My Pages and community pages.
3. Open the folder for the type of layout for which you are creating a web service.
   For example, if you are creating a web service for a layout that applies to search results pages, open the **Search Results Page Layouts** folder.
4. In the **Create Object** drop-down list, click **Web Service — Remote Portlet**.
5. Next to **Remote Server**, click **Browse**.
   The Choose Remote Server dialog box opens.
6. Select the remote server this web service should point to and click **OK**.

   If your adaptive layouts are stored on the computer that hosts the Image Service, you can use the **ImageServer Remote Server**.
7. In the **Portlet URL** box, complete the path to the adaptive layout.
   For example:
   `plumtree/portal/private/pagelayouts/searchresultslayout.html`

The default security for this web service is based on the security of the parent folder. You can change the security when you save this web service (on the **Security** tab page in the Save As dialog box), or by editing this web service (on the **Security** page of the Web Service Editor).

Note:

- If you are creating a web service for a portlet layout, provide at least Select access to any users you want to be able to select the layout in My Pages or community pages.
- If you are creating a web service for a page layout, provide at least Select access to any users you want to be able to select the layout in experience definitions.

Creating an Experience Definition to Display Adaptive Page Layouts

Adaptive layouts let you quickly change the look and feel of areas in the portal user interface using adaptive tags in standard XHTML. Adaptive layouts are displayed in the portal through remote portlet web services. Adaptive page layouts are applied at the experience definition level and affect the entire experience definition. Adaptive portlet layouts are applied at the My Page and community page level and affect only that page.

Before you create an experience definition that displays adaptive page layouts, you must:

- Create any custom adaptive page layouts you want to use
- Create remote portlet web services for the custom adaptive page layouts
- Create the guest user you want to associate with the experience definition
- Create any header and footer portlets you want to use to brand the experience definition

To create an experience definition that displays adaptive page layouts you must have the following rights and privileges:

- Access Administration activity right
- Create Experience Definitions activity right
- At least Edit access to the parent folder (the folder that will store the experience)
- At least Select access to the remote portlet web services for the adaptive page layouts
- At least Select access to the guest user you want to associate with the experience definition
- At least Select access to any header and footer portlets you want to add to the experience definition

1. Click **Administration**.
2. Open the folder in which you want to store the experience definition.

   Tip: You might want to store all of the resources needed by a particular audience of users in the same folder in which you store those users. By securing the folder appropriately and applying experience definition settings to it you can create completely separate and discreet user experiences for each audience of users.

3. In the **Create Object** drop-down list, click **Experience Definition**.
   The Experience Definition Editor opens.
4. On **Experience Definition Features** page, complete the following tasks:

   - *Associating Folders with an Experience Definition* on page 79
   - *Selecting the Portal Menus and Home Page for an Experience Definition* on page 81

5. Click the **Choose Header, Footer & Style** page and complete the following task:

   - *Branding Experience Definitions with Headers and Footers* on page 82

6. Click the **Edit Navigation Options** page and complete the following tasks:

   - Under **Navigation Type**, select **Portlet-Ready Navigation**.
   - *Defining Mandatory Links to Display in an Experience Definition* on page 84

7. Click the **Login Settings** page and complete the following task:

   - *Defining the Guest User Experience for an Experience Definition* on page 85
   - *Disabling Single Sign-On (SSO) for an Experience Definition* on page 86

8. Click the **Adaptive Page Layout Settings** page and complete the following task:

   - *Applying Adaptive Page Layouts* on page 86

9. Click the **Properties and Names** page and complete the following tasks:

   - *Naming and Describing an Object* on page 217

     You can instead enter a name and description when you save this experience definition.

   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219(optional)

The default security for this experience definition is based on the security of the parent folder. You can change the security when you save this experience definition (on the **Security** tab page in the Save As dialog box), or by editing this experience definition (on the **Security** page of the Experience Definition Editor).

Reverting to a Legacy User Interface

If you want to display the user interface used in previous versions of the portal you can do so with settings in the Experience Definition Editor.

1. Open the Experience Definition Editor by creating a new experience definition or editing an existing one.
2. Click the **Adaptive Page Layout Settings** page.
3. Under **Adaptive Page Layout Mode**, clear the **Enable Adaptive Page Layout Mode** box.
4. Perform tasks on the remaining pages as necessary:

   - *Associating Folders with an Experience Definition* on page 79
   - *Selecting the Portal Menus and Home Page for an Experience Definition* on page 81
   - *Branding Experience Definitions with Headers and Footers* on page 82
   - *Selecting a Navigation Scheme for an Experience Definition* on page 82

Important: You must select different header and footer portlets (the layout header and footer portlets will not work in a legacy user interface). If you select a header portlet that does not include navigation, you cannot use **Portlet-Ready Navigation** (or users will not see any navigation).

- *Defining Mandatory Links to Display in an Experience Definition* on page 84
- *Defining the Guest User Experience for an Experience Definition* on page 85
- *Disabling Single Sign-On (SSO) for an Experience Definition* on page 86
- *Naming and Describing an Object* on page 217

  You can instead enter a name and description when you save this experience definition.
- *Localizing the Name and Description for an Object* on page 342 (optional)
- *Managing Object Properties* on page 219 (optional)
- *Setting Security on an Object* on page 221

About Controlling the User Interface with Experience Definitions and Experience Rules

Experience definitions provide multiple user experiences within a single portal. An experience definition defines certain elements of a user experience, such as adaptive page layout settings, branding style, and navigation. An experience rule defines the conditions that, when met, display the associated experience definition to a user.

Experience Definitions

The experience definition specifies the following:

- Which portal menus to display (My Pages, My Communities, Directory)
- What navigation scheme to display
- Which header and footer to display

  Note: The headers and footers can be overridden at the community level.
- Any mandatory links to display
- The default page displayed when a user logs in (such as a My Page, a particular community, or a Knowledge Directory folder)

Users are directed to a particular experience definition in three ways (in the following order):

1. The users satisfy a rule 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.
2. The users are stored in a folder that is associated with the experience definition.
3. If neither of the above conditions are met, users experience the default experience definition for the portal.

Tip: You might want to store all of the resources needed by a particular audience of users in the same folder in which you store those users. By securing the folder appropriately and applying experience definition settings to it you can create completely separate and discreet user experiences for each audience of users.

Experience Rules

When you create an experience rule, you must also place it in rank order in relation to existing rules. The first rule to evaluate to true will be applied. For example, you might create a rule that says that users in the Marketing group see the user interface defined in the Marketing experience definition, and another rule that says that users in the Management group see the user interface defined in the Management experience definition. Since some users may be in both groups, you may decide that you want the Management experience definition to have priority. In this case, you order the two rules so that the Management experience rule is above the Marketing experience rule.

Guest User Experiences

If you want to have different user experiences for different audiences of guest users (users that have not logged in), you might want to create several guest users and assign them different experience definitions. For an example, see the Guest Users section in *About Users* on page 95.

Creating an Experience Definition to Control the User Interface

Experience definitions provide multiple user experiences within a single portal. An experience definition defines certain elements of a user experience, such as adaptive page layout settings, branding style, and navigation. An experience rule defines the conditions that, when met, display the associated experience definition to a user.

Before you create an experience definition you must:

- Create any custom adaptive page layouts you want to use

- Create remote portlet web services for any custom adaptive page layouts
- Create the guest user you want to associate with the experience definition
- Create any header and footer portlets you want to use to brand the experience definition

To create an experience definition you must have the following rights and privileges:

- Access Administration activity right
- Create Experience Definitions activity right
- At least Edit access to the parent folder (the folder that will store the experience)
- If you want to apply adaptive page layouts to the experience definition, at least Select access to the remote portlet web services for the adaptive page layouts
- At least Select access to the guest user you want to associate with the experience definition
- At least Select access to any header and footer portlets you want to add to the experience definition

1. Click **Administration**.
2. Open the folder in which you want to store the experience definition.

   Tip: You might want to store all of the resources needed by a particular audience of users in the same folder in which you store those users. By securing the folder appropriately and applying experience definition settings to it you can create completely separate and discreet user experiences for each audience of users.

3. In the **Create Object** drop-down list, click **Experience Definition**.
   The Experience Definition Editor opens.
4. Complete the tasks on **Experience Definition Features** page:

   - *Associating Folders with an Experience Definition* on page 79
   - *Selecting the Portal Menus and Home Page for an Experience Definition* on page 81

5. Click the **Choose Header, Footer & Style** page and complete the following task:

   - *Branding Experience Definitions with Headers and Footers* on page 82

6. Click the **Edit Navigation Options** page and complete the following tasks:

   - *Selecting a Navigation Scheme for an Experience Definition* on page 82
   - *Defining Mandatory Links to Display in an Experience Definition* on page 84

7. Click the **Login Settings** page and complete the following task:

   - *Defining the Guest User Experience for an Experience Definition* on page 85

- *Disabling Single Sign-On (SSO) for an Experience Definition* on page 86

8. Click the **Adaptive Page Layout Settings** page and complete the following task:
   - *Applying Adaptive Page Layouts* on page 86
9. Click the **Properties and Names** page and complete the following tasks:
   - *Naming and Describing an Object* on page 217

     You can instead enter a name and description when you save this experience definition.
   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219(optional)
10. Click the **Security** page and complete the following task:
    - *Setting Security on an Object* on page 221

Specifying a User Experience for Users in a Folder

You can specify a user experience for users in a folder by associating an experience definition with the folder.

Note: Users will see the associated experience definition only if no other experience rules apply.

You can associate an experience definition with a folder in the Folder Editor or in the Experience Definition Editor.

- *Applying an Experience Definition to a Folder* on page 80
- *Associating Folders with an Experience Definition* on page 79

Note: When a folder has been associated with an experience definition, the icon representing the folder changes to  .

Tip: You might want to store all of the resources needed by a particular audience of users in the same folder in which you store those users. By securing the folder appropriately and applying experience definition settings to it you can create completely separate and discreet user experiences for each audience of users.

Associating Folders with an Experience Definition

You can specify a user experience for users in a folder by associating an experience definition with the folder.

Note: Users will see the associated experience definition only if no other experience rules apply.

1. If the Experience Definition Editor is not already open, open it now.
2. Select the administrative folders you want to associate with this experience definition.
   - To associate an existing folder, click **Add Folder**.
   - To create a new folder, click **Create Folder**.

Note: When a folder has been associated with an experience definition, the icon representing the folder changes to .

Tip: You might want to store all of the resources needed by a particular audience of users in the same folder in which you store those users. By securing the folder appropriately and applying experience definition settings to it you can create completely separate and discreet user experiences for each audience of users.

Applying an Experience Definition to a Folder

You can specify a user experience for users in a folder by associating an experience definition with the folder.

Note: Users will see the associated experience definition only if no other experience rules apply.

1. Click **Administration**.
2. Open the folder to which you want to apply an experience definition in Folder Editor.
3. Click the **Experience Definition Settings** page.
4. In the drop-down list, select the experience definition you want to apply to users stored in this folder.
5. To change the settings for the selected experience definition, click **Edit Profile**.
   This opens the Experience Definition Editor.

Note: When a folder has been associated with an experience definition, the icon representing the folder changes to .

Tip: You might want to store all of the resources needed by a particular audience of users in the same folder in which you store those users. By securing the folder appropriately and applying experience definition settings to it you can create completely separate and discreet user experiences for each audience of users.

Selecting the Portal Menus and Home Page for an Experience Definition

For each experience definition you can specify which portal menus appear and which page users should see when they log in to the portal.

1. If the Experience Definition Editor is not already open, open it now.
2. In the **Enable** column, select the menus (and associated features) you want to include in this experience definition.
   For example, to include the My Pages menu and features, select the box associated with **My Pages**.

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

3. In the **Home** column, select the portal area you want to display when users log in to the portal. For example, to display a particular community when a user logs in, select the button associated with **Communities**.
4. If you selected **Communities** as the home page you must also select a particular community.

   - To choose a home community from existing communities, click **Choose Home Community**.
   - To create a new home community, click **Create Home Community**.

   Note: Users viewing this experience definition must have at least Read access to the community you choose, or they will receive an error after logging in.

5. If you selected the **Knowledge Directory** as the home page you must also select a particular folder.

   - To choose a home folder from existing folders, click **Choose Home Folder**.
   - To create a new home folder, click **Create Home Folder**.

   Note: Users viewing this experience definition must have at least Read access to the folder you choose, or they will receive an error after logging in.

6. If you enabled the **Knowledge Directory**, under **Include these Knowledge Directory Features**, select which related object features you want to display in the Knowledge Directory.

   By default, objects specified as related to a Knowledge Directory folder display to users viewing that folder. To hide a type of related object, clear the associated check box.

Branding Experience Definitions with Headers and Footers

The Choose Header, Footer & Style page of the Experience Definition Editor enables you to add special branding portlets to an experience definition (as well as change the color scheme) to control what certain groups of users see at the top and bottom of portal pages.

You must create the header and footer portlets you want to use before branding the experience definition.

1. If the Experience Definition Editor is not already open, open it now.
2. Click the **Choose Header, Footer & Style** page.
3. In the **Default Style** drop-down list, select a color scheme.

   Note: If you are using adaptive page layouts, the layout will override the style selected here.

4. Under **Add Header**, select the header portlet you want to apply to the experience definition:
   - To add or change the header, click **Add Header**.
   - To remove the header, select it, then click .
5. Under **Add Footer**, select the footer portlet you want to apply to the experience definition:
   - To add or change the header, click **Add Footer**.
   - To remove the footer, select it, then click .

Note: The community template header and footer can be set to override the experience definition header and footer.

Selecting a Navigation Scheme for an Experience Definition

For each experience definition you can specify a default navigation style to define the menu layout and core navigation structure most appropriate for your bandwidth constraints, browser requirements, design needs, deployment size, and end-user expectations.

1. If the Experience Definition Editor is not already open, open it now.
2. Click the **Edit Navigation Options** page.
3. Under **Navigation Type**, choose a navigation scheme.

- **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.
- **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.
- **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.
- **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.
- **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.
- **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.
- **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 you are using adaptive page layouts or when navigation is controlled by portlets (usually header or footer portlets) using navigation

tags. Adaptive page layouts and navigation tags provide developers a faster, easier way to customize navigation than modifying the other available navigation schemes.

Note:

- If you have written your own navigation styles, they should also be available on this page.
- Vertical navigation styles lessen the page width available for portlets on My Pages and community pages.
- If you have selected any navigation option other than Portlet-Ready Navigation, do not use the default adaptive page layouts available with the portal. If you use the default adaptive page layouts with other navigation options, users will see two methods of navigation.
- The experience definition you log into might have a different navigation style than the experience definition you are creating. To make sure that the experience definition you are creating has the appropriate appearance, log in as a user that sees that experience definition.

If you selected **Mandatory Links Only**, you must now define the mandatory links. See *Defining Mandatory Links to Display in an Experience Definition* on page 84.

If you selected **Portlet-Ready Navigation**, you must select the header and footer and/or the adaptive page layout settings that define your navigation. See *Branding Experience Definitions with Headers and Footers* on page 82 and *Applying Adaptive Page Layouts* on page 86.

Defining Mandatory Links to Display in an Experience Definition

For each experience definition you can define links to web pages, experts, documents, and community pages that are displayed to users as part of the navigation.

1. If the Experience Definition Editor is not already open, open it now.
2. Click the **Edit Navigation Options** page.
3. Under **Edit Links**, add and modify links to web pages, experts, documents, and community pages.

   - Click **Add Links** to add links to web pages.
   - Click **Add Experts** to add links to experts.
   - Click **Add Documents** to add links to documents and document folders in the Knowledge Directory.
   - Click **Add Pages** to add links to community pages.

- To remove links, select the links you want to delete and click  .

  To select or clear all link boxes, select or clear the box under **Navigation Link Heading**.
- If these links are shared with another experience definition and you do not want to share them, click  **Separate**.

  If this is a copy of another experience definition or if this experience definition has been copied, you see a warning that the navigation links are shared between the experience definitions. Any changes you make to these links are reflected in the linked experience definitions.
- To change the menu heading that displays to users, type the text in the **Navigation Link Heading** box.

  When you add navigation links, they display in a new menu, similar to the My Pages menu.

Note: You might have access to resources to which members of your experience definition do not have access. If users do not have access to a resource listed as a navigation link, they will not see the link.

Defining the Guest User Experience for an Experience Definition

For each experience definition you can associate a guest user, which lets you define the initial page an unauthenticated user sees when coming to this experience definition. For example, if an unauthenticated user is directed to this experience definition (through application of an experience rule), you can choose to have that user see the My Page layout of the guest user you associate with this experience definition, even if the experience definition is not set to display My Pages. You can also specify what page users see when they log out of an experience definition.

1. If the Experience Definition Editor is not already open, open it now.
2. Click the **Login Settings** page.
3. Under **Guest Settings**, click **Select a Guest User**, and choose a guest user to associate with this experience definition.
4. Under **Login Page Settings**, select what page unauthenticated users should see when they access this experience definition.
   - To display the default My Page for the selected guest user, select **Guest MyPage**.
   - To display the default login page shared across all experience definitions, select **Default Login Page**.

Unauthenticated users viewing this experience definition will be directed to the page you select here if, in the portal configuration file, GuestRedirectToLogin is set to 1, or if they click **Log In**. Otherwise, unauthenticated users see the home page selected for this experience definition.

Note: If you selected **Guest MyPage** make sure that the Portal Login portlet displays on the selected guest user's default My Page so that unauthenticated users can log in.

5. Under **Login Page Settings**, select what page users should see when they log out of this experience definition.
   - To display the default My Page for the selected guest user, select **Guest MyPage**.
   - To display the default login page shared across all experience definitions, select **Default Login Page**.

   Users will be directed to the page you select here if, in the portal configuration file, RedirectOnLogout is set to 1. Otherwise, upon logout, users see the home page specified for this experience definition.

   Note: If you selected **Guest MyPage** make sure that the Portal Login portlet displays on the selected guest user's default My Page so that unauthenticated users can log in.

Disabling Single Sign-On (SSO) for an Experience Definition

You can override portal SSO settings for users in this experience definition. Otherwise the SSO settings in the portal configuration file will apply.

1. If the Experience Definition Editor is not already open, open it now.
2. Click the **Login Settings** page.
3. Under **Disable Single Sign On (SSO)**, specify whether you want to disable SSO for this experience definition.

   Select **Disable SSO Setting** to override portal SSO settings for users in this experience definition. Otherwise the SSO settings in the portal configuration file will apply.

   Note: This check box is unavailable (grayed-out) if the "SSOVendor" setting is 0 in the portal configuration file.

Applying Adaptive Page Layouts

For each experience definition you can specify whether to use adaptive page layouts to display the user interface or use a legacy user interface used in previous versions of the portal.

The tasks described below assume that adaptive page layouts have not been disabled in the portal configuration file and that you have already created the adaptive page layouts as described in the *Adaptive Page Layouts* section of the *AquaLogic User Interaction Development Guide*.

Before you apply adaptive page layouts to an experience definition, you must create remote portlet web services for the adaptive page layouts.

To apply adaptive page layouts you must have at least Select access to the remote portlet web services for the adaptive page layouts.

1. If the Experience Definition Editor is not already open, open it now.
2. Click the **Adaptive Page Layout Settings** page.
3. Under **Adaptive Page Layout Mode**, specify if this experience definition should display adaptive page layouts.

   To enable adaptive page layouts for this experience definition, select **Enable Adaptive Page Layout Mode**.

   Note:

   - This setting is not available if adaptive layouts are disabled in the portal configuration files.
   - This setting controls only adaptive page layouts, not adaptive portlet layouts, which are controlled by the page layouts selected for My Pages and community pages.
   - If you disable adaptive page layouts, this experience definition will display a legacy user interface used in previous versions of the portal.

4. If you enabled adaptive page layouts, under **Layout Chooser for Page Layouts**, specify the layouts you want to display for each page layout type.

   - In the **Base Page Layouts** drop-down list, select the layout for components that are common to each page (header, footer, navigation, content area).

     The layouts listed in this drop-down list correspond to the remote portlet web services stored in the **Page Layouts/Base Page Layouts** administrative folder.
   - In the **Profile Page Layouts** drop-down list, select the layout for components that are common to each user profile page (header, footer, navigation, content area).

     The layouts listed in this drop-down list correspond to the remote portlet web services stored in the **Page Layouts/Profile Page Layouts** administrative folder.
   - In the **Knowledge Directory Layouts** drop-down list, select the layout for the content area of the Directory.

The layouts listed in this drop-down list correspond to the remote portlet web services stored in the **Page Layouts/Knowledge Directory Page Layouts** administrative folder.

Note: The common components of the Directory are specified in the base page layout.

- In the **Search Results Layouts** drop-down list, select the layout for the content area of the search results.

  The layouts listed in this drop-down list correspond to the remote portlet web services stored in the **Page Layouts/Search Results Page Layouts** administrative folder.

  Note: The common components of the search results are specified in the base page layout.

- In the **Portlet Selection Layouts** drop-down list, select what layout to use for the pop-up or fly-out editor used to select portlets.

  The layouts listed in this drop-down list correspond to the remote portlet web services stored in the **Page Layouts/Portlet Selection Page Layouts** administrative folder.

If you have not already done so, you must select **Portlet Ready Navigation** on the **Edit Navigation Options** page of this editor.

About Branding with Header and Footer Portlets

Branding portlets customize the look of your portal 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.

Note: The easiest way to apply branding (and control the look of your user interface) is with adaptive page layouts, which are then applied through experience definitions.

There are several ways to brand your portal:

- There are two example branding portlets that utilize adaptive tags and are installed with the portal: Layout Footer Portlet and Layout Header Portlet. For more information on adaptive tags and the example portlets, see the *Adaptive Page Layouts* section of the *AquaLogic User Interaction Development Guide*.
- You can also create your own custom branding portlets. There are two example branding portlets that are installed with the portal: Classic Footer Portlet and Classic Header Portlet.

For information on creating custom branding portlets, see the *AquaLogic User Interaction Development Guide*.

- With your AquaLogic Interaction license, you can download and install AquaLogic Interaction Publisher, which provides three branding portlet templates that enable you to customize the look and feel of experience definitions and communities: Header Portlet, Footer Portlet, and Content Canvas. You create and configure customized branding portlets from these portlet templates. You can customize the properties, HTML, and default values for the portlet. For information on creating Publisher branding portlets, see the *Administrator Guide for AquaLogic Interaction Publisher* or the Publisher online help.

About Header and Footer Portlet Precedence

You apply header and footer portlets at the community template, community, or experience definition level. The community template settings determine what header and footer are used in the community.

- The community template can force the community to use the experience definition branding.
- The community template can include its own branding that cannot be overridden in the community. This branding overrides the experience definition branding.
- If the community template does not specify any branding restrictions, the community can include its own branding. This branding overrides the experience definition branding.

About 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*.

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.

You can select one of the following navigation schemes for each experience definition you create:

- **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.
- **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.
- **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.
- **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.
- **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.
- **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.
- **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 you are using adaptive page layouts or when navigation is controlled by portlets (usually header or footer portlets) using navigation tags. Adaptive

page layouts and navigation tags provide developers a faster, easier way to customize navigation than modifying the other available navigation schemes.

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.

About Controling the Initial Portal Experience

AquaLogic Interaction includes several features that work together to control users' initial portal experience, such as the user interface and access to content.

Feature	How the Feature is Applied to Users			
	Created Manually	**Self-Registered**	**Imported Through an Authentication Source**	**Created Through Acceptance of an Invitation**
Default Profiles Each user is assigned a default profile at creation. Default profiles define initial My Account settings, such as language, time zone, and portal interface type; the name and number of My Pages; and the layout of the portlets on those My Pages.	Automatically assigned the "Default Profile" created at installation	Automatically assigned the "Default Profile" created at installation	Automatically assigned the default profiles specified in the Authentication Source Editor	Automatically assigned the default profile specified in the Invitation Editor

Feature	How the Feature is Applied to Users			
	Created Manually	**Self-Registered**	**Imported Through an Authentication Source**	**Created Through Acceptance of an Invitation**
Default profiles provide an initial view of the portal, which users can then change to fit their needs.				
Group Membership	All users are automatically added to the Everyone group and can be assigned to groups manually in the User Editor or Group Editor after creation.			
The most efficient way to manage access to content is to assign access privileges to groups. The only way to assign activity rights (which control access to features) is to assign the rights to groups. You can then add new users to the appropriate groups.	Manually assigned in the User Editor during creation	(No additional membership assigned during creation; assigned only to the Everyone group)	Automatically assigned to groups based on the mappings in the Global ACL Sync Map (and any mappings that occur automatically if the authentication source category matches the domain name)	Automatically assigned to groups specified in the Invitation Editor
Mandatory Communities and Portlets	The most efficient way to manage mandatory communities and portlets is to make them mandatory for particular groups. You then add new users to the appropriate groups as mentioned in the previous entry.			
Mandatory communities are communities to				

<table>
<tr><th>Feature</th><th colspan="4">How the Feature is Applied to Users</th></tr>
<tr><th></th><th>Created Manually</th><th>Self-Registered</th><th>Imported Through an Authentication Source</th><th>Created Through Acceptance of an Invitation</th></tr>
<tr><td>which the user cannot unsubscribe. Mandatory portlets are portlets that cannot be removed from a user's My Page.</td><td colspan="4"></td></tr>
<tr><td>Experience Definitions

Experience definitions provide multiple user experiences within a single portal. An experience definition defines certain elements of a user experience, such as adaptive page layout settings, branding style, and navigation.</td><td colspan="4">All users are assigned the experience definition associated with the folder in which the user is stored. You can also use experience rules to assign experience definitions to users.</td></tr>
<tr><td></td><td>Manually create the user in the folder of your choice</td><td>Automatically created in the “Default Experience Definition” folder created at installation</td><td>Automatically created in the folders specified in the Authentication Source Editor</td><td>Automatically created in the folder specified in the Invitation Editor</td></tr>
</table>

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.

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

About Users

Portal users enable you to authenticate the people who access your portal and assign appropriate security for the documents and objects in your portal. Users can be imported from external user repositories, created through the portal, created through invitations, self-registered, or just guests (unauthenticated users).

Default Profiles

Each user is assigned a default profile at creation. Default profiles define initial My Account settings, such as language, time zone, and portal interface type; the name and number of My Pages; and the layout of the portlets on those My Pages. Default profiles provide an initial view of the portal, which users can then change to fit their needs.

Note: Portlet preferences, group memberships, and community memberships are not inherited by users created from default profiles.

Default profiles are defined through special users, created in the Default Profiles folder (accessed through the Default Profiles Utility). These special users cannot log in to the portal. They are solely used to assign settings to new users.

Users Imported From External User Repositories

You can use authentication sources to import users that are already defined in your enterprise in existing user repositories, such as Active Directory or LDAP servers. After users are imported, you can authenticate them with the credentials from those user repositories. You can also import user information (such as name, address, or phone number), which can then be used to populate user profiles or can be passed to content crawlers, remote portlets, or federated searches as user information.

Users Created Through Invitations

You can invite users to your portal through invitations, making it easy for them to create their own accounts and letting you customize their initial portal experiences with content that is of particular interest to them.

Self-Registered Users

Users can create their own accounts through your portal by clicking **Create an account** on the login page. These users are stored in the **Default Experience Definition** portal folder and are included in the AquaLogic Interaction Authentication Source. They are automatically given security privileges based on the "Default Profile" created at installation. Based on this security, users can personalize their views of the portal with My Pages, portlets, and community memberships, and can view portal content.

Note: Your system administrator can disable the **Create an account** functionality.

Guest Users

The portal lets you create multiple guest users. This is useful when you want to have different user experiences for different sets of unauthenticated users. You can accomplish this by creating a guest user for each group of unauthenticated users that you want to see a different user experience. You then associate each guest user with a different experience definition, customize the My Page for each guest user, and use experience rules to direct the guest users to the appropriate experience definition.

For example, you could create one guest user for employees that have not yet logged in to the portal and one for customers visiting your portal. The My Page for the employee guest user would include the login portlet so employees can log in. The My Page for customers might include

information about your company, such as contact numbers and descriptions of your products or services. You would create two experience definitions, associating one guest user with each. Then you would create two experience rules that would direct users to the appropriate experience definition based on the URL they use to access your portal.

Creating Default Profiles to Customize a Users Initial Portal Experience

When new authenticated users are created in the portal, the following settings are based on default profiles: initial My Account settings, name and number of My Pages, and layout of the portlets on those My Pages.

To create a default profile you need the following rights:

- Access Administration activity right
- Access Utilities activity right

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **Default Profiles**.
   The Default Profiles folder opens.
3. In the **Create Object** drop-down list, click **User**.
4. In the **Login Name** box, type a name for this default profile.

   Users created from this default profile will have their own user names and passwords.

   Note:

   - Do not select **This is a guest account**. Instead, to create a guest user, go to a different administrative folder, create a user there, and make that user a guest.
   - Do not add this user to any groups. Group memberships are not inherited by users created from default profiles. You set group membership through invitations or authentication sources.

After you have created a default profile, edit its layout.

Customizing a Default Profile Experience

When new authenticated users are created in the portal, the following settings are based on default profiles: initial My Account settings, name and number of My Pages, and layout of the portlets on those My Pages.

To customize a default profile experience you need the following rights:

- Access Administration activity right
- Access Utilities activity right

1. If you are not already in the Default Profiles folder, click **Administration**, and, in the **Select Utility** drop-down list, click **Default Profiles**.
2. Select the profile that you want to customize.
3. Click  **Edit Profile Layout**.
4. Specify My Account settings, create and delete My Pages, and change the layout of the My Pages.

   Note:

   - Portlet preferences are not inherited by users created from the default profile. Users set their own preferences.
   - Community membership and access to documents and objects are granted through group membership.

After you have customized the default profile, use invitations and authentication sources to assign the profile to new portal users and to assign group membership.

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

You can automatically lock user accounts based on 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

Manually Locking User Accounts

You can manually lock user accounts through the User Editor.

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

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.

You unlock user accounts differently depending on how the account was locked:

- Admin Lock: A portal administrator locked the user account.
- 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.

  Note: 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.

  Note: Users can remove the lock by specifying the correct credentials the next time they log in.

- To remove an Admin Lock or an Automatic Lock with the Release Disabled Logins Utility:
  a) Click **Administration**.
  b) In the **Select Utilities** drop-down list, click **Release Disabled Logins**.
- To remove an Admin Lock or an Automatic Lock with the User Editor:
  a) Click **Administration**.
  b) Navigate to the user whose account you want to unlock and click the user name.
  c) Clear the check box next to **Disable Login**.

- To remove an Agent Locks for all affected users:
  a) Click **Administration**.
  b) Navigate to the authentication source and click its name.
  c) Click **Fully Synchronized Groups** page.
  d) Click **Re-Enable Users**.
  Unlocking these accounts may take a few minutes.

Deleting a User

You should delete users that should no longer have access to your portal.

To delete a user you must have the following rights and privileges:

- Access Administration activity right
- Admin access to the user

- To delete a user:
  a) Click **Administration**.
  b) Navigate to the user.
  c) Select the user you want to delete and click  .
- To delete a user whose account is locked:
  a) Click **Administration**.
  b) In the **Select Utilities** drop-down list, click **Release Disabled Logins**.
  a) Select the user you want to delete and click  .

About Groups

Groups are sets of users, sets of other groups, or both. Groups enable you to more easily control security because you assign each group different activity rights and access privileges. Groups are created in the portal either by adding them individually as portal objects, or by synchronizing with authentication sources (user repositories such as LDAP or Active Directory).

Membership to a group is determined in two ways:

- Members are explicitly defined as specific users and/or other groups on the Group Memberships page.
- Members are dynamically determined based on rules you set up on the Dynamic Membership Rules page.

Dynamic Group Membership

You might want to have users automatically added to or removed from groups based on properties in their user profiles or other group membership. This is called dynamic group membership. For example, you might want to give users access to a community based on their location, title, department, or any other property in their profile. If you have a community for all the branches in Texas, you could set up a rule that states that all employees in Texas are part of the group. If an employee moves to Arizona, and the "State" property in her profile changes, the employee no longer satisfies this rule.

Community Groups

You can create groups inside a community 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.

Community groups are available only within the community. However, you can make a community group available outside of the community by moving the group to a non-community administrative folder.

Roles

A role is not a portal object; it is an association between a group and the activity rights required to perform a job function. For example, the Knowledge Directory administrator role is not an object you define; it relates to administrative responsibilities for those who manage content in the Knowledge Directory.

Before you create portal groups for the purpose of assigning roles, you should familiarize yourself with the definition and scope of the administrative tasks you plan to delegate and the activity rights needed to complete those administrative tasks. Some users will handle many tasks, but those tasks might actually encompass several roles. Before creating a role to cover all these tasks, consider if there are situations where the tasks will be broken down into smaller roles. You can easily assign more than one role to a user.

Groups Created Upon Installation

The following groups are created in the **Portal Resources** folder when you install the portal:

- **Administrators Group**: This group provides full access to everything in the portal: all objects, all utilities, and all portal activities.
- **Everyone**: This group includes all portal users, whether created manually through the administration menu, imported from authentication sources, created through acceptance of an invitation, or created through the Create an Account page.

Example Roles

A role is not a portal object; it is an association between a group and the activity rights required to perform a job function.

This following table describes the activity rights that are defined by default during installation and provides an example map between activity rights and administrative roles. In the example, the role called Content Administrator provides the activity rights required to populate the portal with document records crawled from remote content sources; a separate role called Knowledge Directory Administrator provides the activity rights required to create Knowledge Directory structure. Although some users might fill both roles, others might not. By creating two separate roles, you can assign the roles separately or together.

Role	Activity Rights Needed
Portal Administrator: Manages all areas of the portal	All activity rights
Content Administrator: Populates 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
Knowledge Directory Administrator: Creates Knowledge Directory structure and approves content	Access Smart Sort, Access Unclassified Documents, Access Utilities, Advanced Document Submission, Create Filters, Create Folders, Edit Knowledge Directory, and Self-Selected Experts

Creating and Adding Members to a Group

Groups are sets of users, sets of other groups, or both. Groups enable you to more easily control security because you assign each group different activity rights and access privileges.

To create a group you must have the following rights and privileges:

- Access Administration activity right
- Create Groups activity right
- At least Edit access to the parent folder (the folder that will store the group)
- At least Select access to any groups to which you want to add this group
- At least Select access to any users you want to add to the group

1. Click **Administration**.
2. Open the folder in which you want to store the group.
3. In the **Create Object** drop-down list, click **Group**.
4. Under **Parent Group Memberships**, specify the groups to which this group should be a member:
   - To make this group a member of another group, click **Add Group**, in the Select Groups dialog box, select the groups to which you want to add this group, and click **OK**.
   - To remove a parent group, select it and click .

     To select or clear all of the group boxes, select or clear the box to the left of **Members**.
   - To toggle the order in which the groups are sorted, click **Members**.
5. Under **Group Members**, specify the members of this group:
   - To add members to this group, click **Add User/Group**, in the Select Members dialog box, select the groups and users you want to add to this group, and click **OK**.
   - To remove a member, select it and click .
   - To remove a member, select it and click .

     To select or clear all of the member boxes, select or clear the box to the left of **Members**.
   - To toggle the order in which the members are sorted, click **Members**.

If you want users and groups to be added to this group based on user profile properties or group membership, set dynamic membership rules. If you want members of this group to be able to access administration, create objects, or perform other activities that require special rights, assign activity rights to the group.

Configuring Dynamic Group Membership

You might want to have users automatically added to or removed from groups based on properties in their user profiles or other group membership. This is called dynamic group membership. For example, you might want to give users access to a community based on their location, title, department, or any other property in their profile. If you have a community for all the branches in Texas, you could set up a rule that states that all employees in Texas are part of the group. If an employee moves to Arizona, and the "State" property in her profile changes, the employee no longer satisfies this rule.

Dynamic membership rules are made up of statements that define what must or must not be true to include a user in the group. The statements are collected together in groupings. The grouping defines whether the statements are evaluated with an AND operator (all statements are true) or an OR operator (any statement is true). If some statements should be evaluated with an AND operator and some should be evaluated with an OR operator, you can create separate groupings for the statements. You can also create subgroupings or nested groupings, where one grouping is contained within another grouping. The statements in the lowest-level grouping are evaluated first to define a set of users. Then the statements in the next highest grouping are applied to that set of users to further filter the set of users. The filtering continues up the levels of groupings until all the groupings of statements are evaluated.

1. Open the Group Editor by creating a new group or editing an existing one.
2. Click the **Dynamic Membership Rules** page.
3. Select the operator for the grouping of statements you are about to create:
   - If a user should be added to the group only when all statements in the grouping are true, select **AND**.
   - If a user should be added to the group when any statement in grouping is true, select **OR**.

   Note: The operator you select for a grouping applies to all its statements and subgroupings directly under it.
4. Define each statement in the grouping:
   a) Click **Add Statement**.
   b) In the first drop-down list, select a property.

      This list includes the properties included in the user profile and **Member Of**, which enables you to select a group whose members you want to include or exclude.
   c) In the second drop-down list, select an operator:

- If you selected a user profile property, you can select **Contains** or **Contains No Value**.
- If you selected **Member Of**, you can select **includes** or **excludes**.

d) If you selected **Contains** as the operator, in the text box, enter a value for the property. You can use wildcards.

e) If you selected **Member Of**, select the groups whose members you want to include or exclude. Click , in the Group Chooser dialog box, select a group, and click **OK**.

Note: The Group Chooser dialog box displays only statically defined groups.

- To add more statements, repeat these steps.
- To remove the last statement in a grouping, select the grouping and click **Remove Statement**.

5. If necessary, add more groupings:

- To add another grouping, select the grouping to which you want to add a subgrouping and click **Add Grouping**. Then define the statements for that grouping.

  Note: You cannot add a grouping at the same level as **Grouping 1**.

- To remove a grouping, select the grouping, and click **Remove Grouping**.

  Note:

  - Any groupings and statements in that grouping will also be removed.
  - You cannot remove the top level **Grouping 1**.

6. Click **Preview Members** to see the dynamic members resulting from the rules you defined. Only 1000 members will be displayed.

The dynamic members are updated for this group when you click **Finish**.

The next time you open this group editor, dynamic members are displayed on the **Group Memberships** page.

Dynamic memberships are updated for all groups as part of the **Dynamic Membership Update Agent** job (located in the **Intrinsic Operations** folder). When user profile data changes, the resulting dynamic group membership changes are updated as part of this job.

Assigning Activity Rights to a Group

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

It is not necessary to grant a user the right to create a type of object for that user to manage an object of that type. Management of an object is based solely on a user's access privilege to that object.

1. If the Group Editor is not already open, open it now and display the **Activity Rights** page.
2. Under **Activity Rights**, click  **Add Activity Rights**.
   The Select Activity Rights dialog box opens.
3. Select the activity rights you want to grant to the group and click **OK**.

   For example, if you select Create Jobs, the members of the group will be able to create jobs in the portal.

   To remove activity rights, select the activity right that you want to remove and click  .

Under **Inherited Activity Rights** you see any activity rights granted to the parent groups of this group.

About Importing and Authenticating Users with Authentication Sources

Authentication sources enable you to import users, groups, and group memberships that are already defined in your enterprise in existing user repositories, such as Active Directory or LDAP servers. After users are imported, you can authenticate them with the credentials from those user repositories.

Authentication Providers

An authentication provider is a piece of software that tells the portal how to use the information in the external user repository. BEA provides authentication providers as part of the AquaLogic Interaction Identity Services. The AquaLogic Interaction Identity Service - LDAP is used to import

and authenticate users and group from LDAP servers. The AquaLogic Interaction Identity Service - Active Directory is used to import and authenticate users and groups from Active Directory servers. If your users and groups reside in a custom system, such as a custom database, you can import and authenticate them by writing your own authentication provider using the IDK.

Note:

- Your portal administrator must install the authentication provider before you can create the associated authentication web service. For information on obtaining authentication providers, contact *ALUIsupport@bea.com*. For information on installing authentication providers, refer to the *Installation Guide for AquaLogic Interaction* (available on *edocs.bea.com*) or the documentation that comes with your authentication provider, or contact your portal administrator.
- To learn about developing your own authentication provider, refer to the *BEA AquaLogic User Interaction Development Center*.

Authentication Web Services

Authentication web services enable you to specify general settings for your external user repository, leaving the more detailed settings (like domain specification) to be set in the associated remote authentication sources. This allows you to create different authentication sources to import each domain without having to repeatedly specify all the settings.

Authentication Sources

Authentication sources can import users and/or groups, authenticate imported users, or both import and authenticate. Your security needs determine how many authentication sources to create and what functionality they need. You might be able to create just one authentication source that imports and authenticates all users and groups, but here are a couple examples of when that would not suffice:

- If you want to use single sign-on (SSO), create a synchronization-only authentication source.
- If you want to distinguish users and groups from different domains, create separate synchronization-only authentication sources for each domain, and create an authentication-only authentication source to authenticate users from all domains (assuming they are from the same user repository).

  This enables you to store users and groups imported from different domains in different portal folders or to create separate users or groups with the same name but from different domains.

If you are importing users and groups into the portal, you run a job for the initial import and then continue to run the job periodically to keep the users and groups in the portal synchronized with those in the source user repository.

Note: When you run the job to import users and groups, the portal also creates a group that includes all users imported through the authentication source. This group is named after the authentication source; for example, if your authentication source is called *mySource*, the group would be called *Everyone in mySource*.

How Authentication Works

When you use authentication sources to authenticate portal users, the user credentials are left in the external repository; they are not stored in the portal database. When someone attempts to log in to your portal through an imported user account, the portal confirms the password with the external repository. This means that the user's portal password always matches the password in the external repository. For example, if a user with a portal account imported from Active Directory changes the Active Directory password, the user can immediately log in to the portal with that password. If the user is already logged in to the portal, the user must log in again with the new password, because the portal will no longer be able to recognize the old password.

AquaLogic Interaction Authentication Source

The AquaLogic Interaction Authentication Source is automatically created upon installation. It is the authentication source used for users stored in the portal database (users created upon install, users created manually through the portal, and self-registered users). This authentication source cannot be modified or deleted.

Creating an Authentication Web Service

Authentication web services enable you to specify general settings for your external user repository, leaving the more detailed settings (like domain specification) to be set in the associated remote authentication sources. This allows you to create different authentication sources to import each domain without having to repeatedly specify all the settings.

Before you create an authentication web service, you must:

- Install the authentication provider on the computer that hosts the portal or on another computer
- Create a remote server pointing to the computer that hosts the authentication provider (optional, but recommended)

To create an authentication web service you must have the following rights and privileges:

- Access Administration activity right
- Create Web Service Infrastructure activity right
- At least Edit access to the parent folder (the folder that will store the authentication web service)
- At least Select access to the remote server that the authentication web service will use

1. Click **Administration**.
2. Open the folder in which you want to store the authentication web service.
3. In the **Create Object** drop-down list, click **Web Service — Authentication**.
   The Authentication Web Service Editor opens.
4. On the **Main Settings** page, complete the following task:
   -
5. Click the **HTTP Configuration** page and complete the following task:
   -
6. Click the **Advanced Settings** page and complete the following task:
   -
7. Click the **Authentication Settings** page and complete the following task:
   -
8. Click the **Debug Settings** page and complete the following task:
   -
9. Click the **Properties and Names** page and complete the following tasks:
   - *Naming and Describing an Object* on page 217

     You can instead enter a name and description when you save this authentication web service.
   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219 (optional)

The default security for this authentication web service is based on the security of the parent folder. You can change the security when you save this authentication web service (on the **Security** tab page in the Save As dialog box), or by editing this authentication web service (on the **Security** page of the Authentication Web Service Editor).

Portal administrators with at least Select access to this authentication web service can create authentication sources based on the web service.

Mapping External Document Security to Imported Portal Users with the Global ACL Sync Map

Users imported through an authentication source can automatically be granted access to the content imported by some remote content crawlers. The Global ACL Sync Map shows these content crawlers how to import source document security.

To access the Global ACL Sync Map you must be a member of the Administrators group.

For an example of how importing security works for users imported through an authentication source, see *Example of Importing Security* on page 193.

Creating an Authentication Source to Import and Authenticate Users

You can create a remote authentication source to import and authenticate users and groups from external user repositories.

Before you create an authentication source, you must:

- Install the authentication provider on the computer that hosts the portal or on another computer.
- Create a remote server that points to the computer that hosts the authentication provider.
- Create an authentication web service on which to base the authentication source.
- Create and configure the default profiles you want to apply to imported users.
- Create the folders in which you want to store the imported users.

To create an authentication source you must have the following rights and privileges:

- Access Administration activity right
- Create Authentication Sources activity right
- At least Edit access to the parent folder (the folder that will store the authentication source)
- At least Select access to the authentication web service on which this authentication source will be based
- At least Select access to the default profiles you want to apply to imported users
- At least Select access to the folders in which you want to store the imported users

1. Click **Administration**.
2. Open the folder in which you want to store the authentication source.
3. In the **Create Object** drop-down list, click **Authentication Source - Remote**.
   The Choose Web Service dialog box opens.
4. Select the web service that provides the basic settings for your authentication source and click **OK**.
   The Remote Authentication Source Editor opens.
5. On the **Main Settings** page, complete the following tasks:
   a) *Setting an Authentication Source Category to Distinguish Users and Groups Imported from a Particular Domain* on page 117
   b) *Setting Default Profiles and Target Folders for Imported Users* on page 118
   c) *Setting a Target Folder for Imported Groups* on page 119
6. Click the **Synchronization** page and complete the following tasks:
   a) Under **General Info**, select **Authentication and Synchronization**.
   b) *Specifying Which Users and Groups to Synchronize* on page 119
7. Click the **Fully Synchronized Groups** page and complete the following task:
   - *Specifying What to Do with Users and Groups Deleted from the Source User Repository* on page 122
8. Click the **Set Job** page and complete the following task:
   - *Associating an Object with a Job* on page 300
9. Click the **Properties and Names** page and complete the following tasks:
   - *Naming and Describing an Object* on page 217

     Note: The authentication source name appears in lists of objects from which users will sometimes choose; therefore, the name should clearly convey the purpose of this authentication source.

     You can instead enter a name and description when you save this authentication source.
   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219 (optional)

The default security for this authentication source is based on the security of the parent folder. You can change the security when you save this authentication source (on the **Security** tab page

in the Save As dialog box), or by editing this authentication source (on the **Security** page of the Authentication Source Editor).

Run the job you associated with this authentication source.

If you are importing only partial users or groups or are applying different default profiles to each group of users, after the associated job runs once, return to the Authentication Source Editor and perform any necessary additional tasks.

Importing Users with a Synchronization-Only Authentication Source

You can import users with an authentication source and have them authenticated through an associated authentication partner.

Before you create an authentication source, you must:

- Install the authentication provider on the computer that hosts the portal or on another computer.
- Create a remote server that points to the computer that hosts the authentication provider.
- Create an authentication web service on which to base the authentication source.
- Create and configure the default profiles you want to apply to imported users.
- Create the folders in which you want to store the imported users.
- Create an authentication source that will authenticate users imported with this authentication source.

To create an authentication source you must have the following rights and privileges:

- Access Administration activity right
- Create Authentication Sources activity right
- At least Edit access to the parent folder (the folder that will store the authentication source)
- At least Select access to the authentication web service on which this authentication source will be based
- At least Select access to the authentication source that will authenticate users imported with this authentication source.

1. Click **Administration**.
2. Open the folder in which you want to store the authentication source.
3. In the **Create Object** drop-down list, click **Authentication Source - Remote**.
   The Choose Web Service dialog box opens.

4. Select the web service that provides the basic settings for your authentication source and click **OK**.
   The Remote Authentication Source Editor opens.
5. On the **Main Settings** page, complete the following tasks:
   a) *Setting an Authentication Source Category to Distinguish Users and Groups Imported from a Particular Domain* on page 117
   b) *Setting Default Profiles and Target Folders for Imported Users* on page 118
   c) *Setting a Target Folder for Imported Groups* on page 119
6. Click the **Synchronization** page and complete the following tasks:
   a) Under **General Info**, select **Synchronization with Authentication Partner**.
   b) In the **Authentication Partners** drop-down list, select the authentication source you want to use for authentication.

      Note: If the authentication partner is unavailable, this authentication source will attempt to authenticate users.

   c) *Specifying Which Users and Groups to Synchronize* on page 119
7. Click the **Fully Synchronized Groups** page and complete the following task:
   - *Specifying What to Do with Users and Groups Deleted from the Source User Repository* on page 122
8. Click the **Set Job** page and complete the following task:
   - *Associating an Object with a Job* on page 300
9. Click the **Properties and Names** page and complete the following tasks:
   - *Naming and Describing an Object* on page 217

     Note: The authentication source name appears in lists of objects from which users will sometimes choose; therefore, the name should clearly convey the purpose of this authentication source.

     You can instead enter a name and description when you save this authentication source.
   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219 (optional)

The default security for this authentication source is based on the security of the parent folder. You can change the security when you save this authentication source (on the **Security** tab page

in the Save As dialog box), or by editing this authentication source (on the **Security** page of the Authentication Source Editor).

Run the job you associated with this authentication source.

If you are importing only partial users or groups or are applying different default profiles to each group of users, after the associated job runs once, return to the Authentication Source Editor and perform any necessary additional tasks.

Authenticating Users with an Authentication-Only Authentication Source

If you have more than one authentication source importing users from the same user repository, create an authentication-only authentication source to authenticate your users.

Before you create an authentication source, you must:

- Install the authentication provider on the computer that hosts the portal or on another computer.
- Create a remote server that points to the computer that hosts the authentication provider.
- Create an authentication web service on which to base the authentication source.

To create an authentication source you must have the following rights and privileges:

- Access Administration activity right
- Create Authentication Sources activity right
- At least Edit access to the parent folder (the folder that will store the authentication source)
- At least Select access to the authentication web service on which this authentication source will be based

1. Click **Administration**.
2. Open the folder in which you want to store the authentication source.
3. In the **Create Object** drop-down list, click **Authentication Source - Remote**.
   The Choose Web Service dialog box opens.
4. Select the web service that provides the basic settings for your authentication source and click **OK**.
   The Remote Authentication Source Editor opens.
5. On the Main Settings page, complete the following task:

   - *Setting an Authentication Source Category to Distinguish Users and Groups Imported from a Particular Domain* on page 117

6. Click the **Synchronization** page and , under **General Info**, select **Authentication Only**.
7. Click the **Properties and Names** page and complete the following tasks:

   - *Naming and Describing an Object* on page 217

     Note: The authentication source name appears in lists of objects from which users will sometimes choose; therefore, the name should clearly convey the purpose of this authentication source.

     You can instead enter a name and description when you save this authentication source.

   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219 (optional)

The default security for this authentication source is based on the security of the parent folder. You can change the security when you save this authentication source (on the **Security** tab page in the Save As dialog box), or by editing this authentication source (on the **Security** page of the Authentication Source Editor).

Add this authentication source as the authentication partner for a synchronization-only authentication source.

Importing Users for Single Sign-On (SSO)

You can import users with an authentication source and have them authenticated transparently through single sign-on (SSO).

Before you create an SSO authentication source, you must:

- Install the authentication provider on the computer that hosts the portal or on another computer.
- Create a remote server that points to the computer that hosts the authentication provider.
- Create an authentication web service on which to base the authentication source.
- Create and configure the default profiles you want to apply to imported users.
- Create the folders in which you want to store the imported users.

To create an SSO authentication source you must have the following rights and privileges:

- Access Administration activity right
- Create Authentication Sources activity right
- At least Edit access to the parent folder (the folder that will store the authentication source)
- At least Select access to the authentication web service on which this authentication source will be based

1. Click **Administration**.
2. Open the folder in which you want to store the authentication source.
3. In the **Create Object** drop-down list, click **Authentication Source - Remote**.
   The Choose Web Service dialog box opens.
4. Select the web service that provides the basic settings for your authentication source and click **OK**.
   The Remote Authentication Source Editor opens.
5. On the **Main Settings** page, complete the following tasks:
   a) *Setting an Authentication Source Category to Distinguish Users and Groups Imported from a Particular Domain* on page 117
   b) *Setting Default Profiles and Target Folders for Imported Users* on page 118
   c) *Setting a Target Folder for Imported Groups* on page 119
6. Click the **Synchronization** page and complete the following tasks:
   a) Under **General Info**, select **Synchronization with Authentication Partner**.
   b) In the **Authentication Partners** drop-down list, select **SSO Authentication Source**.
   c) *Specifying Which Users and Groups to Synchronize* on page 119
7. Click the **Fully Synchronized Groups** page and complete the following task:
   - *Specifying What to Do with Users and Groups Deleted from the Source User Repository* on page 122
8. Click the **Set Job** page and complete the following task:
   - *Associating an Object with a Job* on page 300
9. Click the **Properties and Names** page and complete the following tasks:
   - *Naming and Describing an Object* on page 217

     Note: The authentication source name appears in lists of objects from which users will sometimes choose; therefore, the name should clearly convey the purpose of this authentication source.

     You can instead enter a name and description when you save this authentication source.
   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219 (optional)

Run the job you associated with this authentication source.

If you are importing only partial users or groups or are applying different default profiles to each group of users, after the associated job runs once, return to the Authentication Source Editor and perform any necessary additional tasks.

If you have not already done so, you must modify the portal configuration to enable SSO.

Setting an Authentication Source Category to Distinguish Users and Groups Imported from a Particular Domain

On the Main Settings page of the Authentication Source Editor, you set the prefix you want to add to user and group names to distinguish the domain from which they were imported. For example, if you enter *myDomain*, each user name and each group name will be prefixed by the string *myDomain*; myUser becomes *myDomain\myUser* and myGroup becomes *myDomain\myGroup*.

1. If the Authentication Source Editor is not already open, open it now by creating an authentication source.

   Note: You can set the category only during authentication source creation.

2. Under **Category**, in the **Authentication Source Category** box, type the prefix you want to add to user and group names to distinguish that they were imported from this domain.

   Generally, you can set the category to any value you want, but there are a few important considerations:

   - Do not include spaces in the prefix.
   - After you create this authentication source you cannot change the category value.
   - If you are using Windows Integrated Authentication (WIA) as your single sign-on (SSO) authentication provider, your authentication source category must match the domain name.
   - You might want the authentication source category to match the domain name if you are going to import security information. Some content crawlers have the ability to import security information with the imported content, making portal security much easier to maintain. For this to work, the users with access to the imported content need to correspond to portal users, as specified in the Global ACL Sync Map. If the authentication source category matches the name of the source domain, this correspondence is automatic.
   - Multiple authentication sources can use the same category. However, because the prefix is prepended to the user and group names, you need to be certain that the domains involved do not have different users or groups with the same name. That is, if a LizaR user exists on one domain, and a LizaR user exists on another domain, they must be the same user because only one user will be created.

Setting Default Profiles and Target Folders for Imported Users

Specify which default profiles to apply to users imported by an authentication source. A default profile includes portlets, portlet preferences, My Pages, and personalization settings. By assigning a default profile to the imported users, you can control what users see when they first log in to your portal. After that, users can further personalize their views of the portal.

You must have at least Select access to the folder in which you want to store imported groups.

If the Authentication Source Editor is not already open, open it now.

- To apply the same default profile to all users imported by this authentication source, you can specify the following settings when you create the authentication source:
  a) If the Authentication Source Editor is not already open, open it now.
  b) In the **Default Profile** drop-down list, select the default profile to apply to the imported users.
  c) Under **Target Folder**, click **Browse** to select the folder in which to store the imported users.

     If you want to display an experience definition interface to the imported users when they log in, choose a folder to which the experience definition has been applied or apply the experience definition to the chosen folder before you import users.

     By default, users imported by this authentication source are stored in the same folder that stores the authentication source.

- To apply different default profiles to the users in some groups:
  a) Perform a Partial Users Synchronization to import all the groups.
  b) Return to the Authentication Source Editor.
  c) Click  **Add Group**; then, in the Add Group dialog box, select the groups to which you want to apply different default profiles and click **OK**.

     Note: To view the members of a group or edit a group, click the group name.

  d) For each group, perform the following actions:

     1. In the **Default Profile** drop-down list, select the default profile to apply to the imported users.
     2. Under **Target Folder**, click **Browse** to select the folder in which to store the imported users.

If you want to display an experience definition interface to the imported users when they log in, choose a folder to which the experience definition has been applied or apply the experience definition to the chosen folder before you import users.

By default, users imported by this authentication source are stored in the same folder that stores the authentication source.

e) Prioritize the default profiles by changing the order of the groups.

If a user is a member of more than one group in this list, the uppermost default profile is applied. If necessary, move groups up or down in the list.

After you have configured all the settings for this authentication source, you must run a job to import the users and groups.

Setting a Target Folder for Imported Groups

By default, groups imported by an authentication source are stored in the same folder that stores the authentication source, but you can select a different folder if you want.

You must have at least Select access to the folder in which you want to store imported groups.

1. If the Authentication Source Editor is not already open, open it now.
2. Under **New Groups**, click **Browse** to select the folder in which to store the imported groups. The Change Folder dialog box opens.
3. Select the select a folder and click **OK**.

After you have configured all the settings for this authentication source, you must run a job to import the users and groups.

Specifying Which Users and Groups to Synchronize

When you set an authentication source to synchronize users and/or groups from a source user repository, you can specify which users and groups to synchronize.

Note: When you synchronize users/groups, new users/groups are imported into the portal and deleted users/groups are removed from the portal.

1. If the Authentication Source Editor is not already open, open it now.
2. Click the **Synchronization** page.
3. Specify which users and groups to synchronize.

- To import all users and groups from the source domain, select **Full Synchronization**.

  Each time you run the job associated with this authentication source all users and groups will be synchronized with the portal.

- To import the users from selected groups, but not all of the users found on the source domain, perform the following steps:

  1. Select **Partial Users Synchronization**.
  2. Run the job associated with this authentication source.

     All of the groups in the source user repository are imported into the portal, but no users are imported.
  3. Return to the Authentication Source Editor and click the **Fully Synchronized Groups** page.
  4. Select the groups you want to fully synchronize.
  5. Run the job associated with this authentication source again.

     Each time you run the job associated with this authentication source all groups are synchronized, but the only users that are synchronized are the ones that are members of the fully synchronized groups.

- To import all users, but only selected groups, perform the following steps:

  1. Select **Full Synchronization** or **Partial Users Synchronization**.
  2. Run the job associated with this authentication source.
  3. Delete all unwanted groups from the portal.
  4. Return to the Authentication Source Editor and click the **Synchronization** page.
  5. Select **Partial Groups Synchronization**.
  6. Run the job associated with this authentication source again.

     Each time you run the job associated with this authentication source all users are synchronized, but no new groups are imported. Groups are still removed from the portal if they are deleted from the source user repository.

- To import selected users and selected groups, perform the following steps:

  1. Select **Partial Users Synchronization**.
  2. Run the job associated with this authentication source.

     All of the groups on the source domain are imported into the portal, but no users are imported.

3. Delete all unwanted groups from the portal.
4. Return to the Authentication Source Editor and click the **Fully Synchronized Groups** page.
5. Select the groups from which you want to import users.
6. Click the **Synchronization** page.
7. Select **Partial Users and Partial Group Synchronization**.
8. Run the job associated with this authentication source again.

   Each time you run the job associated with this authentication source the only users that are synchronized are the ones that are members of the fully synchronized groups, and no new groups are imported. Groups are still removed from the portal if they are deleted from the source user repository.

- To import no users or groups, choose **No Synchronization**.

4. If users from another authentication source are members of groups from this authentication source or vice versa, select **Import user and group memberships from other authentication sources**.
5. In the **Import batches of** box, type the number of users you want to import at a time.

   The default batch setting is 1000 users. Some databases cannot support a batch of 1000; the most common reason is that the database runs out of space in the rollback segment because it attempts to add all 1000 users within one transaction. This situation terminates the transaction, and no users are imported.

   Note: Raising the import batch number can improve the time it takes to synchronize.

Selecting Groups from Which to Import Users

The Fully Synchronized Groups page of the Authentication Source Editor enables you to choose groups from which you want to import users. The groups that you list on this page are synchronized with the corresponding groups on the source server.

Before you can select groups to fully synchronize, you must import the groups by running the authentication source in Partial Users Synchronization or Partial Users and Partial Group Synchronization mode.

1. If the Authentication Source Editor is not already open, open it now.
2. Click the **Fully Synchronized Groups** page.
3. Select groups from which to import users:

- To add a group, click **Add Group**; then, in the Add Group dialog box, select the groups you want to add and click **OK**.
- To add every group imported by this authentication source, click **Add All Groups**.
- To delete a group, select the group and click .

  To select or clear all of the group boxes, select or clear the box to the left of **Group**.
- To edit a group, click the group name.

Specifying What to Do with Users and Groups Deleted from the Source User Repository

The Fully Synchronized Groups page of the Authentication Source Editor enables you to specify what to do with users and groups deleted from the source user repository. By default the portal users are disabled and groups are moved to a folder for future deletion, but you can change this behavior.

1. If the Authentication Source Editor is not already open, open it now.
2. Click the **Fully Synchronized Groups** page.
3. To delete users rather than disabling them, clear the box next to **Disable users instead of deleting them**.
4. To delete groups rather than moving them to a folder for future deletion, clear the box next to **Defer deletion of groups instead of deleting**.
5. To change the folder in which groups deferred for deletion are stored, click **Browse** and, in the Change Folder dialog box, select the folder and click **OK**.

   By default, groups deferred for deletion are moved to a **Groups to Delete** folder in the same folder that stores the authentication source.

Editing an Authentication Source

To edit an authentication source you must have at least Edit access to it.

1. On the **Main Settings** page, perform the following tasks as necessary:
   - *Setting Default Profiles and Target Folders for Imported Users* on page 118

- *Setting a Target Folder for Imported Groups* on page 119

2. Click the **Synchronization** page and perform the following tasks as necessary:
   a) Under **General Info**, choose whether you want to use this authentication source to authenticate user credentials, import users and groups, or both:
      - To import users and groups and authenticate user credentials, choose **Authentication and Synchronization**. You must also specify what you want to synchronize (step 3).
      - To authenticate user credentials, but not import users and groups, choose **Authentication Only**.
      - To import users and groups, but use an authentication partner to authenticate user credentials, choose **Synchronization with Authentication Partner**. You must also specify the authentication partner (step 2), and what you want to synchronize (step 3).

   b) If you chose **Synchronization with Authentication Partner**, in the **Authentication Partners** drop-down list, choose the authentication source you want to use for authentication (SSO or another authentication source).

      Note: If the authentication partner is unavailable, this authentication source will attempt to authenticate users.

      To use SSO as specified in the portal configuration file, choose **SSO Authentication Source**.

   c) If you chose **Authentication and Synchronization** or **Synchronization with Authentication Partner**, specify what you want to synchronize.

      See *Specifying Which Users and Groups to Synchronize* on page 119.

3. Click the **Fully Synchronized Groups** page and perform the following tasks as necessary:
   - *Selecting Groups from Which to Import Users* on page 121
   - *Specifying What to Do with Users and Groups Deleted from the Source User Repository* on page 122

4. Click the **Set Job** page and perform the following task as necessary:
   - *Associating an Object with a Job* on page 300

5. Click the **Properties and Names** page and perform the followings tasks as necessary:
   - *Naming and Describing an Object* on page 217
   - *Localizing the Name and Description for an Object* on page 342
   - *Managing Object Properties* on page 219

- *Viewing Top Best Bets for an Object* on page 219

6. Click the **Security** page and perform the following tasks as necessary:
   - *Setting Security on an Object* on page 221
7. Click the **Migration History and Status** page and perform the following tasks as necessary:
   - *Approving an Object for Migration* on page 307
   - *Viewing Import History for an Object* on page 222

If this authentication source is set to synchronize users or groups, run the job associated with it.

About Importing User Information with Profile Sources

Profile sources allow you to import user information (such as name, address, or phone number) that is already defined in your enterprise in existing user repositories, such as Active Directory or LDAP servers. The imported user information can be used to populate user profiles or can be passed to content crawlers, remote portlets, or federated searches as user information.

Note:

- You must map the user information to portal properties on the User Information — Property Map (in the User Profile Manager) before you import the user information.
- You must import users through an authentication source before you can import the associated user information.
- You must run a job associated with the profile source to import the user information. You should continue to run the job periodically to keep the user information in the portal synchronized with the information in the source user repository.

Profile Providers

A profile provider is a piece of software that tells the portal how to use the information in the external user repository. BEA provides profile providers as part of the AquaLogic Interaction Identity Services. The AquaLogic Interaction Identity Service - LDAP is used to import user information from LDAP servers. The AquaLogic Interaction Identity Service - Active Directory is used to import user information from Active Directory servers. If your user information resides

in a custom system, such as a custom database, you can import it by writing your own profile provider using the IDK.

Note:

- Your portal administrator must install the profile provider before you can create the associated profile web service. For information on obtaining profile providers, contact *ALUIsupport@bea.com*. For information on installing profile providers, refer to the *Installation Guide for AquaLogic Interaction* (available on *edocs.bea.com*) or the documentation that comes with your profile provider, or contact your portal administrator.
- To learn about developing your own profile provider, refer to the *BEA AquaLogic User Interaction Development Center*.

Profile Web Services

Profile web services enable you to specify general settings for your external user repository, leaving the more detailed settings (like domain specification) to be set in the associated remote profile sources. This allows you to create different profile sources to import information each domain without having to repeatedly specify all the settings.

Viewing User Profiles

User profiles provide information about users, such as address and position. You can view your information or information for other users.

Your portal administrator controls what information you see in a user profile, but the following categories of information are available by default:

- **General Information** includes general contact information such as name, position, and phone number.
- **Folder Expertise** lists the Knowledge Directory folders for which the user is a related expert. Your portal administrator might add a user to a folder as an expert, or, if you have the appropriate permissions, you can add yourself as expert when you are browsing folders in the Knowledge Directory. To open a listed Knowledge Directory folder, click the folder name.
- **Managed Communities** lists the communities that the user has permission to manage. To view a listed community, click the community name.

Note:

- The Folder Expertise and Managed Communities portlets do not display if your portal uses adaptive layouts.

- If information contained in Folder Expertise or Managed Communities is incorrect, contact your portal administrator.

- To view your user profile, in the portal banner, click **My Account ➤ View User Profile** .
- To view another user's profile, search for the user and click the user's name.

  Note: You can view only those properties to which you have access.

Adding Headers and Footers to User Profiles

You can add header and footer portlets to user profiles to control what users see at the top and bottom of the user profile pages.

To add headers and footers you must have the following rights and privileges:

- Access Administration activity right
- Access Utilities activity right
- At least Select access to the header and footer portlets you want to add

1. Click **Administration**.
2. In the **Select Utility** drop-down list, select **User Profile Manager**.
3. Click the **Header and Footer** page.
4. Select the header and footer for the user profile pages.
   a) To add or change the header, under **Community Header**, click **Browse**, then, in the Select a Header dialog box, select the header you want, and click **OK**.
   b) To add or change the footer, under **Community Footer**, click **Browse**, then, in the Select a Footer dialog box, select the footer you want, and click **OK**.
   c) To remove the header, under **Community Header**, click **Remove**.
   d) To remove the footer, under **Community Footer**, click **Remove**.

Editing Your User Profile

You can update your user profile information, such as e-mail address or phone number.

User profile information can be viewed or searched by users or passed to certain portal objects for authentication or use as metadata.

1. In the portal banner, click **My Account**.
2. On the My Account page, click **Edit User Profile**.
3. Type the information you want to provide in the appropriate text boxes.

Note: Your portal administrator may have populated some information automatically and may have set some information to read only.

4. If there are multiple user profile pages listed on the left and you want to change information on another page, click the page name and change your information.
5. When you are done, click **Finish** to save your settings, or click **Cancel** to revert to your previous settings.

Associating User Information with Properties Using the User Information — Property Map

The User Information — Property Map enables you to map user information to user properties in the portal. The information in these user properties can then be displayed in the user's profile, or it can be sent to content crawlers, remote portlets, or federated searches so that users do not have to enter this information on a separate preference page.

To map user information to portal properties you must have the following rights and privileges:

- Access Administration activity right
- Access Utilities activity right
- At least Select access to the properties you want to map

Note: The Full Name attribute is automatically mapped to display name of the user unless you override it on this page.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **User Profile Manager**.
3. Under Edit Object Settings, click **User Information - Property Map**.
4. Add a property. Click **Add**; then, in the **Choose Property** dialog box, select the property you want to add and click **OK**.
5. Map attributes to the property:
   a) Click next to the property name.
   b) In the text box, type the attribute.
      To map the property to multiple attributes, separate the attribute names with commas (,).
6. Repeat Steps 4 and 5 to map additional properties.

To remove properties, select the property you want to remove and click ✕ .

After mapping user information to portal properties, you need to import the user information through profile sources or have users manually enter the information by editing their user profiles.

Creating a Profile Web Service

Profile web services enable you to specify general settings for your external user repository, leaving the more detailed settings (like domain specification) to be set in the associated remote profile sources. This allows you to create different profile sources to import information each domain without having to repeatedly specify all the settings.

Before you create a profile web service, you must:

- Install the profile provider on the computer that hosts the portal or on another computer
- Create a remote server pointing to the computer that hosts the profile provider (optional, but recommended)

To create a profile web service you must have the following rights and privileges:

- Access Administration activity right
- Create Web Service Infrastructure activity right
- At least Edit access to the parent folder (the folder that will store the profile web service)
- At least Select access to the remote server that the profile web service will use

1. Click **Administration**.
2. Open the folder in which you want to store the profile web service.
3. In the **Create Object** drop-down list, click **Web Service — Profile**.
   The Profile Web Service Editor opens.
4. On the **Main Settings** page, complete the following task:
   •
5. Click the **HTTP Configuration** page and complete the following task:
   •
6. Click the **Advanced Settings** page and complete the following task:
   •

7. Click the **Authentication Settings** page and complete the following task:
   -
8. Click the **Debug Settings** page and complete the following task:
   -
9. Click the **Properties and Names** page and complete the following tasks:
   - *Naming and Describing an Object* on page 217

     You can instead enter a name and description when you save this profile web service.
   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219 (optional)

The default security for this profile web service is based on the security of the parent folder. You can change the security when you save this profile web service (on the **Security** tab page in the Save As dialog box), or by editing this profile web service (on the **Security** page of the Profile Web Service Editor).

Portal administrators with at least Select access to this profile web service can create profile sources based on the web service.

Importing User Information from External Repositories with Remote Profile Sources

Profile sources allow you to import user information (such as name, address, or phone number) that is already defined in your enterprise in existing user repositories, such as Active Directory or LDAP servers. The imported user information can be used to populate user profiles or can be passed to content crawlers, remote portlets, or federated searches as user information.

Before you create a remote profile source, you must:

- Import users with an authentication source.
- If necessary, create portal properties for the attributes you want to import.
- Associate the portal properties with the user object through the Global Object Property Map.
- Map user attributes from the source user repository to portal properties with the User Information Property Map.
- Install the profile provider on the computer that hosts the portal or on another computer.
- Create a remote server that points to the computer that hosts the profile provider.

- Create a profile web service on which to base the profile source.

To create a profile source you must have the following rights and privileges:

- Access Administration activity right
- Create Profile Sources activity right
- At least Edit access to the parent folder (the folder that will store the authentication source)
- At least Select access to the profile web service on which this profile source will be based

1. Click **Administration**.
2. Open the folder in which you want to store the profile source.
3. In the **Create Object** drop-down list, click **Profile Source - Remote**.
   The Choose Web Service dialog box opens.
4. Select the web service that provides the basic settings for your profile source and click **OK**.
   The Remote Portlet Editor opens, displaying the Main Settings page.
5. On the **Main Settings** page, complete the following tasks:
   - *Selecting a Unique Key for a Profile Source* on page 131
   - *Selecting the Users and Groups for Which to Import Profile Information* on page 131
6. Click the **Property Map** page and complete the following task:
   - *Mapping Source User Attributes to Portal PropertiesYou can select the users and groups for which user information should be imported.*
7. Click the **Set Job** page and complete the following task:
   - *Associating an Object with a Job* on page 300
8. Click the **Properties and Names** page and complete the following tasks:
   - *Naming and Describing an Object* on page 217

     You can instead enter a name and description when you save this profile source.
   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219 (optional)

The default security for this profile source is based on the security of the parent folder. You can change the security when you save this profile source (on the **Security** tab page in the Save As dialog box), or by editing this profile source (on the **Security** page of the Profile Source Editor).

Run the job associated with this profile source.

Selecting a Unique Key for a Profile Source

Each profile source must include a unique key that is used to identify the user to the profile provider.

1. If the Profile Source Editor is not already open, open it now and display the **Main Settings** page.
2. Under **Profile Unique Key**, select the key that will be used to identify the user to the profile provider.

   - **Remote Unique Name** — This is the default. The user's imported unique name will be sent to the remote provider to identify the user. Common examples are the GUID or User Name.
   - **Remote Authentication Name** — The user's imported authentication name will be sent to the remote provider. In most cases, this is the same as the unique name.
   - **User Property Value** — The value of a property associated with each user will be sent to identify this user. Typically, this value is imported by another profile source.

     If you select this option, you must also select which property to use: click **Choose Property**, in the Choose Property dialog box, select a property and click **OK**.

     To change the selected property, click the property name.

Selecting the Users and Groups for Which to Import Profile Information

You can select the users and groups for which user information should be imported.

1. If the Profile Source Editor is not already open, open it now and display the **Main Settings** page.
2. Under **Profile Source Membership**, select the users and groups for which user information should be imported.

   - To add users or groups, click **Add Users/Groups**; then, in the Profile Source Membership dialog box, select the users and groups you want to add and click **OK**.
   - To remove a user or group, select the user or group and click .

     To select or clear all of the user and group boxes, select or clear the box to the left of **Users/Groups**.
   - To toggle the order in which the folders are sorted, click **Users/Groups**.

Clearing User Information Imported by a Profile Source

You can delete all user information previously imported by a profile source. This is useful when you add a new user property and want to look it up and update it for all users, or when you change a property from read-write to read-only and want to overwrite previous user modifications.

To delete user information imported by a profile source you must have the following rights and privileges:

- Access Administration activity right
- At least Edit access to the profile source

1. Click **Administration**.
2. Navigate to and open the profile source.
   The Remote Portlet Editor opens, displaying the Main Settings page.
3. Click **Clear History**.

Run the job associated with the profile source to import the user information again.

About 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. After you create an invitation, you generate an invitation link to send to invitees. The invitation link expires after a specified number of users is created from the link or after the specified date. You can generate multiple invitation links for one invitation, each with different expiration settings.

To accept the invitation, the user clicks the link included in the e-mail and follows the directions to create a new user and 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 stored in the folder you specify in the invitation and are included in the AquaLogic Interaction Authentication Source. They are automatically given security

privileges based on the default profile you specify in the invitation. Based on this security, users can personalize their views of the portal with My Pages, portlets, and community memberships, and can view portal content.

Inviting Users to Your Portal

Before you create an invitation, you must:

- Create the default profile you want to apply to the users who accept the invitation.
- Create the folder in which you want to store the users who accept the invitation.

To create an invitation you must have the following rights and privileges:

- Access Administration activity right
- Create Invitations activity right
- At least Edit access to the parent folder (the folder that will store the invitation)

1. Click **Administration**.
2. Open the folder in which you want to store the invitation.
3. In the **Create Object** drop-down list, click **Invitation**.
   The Invitation Editor opens, displaying the Main Settings page.
4. Select a folder in which to store the users who accept this invitation. Click **Browse**; then, in the **Select a Folder** dialog box, choose a folder and click **OK**.

   If you want to display a particular experience definition interface to users when they log in, choose a folder to which the experience definition has been applied or apply the experience definition to the chosen folder before you send the invitation.
5. In the **Default User Image** drop-down list, select the default profile to apply to users who accept the invitation.

   The default profile defines the user's initial view of the portal.
6. Select the groups to which you want to add users who accept the invitation.

   - To add invitees to a group, click  **Add Group**; then, in the **Select Groups** dialog box, select the groups you want to add and click **OK**.
   - To remove a group from the list, select the group and click .

     To select or clear all of the group check boxes, select or clear the box to the left of **Group Name**.
   - To toggle the order in which the groups are sorted, click **Group Name**.

After creating the invitation, you need to generate an invitation link and e-mail it to your invitees.

Sending Invitations

To send an invitation, you generate a link to e-mail to recipients. Recipients who follow this link are prompted to create a new account in your portal and can then begin customizing their views of your portal and exploring its contents.

Before you send an invitation, you must:

- Create the invitation.

To send an invitation you must have the following rights and privileges:

- Access Administration activity right
- At least Edit access to the invitation

You can create

1. Click **Administration**.
2. Open the folder in which the invitation is stored.
3. Select the invitation and click **Send Invitation**.
   The Send Invitation page opens.
4. If you have not already done so, create an invitation link. Click **Create New Invitation Link**.

   If you have already created an invitation link with the expiration settings you want to use, skip to Step 6.
5. In the Create New Invitation Link dialog box, specify settings to prevent this link from being circulated and allowing unintended users access to secured content in your portal.
   a) In the **Name** box, type a name for this link that makes clear to you and other portal administrators what this link is for.
   b) In the **Number of Invitations** box, type the maximum number of users that can be created from this link.
   c) In the **Expiration Date** box, type the date after which this link displays an error and will not allow users to create a portal account.

      To choose the date from a calendar, click .
   d) To create the link, click **Finish**.

6. To display the invitation link, click the link name.
7. Copy and paste the invitation link into an e-mail, modify the message as desired, and send it to your invitees.

   Note: The only way to cancel an invitation is to delete the invitation, so be sure your invitation is correct before you e-mail it to anyone.

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.

To configure user activity auditing and audit user activity you must have the following rights:

- Access Administration activity right
- Access Utilities activity right

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

You can specify what types of events should be logged.

To access the Audit Manager you must be a member of the Administrators Group.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **Audit Manager**.
3. Under **Message Types**, specify what types of events should be logged:

Message Type	Description
Item Change	Creates an entry every time an object is edited.
Item Deletion	Creates an entry every time an object is deleted.
Locked Account	Creates an entry every time a user account is locked after a number of failed login attempts.
Security Change	Creates an entry every time an object's security is edited.

Message Type	Description
User Login	Creates an entry every time a user successfully logs in to the portal.
Global System Change	Creates an entry every time an edit is made to the Global ACL Sync Map, the Global Property Map, the Global Content Type Map, the Global Object Property Map, or the User Information Property Map; every time job folders or Automation Services are registered; and every time global system settings are changed through the various portal utilities.

Querying User Activity Audit Information

You can query the user activity logs.

To access the Audit Manager you must be a member of the Administrators Group.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **Audit Manager**.
3. Click the **Create Audit Query** page.
4. Under **Search Criteria**, limit the information returned by your query:

   Note: If you do not specify any information on this page, your query returns a description of every audit record that is stored in the database.

   - To limit your query to a particular type of object, in the **Item Type** drop-down list, choose the object.

     For example, you might want to see only audit messages referring to modifications of content crawlers.

   - To limit your query to objects of a particular name, in the **Item Name** box, type the text you want to search for and, in the drop-down list, choose whether you want your search for approximate or exact matches.

     If you search for approximate matches, the portal returns items that include your text in any part of the name; if you search for exact matches, the portal returns only those items in which the item name equals the text you specify. For example, you could request only

audit messages referring to actions on Sales content crawlers or Sales portlets by entering `Sales` in the text box and choosing **Approximate**.

- To limit your query to actions performed by a particular user, in the **Username** box, type the text you want to search for and, in the drop-down list, choose whether you want to search for approximate or exact matches.
- To limit your query to actions performed on a particular portal server, in the **Server Name** box, type the text you want to search for and, in the drop-down list, choose whether you want to search for approximate or exact matches.

  For example, you could retrieve messages for all jobs run on the Automation Service named PortalJobs.
- To limit your query to audit messages containing a particular word, in the **Word in Message** box, type the text you want to search for.

  For example, to limit your query to all messages relating to a particular group, type the group name in this box.
- To limit your query to particular types of messages, choose the types.

Message Type	Description
Item Change	Entries corresponding to every time an object is edited.
Item Deletion	Entries corresponding to every time an object is deleted.
Locked Account	Entries corresponding to every time a user account is locked after a number of failed login attempts.
Security Change	Entries corresponding to every time an object's security is edited.
User Login	Entries corresponding to every time a user successfully logs in to the portal.
Global System Change	Entries corresponding to every time an edit is made to the Global ACL Sync Map, the Global Property Map, the Global Document Type Map, the Global Object Property Map, or the User Information Property Map; every time job folders or Automation Services are registered; and every time global system settings are changed through the various portal utilities.

5. To limit your query to a particular period, in the **Time Interval** boxes, enter the starting and ending date and time you want to search.

6. Select the order in which you want to sort audit messages.

   By default, the most recent audit messages are displayed first. To change the sort to display the oldest audit messages first, choose **Oldest to newest**.

7. In the **Results per page** box, type the maximum number of messages to display per page.
8. Click the **Run Query** page.

User Activity Audit Query Results

When you run an audit query, the results display on the Run Query page of the Audit Manager.

Column	Description
Item Type	Displays the type of object that was modified: for example, Content Crawler, Portlet, or User.
Item Name	Displays the name of the object that was modified: for example, the Meeting Minutes Content Crawler.
User	Displays the name of the user who performed the action on the object.
Server	Displays the server from which the object was modified.
Message Type	Displays the type of action performed on the object: for example, User Login or Item Change.
Time	Displays the date and time the object was modified.
Message	Displays the text of the message.

Archiving Audit Messages

You can specify how and when to archive audit messages.

To access the Audit Manager you must be a member of the Administrators Group.

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. To change the frequency, edit the Audit Log Management Job.

1. Click **Administration**.

2. In the **Select Utility** drop-down list, click **Audit Manager**.
3. Under **Archiving Agent**, specify the settings for your auditing archive:
   a) In the **Network path of archive files** box, type the path to the folder in which you want to store audit archive files.
   b) In the **Days to keep messages in database** box, type the number corresponding to how many days worth of messages you want to store in the portal database.

      Only messages in the portal database are available for audit query. After the specified amount of time, messages are moved from the database into the archive files.
   c) In the **Days to keep messages in files** box, type the number corresponding to the number of days you want to store the message files.

      After the specified period, messages are deleted from these files and no longer available.

Deleting Audit Messages and Archives

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

To access the Audit Manager you must be a member of the Administrators Group.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **Audit Manager**.
3. Under **Delete Messages**, specify which messages you want to delete from the portal database (they are not moved into the audit archive) and which archives you want to delete from your file system:
   a) In the **Delete Messages and Archives prior to** box, type the date for which you want to delete messages and archives.

      Any messages and archives with this date or an earlier date are deleted.
   b) In the **Message types to delete** section, choose the types of messages that you want to delete from the database.

Message Type	Description
Item Change	Entries corresponding to every time an object is edited.
Item Deletion	Entries corresponding to every time an object is deleted.
Locked Account	Entries corresponding to every time a user account is locked after a number of failed login attempts.

Message Type	Description
Security Change	Entries corresponding to every time an object's security is edited.
User Login	Entries corresponding to every time a user successfully logs in to the portal.
Global System Change	Entries corresponding to every time an edit is made to the Global ACL Sync Map, the Global Property Map, the Global Document Type Map, the Global Object Property Map, or the User Information Property Map; every time job folders or Automation Services are registered; and every time global system settings are changed through the various portal utilities.

c) If you want to delete these messages and archives when you click **Finish**, select **Yes** next to **Delete Messages and Archives when 'Finish' is clicked**.

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.

About the Portal Knowledge Directory

The Knowledge Directory is similar to a file system tree in that documents are organized in folders and subfolders. A folder can contain documents uploaded by users or imported by content crawlers, as well as links to people, portlets, and communities. If your administrator has given you permission, you might also be allowed to add documents to the Knowledge Directory, or submit yourself as an expert on a particular topic.

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*. For example, you probably want to organize the Knowledge Directory in a way that allows you to easily delegate administrative responsibility for the content and facilitate managed access with access control lists (ACLs).

When you open the Directory, you see the folders and subfolders to which you have at least Read access. After you have opened a Directory folder, you see additional features: documents, document display options, subfolders, and related objects.

Documents

On the left, you see the documents to which you have at least Read access. Each document includes an icon to signify what type of document it is (for example: web page, PDF, MS Word document), the document name, the document description, when the document was last modified, a link to view additional document properties, and a link that displays the URL to this document (enabling you to e-mail a link to the document). At the bottom of the list of documents, you see page numbers indicating how many pages of documents exist in this folder.

Document Display Options

At the top of the list of documents, you see drop-down lists that let you change how documents are sorted, how many documents are displayed per page, and filter what types of documents are displayed.

Subfolders and Related Objects

On the right, you see the subfolders in this folder, and any objects that the folder administrator has specified as related to this folder.

Note: You see only those folders and objects to which you have at least Read access.

- Under **Subfolders**, you see the subfolders in this folder.
- Under **Related Communities**, you see communities that have information related to the documents in this folder.
- Under **Related Folders**, you see other Directory folders that have information related to the documents in this folder.
- Under **Related Portlets**, you see portlets that have information or functionality related to the documents in this folder.
- Under **Related Experts**, you see the users that are familiar with the documents in this folder (for example, an expert might have written one of the documents in the folder).
- Under **Related Content Managers**, you see the users that manage the documents in this folder and the content sources and content crawlers associated with this folder.

Browsing Documents in the Portal Knowledge Directory

The Knowledge Directory is similar to a file system tree in that documents are organized in folders and subfolders. A folder can contain documents uploaded by users or imported by content crawlers, as well as links to people, portlets, and communities. If your administrator has given you

permission, you might also be allowed to add documents to the Knowledge Directory, or submit yourself as an expert on a particular topic.

When you open the Directory, you see the folders and subfolders to which you have at least Read access.

- To edit the directory (add or edit folders), click **Edit Directory**.

  Note: You must have the Edit Knowledge Directory activity right to see this button. You must also have at least Edit access to a folder or document to be able to edit the folder or document.

- To open a folder or subfolder, click its name.

  Note: If the folder includes a description, it appears as a tooltip. To view the description, place your mouse over the folder name.

  After you have opened a Directory folder, you see the additional features described next.
- To open a document, click its name.
- To view the properties of a document, click the **Properties** link under the document description.
- To tell other portal users about a document:
  a) Under the document description, click **Send Document Link**.
  b) In the **Document Link** dialog box, copy the text, then click **Close**.
  c) In your e-mail application, paste the text into an e-mail message and send it.

  When other portal users click the URL in your e-mail, the document opens. If a user does not have permission to see the document, an error message is displayed.

- To submit a document to the portal, click **Submit Documents**.

  Note: You must have at least Edit access to the folder and at least Select access to the content source that provides access to the document to be able to submit a document.

- To view another page of items, at the bottom of the list of documents, click a page number or click **Next >>**.
- To change the sort order of documents between ascending and descending, in the **Sort** by drop-down list, select the desired option: **Document Name Ascending** or **Document Name Descending**.
- To change the number of documents that are displayed per page, in the **Items per page** drop-down list, select the desired number.

  By default, 20 items are shown per page.

- To filter the documents by document type (for example, MS Word documents or PDF documents), in the **Show only item type** drop-down list, select the desired document type.
- To open a subfolder, under **Subfolders**, click the subfolder name.

  Note: Beneath the banner, you see the hierarchy for the folder you are viewing (sometimes referred to as a breadcrumb trail). To move quickly to one of these folders, click the folder's name.

- To create a subfolder in this folder, under **Subfolders**, click **Create Folder**. In the **Create Document Folder** dialog box, type a name and description for the folder, and click **OK**.
- To view a related community, under **Related Communities**, click the community name.

  Note: If you have at least Select access to the community, you can join the community.

- To open a related folder, under **Related Folders**, click the folder name.
- To preview a related portlet, under **Related Portlets**, click the portlet name.

  Note: If you have at least Select access to the portlet, from the portlet preview page, you can add the portlet to one of your My Pages.

- To view the user profile for a related expert, under **Related Experts**, click the user's name.

  Note: If you have the Self-Selected Experts activity right, and are not already listed as an expert, click **Add Me** to add yourself as an expert on the folder's topic.

- To view the user profile for a related content manager, under **Related Content Managers**, click the user's name.

Editing Documents in the Portal Knowledge Directory

The Knowledge Directory is similar to a file system tree in that documents are organized in folders and subfolders. A folder can contain documents uploaded by users or imported by content crawlers, as well as links to people, portlets, and communities. If your administrator has given you permission, you might also be allowed to add documents to the Knowledge Directory, or submit yourself as an expert on a particular topic.

When you open the Directory, you see the folders and subfolders to which you have at least Read access.

- To edit the directory (add or edit folders), click **Edit Directory**.

Note: You must have the Edit Knowledge Directory activity right to see this button. You must also have at least Edit access to a folder or document to be able to edit the folder or document.

- To open a folder or subfolder, click its name.

  Note: If the folder includes a description, it appears as a tooltip. To view the description, place your mouse over the folder name.

  After you have opened a Directory folder, you see the additional features described next.
- To open a document, click its name.
- To view the properties of a document, click the **Properties** link under the document description.
- To tell other portal users about a document:

  a) Under the document description, click **Send Document Link**.
  b) In the **Document Link** dialog box, copy the text, then click **Close**.
  c) In your e-mail application, paste the text into an e-mail message and send it.

  When other portal users click the URL in your e-mail, the document opens. If a user does not have permission to see the document, an error message is displayed.

- To submit a document to the portal, click  **Submit Documents**.

  Note: You must have at least Edit access to the folder and at least Select access to the content source that provides access to the document to be able to submit a document.

- To view another page of items, at the bottom of the list of documents, click a page number or click **Next >>**.
- To change the sort order of documents between ascending and descending, in the **Sort** by drop-down list, select the desired option: **Document Name Ascending** or **Document Name Descending**.
- To change the number of documents that are displayed per page, in the **Items per page** drop-down list, select the desired number.

  By default, 20 items are shown per page.
- To filter the documents by document type (for example, MS Word documents or PDF documents), in the **Show only item type** drop-down list, select the desired document type.
- To open a subfolder, under **Subfolders**, click the subfolder name.

  Note: Beneath the banner, you see the hierarchy for the folder you are viewing (sometimes referred to as a breadcrumb trail). To move quickly to one of these folders, click the folder's name.

- To create a subfolder in this folder, under **Subfolders**, click **Create Folder**. In the **Create Document Folder** dialog box, type a name and description for the folder, and click **OK**.
- To view a related community, under **Related Communities**, click the community name.

  Note: If you have at least Select access to the community, you can join the community.
- To open a related folder, under **Related Folders**, click the folder name.
- To preview a related portlet, under **Related Portlets**, click the portlet name.

  Note: If you have at least Select access to the portlet, from the portlet preview page, you can add the portlet to one of your My Pages.
- To view the user profile for a related expert, under **Related Experts**, click the user's name.

  Note: If you have the Self-Selected Experts activity right, and are not already listed as an expert, click **Add Me** to add yourself as an expert on the folder's topic.
- To view the user profile for a related content manager, under **Related Content Managers**, click the user's name.

Folder Toolbar

You can use the buttons in the folder toolbar to perform actions on the folder you are viewing.

Button	Action
	Expand all the object types so you can view all objects stored in this folder.
	Edit the folder you are viewing. **Note:** You see this button only if you have at least Edit access to the folder.
Up	Navigate up to the parent folder. **Note:** You see this button only when you are viewing a subfolder.

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 access the Knowledge Directory Preferences Utility you must be a member of the Administrators group.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **Knowledge Directory Preferences**.
3. In the **Subfolder Description type** list, choose the type of subfolder description to display in the Knowledge Directory:

Option	Description
none	Displays no subfolder description
abbreviated	Displays only the first 100 characters of the folder description
full	Displays the full subfolder description

4. In the **Maximum number of subfolders to display** list, choose the number of subfolders to display under the current folder.
5. In the **Number of subfolder columns** list, choose a number of columns to display subfolders.

   Note:

   - Documents are always displayed in a single column.
   - This setting does not apply in adaptive page layout mode.

6. In the **Number of documents to show per page** box, type a number.
7. In the **Document Description type** list, choose the type of document description to display in the Knowledge Directory:

Option	Description
none	Displays no document description
abbreviated	Displays only the first 100 characters of the document description
full	Displays the full document description

8. In the **Related Resources placement** list, choose the desired placement, relative to folders and documents: **Left**, **Right**, **Top**, or **Bottom**.

   Note: Related resources are specified on the Related Resources page of the Folder Editor.

9. In the **Browsing Source** list, choose the source of the folder information that displays when browsing the Knowledge Directory:

Option	Description
Search	Uses the portal Search Service to generate the list of folder contents
Database	Queries the portal database

   Note: If you have a large collection of documents, you can improve browsing performance by choosing **Search**.

10. In the **Default Document Submission Content Type** list, choose the default content type, which is used when you submit a document that is not mapped to any content type.

    If you do not want to specify a default, choose **None**.

11. Under **Browsing Column Properties**, select the properties you want to display as custom columns when browsing documents the Knowledge Directory:

    - To add a property, click **Add Property**, then, in the drop-down list that appears, select the desired property.

      Note: Only numeric and date properties can be selected as custom column properties.

    - To delete properties, select the properties you want to delete and click .
    - To change the order in which properties display use the icons to the right of the properties:
      - To move a property to the top of this list, click .
      - To move a property up one space in this list, click .
      - To move a property down one space in this list, click .
      - To move a property to the bottom of this list, click .

      The order in which properties appear on this page is the order in which the columns appear in the Knowledge Directory.

Troubleshooting Knowledge Directory Issues

If you get a timeout error when applying a security change to all child objects in a Knowledge Directory hierarchy, you can use the following workaround.

This issue is most likely caused by trying to apply changes to many levels of nested subfolders with a large number of child objects (other folders or documents). In this situation, you can perform the following steps to work around the issue:

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.

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.

Complete the following tasks to enable managed discovery of enterprise content through the portal:

- For all file types you plan to support in your portal, configure document properties to store document metadata and to enable document filters used by the Knowledge Directory, content crawlers, the Smart Sort utility, and the Search Service.
- 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.
- Allow users at least Edit access to the folders in the Knowledge Directory to which you want them to be able to upload document records.
- Configure portlets that users can add to their My Pages.
- Create communities that users can add to their My Communities list.
- Configure content crawlers and crawl jobs to create links to back-end content sources, such as internet locations, file system locations, Documentum Content Servers, Exchange Servers, Lotus Notes Servers, or other IMAP-compliant servers.
- Run a Search Update job to index these documents so that they can be discovered with the search.

Permissions Required for Accessing, Crawling, and Submitting Documents

There are several kinds of permissions a user needs to view, submit, or crawl documents.

Action	Permissions Needed
Access documents imported into the portal	• Read access to the document link in the Knowledge Directory • Read access to the Knowledge Directory folder in which the link is stored • Read access to the content source used to import the document • If the document is not gatewayed, access to the document in the source repository
Crawl documents into the portal	• Edit access to the Knowledge Directory folder into which they are crawling documents

Action	Permissions Needed
	• Edit access to the administrative folder in which they are creating the content crawler • Select access to the content source • Access Administration activity right • Create Content Crawlers activity right • Select access to a job that can run the content crawler or Create Jobs activity right plus Edit access to an administrative folder that is registered to an Automation Service
Submit a document into the portal	• Edit access to the Knowledge Directory folder into which they are submitting a document • Select access to a content source that supports document submission • If the associated content web service does not support browsing, knowledge of the path to the document

Note: If you have content sources that access sensitive information, be aware that users that have access to the content source and have the additional permissions listed in the table could access anything that the user that the content source impersonates can access. For this reason, you might want to create multiple content sources that access the same repository but that use different authentication information and for which you allow different users access.

Using Simple Submission to Submit or Upload Documents to the Portal Knowledge Directory

With the proper permissions, you can submit documents to the Knowledge Directory. Depending on your portal configuration, you might also be able to upload a file to the Knowledge Directory. When you upload a file, it is copied from the remote repository into the portal's document repository and a pointer is created to that copied file.

Before you submit or upload a document to the Knowledge Directory:

- Make sure the language of the document you are submitting matches the language specified as your locale on the **Edit Locale Settings** page of **My Account**. For example, if your default locale is Japanese, the document you are submitting must also be in Japanese. If you are

submitting a document in a language different from your default locale, either change your default locale to match the language of the document before you submit it, or, if you have permission to edit the Knowledge Directory, you can use Remote Document or Web Document submission to select a different language.

- If you are using Netscape 7.1, make sure the ANSI character set of your client matches that of the name of the file that is being uploaded. The ANSI character set is determined by the default system locale configuration of your client. For instructions on setting the default locale, refer to your operating system documentation.

To submit or upload a document to the Knowledge Directory you must have the following rights and privileges:

- At least Edit access to the parent folder (the folder that will store the document)
- At least Select access to the content source that provides access to the location where the document is stored

Note: If you want to use a content type other than the default content type associated with the content source, if you want to submit a document to more than one Knowledge Directory folder, or if you want to submit a document in a language different from your default locale, use Remote Document or Web Document submission, available when editing the Knowledge Directory.

1. Click **Directory**.
2. Open the folder in which you want to place the document.
3. Click **Submit Documents**.
   The Submit a Document dialog box opens.

   If you are editing the Knowledge Directory, you open the Submit a Document dialog box by selecting **Simple Submit** in the **Submit Document** drop-down list on the right.
4. In the **Document source** drop-down list, accept the default document source or select another.

   The document source tells the portal how to find the document you are submitting.

   Note: If you are uploading a file, you must select **Content Upload**.
5. Specify a file by performing one of the following actions:
   - If you are submitting a web document, in the **URL** text box, type the document's URL.
   - If you are submitting or uploading a file, specify a file by performing one of the following actions:
     - Type the UNC path to the document in the **File path** text box.

If you are leaving the file in the remote location, you must type a network path (for example, `\\myComputer\myFolder\myFile.txt`). If you are uploading the file, the path can be a local path (for example, `C:\myFolder\myFile.txt`) or a network path.

- Click **Browse** to navigate to the location of the file you want to submit.

  If you are leaving the file in the remote location, you must supply a network path to the file, and therefore, you cannot browse your local drives; you must browse the network to your computer and then to the location of the file. If you are uploading the file, you can browse to local drives or network drives.

  Note:

  - Depending on how the administrator configured the content source, the **Browse** button might not appear, therefore you might not be able to browse to the file. If you do not see a **Browse** button, type the path in the **File path** text box.
  - If the **Browse** button does display but you cannot browse to the folder where the file you want to submit is located, the content source you chose might not have the necessary privileges to access the file location. Click **Cancel** and resubmit the file using a different content source.

6. If desired, override the default name or description.

   - To override the default name, select **Use this name** and, in the text box, type the name.
   - To override the default description, select **Use this description** and, in the text box, type the description.

Once the folder administrator (one who has Admin access to the folder) approves your submission, links to the document you submitted or uploaded appear in the Knowledge Directory.

Using Advanced Submission to Submit or Upload Documents to the Portal Knowledge Directory

With the proper permissions, you can use advanced submission to submit documents to the Knowledge Directory. Advanced submission enables you to select a content type other than the default content type associated with the content source, submit a document to more than one Knowledge Directory folder, or submit a document in a language different from your default locale. Depending on your portal configuration, you might also be able to upload a file to the

Knowledge Directory. When you upload a file, it is copied from the remote repository into the portal's document repository and a pointer is created to that copied file.

To use advanced submission to submit or upload a document to the Knowledge Directory you must have the following rights and privileges:

- Edit Knowledge Directory activity right
- At least Edit access to the parent folder (the folder that will store the document)
- At least Select access to the content source that provides access to the location where the document is stored

Note: If you are using Netscape 7.1, make sure the ANSI character set of your client matches that of the name of the file that is being uploaded. The ANSI character set is determined by the default system locale configuration of your client. For instructions on setting the default locale, refer to your operating system documentation.

1. Click **Directory**.
2. Click **Edit Directory**.
3. Open the folder in which you want to place the document.
4. In the **Submit Document** drop-down list on the right, select **Remote Document**.
   The Choose a Content Source dialog box opens.
5. Select the content source that provides access to the content you want to submit or upload and click **OK**.

   Note: If you are uploading a file, you must select **Content Upload**.

6. Specify a file by performing one of the following actions:

   - Type the UNC path to the document in the **File path** text box.

     If you are leaving the file in the remote location, you must type a network path (for example, `\\myComputer\myFolder\myFile.txt`). If you are uploading the file, the path can be a local path (for example, `C:\myFolder\myFile.txt`) or a network path.

   - Click **Browse** to navigate to the location of the file you want to submit.

     If you are leaving the file in the remote location, you must supply a network path to the file, and therefore, you cannot browse your local drives; you must browse the network to your computer and then to the location of the file. If you are uploading the file, you can browse to local drives or network drives.

     Note:

- Depending on how the administrator configured the content source, the **Browse** button might not appear, therefore you might not be able to browse to the file. If you do not see a **Browse** button, type the path in the **File path** text box.
- If the **Browse** button does display but you cannot browse to the folder where the file you want to submit is located, the content source you chose might not have the necessary privileges to access the file location. Click **Cancel** and resubmit the file using a different content source.

7. If you want to override the default name or description, click the **Name and Description** page and edit the values.
   - To override the default name, edit the value in the **Name** box.
   - To override the default description, edit the value in the **Description** box.

Once the folder administrator (one who has Admin access to the folder) approves your submission, links to the document you submitted or uploaded appear in the Knowledge Directory.

Using Advanced Submission to Submit Web Documents to the Portal Knowledge Directory

With the proper permissions, you can use advanced submission to submit documents to the Knowledge Directory. Advanced submission enables you to select a content type other than the default content type associated with the content source, submit a document to more than one Knowledge Directory folder, or submit a document in a language different from your default locale.

To use advanced submission to submit a document to the Knowledge Directory you must have the following rights and privileges:

- Edit Knowledge Directory activity right
- At least Edit access to the parent folder (the folder that will store the document)
- At least Select access to the content source that provides access to the location where the document is stored

Note: If you are using Netscape 7.1, make sure the ANSI character set of your client matches that of the name of the file that is being uploaded. The ANSI character set is determined by the default system locale configuration of your client. For instructions on setting the default locale, refer to your operating system documentation.

1. Click **Directory**.

2. Click **Edit Directory**.
3. Open the folder in which you want to place the document.
4. In the **Submit Document** drop-down list on the right, select **Web Document**.
   The Choose a Content Source dialog box opens.
5. Select the content source that provides access to the content you want to submit and click **OK**.

   Note: If you are submitting an unsecured web document, you can select **World Wide Web**.

6. In the **URL** text box, type the document's URL.
7. Under **Choose Content Type**, select the content type to apply to this document.
   - To use the folder's default content type, leave **Default Content Type** selected.
   - To choose a different content type, select **This content type**, click **Change**, in the dialog box, select the content type you want to use, and click **OK**.
8. Under **Choose Knowledge Directory Folders**, specify into which folders you want to submit this document.
   - To add a folder, click **Add Folder**.
   - To remove folders, select the folders you want to delete and click .
   - To change the order of the names in the list from ascending to descending alphabetical order (or vice versa), click **Folder Names**.
9. Under **Document Content Language**, choose the language used for the majority of the document's content.

   The language you choose is the language by which the document is indexed. The search engine uses the language when searching.
10. If you want to override the default name or description, click the **Name and Description** page and edit the values.
    - To override the default name, edit the value in the **Name** box.
    - To override the default description, edit the value in the **Description** box.

Once the folder administrator (one who has Admin access to the folder) approves your submission, links to the document you submitted or uploaded appear in the Knowledge Directory.

About Document and Object Properties

Properties provide information about, as well as a way to search for, documents and objects in your portal. For example, you might want to create an Author property so users can find all the documents or objects created by a particular user.

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.

Creating a Property to Store Object Metadata

Properties provide information about, as well as a way to search for, documents and objects in your portal. For example, you might want to create an Author property so users can find all the documents or objects created by a particular user.

To create a property you must have the following rights and privileges:

- Access Administration activity right
- Create Properties activity right
- At least Edit access to the parent folder (the folder that will store the property)

1. Click **Administration**.
2. Open the folder in which you want to store the property.
3. In the **Create Object** drop-down list, click **Property**.
   The Property Editor opens, displaying the **Main Settings** page.
4. In the **Property Type** drop-down list, choose what kind of information this property stores.

   - **Textual** stores text values.
   - **Simple Number** stores whole numbers.
   - **Floating Point Number** stores numbers that include decimal points.
   - **Date** stores date values.
   - **Reference** stores a reference to an administrative object in the portal.

     After choosing this option, in the second drop-down list, choose what type of administrative object this property references.
   - **Encrypted Text** stores encrypted text values.

Note: After you save this property, the property type cannot be changed.

5. If this property stores a web address, select **Treat this property like an URL**.

   If you choose this option, users can click-through the property, so the values for this property must always be URLs.

   Note: This option is only available for textual properties.

6. If this property applies to documents imported into the portal, select **This property is supported for use with documents**.

   Note: This option is not available for reference properties.

7. If this property is generated automatically and you want to store the value in the database but not display it on the **Properties and Names** page of object editors, clear the **This Property is visible in the UI** check box.
8. If you do not want users to be able to edit the values for this property, select **Read Only**.
9. If you want users to be able to search for objects based on the values for this property, select **Searchable**.

   For example, if you specify that the Author property is searchable, users can search for all the objects created by a particular person.

10. If you want to require that users specify a value for this property before they can save the associated object, select **Make this property mandatory**.

    Note: This option is not available if this property is set to Read Only.

11. If you want to allow users to specify more than one value for this property, select **Multiple values can be selected for this Property**.
12. In the **Property Chooser Type** drop-down list, specify what format should be used for value selection.

    - **None** displays a text box in which users can type property values.
    - **Managed Dropdown** displays a drop-down list of values you specify from which users can choose.

      - To create the values users can choose from, click **Add Value** and type a value in the text box.
      - To remove a value, select the value and click .

To select or clear all value check boxes, select or clear the box to the left of **Value Names**.

- **Unmanaged Dropdown** displays a drop-down list populated with the values from a database table from which users can choose.

  1. In the **Database Table Name** box, type the name of the table from which you want to populate your list.
  2. In the **Pick Column** box, specify the column from which you want to populate the list.
  3. In the **Sort Column** box, type the name of the column upon which the values are sorted.

- **Tree** displays a hierarchical list populated with the values from a database table from which users can choose.

  1. In the **Database Table Name** box, type the name of the table from which you want to populate the list.
  2. In the **Pick Column** box, specify the column from which you want to populate the list.
  3. In the **Sort Column** box, type the name of the column upon which the values are sorted.

  - To add a column, click  **Add Value** and enter the **Pick Column** and **Sort Column** values.
  - To remove a column, select it and click  .
  - To select or clear all the column check boxes, select or clear the box to the left of **Pick Column**.

Note: This option is not available for date or reference properties.

13. Click the **Names and Descriptions** page and complete the following tasks:

- *Naming and Describing an Object* on page 217

  You can instead enter a name and description when you save this property.
- *Localizing the Name and Description for an Object* on page 342 (optional)

If you set this property to be searchable, you need to rebuild the search index.

Mapping Source Document Attributes to Portal Properties Using the Global Document Property Map

The Global Document Property Map provides default mappings for properties common to the documents in your portal. When users import documents into the portal (either manually or through a content crawler), property values can be extracted from the source documents according to the property mappings you specify in the associated content types and the property mappings from the Global Document Property Map.

To access the Global Document Property Map you must be a member of the Administrators Group.

When a user imports 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 crawler's content type settings.
2. The portal populates property values based on the property mappings in the content type.
3. If there are additional mapped properties in the Global Content Type Map (not included in the content type's property mappings), the portal populates the property values, based on those property mappings.

Therefore, you can map common properties in the Global Document Property Map and specify only special mappings, default values, and override values in content types.

Note: Some property mappings are set when the portal is installed so that the portal can produce some general metadata even if you do not create any specialized mappings.

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.

   Note: Some property mappings are set when the portal is installed so that content crawlers can produce some general metadata even if you do not create any specialized mappings.

   - To add a property mapping, click  **Add Property**; then, in the Add Property dialog box, select the properties you want to add and click **OK**.
   - To associate an attribute, click the property name and, in the text box, type the associated attributes, separated by commas (,).

     The first attribute with a value populates the property.

- To delete a mapping, select the mapping and click  .

  To select or clear all of the mapping check boxes, select or clear the box to the left of **Property Name**.
- To toggle the order in which the properties are sorted, click **Property Name**.

Note: You can map any attribute from the source document. For information on source document attributes, refer to the documentation for the third-party software.

HTML Metadata Handling

Generally, you will be able to determine what source document attributes can be mapped to portal properties, but it might not be as clear in HTML documents.

This table shows the names of the attributes that are returned by the HTML accessor. You can map the attribute names to portal properties.

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

HTML Metadata	Name of Attribute Returned by HTML Accessor	Default Mapping or Mapping Suggestion
<TITLE> Tag	Title	Title (default)
<META> Tag	The attribute name is the NAME value. Example: `<META NAME="creation_date" CONTENT="18-Jan-2004">` The attribute that would be extracted from the example would be named “creation_date”	Using the example, you could map creation_date to the Created property.
Headline Tags	The attribute name is the name of the tag followed by an ordinal, one-based index in parentheses. The Accessor returns a value for each headline tag (<H1>, <H2>, <H3>, <H4>, <H5>, and <H6>) and each bold tag (<B>).	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).

HTML Metadata	Name of Attribute Returned by HTML Accessor	Default Mapping or Mapping Suggestion

	Example:
	`<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`	

HTML Comments	It is 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 the following source document attribute-value pairs:
	`Writer jm`
	`AP md`
	`Copy editor mr`
	`Web editor ad`
		Using the example, you could map Writer to the portal property Author.

Parent URL	Documents imported via a web content crawl return an attribute named Parent URL with the value of the URL of the parent page that contains a link to the document.	URL (default)

HTML Metadata	Name of Attribute Returned by HTML Accessor	Default Mapping or Mapping Suggestion
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.

HTML Metadata	Name of Attribute Returned by HTML Accessor	Default Mapping or Mapping Suggestion
		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.

Associating Properties with Portal Objects Using the Global Object Property Map

The Global Object Property Map displays all the types of portal objects with which you can associate properties. When users create a portal object, they can specify values for the associated properties on the Properties and Names page of the object's editor.

To access the Global Object Property Map you must be a member of the Administrators group.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **Global Object Property Map**.
3. To associate properties, click  ; then, in the Choose Property dialog box, select the properties you want to associate with the object, and click **OK**.

To toggle the order in which the objects are sorted, click **Objects**.

Associating User Information with Properties Using the User Information — Property Map

The User Information — Property Map enables you to map user information to user properties in the portal. The information in these user properties can then be displayed in the user's profile, or it can be sent to content crawlers, remote portlets, or federated searches so that users do not have to enter this information on a separate preference page.

To map user information to portal properties you must have the following rights and privileges:

- Access Administration activity right
- Access Utilities activity right
- At least Select access to the properties you want to map

Note: The Full Name attribute is automatically mapped to display name of the user unless you override it on this page.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **User Profile Manager**.
3. Under Edit Object Settings, click **User Information - Property Map**.
4. Add a property. Click **Add**; then, in the **Choose Property** dialog box, select the property you want to add and click **OK**.
5. Map attributes to the property:
   a) Click next to the property name.
   b) In the text box, type the attribute.
      To map the property to multiple attributes, separate the attribute names with commas (,).
6. Repeat Steps 4 and 5 to map additional properties.

   To remove properties, select the property you want to remove and click .

After mapping user information to portal properties, you need to import the user information through profile sources or have users manually enter the information by editing their user profiles.

Managing Object Properties

You can add or edit properties for an object.

1. Open the object's editor by creating a new object or editing an existing object.
2. Under Object Properties, change the properties and values:
   - To add or delete properties for all objects of this type, click **Open Properties Map**.

     This displays the Global Object Property Map.
   - To enter values for properties, type the value in the box to the right of the property.

   Note: The Object Properties section does not display if you are creating or editing a property.

About Controlling Document Placement with Filters

Filters control what content goes into which folder when crawling in documents or using Smart Sort to filter content into new folders. A filter sets conditions that document links must pass in order to be sorted 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 content of documents. Statements can operate on the basic fields or any other additional document properties. Statements define what must or must not be true to allow the document to pass the filter. The statements are collected together in groupings. The grouping defines whether the statements are evaluated with an AND operator (all statements are true) or an OR operator (any statement is true). If some statements should be evaluated with an AND operator and some should be evaluated with an OR operator, you can create separate groupings for the statements. You can also create subgroupings or nested groupings, where one grouping is contained within another grouping. The statements in the lowest-level grouping are evaluated first to define a set of results. Then the statements in the next highest grouping are applied to that set of results to further filter the results. The filtering continues up the levels of groupings until all the groupings of statements are evaluated.

Creating Filters to Control the Placement of Documents

To create a filter you must have the following rights and privileges:

- Access Administration activity right
- Create Filters activity right
- At least Edit access to the parent folder (the folder that will store the filter)

1. Click **Administration**.
2. Open the folder in which you want to store the filter.
3. In the **Create Object** drop-down list, click **Filter**.
   The Filter Editor opens.
4. On **Main Settings** page, complete the following task:
   - *Defining Filter Conditions* on page 167
5. Click the **Properties and Names** page and complete the following tasks:
   - *Naming and Describing an Object* on page 217

     You can instead enter a name and description when you save this filter.
   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219(optional)

Add the filter to folders.

Defining Filter Conditions

A filter is a combination of a basic fields search and statements. The basic fields search operates on the name, description, and content of documents. Statements can operate on the basic fields or any other additional document properties. Statements define what must or must not be true to allow the document to pass the filter. The statements are collected together in groupings. The grouping defines whether the statements are evaluated with an AND operator (all statements are true) or an OR operator (any statement is true). If some statements should be evaluated with an AND operator and some should be evaluated with an OR operator, you can create separate groupings for the statements. You can also create subgroupings or nested groupings, where one grouping is contained within another grouping. The statements in the lowest-level grouping are evaluated first to define a set of results. Then the statements in the next highest grouping are applied to that set of results to further filter the results. The filtering continues up the levels of groupings until all the groupings of statements are evaluated.

A filter needs at least a basic fields search or a statement.

1. If the Filter Editor is not already open, open it now and display the **Main Settings** page.
2. To search the name, description, and content values, type the text you want to search for in the **Basic fields search** text box.

   You can use the text search rules described in *Using Text Search Rules* on page 364.
3. Select the operator for the grouping of statements you are about to create:

   - If a document should pass the filter only when all statements in the grouping are true, select **AND**.
   - If a document should pass the filter when any statement in grouping is true, select **OR**.

   Note: The operator you select for a grouping applies to all its statements and subgroupings directly under it.
4. Define each statement in the grouping:

   a) Click **Add Statement**.

   b) In the first drop-down list, select the searchable property for which you want to filter the values.

   c) In the second drop-down list, select the operator to apply to this condition.

   This list will vary depending on the property selected:

   - For any text property, you can search for a value that contains your search string, or you can search for properties that have never had a value (**Contains No Value**).

     Note: If the property contained a value at some point, but the value has been deleted, the property will not match the Contains No Value condition.
   - For any date property, you can search for a value that comes after, comes before, is, or is not the date and time you enter in the boxes. You can also search for a value within the last number of minutes, hours, days, or weeks that you enter in the box.
   - For any number property, you can search for a value that is greater than, is less than, is, is not, is greater than or equal to, or is less than or equal to the number you enter in the text box.

   d) In the box (or boxes), specify the value the property must meet.

   Note: If you are searching for a text property, you can use the text search rules described in *Using Text Search Rules* on page 364.

To remove the last statement in a grouping, select the grouping, and click **Remove Statement**.

5. If necessary, add more statements by repeating Step 4.
6. If necessary, add more groupings:
   - To add another grouping, select the grouping to which you want to add a subgrouping, click **Add Grouping**, then define the statements for that grouping (as described in Step 4).

     Note: You cannot add a grouping at the same level as **Grouping 1**.
   - To remove a grouping, select the grouping, and click **Remove Grouping**.

     Note:

     - Any groupings and statements in that grouping will also be removed.
     - You cannot remove **Grouping 1**.

7. To verify that the filter works, click **Test Filter**.
   The results of the test appear in **Filter Test Report**.

You might want to test your filters before you use them extensively. See *Testing Filters* on page 169.

Testing Filters

You might want to test your filters before you use them extensively.

- To test filters, crawl content into a test folder; then perform one of the following tests:
  - Run an advanced search on the folder using the same criteria you used for your filter.

    If your advanced search returns the content you expected, you can apply the filter to the appropriate folder, confident that the filter will allow the proper documents into the folder.
  - Add the filters to subfolders and perform a smart sort to sort the content into the subfolders according to the filters.

    If the subfolders contain the content you expected, the filters work correctly.

Applying a Filter to a Folder

After you create a filter, you assign it to folders to control what content goes into the folder when crawling in documents or using Smart Sort to filter content into new folders.

Note: When users submit documents, the filters do not apply.

1. Click **Directory**.
2. Click **Edit Directory**.
3. Open the Folder Editor for the folder to which you want to apply a filter.
   - To edit the root folder (or folder you are in), in the action toolbar in the upper-right of the Edit Directory page, click .
   - To edit a subfolder, click to the right of the folder name.
4. Under **Filter Settings**, select **Links that pass** and choose whether documents must pass all filters or at least one filter to sort into this folder.
5. Under **Filters**, specify the filters that documents must pass to sort into this folder.
   - To add a filter, click **Add Filter**, select filters, and click **OK**.
   - To create a filter, click **Create Filter**.
   - To remove filters, select the filters you want to remove and click .
   - To toggle the names in the list between ascending and descending alphabetical order, click **Filter Names**.

About Content Types

Content types specify several options — the source content format (such as Microsoft Office, web page, or Lotus Notes document), whether the text of the content should be indexed for searching, and how to populate values for document properties. You should create a separate content type for each unique combination of these options. For example, if departments use different Microsoft

Word attributes for document descriptions, you might have to create one content type that pulls the description from the Subject attribute and one that pulls it from the Comments attribute.

Mapping Document Metadata to Portal Properties with Content Types

Content types specify several options — the source content format (such as Microsoft Office, web page, or Lotus Notes document), whether the text of the content should be indexed for searching, and how to populate values for document properties. You should create a separate content type for each unique combination of these options. For example, if departments use different Microsoft Word attributes for document descriptions, you might have to create one content type that pulls the description from the Subject attribute and one that pulls it from the Comments attribute.

To create a content type you need the following rights and privileges:

- Access Administration activity right
- Create Content Types activity right
- At least Edit access to the parent folder (the folder that will store the content type)

1. Click **Administration**.
2. Open the folder in which you want to store the content type.
3. In the **Create Object** drop-down list, choose **Content Type**.
   The Content Type Editor opens.
4. In the **Document Accessor** drop-down list, choose the accessor associated with the type of document for which you are creating this content type.

   The accessor determines how the portal extracts information from these documents. Use the File Accessor to extract general information (such as name and description) from any type of document. Use other accessors to extract more detailed information; for example, the HTML Accessor can extract title and author values.
5. If these documents include textual content and you want that content to be searchable, select **Index documents of this type for Search**.

   For example, you would not want to index .zip files, but you probably would want to index .txt files.
6. Specify property mappings.

   - To add a property mapping, click  **Add Property**; then, in the Select Properties dialog box, select the properties you want to map and click **OK**.

- To change the source document attributes that are associated with a property or to add a default or override value, click  to the far-right of the property and type the settings in the appropriate text boxes.
  - If you want the values for this property to be extracted from the source document attributes, in the **Mapped Attributes** box, type the associated attributes, separated by commas (,).

    Content crawlers search the attributes in the order in which you list them here; if there is no value for the first listed attribute, the content crawler looks for the second listed attribute, and so on.
  - If you want this property to have a value even when the source document does not have a value for any mapped attribute, in the **Default Value** box, type the value you want this property to be given.
  - If you want this property to have the same value for all documents, in the **Override Value** box, type the value you want this property to be given.

    Note: If you set an override value, the attribute mappings and default value for this property are ignored.
  - To save your settings, click  .
- To remove a mapping, select it and click  .

  To select or clear all of the mapping check boxes, select or clear the box to the left of **Properties**.
- To toggle the order in which the properties are sorted, click **Properties**.

Mapping Content Types to Imported Content Using the Global Content Type Map

The Global Content Type Map enables you to map file extensions (for example, .doc, .txt, .html) to content types, to define which content types are applied to imported content (whether imported by a content crawler or uploaded by a user).

To access the Global Content Type Map you must be a member of the Administrators group.

The initial content types and mappings enable you to import any type of content into the portal, but you will probably want to create custom content types and mappings to import metadata specific to your company's needs.

Note:

- Users with the Advanced Document Submission activity right can use remote document submission or web document submission and can override the default content type specified in the Global Content Type Map.
- When users create content crawlers, the **Content Type** page displays the default mappings specified in the Global Content Type Map. They can then override these mappings to fit the needs of the individual content crawler.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, choose **Global Content Type Map**.
3. Configure identifiers for content types.

   You probably want to index full-text and import specific metadata from the majority of content you import into the portal. However, there are some file types that do not include much metadata and cannot be full-text indexed (for example, .exe or .zip files). You can map these file types to the **Non Indexed Files** content type. Initially, the Global Content Type Map specifies that any file extension that is not mapped uses the **Non Indexed Files** content type (the last identifier in the list is *, which includes all file extensions).

   The * identifier mapping allows any type of file to be imported into the portal.

   - If you want to limit the types of files that can be imported into the portal:
     - Remove the mappings for any file types you want to exclude from the portal.
     - Add mappings for any non-indexed files you want to include in the portal (for example, .zip files).
     - Remove the * mapping.
   - If you do not want to limit the types of files that can be imported into the portal, keep the * mapping at the bottom of the list so that it is not applied to a file type that has a mapping.

Prioritizing a List of Objects

You can change the order of objects

- To move a group to the top of the list, click ⤒.

- To move a group up one space in the list, click ▲.
- To move a group down one space in the list, click ▼.
- To move a group to the bottom of the list, click ⤓.

About Content Sources

Content sources provide access to external content repositories, enabling users to submit documents and content managers to create content crawlers to import documents into the Knowledge Directory. Each content source is configured to access a particular document repository with specific authentication. 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.

There are two types of content sources: web content sources and remote content sources. Web content sources provide access to web sites. Remote content sources provide access to external content repositories, such as a Windows NT file system, Documentum, Microsoft Exchange, or Lotus Notes.

Note: If you delete a content source from which documents have been imported into the portal, the links to the documents will still exist, but users will no longer be able to access these documents.

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 from that content source.

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 at least Select 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, as described previously in Content Source Histories.

Content Sources Available with the Portal

Some content sources (and their necessary content web services and remote servers) are automatically created in the **Portal Resources** folder when you install the portal. There are also content sources that are available with the portal installation, but require additional steps to complete installation. For information on the additional installation steps, refer to the *Installation Guide for AquaLogic Interaction*.

- **World Wide Web:** This content source provides access to any unsecured web site.
- **Content Upload:** This content source lets users upload a document from an internal network. You should upload a document if it is not normally accessible by the users you want to see it. For example, if the document is located on your computer and your computer is not accessible by other users, you should upload the document. Additionally, if you are running an extranet, where users may not typically have access to your internal network, you should upload documents you want to make accessible externally.

Creating a Content Web Service

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

Before you create a content web service, you must:

- Install the content provider on the computer that hosts the portal or on another computer
- Create a remote server pointing to the computer that hosts the content provider (optional, but recommended)

To create a content web service you must have the following rights and privileges:

- Access Administration activity right
- Create Web Service Infrastructure activity right
- At least Edit access to the parent folder (the folder that will store the content web service)
- At least Select access to the remote server that the content web service will use

1. Click **Administration**.
2. Open the folder in which you want to store the content web service.
3. In the **Create Object** drop-down list, click **Web Service — Content**.
   The Content Web Service Editor opens.
4. On the **Main Settings** page, complete the following task:
   -
5. Click the **HTTP Configuration** page and complete the following task:
   -
6. Click the **Advanced URL Settings** page and complete the following task:
   -
7. Click the **Advanced Settings** page and complete the following task:
   -
8. Click the **Authentication Settings** page and complete the following task:
   -
9. Click the **Preferences** page and complete the following task:
   -
10. Click the **User Information** page and complete the following task:
    -
11. Click the **Debug Settings** page and complete the following task:
    -
12. Click the **Properties and Names** page and complete the following tasks:

- *Naming and Describing an Object* on page 217

  You can instead enter a name and description when you save this content web service.
- *Localizing the Name and Description for an Object* on page 342 (optional)
- *Managing Object Properties* on page 219 (optional)

The default security for this content web service is based on the security of the parent folder. You can change the security when you save this content web service (on the **Security** tab page in the Save As dialog box), or by editing this content web service (on the **Security** page of the Content Web Service Editor).

Portal administrators with at least Select access to this content web service can create content sources based on the web service.

Providing Access to External Content with Remote Content Sources

Content sources provide access to external content repositories, enabling users to submit documents and content managers to create content crawlers to import documents into the Knowledge Directory. Each content source is configured to access a particular document repository with specific authentication. 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.

Before you create a remote content source, you must:

- Install the crawl provider on the computer that hosts the portal or on another computer
- Create a remote server pointing to the computer that hosts the crawl provider (optional, but recommended)
- Create a crawler web service on which to base this content source

To create a remote content source you must have the following rights and privileges:

- Access Administration activity right
- Create Content Sources activity right
- At least Edit access to the parent folder (the folder that will store the content source)

1. Click **Administration**.
2. Open the folder in which you want to store the content source.
3. In the **Create Object** drop-down list, click **Content Source - Remote**.

The Choose Web Service dialog box opens.

4. Select the web service that provides the basic settings for your content source and click **OK**. The Remote Content Source Editor opens, displaying the Main Settings page.
5. Complete the tasks on **Main Settings** page:

   - If necessary, edit the content web service associated with this content source by clicking the web service name.
   - *Gatewaying Imported Content* on page 183

   Note: Depending on what type of remote content source you are creating, you might see additional settings and additional pages.

6. Click the **Properties and Names** page and complete the following tasks:

   - *Naming and Describing an Object* on page 217

     You can instead enter a name and description when you save this content source.

   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219(optional)

Users with at least Select access to this content source can now submit documents from this content source or create content crawlers that will access this content source.

Providing Access to Web Content with Web Content Sources

Content sources provide access to external content repositories, enabling users to submit documents and content managers to create content crawlers to import documents into the Knowledge Directory. Each content source is configured to access a particular document repository with specific authentication. 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.

Note: The World Wide Web content source, created upon install, provides access to any unsecured web site.

To create a web content source you must have the following rights and privileges:

- Access Administration activity right
- Create Content Sources activity right
- At least Edit access to the parent folder (the folder that will store the content source)

1. Click **Administration**.

2. Open the folder in which you want to store the content source.
3. In the **Create Object** drop-down list, click **Content Source - WWW**.
   The Content Source Editor opens.
4. Complete the tasks on **Main Settings** page:
   - *Providing Access to Web Content by Impersonating a User* on page 181
   - *Gatewaying Imported Content* on page 183
   - *Selecting a Web Service for Gatewayed Content* on page 180 (only necessary if you selected to gateway content)
5. Click the **Proxy Server Configuration** page and complete the following task:
   - *Providing Access to Web Content Through a Proxy Server* on page 179
6. Click the **Login Form Settings** page and complete the following task:
   - *Providing Access to Web Content Through a Login Form* on page 181
7. Click the **Cookie Information** page and complete the following task:
   - *Providing Access to Web Content Through Cookies* on page 182
8. Click the **Header Information** page and complete the following task:
   - *Providing Access to Web Content Through Header Information* on page 183
9. Click the **Properties and Names** page and complete the following tasks:
   - *Naming and Describing an Object* on page 217

     You can instead enter a name and description when you save this content source.
   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219(optional)

Users with at least Select access to this content source can now submit documents from this content source or create content crawlers that will access this content source.

Providing Access to Web Content Through a Proxy Server

If you use a proxy server to access the internet, you can specify the proxy server settings on the Proxy Server Configuration page of the Content Source Editor.

1. Open the Content Source Editor by creating a new web content source or editing an existing one.
2. Click the **Proxy Server Configuration** page.
3. In the **Address** box, type the name of your proxy server.
4. In the **Port** box, type the port number for your proxy server.
5. If this proxy server requires security information:
   a) In the **User name** box, type the name of the user you want the portal to impersonate to access this proxy server.
   b) In the **Password** box, type the password for the user you specified.
   c) In the **Confirm** box, type the password again.
6. If you do not require the proxy server to access computers hosted on your local network, select **Bypass proxy server for local addresses**.
7. If there are other sites that do not require the proxy server, in the **Do not use for addresses beginning with** box, type the base URLs of these web sites.

   Separate multiple URLs with semicolons (;).

Selecting a Web Service for Gatewayed Content

If you selected to gateway the content from this content source, you need to select a web service to associate with this content source. You can specify that information on the Main Settings page of the Web Content Source Editor.

1. Open the Content Source Editor by creating a new web content source or editing an existing one.
   The **Main Settings** page opens.
2. Under **Web Service**, associate a content web service with this content source:

   This section appears only if you selected to gateway content.

   - To associate an existing content web service, click **Browse**; then, in the **Choose Web Service** dialog box, choose a content web service and click **OK**.
   - To remove the association, click **Remove**.
   - To edit the associated content web service, click its name.

Providing Access to Web Content by Impersonating a User

If the web site accessed by this content source requires a specific user name and password to access the site, you can specify that information on the Main Settings page of the Web Content Source Editor.

1. Open the Content Source Editor by creating a new web content source or editing an existing one.
   The **Main Settings** page opens.
2. Under **Target Site Security**, specify the security information required to access this web site:
   a) In the **User name** box, type the name of the user that this portal will impersonate to access content from this web site.
   b) In the **Password** box, type the password for the user you entered in Step a.
   c) In the **Confirm password** box, type the same password you entered in Step b.

Providing Access to Web Content Through a Login Form

If the web site accessed by this content source requires users to complete a form to access the site, you can specify the login form settings on the Login Form Settings page of the Content Source Editor.

1. Open the Content Source Editor by creating a new web content source or editing an existing one.
2. Click the **Login Form Settings** page.
3. In the **Login URL** box, type the URL to the login form that needs to be completed.
4. In the **Post URL** box, type the URL to which this login form posts data.
   To find the URL, search the form's source HTML for the <FORM> tag; the ACTION attribute contains the URL to which the form posts.
5. Under **Form Fields**, specify the information needed to gain access to this site:
   To determine this information, you can either contact the person who wrote the form or search the form's source HTML for each <INPUT> tag.
   - To add information for an <INPUT> tag:
     1. Click  **Add**.
     2. In the **Name** box, type the text after "name=" from the <INPUT> tag.

For example, if the form includes "<INPUT type="password" name="Password" size="10">, type `Password`.

3. In the **Value** box, type the text you would normally type in the form field.

   Using the example from the previous step, you would type the password needed to access the site.

- To remove a name/value pair, select the name/value and click ✕.

  To select or clear all of the name/value pair boxes, select or clear the box to the left of **Name**.

Providing Access to Web Content Through Cookies

If the web site accessed by this content source requires information to be sent in the form of cookies, you can specify the cookie settings on the Cookie Information page of the Content Source Editor.

1. Open the Content Source Editor by creating a new web content source or editing an existing one.
2. Click the **Cookie Information** page.
3. Determine what cookie information you need to send through one of the following methods:
   - Contact the person who wrote the form.
   - Viewing the cookies through your internet browser:
     1. Set your internet browser to prompt you before accepting cookies.

        Refer to your browser's online help for instructions.
     2. Navigate to the web site this content source will access.
     3. When prompted to accept a cookie, view the cookie information.

        For each cookie you receive, make note of the name and data values and the base URL from which it was sent.

4. Under **Cookies**, specify the cookie information needed to gain access to this site:
   - To add information for a cookie:
     1. Click **Add**.

2. In the **Name** box, type the text displaying in the Name field for the cookie.
3. In the **Value** box, type the text displaying in the Data field for the cookie.
4. In the **Cookie URL** box, type the base URL from which the cookie was sent.

   For example, if you need a cookie to access all areas of a web site, you might type `http://www.mysite.com`, but if the cookie is needed to access only a particular area of a web site, you might type `http://www.mysite.com/securedcontent`.

- To remove a cookie, select the cookie and click ✕.

  To select or clear all of the cookie boxes, select or clear the box to the left of **Name**.

Providing Access to Web Content Through Header Information

If the web site accessed by this content source requires header information to access the site, you can specify the header information on the Header Information page of the Content Source Editor.

1. Open the Content Source Editor by creating a new web content source or editing an existing one.
2. Click the **Header Information** page.
3. Paste the required header information into the text box if one of the following is true:
   - The web site accessed by this content source only responds to requests with specific information in the included HTTP header.
   - Your proxy server sends requests beyond your firewall only if there is specific information in the header.

Gatewaying Imported Content

When users click a link to an imported document, they can either be directed to the actual location of the source document or the content can be gatewayed, and the user will be redirected to a URL (generated from the settings in your portal configuration file) that, in turn, displays the document. Gatewaying content allows users to view documents they might not otherwise be able to view, either due to security on the source repository or firewall restrictions. You configure gateways settings on the Main Settings page of the Content Source Editor.

1. Open the Content Source Editor by creating a new web content source or editing an existing one.

The **Main Settings** page opens.

2. Under **URL Type**, specify what happens when users follow document links:

   - If you want to direct users to the actual location of the document, choose **Does not use the Gateway to open documents**. Be warned, however, that with this option, even users with access to this content source's documents will not be able to open the documents if the documents are not publicly available and the users are not connected to your network.
   - If you want to redirect users to a URL (generated by the settings in your portal configuration file) that, in turn, displays the document, choose **Uses the gateway to open documents**.

     Note:

     - If you want your users to be able to view documents even when they are not connected to your network, you should choose this option.
     - If the associated content web service supports content upload (specified on the Advanced Settings page of the Content Web Service Editor), you must use the gateway or content uploads will fail.

   By default web content sources do not gateway content, whereas remote content sources do.

About Importing Content with Content Crawlers

Content crawlers enable you to import content into the portal. Web content crawlers enable you to import content from web sites. Remote content crawlers enable you to import content from external content repositories such as a Windows NT file system, Documentum, Microsoft Exchange, or Lotus Notes.

Crawl Providers

A crawl provider is a piece of software that tells the portal how to use the information in the external content repository. BEA provides several crawl providers: World Wide Web (WWW) (installed with the portal software), Windows NT File (included with the portal software), Documentum, Microsoft Exchange, and Lotus Notes. If your content resides in a custom system, such as a custom database, you can import it by writing your own crawl provider using the IDK.

Note:

- Your portal administrator must install the crawl provider before you can create the associated content web service. For information on obtaining crawl providers, contact *ALUIsupport@bea.com*. For information on installing crawl providers, refer to the *Installation Guide for AquaLogic Interaction* (available on *edocs.bea.com*) or the documentation that comes with your crawl provider, or contact your portal administrator.
- To learn about developing your own crawl provider, refer to the *BEA AquaLogic User Interaction Development Center*.
- For a summary of AquaLogic Interaction crawl providers, as well as guidelines on best practices for deploying content crawlers, see the *Deployment Guide for BEA AquaLogic User Interaction*.

Content Web Services

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

Content Sources

Content sources provide access to external content repositories, enabling users to submit documents and content managers to create content crawlers to import documents into the Knowledge Directory. Each content source is configured to access a particular document repository with specific authentication. 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.

Best Practices

- 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 web 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*.

Importing Content from External Document Repositories with Remote Content Crawlers

You can create a remote content crawler to import content (and security) from external document repositories.

Before you create a remote content crawler, you must:

- Install the content provider on the computer that hosts the portal or on another computer.
- Create a remote server.
- Create a content web service.
- Create a content source.
- Create the folders in which you want to store the imported content.
- Create and apply any filters to the folders to control the sorting of imported content.
- Create any users and groups to which you want to grant access to the imported content.

To create a remote content crawler you must have the following rights and privileges:

- Access Administration activity right
- Create Content Crawlers activity right
- At least Edit access to the parent folder (the folder that will store the content crawler)
- At least Select access to the content web service on which this content crawler will be based
- At least Select access to the folders in which you want to store the imported content

- At least Select access to the users and groups to which you want to grant access to the imported content

1. Click **Administration**.
2. Open the folder in which you want to store the content crawler.
3. In the **Create Object** drop-down list, click **Content Crawler —Remote**.
   The Choose Content Source dialog box opens.
4. Select the content source that provides access to the content you want to crawl and click **OK**.
   The Remote Content Crawler Editor opens.
5. Complete the following tasks on the **Main Settings** page:
   - Setting a Target Folder for Imported Content
   - Mirroring the Source Folder Structure
   - Automatically Approving Imported Content
   - Importing Content Security
   - Granting Access to Imported Content
6. Click the **Document Settings** page and complete the following tasks:
   - Specifying When Imported Documents Expire
   - Specifying Refresh Settings for Imported Links and Property Values
7. Click the **Content Type** page and complete the following task:
   - Assigning Content Types to Imported Content
8. Click the **Advanced Settings** page and complete the following tasks:
   - Selecting a Language for Imported Content
   - Specifying What to Do with Rejected Documents
   - Specifying What to Do On Subsequent Crawls
   - Marking Imported Content with a Content Crawler Tag
   - Specifying Maximum Threads Settings
9. Click the **Set Job** page and complete the following task:
   - *Associating an Object with a Job* on page 300
10. Click the **Properties and Names** page and complete the following tasks:
    - *Naming and Describing an Object* on page 217

You can instead enter a name and description when you save this content crawler.

- *Localizing the Name and Description for an Object* on page 342 (optional)
- *Managing Object Properties* on page 219(optional)

The default security for this content crawler is based on the security of the parent folder. You can change the security when you save this content crawler (on the Security tab page in the Save As dialog box), or by editing this content crawler (on the Security page of the Content Crawler Editor).

To import content, run the job you associated with this content crawler.

Importing Web Content with Web Content Crawlers

You can create a web content crawler to import content from web sites and RSS feeds.

Before you create a web content crawler, you must:

- Create a content source, if necessary, to access secured content.
- Create the folders in which you want to store the imported content.
- Create and apply any filters to the folders to control the sorting of imported content.
- Create any users and groups to which you want to grant access to the imported content.

To create a web content crawler you must have the following rights and privileges:

- Access Administration activity right
- Create Content Crawlers activity right
- At least Edit access to the parent folder (the folder that will store the content crawler)
- At least Select access to the content web service on which this content crawler will be based
- At least Select access to the folders in which you want to store the imported content
- At least Select access to the users and groups to which you want to grant access to the imported content

1. Click **Administration**.
2. Open the folder in which you want to store the content crawler.
3. In the **Create Object** drop-down list, click **Content Crawler — WWW**.
   The Choose Content Source dialog box opens.
4. Select the content source that provides access to the content you want to crawl and click **OK**.
   The Web Content Crawler Editor opens.
5. Complete the following tasks on the **Main Settings** page:

- Defining Where and How Far to Crawl
- Setting a Target Folder for Imported Content
- Automatically Approving Imported Content
- Granting Access to Imported Content

6. Click the **Web Page Exclusions** page and complete the following task:
   - Avoiding Importing Unwanted Web Content
7. Click the **Target Settings** page and complete the following task:
   - Specifying a Time-Out Period for a Web Content Crawler
8. Click the **Document Settings** page and complete the following tasks:
   - Specifying When Imported Documents Expire
   - Specifying Refresh Settings for Imported Links and Property Values
9. Click the **Content Type** page and complete the following task:
   - Assigning Content Types to Imported Content
10. Click the **Advanced Settings** page and complete the following tasks:
    - Selecting a Language for Imported Content
    - Specifying What to Do with Rejected Documents
    - Specifying What to Do On Subsequent Crawls
    - Marking Imported Content with a Content Crawler Tag
    - Specifying Maximum Threads Settings
11. Click the **Set Job** page and complete the following task:
    - *Associating an Object with a Job* on page 300
12. Click the **Properties and Names** page and complete the following tasks:
    - *Naming and Describing an Object* on page 217

      You can instead enter a name and description when you save this content crawler.
    - *Localizing the Name and Description for an Object* on page 342 (optional)
    - *Managing Object Properties* on page 219(optional)

The default security for this content crawler is based on the security of the parent folder. You can change the security when you save this content crawler (on the Security tab page in the Save As dialog box), or by editing this content crawler (on the Security page of the Content Crawler Editor).

To import content, run the job you associated with this content crawler.

Refreshing Content from Content Crawlers

You can refresh metadata and import new content from content crawlers that have previously imported content.

If you are editing an existing content crawler, you see the section Importing Documents. Under Importing Documents, specify whether to import only new documents. By default, the content crawler attempts to import only new documents (those that have not been previously imported by this content crawler or other content crawlers that access this same content source). You can change the content crawler setting to import multiple copies of each document, which might be useful while testing your content crawlers. You can also specify whether the content metadata should be updated.

1. To import only new documents, select **Import only new links**.
   New options display.

   If you want to import all content again the next time this content crawler runs, leave the option unselected and skip the rest of the steps.

2. Specify what new links means:

   - To import only those documents that have not been previously imported by this content crawler, choose **by this Content Crawler**.
   - To import only those documents that have not been imported from the associated content source (either by this content crawler, another content crawler, or manually by a user), choose **from this Content Source**.

   Note: The option you choose here also applies to the rejection history and deletion history. For example, if you select **from this Content Source**, the rejection history includes content rejected by any content crawler that has crawled the content source.

3. To refresh the previously imported documents as specified on the Document Settings page, select **refresh them**.

   Generally, refreshing documents is the job of the Document Refresh Agent; refreshing documents slows the content crawler down. However, if you changed the document settings

for this content crawler or changed the property mappings in the associated content types, refreshing documents updates these settings for the previously imported documents.

Note: If you are crawling an RSS feed, the **refresh them** option refreshes the properties (such as the title and description) with the values from the target documents, not the RSS feed. If you want to retain the properties from the RSS feed, do not select **refresh them**.

4. If you created additional folders or applied different filters to destination folders, select **try to sort them into additional folders** to sort the previously imported documents into new Knowledge Directory folders.

   Another content crawler might have imported documents from the same content source but into different folders than the destination folders specified for this content crawler. Make sure you really want to re-sort those documents into the destination folders specified for this content crawler.

5. To re-import documents that were previously deleted (manually, due to expiration, or due to missing source documents), select **regenerate deleted links**.

   Note: This might re-import documents that were at one time deemed inappropriate for your portal.

If absolutely necessary, you can delete the history of documents that have been deleted from the portal. Remember that the deletion history is defined by what you specified as new documents in Step 2.

- If you chose **by this Content Crawler**, the history includes all documents imported by this content crawler that have been deleted.
- If you chose **from this Content Source**, the history includes all documents imported from this content source that have been deleted. Therefore, you are deleting the history for all content crawlers that import documents from this content source.

If you are still sure that you must delete the record of documents deleted from the portal, click **Clear Deletion History**.

Testing a Content Crawler

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.

Create a test folder and 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.

1. 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.

2. 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.

3. 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 **Search Only in this Folder**. The links that are returned by your search are for the documents that will pass your filter.

Troubleshooting the Results of a Crawl

There are several things you can troubleshoot if your content crawler does not import the expected content.

- Make sure your folder filters are correctly filtering content.

  To learn about testing your filters, see *Testing Filters* on page 169.

- Make sure your content crawler did not place unwanted content into the target folder.

If a document does not filter into any subfolders, your content crawler might place the document in the target folder. This is determined by a setting on the Main Settings page of the Folder Editor.

- Make sure the content crawler did not place content into the Unclassified Documents folder.

  If a document cannot be placed in any target folders or subfolders, your content crawler might place the document in the Unclassified Documents folder. This is determined by a setting on the Advanced Settings page of the Content Crawler Editor.

  If you have the correct permissions, you can view the Unclassified Documents folder when you are editing the Knowledge Directory or by clicking **Administration**, then, in the **Select Utility** drop-down list, select **Access Unclassified Documents**.
- Make sure you have at least Edit access to the target folder.
- For web content crawlers, make sure the robot exclusion protocols or any exclusions or inclusions are not keeping your content crawler from importing the expected content.

  This is determined by a setting on the Web Page Exclusions page of the Content Crawler Editor.
- Make sure the authentication information specified in the associated content source allows the portal to access content.
- Review the job history for additional information.

Example of Importing Security

Assume that you create an authentication source called *myAuthSource* importing users and groups into the portal from a source domain called *myDomain*. This authentication source uses the category *Employees*. Therefore, the text "Employees\" is prepended to each user's name and each group's name to distinguish these users and groups from those imported through other authentication sources. For example, if you have a user *myDomain\Mary* in the source domain, the user is imported into the portal as *Employees\Mary*.

Every authentication source automatically creates a group that includes all the users imported through that authentication source. In this example, because the authentication source is called *myAuthSource*, the group that includes all imported users is called *Everyone in myAuthSource*.

Suppose that you want to import content from a Lotus Notes system called *myNotes*, which includes users and groups equivalent to those found in the myDomain domain. Because you have already imported these groups and users into the portal, your Notes content crawler can import Notes security information along with each Notes document. The groups in the Notes system do not have to have the same names as their corresponding groups in the myDomain domain or in the

portal; the important thing is that there are Notes groups that have equivalent portal groups. If there are Notes groups that do not have equivalent groups in the portal, your Notes content crawler will ignore security settings referring to such groups.

When your Notes content crawler finds a document, it creates a list of the Notes groups that have access to it. This list is called an ACL (Access Control List). The ACLs created for Notes documents do not contain entries for specific Notes users, only for Notes groups. (Notes content crawlers only grant access to portal groups. Windows File content crawlers do grant access to portal users.) Each ACL entry is written as {Notes Server Name}\{Notes Group Name}. In this example, the content crawler creates an ACL with the single entry *myNotes\Engineering*, because this is the only Notes group that has access to that document.

The content crawler then refers to the Global ACL Sync Map to determine which portal group corresponds to this Notes group. This is a two-stage process:

1. Knowing that you would import documents and security through Notes content crawlers, on the Prefix: Domain Map page, you mapped the myAuthSource category *Employees* to the source domain *myNotes*. Guided by this entry, your content crawler modifies the ACL entry from *myNotes\Engineering* to *Employees\Engineering*.
2. Knowing that your Notes system uses a different group name than your myDomain domain, on the Portal: External Group Map page, you mapped the Notes system group *Engineering* to the myDomain group, now, the portal group, *Developers*. Guided by this entry, your content crawler modifies the ACL entry from *Employees\Engineering* to *Employees\Developers*.

As a result, all the users in the portal group *Developers* are automatically granted access to the document.

Destination Folder Flow Chart

This flow chart shows how a content crawler determines into which folders to import content. The process starts in the upper-left corner. The content crawler goes through this process for the destination folder you select and then continues down the levels of subfolders, if necessary. The content crawler repeats this process for each destination folder you add to the Main Settings page of the Content Crawler Editor.

If the content crawler is set to ignore the filters of destination folders, the first step in this flow chart is treated as if the document passes the filters for the folder. Be aware that only the filters of the destination folders will be ignored; if the destination folder has any *subfolders* with filters, these subfolder filters will not be ignored.

Note: If the document does not pass the filters of the destination folder, the content crawler checks to see if the destination folder has a default folder. It is only for subfolders of the destination folder that the content crawler checks to see if the *parent folder* has a default folder.

Metadata Imported by Content Crawlers

Content crawlers index the full document text, but some content crawlers can import additional metadata.

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

Creating a Snapshot Query to Display the Results of a Search in a Portlet or an E-mail Message

Snapshot portlets enable you to display the results of a search 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 you must have the following rights and privileges:

- Access Administration activity right
- Create Snapshot Queries activity right
- At least Edit access to the parent folder (the folder that will store the snapshot query)
- At least Select access to any properties by which you want to filter your results

- At least Select access to any Knowledge Directory or administrative folders to which you want to restrict your results

1. Click **Administration**.
2. Open the folder in which you want to store the snapshot query.
3. In the **Create Object** drop-down list, click **Snapshot Query**.
   The Snapshot Query Editor opens.
4. On the **Construct Snapshot Query** page, complete the following tasks:
   - *Defining Snapshot Query Conditions* on page 198
   - *Limiting a Snapshot Query* on page 200
5. Click the **Format Snapshot Query Result** page and complete the following task:
   - *Formatting the Results of a Snapshot Query* on page 201
6. Click the **Preview Snapshot Query Result** page and complete the following task:
   - *Previewing the Results of a Snapshot Query* on page 202
7. Click the **Snapshot Portlet List** page and complete the following task:
   - *Creating a Snapshot Portlets to Display the Results of a Snapshot Query* on page 204
8. Click the **Properties and Names** page and complete the following tasks:
   - *Naming and Describing an Object* on page 217

     You can instead enter a name and description when you save this <object>.
   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219 (optional)
9.

If you did not create a snapshot portlet on the **Snapshot Portlet List** page, a snapshot portlet is automatically created when you save this snapshot query.

The default security for this snapshot query is based on the security of the parent folder. You can change the security when you save this snapshot query (on the **Security** tab page in the Save As dialog box), or by editing this snapshot query (on the **Security** page of the Snapshot Query Editor).

Defining Snapshot Query Conditions

A snapshot query is a combination of a basic fields search and statements. The basic fields search operates on the name, description, and content of documents and objects. Statements can operate on the basic fields or any other additional document or object properties. Statements define what must or must not be true to return the document or object in the results. The statements are collected together in groupings. The grouping defines whether the statements are evaluated with an AND operator (all statements are true) or an OR operator (any statement is true). If some statements should be evaluated with an AND operator and some should be evaluated with an OR operator, you can create separate groupings for the statements. You can also create subgroupings or nested groupings, where one grouping is contained within another grouping. The statements in the lowest-level grouping are evaluated first to define a set of results. Then the statements in the next highest grouping are applied to that set of results to further filter the results. The filtering continues up the levels of groupings until all the groupings of statements are evaluated.

A snapshot query needs at least a basic fields search or a statement.

1. If the Snapshot Query Editor is not already open, open it now and display the **Construct Snapshot Query** page.
2. To search the name, description, and content values, type the text you want to search for in the **Basic fields search** text box.

   You can use the text search rules described in *Using Text Search Rules* on page 364.
3. Select the operator for the grouping of statements you are about to create:

   - If a document or object should be returned only when all statements in the grouping are true, select **AND**.
   - If a document or object should be returned when any statement in grouping is true, select **OR**.

   Note: The operator you select for a grouping applies to all its statements and subgroupings directly under it.
4. Define each statement in the grouping:

   a) Click **Add Statement**.

   b) In the first drop-down list, select the searchable property for which you want to filter the values.

   c) In the second drop-down list, select the operator to apply to this condition.

   This list will vary depending on the property selected:

- For any text property, you can search for a value that contains your search string, or you can search for properties that have never had a value (**Contains No Value**).

  Note: If the property contained a value at some point, but the value has been deleted, the property will not match the Contains No Value condition.

- For any date property, you can search for a value that comes after, comes before, is, or is not the date and time you enter in the boxes. You can also search for a value within the last number of minutes, hours, days, or weeks that you enter in the box.
- For any number property, you can search for a value that is greater than, is less than, is, is not, is greater than or equal to, or is less than or equal to the number you enter in the text box.

d) In the box (or boxes), specify the value the property must meet.

Note: If you are searching for a text property, you can use the text search rules described in *Using Text Search Rules* on page 364.

To remove the last statement in a grouping, select the grouping, and click **Remove Statement**.

5. If necessary, add more statements by repeating Step 4.
6. If necessary, add more groupings:

- To add another grouping, select the grouping to which you want to add a subgrouping, click **Add Grouping**, then define the statements for that grouping (as described in Step 4).

  Note: You cannot add a grouping at the same level as **Grouping 1**.

- To remove a grouping, select the grouping, and click **Remove Grouping**.

  Note:

  - Any groupings and statements in that grouping will also be removed.
  - You cannot remove **Grouping 1**.

You might also want to limit your search by language, object types, or folders.

Limiting a Snapshot Query

You can limit your snapshot query to specific languages, portal repositories, objects, folders, projects, or portlets.

1. If the Snapshot Query Editor is not already open, open it now and display the **Construct Snapshot Query** page.
2. To limit your search to a specific language, under **Limit Search by Document Language**, select a language from the **Specify Language** drop-down list.
3. Under **Specify Range of Search**, select which of the repositories to search: Knowledge Directory, Portal (administrative objects), Collaboration, Publisher.

   You will see check boxes for Collaboration or Publisher only if you have these products installed.
4. Next to **Repository General Settings**, choose whether to search folders or documents from any source within the repositories you selected, and whether to search all subfolders within those repositories.
5. If you selected the Knowledge Directory as one of the repositories to search, specify settings under **Knowledge Directory Search Settings**.

   - Next to **Search Results Contain**, select **Knowledge Directory Folders**, and/or **Knowledge Directory Documents**.

     Note: You must select at least one of these options.

   - To restrict the search to selected folders, click  **Add Document Folder**, in the **Add Document Folder** dialog box, select the folders to which you want to restrict your search and click **OK**.
   - To remove a folder, select it and click  .
6. If you selected the portal as one of the repositories to search, specify settings under **Portal Search Settings**.

   - Next to **Search Results Contain**, select the portal administrative object types to include in the search, such as portlets or communities.

     Note: To select or clear all the object types, select or clear the box next to **All Types**.

- To restrict the search to selected folders, click  **Add Administrative Folder**, in the **Add Administrative Folder** dialog box, select the folders to which you want to restrict your search and click **OK**.
- To remove a folder, select it and click  .

7. If you selected Collaboration as one of the repositories to search, under **Restrict to Selected Collaboration Projects**, click **Add Project** to select one or more Collaboration projects to search, then click **Finish**.
8. If you selected Publisher as one of the repositories to search, under **Restrict to Selected Publisher Portlets**, click **Add Publisher Portlet** to select one or more Publisher portlets to search, then click **Finish**.

Next, specify the format for your results.

Formatting the Results of a Snapshot Query

You can define how your snapshot query results appear. By default, results are listed in order of relevance; that is, those results that most closely match your query are listed first. You can change the order in which results are displayed, limit the number of items returned, specify a style in which the snapshot portlet will be displayed, and e-mail results to users.

1. If the Snapshot Query Editor is not already open, open it now and display the **Format Snapshot Query Result** page.
2. In the **Maximum items displayed** field, type the number of items that should appear on a page.
3. In the **Order results by** drop-down list, select the property type by which you want to sort results.

   For example, you can sort your search results by Content Type ID or Object Last Modified.
4. To select the available fields for display on search results, under **Query Return Fields**, click  **Add Query Fields**, select the fields you want to add, and click **OK**.

   The fields you add here can be selected in the administrative preferences of snapshot query portlets associated with this snapshot query. Selecting all or a subset of these fields in the administrative preferences of a particular snapshot query portlet determines what end users see in results appearing in that portlet.
5. If you want the content snapshot portlet to appear with a subscribe button that enables users to receive e-mail about search results, select**Enable e-mail subscriptions**.

Note:

- You must configure an external operation to send e-mail notifications for this snapshot query. See *E-mailing the Results of a Snapshot Query* on page 202.
- Users receive the e-mail only if their e-mail addresses are available in their user profiles.

Next, preview your results.

Previewing the Results of a Snapshot Query

You can preview the results of a snapshot query before you save it.

If the Snapshot Query Editor is not already open, open it now and display the **Preview Snapshot Query Result** page.
The fields displayed in these results are the ones you added on the **Format Snapshot Query Result** page, under **Query Return Fields**. However, for each snapshot portlet associated with this query, you can select all or a subset of the available query return fields in the portlet's administrative preferences.

Next, create a snapshot portlet on the **Snapshot Portlet List** page or save the snapshot query to automatically create a snapshot portlet.

E-mailing the Results of a Snapshot Query

You can e-mail the results of a snapshot query to users by creating an external operation and editing `SavedSearchMailer.bat` or `SavedSearchMailer.sh`.

Before you create an external operation to e-mail the results of a snapshot query, you must create the snapshot query and snapshot portlet.

To create an external operation you must have the following rights and privileges:

- Access Administration activity right
- At least Edit access to the Snapshot Query Mailer external operation
- At least Edit access to the parent folder (the folder that will store the external operation)
- At least Select access to the job that will run this external operation or Create Jobs activity right to create a job to run this external operation

1. If you have not already done so, edit `SavedSearchMailer.bat` or `SavedSearchMailer.sh` to specify the settings for your mail server and customize the e-mail values.

The saved search mailer file is located on the computer that hosts the Automation Service, iin *Install_Dir*\scripts (for example, C:\bea\alui\ptportal\6.5\scripts or /opt/bea/alui/ptportal/6.5/scripts).

You must replace the following argument values:

Argument	Description
SENDER	The name you want to display as the From value in the automated e-mails
MAIL_SERVER	Your SMTP mail server
REPLYTO	The e-mail address that users can reply to from the automated e-mails

Optionally, you can replace the following argument values:

Argument	Description
USER	The name of the user you want to send the automated e-mails
PWD	The password for the user that will send the automated e-mails
MIMETYPE	The MIME type you want to use for the automated e-mails
SUBJECT	The text you want to display in the automated e-mail subject line By default the subject includes the name of the snapshot query (represented by <search_name>) and the name of the user receiving the results (represented by <name>).
BODY_HEADER	The text you want to display at the top of the automated e-mail body
BODY_SEPARATOR	Any code you want to use to generate a separation between the header and the results
BODY_FOOTER	The text you want to display at the bottom of the automated e-mail body

2. In the portal, click **Administration**.
3. Open the snapshot portlet for which you want to e-mail results.
4. Click the **Properties and Names** page.
5. Copy or make note of the **Object ID**, then close the snapshot portlet.
6. Open the **Intrinsic Operations** folder.
7. Select the **Snapshot Query Mailer** external operation and click  .

8. In the Target Folder dialog box, select the folder in which you want to store your new external operation and click **OK**.
9. Open the copy of the snapshot query mailer you just created.
10. Replace `200` in the arguments with the object ID of the snapshot query you want to e-mail.
11. Click the **Set Job** page and complete the following task:
    - *Associating an Object with a Job* on page 300
12. Click the **Properties and Names** page and rename the external operation.

Set the job to run on a regular basis.

Creating a Snapshot Portlets to Display the Results of a Snapshot Query

You can create a snapshot portlet to display the results of a snapshot query on a portal page.

- To create a content snapshot portlet associated with this snapshot query, click  **Create Content Snapshot Portlet**.
  The portlet appears under **Portlet List**, and is added to the same administrative folder as this snapshot query.

  Note:

  - If you create a content snapshot portlet manually (rather than having one automatically created when saving the snapshot query), the name of the portlet will be New Snapshot Query.
  - If you do not manually create a content snapshot portlet on this page, one will be created automatically when you save the snapshot query; the portlet will have the same name as the snapshot query.
  - To delete a snapshot portlet, you must delete it from the administrative folder that contains its associated snapshot query.

- To change the name of a snapshot portlet:
  a) Click the portlet name.
  The Portlet Editor opens.
  b) Click the **Properties and Names** page.
  c) Edit the name.

- To select the fields displayed in the results and the fields that users can search on.for a snapshot portlet, edit the administrative preferences:
  a) Click the portlet name.
     The Portlet Editor opens.
  b) Click the **Edit** button next to **Configure this Portlet**.
  c) Edit the preferences.

Users with at least Select access to this portlet can now add this portlet to their My Pages. Community administrators with at least Select access to this portlet can now add this portlet to their communities.

Managing Administrative Objects and Portal Utilities

The Administrative Objects Directory allows you to create and manage administrative objects and access portal utilities. It displays all objects to which you have at least Read access.

To display the Administrative Objects Directory click **Administration**.

Note: To access the Administrative Objects Directory you must have the Access Administration activity right.

- To view an administrative object, click its name. For details on viewing objects, see *Viewing Objects* on page 210.
- To search for objects, use one of the following methods:
  - *Searching for Objects in the Administrative Objects Directory* on page 210
  - *Searching for Objects or Documents Using Advanced Search* on page 211
- To create an object, open the folder in which you want to store the object, then, in the **Create Object** drop-down list, select the type of object you want to create.

  The list displays only those objects you have permission to create in the folder you are viewing.
- To access portal utilities, in the **Select Utility** drop-down list, select the utility you want to use.

  The list displays only those utilities to which you have access.
- To perform actions on the folder you are viewing, use the buttons in the folder toolbar (to the right of the folder title).

  For a description of these buttons, see *Folder Toolbar* on page 146.

- To perform actions on the objects in the folder, use the buttons in the action toolbar (to the right of the **Create Object** and **Select Utility** drop-down lists).

  For a description of these buttons, see *Action Toolbar* on page 215.

Creating or Editing an Administrative Folder

Administrative folders provide a hierarchical structure that make it easy to organize portal objects and manage security.

Tip: You might want to store all of the resources needed by a particular audience of users in the same folder in which you store those users. By securing the folder appropriately and applying experience definition settings to it you can create completely separate and discreet user experiences for each audience of users.

To create an administrative folder you must have the following rights and privileges:

- Access Administration activity right
- Create Admin Folders activity right
- At least Edit access to the parent folder (the folder in which you are creating the new folder)

1. Click **Administration**.
2. If necessary, open the folder in which you want to store the new folder.
3. In the **Create Object** drop-down list, click **Administrative Folder**.
   The Create Administrative Folder dialog box opens.
4. In the **Name** box, type a name for the folder.

   This name appears in lists of objects from which users will sometimes choose; therefore, the name should clearly convey the purpose of this folder.
5. In the **Description** box, type a description for the folder.

   This description appears in the Administrative Objects Directory to provide other administrators further details on the purpose of this folder.
6. Click **OK**.

You can perform additional tasks when you edit this folder:

- *Applying an Experience Definition to a Folder* on page 80
- *Localizing the Name and Description for an Object* on page 342

- *Managing Object Properties* on page 219
- *Viewing Top Best Bets for an Object* on page 219
- *Setting Security on an Object* on page 221
- *Approving an Object for Migration* on page 307
- *Viewing Import History for an Object* on page 222

Editing an Administrative Folder

Administrative folders provide a hierarchical structure that make it easy to organize portal objects and manage security.

To edit an administrative folder you must have at least Edit access to the folder.

1. Click **Administration**.
2. Select the folder you want to edit and click **Edit Subfolder**.
3. Perform tasks as necessary on the **Experience Definition Settings** page:
   - *Applying an Experience Definition to a Folder* on page 80
4. Perform tasks as necessary on the **Properties and Names** page:
   - *Naming and Describing an Object* on page 217
   - *Localizing the Name and Description for an Object* on page 342
   - *Managing Object Properties* on page 219
   - *Viewing Top Best Bets for an Object* on page 219
5. Perform tasks as necessary on the **Security** page:
   - *Setting Security on an Object* on page 221
6. Perform tasks as necessary on the **Migration History and Status** page:
   - *Approving an Object for Migration* on page 307
   - *Viewing Import History for an Object* on page 222

Viewing Objects

You can view objects through the Administrative Objects Directory and through some of the portal utilities.

Note: To view an object you must have the Access Administration activity right and at least Read access to the object.

To display the Administrative Objects Directory and access portal utilities, click **Administration**.

- To view the contents of a folder, click its name.
- To view all objects of a particular type stored in the folder, click the object type (for example, crawler, portlet, or user). To hide these objects, click the object type again.
- To expand all the object types so you can view all objects stored in this folder, click  .
- To change the column used for sorting or to toggle the sort order between ascending and descending, click the column name.

  You see an icon (▼ or ▲) to the right of the column name by which the objects are sorted.
- While viewing a subfolder, to navigate up to the parent folder, click  **Up**, or click the parent folder name in folder path.
- By default, only the first 50 objects display. To view the next 50 objects, click **Next >>**, or to view another set of 50 objects, click a number range.
- To view default profiles, in the **Select Utilities** drop-down list, select **Default Profiles**.

Searching for Objects in the Administrative Objects Directory

You can search for objects in the Administrative Objects Directory.

1. Click **Administration**.
2. Specify your criteria in the **Object Search** boxes:

- To limit your search to particular types of objects (for example, crawler, portlet, or user), in the first drop-down list, select the object type.
- To limit your search to particular folders, in the second drop-down list, choose whether to search all administrative folders, only the folder you are currently viewing, or the folder you are viewing plus its subfolders.
- In the text box, type the text you want to search for and click  .

Searching for Objects or Documents Using Advanced Search

You can perform an advanced search, using metadata properties and location, to find objects or documents.

In the portal banner or in the Administrative Objects Directory, click **Advanced Search**.

- To search for text in the name or description of an object or document, type the text you want to search for in the **Search** for text box.

  You can use the text search rules.

- To search for property values, click **Add Criteria**, and specify the property criteria in the boxes that appear:

  a) In the first drop-down list (property), select the searchable property for which you want to filter the values.

  b) In the second drop-down list (operator), select the operator to apply to this condition.

  This list will vary depending on the property selected:

  - For any text property you can search for a value that contains your search string (Contains), or you can search for properties that are blank (Contains No Value).

    Note: To exclude results with particular values, select **Contains**, then type "not" followed by words that you want to exclude from your search. For example, if you want to search for documents about retirement benefits, excluding pension plans, then type "retirement benefits" in the Search for text box, select Contains from the drop-down list, and type "not pension plan" in the text box.

- For any date property you can search for a value that comes after, comes before, is, or is not the date and time you choose or for a value that occurs in the last number of minutes, hours, days, or weeks that you specify.
- For any number property you can search for a value that is greater than, is less than, is, is not, is greater than or equal to, or is less than or equal to the number you enter in the text box.

c) In the **Value** text box, enter the value the property must have, or not have, depending on which operator you selected.

Note: If you are searching for a text property, you can use the text search rules.

- To remove a property condition, select the condition and click ✕ (next to **Add Criteria**).
- Specify how you want your search criteria handled:
  - To meet all the conditions you define, select **All Criteria**.

    Selecting All Criteria is equivalent to using AND.
  - If you want your search results to meet at least one of the conditions you define, select **Any Criterion**.

    Selecting Any Criterion is equivalent to using OR.
- To restrict your search to specific Knowledge Directory folders, click **Add Document Folder**. In the Select folder for search dialog box, select the folders you want to search and click **OK**.
- To restrict your search to specific Administrative Objects Directory folders, click **Add Administration Folder**. In the Select folder for search dialog box, select the folders you want to search and click **OK**.
- To remove folders from your list, select the folders and click ✕ (next to **Add Administration Folder**).

  To select or clear all folder boxes, select or clear the box next to **Folder Names**.
- Specify whether you want to include subfolders.

  By default, the portal searches subfolders. To exclude subfolders from your search, clear the box next to **Include subfolders**.
- To specify the number of results to display on a page, in the **Results per page** drop-down list, choose a value.

- To restrict your search to a specific language, in the **This language only** drop-down list, choose a language.
- To limit your search to specific object types, in the **Result Types** list, select the object types you want to search.

  To select or clear all object type boxes, select or clear the box next to **Object Type**.
- To set all search conditions back to the defaults, click **Clear**.
- To perform your search, click **Search**.

Complex Property Search Example

You can use multiple property criteria to define complex property searches. For example, if you want to find documents published after a certain date by a specific branch of a company, you could set the property criteria to the following values:

- First Criterion:
  - Property = Object Created
  - Operator = Comes After
  - Value = December 30, 2003

  Your search results would be limited to objects created after December 30, 2003.
- Second Criterion:
  - Property = Company
  - Operator = Contains
  - Value = Company A

  Your search results would be limited to objects where the company property contains Company A.
- Third Criterion:
  - Property = Address
  - Operator = Contains
  - Value = San Francisco

  Your search results would be limited to objects that contain San Francisco in the address.

- You would also want to limit your Result Types to Documents, so that only documents were returned in your results.

Opening an Object Editor

Many administrative tasks start with opening an object editor. You open an editor by creating a new object or editing an existing object.

Note: To edit an object you must have the Access Administration activity right, and you must have at least Edit access to the object.

- To create a new object see *Creating an Object* on page 214.
- To edit an object from Administration see *Editing an Object from Administration* on page 215.
- To edit an object from the search results page perform a search and click the name of the object you want to edit.
- To edit a community while viewing it, in the **My Communities** menu, select the community you wan to edit, then click **Edit Page**.

Creating an Object

You can create objects such as folders, portlets, and users in the Administrative Objects Directory.

Note: To create an object you must have the Access Administration activity right and the activity right associated with creating that type of object. You must also have at least Edit access to the parent folder (the folder that will store the object).

1. Click **Administration**.
2. If necessary, open the folder in which you want to store the object.

   You can create folders in the root of the Administrative Objects Directory, but other objects must be created in a folder.

   Note: You can create default profiles only from the Default Profiles Utility: in the **Select Utilities** drop-down list, select **Default Profiles**.
3. In the **Create Object** drop-down list, select the type of object you want to create.

The list displays only those objects you have permission to create in the folder.

Depending on the type of object you are creating, either the object editor opens, or a dialog box opens prompting you to choose a content source, template, or web service.

4. If necessary, select the appropriate content source, template, or web service and click **OK**.

   A remote authentication source, a remote content source, an outgoing federated search, or a profile source requires a web service. A portlet requires a template or web service. A content crawler requires a content source.

   The object editor opens.

5. Complete the editor.

Editing an Object from Administration

You can edit objects from the Administrative Objects Directory.

Note: You must have the Access Administration activity right and at least Edit access to the object you want to edit.

1. Click **Administration**.
2. Navigate to the object you want to edit.
3. Open the object editor:
   - To open the Folder Editor, select the folder, and click **Edit Subfolder**.
   - To open the Community Editor, click the community name, and click .
   - To open any other object's editor, click the name of the object.

Action Toolbar

You can use the buttons in the action toolbar to perform actions on the objects in a folder.

Button	Action
	Move an object to a different folder. Select one or more objects and click . In the Target Folder dialog box, expand the folders as necessary, select a folder, and click **OK**.

Button	Action
	Copy an object to a different folder. Select one or more objects and click  . In the Target Folder dialog box, expand the folders as necessary, select a folder, and click **OK**.
	Delete an object. Select one or more objects and click  . In the confirmation dialog box, click **Apply Now**.
	Modify the security settings on one or more objects. Select the objects and click  .
	Add one or more portlets or communities to one or more groups. You can use this bulk add operation to add portlets to users' My Pages or subscribe users to communities based on group membership. Select the portlets or communities and click  .
	Marks an object for migration to another portal. Select one or more objects and click  . In the Script Prompt dialog box, describe why you want this object migrated and click **OK**.
 **Edit Subfolder**	Edit a subfolder. Select a subfolder and click  **Edit Subfolder**. **Note:** You can only edit folders to which you have at least Edit access.
 **Edit Profile Layout**	Edit the portal layout for guest users and default profiles. Select a guest user or default profile and click  **Edit Profile Layout** to log in as that user and change the user's portal layout. **Note:** You see this button only when you are viewing users.
 **Send Invitation**	Send an invitation or access a previously generated invitation link. Select the invitation and click  **Send Invitation**. **Note:** You see this button only when you are viewing invitations.

Button	Action
Run Once	Set a job to run one time, immediately, without changing the schedule. Select a job and click **Run Once**. **Note:** You see this button only when you are viewing jobs.
Enable **Disable**	**Enable** makes a user or web service available for use. **Disable** makes a user or web service unavailable for use. Select a user or web service and click **Enable** or **Disable**. **Note:** You see these buttons only when you are viewing users or web services.

Naming and Describing an Object

You can add or edit a name and description for an object.

1. Open the object's editor by creating a new object or editing an existing object.
2. In the **Name** box, type a name for this object.

   The name should clearly convey what this object is or what it can be used for.

   If you are naming an authentication-only authentication source, this name appears in the Authentication Partners list when you create the associated synchronization-only authentication source.
3. In the **Description** text box, type a description for this object.

   The description should provide additional detail to convey what this object is or what it can be used for.

   If you are naming an authentication source that synchronizes users, this description appears in the Authentication Source drop-down list when users log in. Users imported from an external source must choose the appropriate authentication source to log in to the portal.
4. If your portal administrator did not set a mandatory object language, in the **Primary Language** drop-down list, select the language for the name and description you entered.

If your portal administrator did set a mandatory object language, you see the mandatory language instead of a drop-down list. You cannot change this setting. The name and description you entered must be in the mandatory language.

If a localized name and description is not available in a user's selected language, the user will see the name and description in the specified primary language.

If you want to add names and descriptions for other languages, see *Localizing Object Names and Descriptions* on page 341.

Localizing the Name and Description for an Object

You can localize object names and descriptions, so that users see the names and descriptions in their chosen language. For example, if you have an object called "Engineering," you could add "Ingénierie" as the localized entry for French. A user viewing the French user interface would see the Ingénierie as the object name, as well as any other names and descriptions localized into French.

Note:

- You can localize names and descriptions into only the languages supported by the portal.
- You cannot localize names and descriptions for users.
- You can localize the names and descriptions for all objects at the same time using the Localization Manager.

1. Open the object's editor by creating a new object or editing an existing object.
2. Select **Supports Localized Names**.
   The **Localized Names and Descriptions** section appears.
3. Add or edit the localized names and descriptions:
   - To add an entry for a language, click **New Localized Name**, then, in the Name and Description dialog box, enter the localized name and/or description, select the appropriate language, and click **Finish**.
   - To edit an existing entry, click the entry you want to change, then, in the Name and Description dialog box, edit the entry as necessary, and click **Finish**.
   - To remove existing entries, select the entries you want to remove and click .

     To select or clear all entries, select or clear the check box to the left of **Name**.

Viewing Top Best Bets for an Object

The Properties and Names page displays the top best bet terms set for this object. When users do a top best bet search on these terms, they go directly to this object instead of seeing the normal search results.

- To link to the end user view of this object, click the link under **URL**.

  You can use this URL to go directly to a top best bet. This can be useful if you want to direct users to an object or document related to a particular issue, but the object or document changes frequently. For example, you might want to direct customers to your current privacy statement, but you need to keep copies of older privacy statements in your portal for internal reference. You could create a top best bet that points to the current privacy statement and add a link to that top best bet on your customer account page. When your privacy statement is updated, you can change the top best bet without having to change any links you made to the privacy statement.

Managing Object Properties

You can add or edit properties for an object.

1. Open the object's editor by creating a new object or editing an existing object.
2. Under Object Properties, change the properties and values:

   - To add or delete properties for all objects of this type, click **Open Properties Map**.

     This displays the Global Object Property Map.

   - To enter values for properties, type the value in the box to the right of the property.

   Note: The Object Properties section does not display if you are creating or editing a property.

Associating an Object with a Job

On the Set Job page, you can associate an object with a new or existing job. You run jobs to import users with authentication sources, import content with content crawlers, run external operations, and import user information with profile sources.

Before you can run jobs, you must:

- 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 300.
- Register the Automation Service with the portal, as described in *Registering Automation Services* on page 298.
- Assign administrative folders to the registered Automation Services, as described in *Registering Job Folders to Run Jobs* on page 299.

1. Open the object's editor by creating a new object or editing an existing object.
2. Click the **Set Job** page.
3. Associate the object with one or more jobs:
   - To run this object with an existing job, click **Add Job**; then, in the Choose Jobs dialog box, select the jobs you want to add this object to and click **OK**.
   - To create a new job to run this object, click **Create Job**, then, in the Job Editor, schedule your job and click **Finish**.
   - To remove a job, select the job and click .

     To select or clear all of the job check boxes, select or clear the check box to the left of **Job Name**.
   - To edit a job, click the job name.

     If you added this object to an existing job, you might want to verify that the job is scheduled to run.
   - To change the order in which the jobs are sorted, click **Job Name**.

Changing the Owner of an Object

The owner of an object is displayed on the Set Job page. If you are a portal administrator, you can change the owner.

You might want to change object owner for several reasons:

- If the owner is deleted from the portal, you must assign an existing portal user as the object owner before you can run the job.
- When a job runs, it might require access to portal objects that the current owner does not have access to. For example, a content crawler needs access to the folders into which it imports content. You might need to change the owner to provide the proper access.

1. Open the object's editor by creating a new object or editing an existing object.
2. Click the **Set Job** page.
3. To change the owner, click **Change Owner**; then, in the Choose User dialog box, choose the user whom you want to make the object owner and click **OK**.

   Note: If you are not a portal administrator, you see the name of the owner, but you cannot change it.

Setting Security on an Object

By default, a new object inherits the security of the parent folder, but you can override the inherited security.

Note: You cannot override the inherited security for users; user security is always the same as the folder in which the user is stored. If you do not want a user to be returned in some users' searches, make sure those users are not allowed access to the folder in which the user is stored.

1. Open the object's editor by creating a new object or editing an existing object.
2. Click the **Security** page.
3. Specify which users and groups can access this object and what type of access they have:

   - To allow additional users or groups access to this object, click  **Add Users/Groups**.

- To specify the type of access a user or group has, in the drop-down list under the **Privilege** column, select the access type.

  For a description of the available privileges, see *About Access Privileges* on page 60.

  Note: If a user is a member of more than one group included in the list, or if they are included as an individual user and as part of a group, that user gets the highest access available to her for this object. For example, if a user is part of the Everyone group (which has Read access) and the Administrators Group (which has Admin access), that user gets the higher privilege to the community: Admin.

- To delete a user or group, select the user or group and click ✕.

  To select or clear all of the user and group check boxes, select or clear the check box to the left of **Users/Groups**.
- To see what users are included in a group, click the group name.
- To change the column used for sorting or to toggle the sort order between ascending and descending, click the column name.

  You see an icon (▾ or ▴) to the right of the column name by which the objects are sorted.

Viewing Import History for an Object

You can see if an object was imported from another portal.

1. Open the object's editor by creating a new object or editing an existing object.
2. Click the **Migration History and Status** page.

If this object was imported from another portal, information displays under **Import History**:

Column	Description
Migration Date	Displays the date the object was copied into this portal.
Migration Comment	Displays the comment entered by the user who requested migration of the object.

Column	Description
Source Portal UUID	Displays the unique identifier of the originating portal.

Requesting That an Object Be Migrated

You can request that an object be added to a migration package to be exported to another portal.

Note: You must have at least Edit access to the object for which you want to request migration.

1. Search for the object or click **Administration** and navigate to the object.
2. Select the object and click  .
3. In the Script Prompt dialog box, describe why you want this object migrated and click **OK**.

To view the status of your migration request, open the object's editor and click the **Migration History and Status** page. Under **Migration Status**, you see whether your request is waiting for approval, has been approved, or has been rejected, as well as your comments and any comments from the portal administrator.

Approving an Object for Migration

When users want an object to be migrated, they submit a migration request. A portal administrator can then approve the request, and the object is added to the migration package.

1. Open the object's editor by creating a new object or editing an existing object.
2. Click the **Migration History and Status** page.
   Under **Migration Status**, you see whether this object has been requested for migration, and, if so, whether it is waiting for approval, has been approved, or has been rejected.
3. If you are a member of the Administrators group, and you want to add this object to the migration package to be migrated to another portal, select **Approve this object for migration**.

   Note: Users who are not members of the Administrators group do not see this option.

After approving objects for migration, you can use the Migration - Export Utility to create a migration package.

Creating Remote Servers

Remote servers group together web services that are installed on the same computer and require the same type of authentication. With a remote server, you enter the base URL and authentication settings just once for multiple web services, and, if you need to move the web services, you just need to change the remote server settings.

To create a remote server you must have the following rights and privileges:

- Access Administration activity right
- Create Remote Server activity right
- At least Edit access to the parent folder (the folder that will store the remote server)

1. Click **Administration**.
2. Open the folder in which you want to store the remote server.
3. In the **Create Object** drop-down list, click **Remote Server**.
   The Remote Server Editor opens.
4. On the **Main Settings** page, complete the following tasks:
   - *Specifying the Location and Authentication Settings for a Remote Server* on page 66
5. Click the **Properties and Names** page and complete the following tasks:
   - *Naming and Describing an Object* on page 217

     You can instead enter a name and description when you save this remote server.
   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219 (optional)

The default security for this remote server is based on the security of the parent folder. You can change the security when you save this remote server (on the **Security** tab page in the Save As dialog box), or by editing this remote server (on the **Security** page of the Remote Server Editor).

About Extending Portal Services with Portlets

Portlets provide portal users customized tools and services as well as information. Portlets let you to integrate applications, tools, and services into your portal, while taking advantage of portal security, caching, and customization. Users can then add these portlets to their My Pages or to community pages.

Portlets

Portlets can be intrinsic or remote. An intrinsic portlet consists of one or more sets of code that are located on the portal computer. Your portal administrator needs to install this code in the correct location before an intrinsic portlet can be created. A remote portlet is a portlet hosted by a separate remote server. When a user displays a My Page or community page that includes a remote portlet, the portal contacts the appropriate remote server to obtain updated portlet content.

Some portlets can be placed only in certain areas of the page:

- Header portlets can be added to communities, community templates, and experience definitions to change the branding of these objects by replacing a 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 (specified in the page layout). Content canvas portlets can display across the entire width of the page or across one or two columns. You cannot add more than one content canvas portlet per page.

Portlet Web Services

Portlet web services allow you to specify *functional* settings for your portlets, leaving the *display* settings to be set in each associated portlet. There are intrinsic portlet web services and remote portlet web services.

An intrinsic portlet web service references one or more sets of code that are located on the portal computer. Your portal administrator must install this code in the correct location before you can create the associated intrinsic portlet web service.

A remote portlet web service references services hosted by a separate remote server. These services can be hosted by a web site or can be provided by code on a remote server. If the code is hosted by a remote server, your portal administrator must install this code before you can create the associated remote portlet web service. When a user displays a My Page or community page that includes a remote portlet, the portal contacts the appropriate remote server to obtain updated portlet content.

Portlet Templates

Portlet templates allow you to create multiple instances of a portlet, each 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 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.

Portlet Bundles

Portlet bundles are groups of related portlets, packaged together for easy inclusion on My Pages or community pages. 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.

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* or the documentation provided with the portlet software.

Portlets Available with the Portal

Some portlets (and their necessary portlet web services and remote servers) are automatically created in the **Portal Resources** folder when you install the portal. There are also portlets that are available with the portal installation, but require additional steps to complete installation. For information on the additional installation steps, refer to the *Installation Guide for AquaLogic Interaction*.

Note: You can also create your own portlets, 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*.

The following navigation portlet can be used with the Portlet-Ready Navigation scheme (set in an experience definition) to provide custom navigation for your portal:

- **Navigation Tags Header Portlet:** This portlet is provided as an example of a custom header that includes navigation tags; you can customize it and use it in communities or experience definitions. This portlet is stored in the Portal Resources folder.

  Note: 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.

The following branding portlets enable you to add custom branding to your portal pages:

- **Classic Footer Portlet:** This portlet is provided as an example of a custom footer that you can customize and use in communities or experience definitions.
- **Classic Header Portlet:** This portlet is provided as an example of a custom header that you can customize and use in communities or experience definitions.
- **Layout Footer Portlet:** This portlet is provided as an example of a custom footer that uses adaptive tags; you can customize it and use it in communities or experience definitions.

- **Layout Header Portlet:** This portlet is provided as an example of a custom header that uses adaptive tags; you can customize it and use it in communities or experience definitions.

The following login portlets can be added to guest default profiles so users can log in to the portal:

- **Portal Login:** This portlet allows users to log in to the portal. You probably want to add this to all your guest users' home pages so that users can log in from the default page displayed when they navigate to your portal.
- **Tag Login Portlet:** This portlet is provided as an example of a custom login portlet that uses adaptive tags; you can customize it and add it to your guest users' home pages so that users can log in from the default page displayed when they navigate to your portal. This portlet is stored in the Portal Resources folder. For information on adaptive tags, see the *Adaptive Page Layouts* section of the *AquaLogic User Interaction Development Guide*.

The following user profile portlets are included on the user profile page by default:

- **Folder Expertise:** This portlet displays the folders for which the user is an expert. Portal administrators can add users to a folder as an expert through the Related Resources page of the Folder Editor, or, if users have the Self-Selected Experts activity right, they can add themselves as experts when they are browsing folders in the Knowledge Directory. This portlet is stored in the Portal Resources folder and is displayed on the user profile page by default.
- **General Information:** This portlet displays user profile information such as name and address, but it is configurable by the portal administrator to display any information. If your portal displays a legacy layout (rather than adaptive layouts), this portlet is displayed on the user profile page by default. This portlet is stored in the Portal Resources folder.
- **Managed Communities:** This portlet displays the communities to which the user has Edit or Admin access. If your portal displays a legacy layout (rather than adaptive layouts), this portlet is displayed on the user profile page by default. This portlet is stored in the Portal Resources folder.

The following portlets are ready to be added to My Pages and community pages:

- **Job Histories Intrinsic Portlet:** This portlet displays the same job history information that is displayed on the Job History page of the Automation Service Manager. This portlet is stored in the Portal Resources folder.
- **Portal Search:** This portlet lets users search your portal and access their saved searches. Users might want to add this to their home page for easy access to their saved searches. This portlet is stored in the Portal Resources folder.
- **RSS Reader Portlet:** This portlet lets users specify an RSS or ATOM feed to display on a My Page. This portlet is stored in the Portal Resources/RSS Reader folder, but is available only if the portal administrator installed the Remote Portlet Service and imported the RSS Reader migration package.

- **RSS Community Reader Portlet:** This portlet lets community managers specify an RSS or ATOM feed to display on a community page. This portlet is stored in the Portal Resources/RSS Reader folder, but is available only if the portal administrator installed the Remote Portlet Service and imported the RSS Reader migration package.
- **User Activities:** This portlet displays a user's status history and any other recent activities that are submitted by other applications. This portlet is stored in the Activity Service folder, but is available only if the portal administrator installed the Remote Portlet Service and imported the Activity Service migration package.

  To view another user's activities, open the user's profile and look at the User Activities portlet displayed in the profile. To subscribe to e-mail notification or an RSS feed of the user's activity, click the appropriate button at the bottom of the user's User Activities portlet.
- **User Status:** This portlet lets users post their current status. This portlet is stored in the Activity Service folder, but is available only if the portal administrator installed the Remote Portlet Service and imported the Activity Service migration package.

The following portlet templates (and any necessary portlet web services and remote servers) are created when you install the portal:

- **Community Links Portlet Template:** This template is used by the portal to create portlets that display the links saved in a Community Knowledge Directory folder. This portlet template is stored in the Portal Resources folder.
- **Content Snapshots:** This template is used by the portal to create portlets that display the results of a Snapshot Query. This portlet template is stored in the Portal Resources folder.

Providing Access to Existing Web Applications with Portlets

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.

1. (Recommended) Create a lockbox in the portal for the existing application, and have users supply their login credentials for that lockbox.
2. Create a remote server in the portal for the existing application.
3. Create a remote portlet web service in the portal to associate with a portlet you will create for the existing application.

If you created a lockbox, use it to supply the user credentials for authenticating to this application.

4. Create a portlet based on the web service you created.
5. Add the portlet to My Pages or communities.

Creating an Intrinsic Portlet Web Service

Portlet web services allow you to specify *functional* settings for your portlets, leaving the *display* settings to be set in each associated portlet. An intrinsic portlet web service references one or more sets of code that are located on the portal computer.

Before you create an intrinsic portlet web service, you must:

- Install the portlet code on the computer that hosts the portal
- Create a remote server pointing to the computer that hosts the portlet code (optional, but recommended)

To create an intrinsic portlet web service you must have the following rights and privileges:

- Access Administration activity right
- Create Web Service Infrastructure activity right
- At least Edit access to the parent folder (the folder that will store the intrinsic portlet web service)
- At least Select access to the remote server that the intrinsic portlet web service will use

1. Click **Administration**.
2. Open the folder in which you want to store the intrinsic portlet web service.
3. In the **Create Object** drop-down list, click **Web Service — Intrinsic Portlet**.
   The Intrinsic Portlet Web Service Editor opens.
4. On the **Main Settings** page, complete the following task:
   -
5. Click the **Advanced Settings** page and complete the following task:
   -
6. Click the **Preferences** page and complete the following task:

-

7. Click the **Alternative Browsing Devices** page and complete the following task:
   -
8. Click the **Properties and Names** page and complete the following tasks:
   - *Naming and Describing an Object* on page 217

     You can instead enter a name and description when you save this intrinsic portlet web service.
   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219 (optional)

The default security for this intrinsic portlet web service is based on the security of the parent folder. You can change the security when you save this intrinsic portlet web service (on the **Security** tab page in the Save As dialog box), or by editing this intrinsic portlet web service (on the **Security** page of the Intrinsic Portlet Web Service Editor).

Portal administrators with at least Select access to this intrinsic portlet web service can create portlets or portlet templates based on the web service.

Creating a Remote Portlet Web Service

Portlet web services allow you to specify *functional* settings for your portlets, leaving the *display* settings to be set in each associated portlet. A remote portlet web service references services hosted by a separate remote server.

Before you create a remote portlet web service, you must:

- Install the portlet code on the computer that hosts the portal or on another computer
- Create a remote server pointing to the computer that hosts the portlet code (optional, but recommended)

To create a remote portlet web service you must have the following rights and privileges:

- Access Administration activity right
- Create Web Service Infrastructure activity right

- At least Edit access to the parent folder (the folder that will store the remote portlet web service)
- At least Select access to the remote server that the remote portlet web service will use

1. Click **Administration**.
2. Open the folder in which you want to store the remote portlet web service.
3. In the **Create Object** drop-down list, click **Web Service — Remote Portlet**.
   The Remote Portlet Web Service Editor opens.
4. On the **Main Settings** page, complete the following task:
   -
5. Click the **HTTP Configuration** page and complete the following task:
   -
6. Click the **Advanced URL Settings** page and complete the following task:
   -
7. Click the **Advanced Settings** page and complete the following task:
   -
8. Click the **Authentication Settings** page and complete the following task:
   -
9. Click the **Preferences** page and complete the following task:
   -
10. Click the **User Information** page and complete the following task:
    -
11. Click the **Debug Settings** page and complete the following task:
    -
12. Click the **Alternative Browsing Devices** page and complete the following task:
    -
13. Click the **Properties and Names** page and complete the following tasks:

- *Naming and Describing an Object* on page 217

  You can instead enter a name and description when you save this remote portlet web service.
- *Localizing the Name and Description for an Object* on page 342 (optional)
- *Managing Object Properties* on page 219 (optional)

The default security for this remote portlet web service is based on the security of the parent folder. You can change the security when you save this remote portlet web service (on the **Security** tab page in the Save As dialog box), or by editing this remote portlet web service (on the **Security** page of the Remote Portlet Web Service Editor).

Portal administrators with at least Select access to this remote portlet web service can create portlets or portlet templates based on the web service.

Providing Custom Tools and Services with Portlets

Portlets provide portal users customized tools and services as well as information. Portlets let you to integrate applications, tools, and services into your portal, while taking advantage of portal security, caching, and customization. Users can then add these portlets to their My Pages or to community pages.

Before you create a portlet, you must:

- Install the portlet code on the computer that hosts the portal or, if your portlet does not rely on any portal code, you can instead install it on another computer
- If you installed the portlet code on a computer other than the one that hosts the portal, create a remote server to point to the remote computer
- Create a portlet web service on which to base your portlet
- Optionally, create a portlet template on which to base your portlet

Note: For information on installing portlet code, refer to the *Installation Guide for AquaLogic Interaction* (available on *edocs.bea.com*) or the documentation that comes with your portlet, or contact your portal administrator.

To create a portlet you must have the following rights and privileges:

- Access Administration activity right

- Create Portlets activity right
- At least Edit access to the parent folder (the folder that will store the portlet)

1. Click **Administration**.
2. Open the folder in which you want to store the portlet.
3. In the **Create Object** drop-down list, click **Portlet**.
   The Choose Template or Web Service dialog box opens.
4. Select the template or web service that provides the basic settings for your portlet and click **OK**.

   Use a template when possible. When you use a template, your portlet inherits the template's web service as well as its default settings. Some web services that are designed to work with templates might not work correctly if you bypass the template and make a new portlet directly from the web service object.

   The Portlet Editor opens.
5. On the **Main Settings** page, complete the following tasks:

   - *Editing the Administrative Preferences for a Portlet* on page 239
   - *Specifying the Size, Type, and Orientation for a Portlet* on page 234

6. Click the **Properties and Names** page and complete the following tasks:

   - *Naming and Describing an Object* on page 217

     You can instead enter a name and description when you save this portlet.
   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219 (optional)

The default security for this portlet is based on the security of the parent folder. You can change the security when you save this portlet (on the **Security** tab page in the Save As dialog box), or by editing this portlet (on the **Security** page of the Portlet Editor).

Specifying the Size, Type, and Orientation for a Portlet

You can specify the size, type, and orientation for a portlet on the Main Settings page of the Portlet Editor.

1. If the Portlet Editor is not already open, open it now and display the **Main Settings** page.
2. Specify what type of portlet this is.

- **Narrow**: 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.

  If you created this portlet from a portlet template that creates narrow portlets or if you are editing an existing a narrow portlet, you can change it to a Wide portlet but not to a header, footer, or content canvas portlet.
- **Wide**: Wide portlets can be added only to wide columns.

  If you created this portlet from a portlet template that creates wide portlets or if you are editing an existing a wide portlet, you can change it to a narrow portlet but not to a header, footer, or content canvas portlet.
- **Header**: Header portlets can be added to communities, community templates, and experience definitions to change the branding of these objects by replacing a banner at the top of the page (so that it differs from the top banner displayed by the main portal).

  You cannot change this setting if you created this portlet from a portlet template that creates header portlets or if you are editing an existing header portlet.
- **Footer**: 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).

  You cannot change this setting if you created this portlet from a portlet template that creates footer portlets or if you are editing an existing footer portlet.
- **Content Canvas**: Content canvas portlets can be added below the top banner on community pages that include a content canvas space (specified in the page layout). Content canvas portlets can display across the entire width of the page or across one or two columns. You cannot add more than one content canvas portlet per page.

  You cannot change this setting if you created this portlet from a portlet template that creates content canvas portlets or if you are editing an existing content canvas portlet.

3. If this is a narrow or wide portlet, specify whether this portlet is a community-only portlet.
   - If you want to allow users to add this portlet to My Pages or community pages, choose **For My Pages or Community pages**.
   - If you want to allow users to add this portlet only to community pages, choose **For Community pages only**.
4. If this is a narrow or wide portlet and you do not want to display the title of this portlet when it is added to a page, select **Suppress Portlet's title bar**.

Note: If this portlet includes preferences or help, suppressing the title bar will make these features unavailable in the portlet.

Caching Portlet Content

You might occasionally want to run a job to cache portlet content (for example, if the portlet takes a couple minutes to render). When the job runs, it creates a snapshot of the portlet content (in the form of a static HTML file) that can be displayed on a web site. The file is stored in the shared files directory (for example, `C:\bea\ALUI\ptportal\6.5`) in `\StagedContent\Portlets\<portletID>\Main.html`. You can then create another portlet that simply displays the static HTML.

Note: The shared files directory path is set on the **Portal URL Manager** page of the Portal Settings Utility.

To run a portlet as a job you must have the following rights and privileges:

- Access Administration activity right
- Create Jobs activity right
- At least Edit access to the parent folder (the folder that will store the job)
- At least Select access to the portlet

Note:

- Because intrinsic portlets rely on the portal application, you cannot run an intrinsic portlet as a job.
- Because the content produced is static you should only run portlets that present information that is valuable when updated on a periodic basis. For example, a report portlet would be ideal to run as a job, while more interactive portlets, like application interfaces would not be appropriate.
- If the portlet includes preferences, the preferences for the user that creates the job will be used.

1. Click **Administration**.
2. Open the folder in which you want to store the portlet job.

   Note: In order for the job to run, the folder must be registered with an Automation Service.

3. In the **Create Object** drop-down list, select **Job**.
4. On the Main Settings page, click **Add Operation**.
5. Select the portlets you want to run with this job, and click **OK**.

6. Under **Schedule**, select the frequency with which you want this job to run.

Setting Security for a Portlet

By default, a new portlet inherits the security of the parent folder, but you can change the security of each individual portlet.

1. Open the Portlet Editor by creating a new portlet or editing an existing one.
2. Click the **Security** page.
3. Specify which users and groups can access this portlet and what type of access they have:

   - To allow additional users or groups access to this portlet, click **Add Users/Groups**.
   - If this portlet can be added to My Pages and is not a header, footer, or content canvas portlet, you can force users or groups to include this portlet on their default My Pages. To do so, in the **Mandatory** drop-down list, click **Mandatory**.

     Note: Users and groups for which this portlet is mandatory will not be able to remove this portlet from their My Pages.

   - To specify the type of access a user or group has, in the drop-down list under the **Privilege** column, select the access type.

     For a description of the available privileges, see *About Access Privileges* on page 60.

     Note: If a user is a member of more than one group included in the list, or if they are included as an individual user and as part of a group, that user gets the highest access available to her for this object. For example, if a user is part of the Everyone group (which has Read access) and the Administrators Group (which has Admin access), that user gets the higher privilege to the community: Admin.

   - To delete a user or group, select the user or group and click .

     To select or clear all of the user and group check boxes, select or clear the check box to the left of **Users/Groups**.
   - To see what users are included in a group, click the group name.
   - To change the column used for sorting or to toggle the sort order between ascending and descending, click the column name.

     You see an icon (or) to the right of the column name by which the objects are sorted.

- If you chose Mandatory for any user or group, in the **Mandatory Portlet Priority** drop-down list, set this portlet's priority.

  The priority determines the portlet's placement on the My Page; portlets with higher priority display closer to the upper-left of the My Page then portlets with lower priority.

Portlet Preferences

Portlets can includes several different types of preferences.

Preference Type	Description	Who Can Set Them and Where
Administrative Preferences	These preferences affect everyone's view of the portlet. For example, setting which e-mail server an e-mail portlet should connect to.	They are set by the portlet creator on the Main Settings page of the Portlet Editor, or by users with administrative rights from **My Pages ➤ Edit Portlet Preferences** or by clicking the edit icon in a portlet's title bar.
Personal Preferences	These preferences affect that user's view of the portlet. For example, setting how many e-mails are displayed in an e-mail portlet .	They are set by the user from **My Pages ➤ Edit Portlet Preferences** or **My Communities ➤ Edit Portlet Preferences** .
Community Preferences	These preferences affect everyone's view of portlets in that community. For example, setting a specific public e-mail folder to display in an e-mail portlet, and setting 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. 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.

Preference Type	Description	Who Can Set Them and Where
Portlet Template Preferences	These preferences affect the portlet template itself and all portlets created from that template. For example, specifying which portlet web service a portlet uses.	These preferences are set by the portlet template creator on the **Main Settings** page of the Portlet Template Editor. 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.

Editing the Administrative Preferences for a Portlet

You can configure the administrative preferences for a portlet on the Main Settings page of the Portlet Editor.

1. If the Portlet Editor is not already open, open it now and display the **Main Settings** page.
2. If the associated web service includes administrative preferences (specified on the **Preferences** page of the Portlet Web Service Editor), click **Edit** to edit the preferences.

Managing User Credentials for External Applications Through the Credential Vault Manager

You can provide secure portal access to existing web applications by setting up lockboxes to store user credentials. For example, you might want to provide portal access to a secured employee benefits system. Users can enter their user authentication information through the Password Manager on the My Account page and not have to enter the information each time they access the secured application through the portal.

To access the Credential Vault Manager you must be a member of the Administrators Group.

You can create a lockbox for each secured application the user needs to access through the portal.

1. Click **Administration**.

2. In the **Select Utility** drop-down list, click **Credential Vault Manager**.
3. Create a lockbox for each secured application you will provide access to through the portal.
   - To create a new lockbox, click **New Lockbox**.

     The Lockbox Editor opens.
   - To edit an existing lockbox, click its name.

     The Lockbox Editor opens.
   - To delete a lockbox, select it and click .

After setting up lockboxes:

- Users need to enter their login credentials through the Password Manager on the My Account page.
- To send credentials in portlet headers, using RSA public key/private key encryption, after setting up a lockbox, you must associate the lockbox with the remote portlet web service (on the Authentication Settings page), enter the public key for RSA encryption in the remote server (on the Main Settings page), and use the IDK to provide the private key for RSA encryption (see BEA AquaLogic User Interaction Development Center for information).

Creating or Editing a Lockbox to Store User Credentials for External Applications

Create a lockbox for each secured application the user needs to access through the portal.

To access the Credential Vault Manager you must be a member of the Administrators Group.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **Credential Vault Manager**.
3. Click **New Lockbox** or click an existing lockbox to edit it.
4. In the **Name** box, type a name for the lockbox.

   Users will see this name in a list of their external accounts when they click **Password Manager** on the **My Account** page. The name should clearly identify the external system for which users will enter their login credentials.
5. In the **Description** text box, type a description for this lockbox.

This description displays in the Administrative Objects Directory to help other administrators understand what this object is.

6. If your portal administrator did not set a mandatory object language, in the **Primary Language** drop-down list, select the language for the name and description you entered.

   If your portal administrator did set a mandatory object language, you see the mandatory language instead of a drop-down list. You cannot change this setting. The name and description you entered must be in the mandatory language.

   If a localized name and description is not available in a user's selected language, the user will see the name and description in the specified primary language.

7. If you want to add localized names and descriptions:
   a) Select **Supports Localized Names**.
      The **Localized Names and Descriptions** section appears.
   b) Add or edit the localized names and descriptions:

   - To add an entry for a language, click **New Localized Name**, then, in the Name and Description dialog box, enter the localized name and/or description, select the appropriate language, and click **Finish**.
   - To edit an existing entry, click the entry you want to change, then, in the Name and Description dialog box, edit the entry as necessary, and click **Finish**.
   - To remove existing entries, select the entries you want to remove and click .

     To select or clear all entries, select or clear the check box to the left of **Name**.

8. Under **Lockbox Properties**, enter names for the user name and password properties for this lockbox.

   End users will see these names in the Password Manager when entering their login credentials for the external system corresponding to this lockbox.

   These properties will be created when you save the lockbox. After you have saved the lockbox, these properties appear as links. Click the links to edit the properties.

About Providing Content and Services to a Group of Users through Communities

Communities are sites within a portal designed for a specific audience or task, such as collaborative projects. The pages, portlets, layout, community preferences, and subcommunities within a community are determined by the community administrator. Although the community administrators determine which portlets are displayed in a community, a portlet itself might allow community members to change the content within each portlet.

You might have 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 might have a separate community containing project milestones, regulatory compliance requirements, and technical specifications.

You are automatically subscribed to communities based on your group membership. You can also join communities on your own. Some community subscriptions might be mandatory, but you can unsubscribe from those that are not. The communities you are subscribed to appear in the **My Communities** menu. Some mandatory communities might also appear as tabs in the menu area.

Community Menus

The community might include the following menus:

- The community menu displays all the community pages, and—if enabled by the administrator—the Community Knowledge Directory. The community pages display portlets. The Community Knowledge Directory displays the members of the community, any subcommunities of the community, and any other folders and contents the community administrator added.
- **Subcommunities** displays any subcommunities within the current community.

- **Related Communities** displays any communities that are stored in the same administrative folder as the current community.

  This menu only appears if you have access to related communities.

Note: Your portal administrator might use a navigation scheme with customized menu options.

Community Templates

Each community is based on a community template. Community templates define the basic structure for the resulting communities, such as which page templates to include and, optionally, a header or footer for the community. A single community template can be used by many different communities, allowing you to keep similar types of communities looking analogous. For example, you might want all communities based on departments to look similar and contain similar content, but you might want communities based on projects to look different.

Page Templates

Each community page is based on a page template. Page templates define the basic structure for the resulting community pages, such as the column layout and which portlets to include. 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.

Template Inheritance

When you create a community or a community page, you can decide whether to inherit the underlying community template or page template. Inheriting a template has several effects on the resulting communities and pages:

- You cannot remove objects that are included as part of the template. For example, you cannot remove pages that came from the community template and you cannot remove portlets that came from the page template.
- Any changes made to the template are inherited by the resulting communities or community pages. For example, if a page template is removed from a community template, the page is removed from any communities that inherited the template; if a portlet is removed from a page template, the portlet is removed from any pages that inherited the template.

Note: If you inherit a community template, you also inherit the included page templates.

Subcommunities

Subcommunities (along with community 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.

Community Groups and Community Portlets

With the appropriate activity rights, you can create groups and portlets inside a community without affecting portal groups or portlets. For example, you might have a group that is responsible for maintaining schedules in a specific community without making that group a portal group. Or you may create a community links portlet inside a community for the convenience of community members.

Note: Community groups and community portlets are available only to the community in which they are associated. If you want to use them outside the community, you can move them from the community folder to another administrative folder.

Community Knowledge Directory

The Community Knowledge Directory, if enabled, displays community resources in an organizational structure that is relevant to the community (as opposed to the broader portal audience). It includes a list of community members, displayed in the **Members** folder, and a list of subcommunities, displayed in the **Subcommunities** folder. Community administrators can also create folders that contain links to relevant web pages, community experts, portal documents, or community pages.

Note: Communities and subcommunities have separate Community Knowledge Directories.

Community Links Portlets

A community links portlet displays a snapshot of the links in a single Community Knowledge Directory folder. You can then add the portlet to a page in the community or invite users to add the portlet to their My Pages to provide quick access to the community resources. With the proper access privileges, you can also use the portlet to add or delete content from the associated Community Knowledge Directory folder.

Creating a Community Template

Community templates define the basic structure for the resulting communities, such as which page templates to include and, optionally, a header or footer for the community.

Before you create a community template, you must:

- Create any page templates you want to add to this community template
- Create any header and footer portlets you want to add to the community template

To create a community template you must have the following rights and privileges:

- Access Administration activity right
- Create Community Infrastructure activity right
- At least Edit access to the parent folder (the folder that will store the community template)
- At least Select access to any page templates you want to add to this community template
- At least Select access to any header and footer portlets you want to add to the community template

1. Click **Administration**.
2. Open the folder in which you want to store the community template.
3. In the **Create Object** drop-down list, click **Community Template**.
   The Community Template Editor opens.
4. On the **Main Settings** page, complete the following task:
   - *Adding Page Templates to a Community Template* on page 247
5. Click the **Header and Footer** page and complete the following task:
   - *Adding Headers and Footers to Community Templates* on page 248
6. Click the **Properties and Names** page and complete the following tasks:
   - *Naming and Describing an Object* on page 217

     You can instead enter a name and description when you save this <object>.
   - *Localizing the Name and Description for an Object* on page 342 (optional)
   - *Managing Object Properties* on page 219 (optional)

The default security for this community template is based on the security of the parent folder. You can change the security when you save this community template (on the **Security** tab page in the Save As dialog box), or by editing this community template (on the **Security** page of the Community Template Editor).

Adding Page Templates to a Community Template

Page templates define the basic structure for the resulting community pages, such as the column layout and which portlets to include.

To add page templates to a community template you need at least Select access to the page templates.

If the Community Template Editor is not already open, open it now and display the **Main Settings** page.

- To add a page template, click **Add Page Templates**, then, in the Add Page Templates dialog box, select the page templates you want to add to this community template and click **OK**.
- To remove a page template, select the template and click .

  To select or clear all of the page template check boxes, select or clear the box to the left of **Page Template Names**.

  Note: If you remove a page template, the associated page will be removed from any communities that inherit the changes in this template. It might take up to 15 minutes for this occur.
- To change the order in which the pages will appear in the communities created from this template, use the arrow icons to the right of the page templates.
  - To move a page to the top of this list, click .
  - To move a page up one space in this list, click .
  - To move a page down one space in this list, click .
  - To move a page to the bottom of this list, click .

  Note: If you change the order of the pages, and communities have previously been created from this community template, the page order will change in all of the communities derived from this community template.

Adding Headers and Footers to Community Templates

You can add header and footer portlets to community templates to control what community members see at the top and bottom of the pages in the associated communities.

To add headers and footers you must have the following privileges:

- At least Edit access to the community template
- At least Select access to the header and footer portlets you want to add

If the Community Template Editor is not already open, open it now and display the **Header and Footer** page.

- To add or change the header, under **Community Header**, click **Browse**, then, in the Select a Header dialog box, select the header you want, and click **OK**.
- To add or change the footer, under **Community Footer**, click **Browse**, then, in the Select a Footer dialog box, select the footer you want, and click **OK**.
- To remove the header, under **Community Header**, click **Remove**.
- To remove the footer, under **Community Footer**, click **Remove**.
- To use experience definition headers and footers, select **Force community to use header and footer from experience definition**.

  This forces communities created from this template to use the header and footer from the experience definition rather than from this template or from the community itself. Because users might be assigned to different experience definitions with different headers and footers, if you select this option, communities created from this template might display different headers and footers to different users.

Providing Content and Services to a Group of Users through a Community

Communities are sites within a portal designed for a specific audience or task, such as collaborative projects. You might have 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 might have a separate community containing project milestones, regulatory compliance requirements, and technical specifications.

Before you create a community, you must:

- Create the community template on which this community will be based
- Create the page templates on which your community pages will be based
- Create any portlets you want to add to the community pages

To create a community you must have the following rights and privileges:

- Access Administration activity right
- Create Communities activity right
- At least Edit access to the parent folder (the folder that will store the community)
- At least Select access to the community template on which this community will be based
- At least Select access to any page templates you will use to create community pages
- At least Select access to any portlets you want to add to the community pages

1. Click **Administration**.
2. Open the folder in which you want to store the community.
3. In the **Create Object** drop-down list, click **Community**.
   The Community Editor opens.
4. On the **Community Pages** page, complete the following tasks:
   - <xrefs to tasks that can be performed on this page>
5. Click the **Header and Footer** page and complete the following tasks:
   -
6. Click the **Subcommunities** page and complete the following tasks:
   -
7. Click the **This Community's Groups** page and complete the following tasks:
   -
8. Click the **Header and Footer** page and complete the following tasks:
   -
9. Click the **Portlet Preferences** page and complete the following tasks:
   -

10. Click the **This Community's Portlets** page and complete the following tasks:
    -
11. Click the **Properties and Names** page and complete the following tasks:
    - *Naming and Describing an Object* on page 217

      You can instead enter a name and description when you save this community.
    - *Localizing the Name and Description for an Object* on page 342 (optional)
    - *Managing Object Properties* on page 219(optional)

The default security for this community is based on the security of the parent folder. You can change the security when you save this community (on the **Security** tab page in the Save As dialog box), or by editing this community (on the **Security** page of the Community Editor).

Applying a Community Template to a Community

Each community is based on a community template. Community templates define the basic structure for the resulting communities, such as which page templates to include and, optionally, a header or footer for the community. A single community template can be used by many different communities, allowing you to keep similar types of communities looking analogous. For example, you might want all communities based on departments to look similar and contain similar content, but you might want communities based on projects to look different.

1. If the Community Editor is not already open, open it now.
2. In the **Community Templates** section, click **Select Community Template**.

   Note: If you are editing an existing community, the button says Change Community Template. Before changing the template, read the important notes below.

   The Community Templates dialog box opens.
3. Select a template and click **OK**.
4. If you do not want this community to inherit future changes to the template, clear the box next to **Inherit the Template**.

   If you select to inherit changes, any change applied to the community template affects the community. For example, if a page is removed from the community template, the page will be removed from this community as well. Additionally, if you inherit changes, you cannot delete pages associated with the template, but you can add new pages and change the order of the pages.

5. Click **OK**.

Important: After a community is created, you can select a different community template to use. When selecting a different community template, note the following:

- Any pages from the old community template that are not part of the new community template will be removed.
- If you have set special headers and footers for your community, switching to a community template that enforces a header or footer will remove your header or footer.

Setting the Community Home Page and Ordering Community Pages

The order in which pages are displayed in the Community Pages list is the order in which the page links will display to users. By default, the first page you add to a community, whether directly or via a community template, will be the home page of your community.

If the Community Editor is not already open, open it now.

- To change the home page, move the desired page to the top of the **Community Pages** list by clicking ⤒ to the far right of the page name.
- To move a page up one space in this list, click ▲.
- To move a page down one space in this list, click ▼.
- To move a page to the bottom of this list, click ⤓.

Adding Headers and Footers to Communities

You can add header and footer portlets to communities to control what community members see at the top and bottom of the pages in the community.

To add headers and footers you must have the following privileges:

- At least Edit access to the community
- At least Select access to the header and footer portlets you want to add

If the Community Editor is not already open, open it now and display the **Header and Footer** page.

- To add or change the header, under **Community Header**, click **Browse**, then, in the Select a Header dialog box, select the header you want, and click **OK**.
- To add or change the footer, under **Community Footer**, click **Browse**, then, in the Select a Footer dialog box, select the footer you want, and click **OK**.
- To remove the header, under **Community Header**, click **Remove**.
- To remove the footer, under **Community Footer**, click **Remove**.

Setting Security on a Community

By default, a new community inherits the security of the parent folder, but you can change this security.

1. Open the Community Editor by creating a new community or editing an existing community.
2. Click the **Security** page.
3. Specify which users and groups can access this community and what type of access they have:
   - To allow additional users or groups access to this community, click **Add Users/Groups**.
   - To specify whether this community is mandatory, select an option in the **Mandatory** drop-down list:

     By default communities are **Not Mandatory**.

     - To force users or groups to be members of this community, select **Mandatory**.
     - To force users or groups to be members of this community and add a tab to the portal banner for this community, select **Mandatory with Tab**.

     Note: Users and groups for which this community is mandatory will not be able to unsubscribe from this community, that is, this community will always be available in their My Communities menu.

   - To specify the type of access a user or group has, in the drop-down list under the **Privilege** column, select the access type.

     For a description of the available privileges, see *About Access Privileges* on page 60.

     Note: If a user is a member of more than one group included in the list, or if they are included as an individual user and as part of a group, that user gets the highest access available to her for this object. For example, if a user is part of the Everyone group (which has Read access) and the Administrators Group (which has Admin access), that user gets the higher privilege to the community: Admin.

- To delete a user or group, select the user or group and click ✕.

  To select or clear all of the user and group check boxes, select or clear the check box to the left of **Users/Groups**.
- To see what users are included in a group, click the group name.
- To change the column used for sorting or to toggle the sort order between ascending and descending, click the column name.

  You see an icon (▾ or ▴) to the right of the column name by which the objects are sorted.
- If you chose Mandatory with Tab for any user or group, in the **Mandatory Tab Priority** drop-down list, set this community tab's priority.

  The priority determines the order in which tabs display in the portal banner: tabs with higher priority display before tabs with lower priority.

Creating a Community Page Template

Page templates define the basic structure for the resulting community pages, such as the column layout and which portlets to include. 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.

Note: To create a page template you must have the following rights and privileges:

- Access Administration activity right
- Create Community Infrastructure activity right
- At least Edit access to the parent folder (the folder that will store the page template)
- At least Select access to any portlets you want to add to the page

1. Click **Administration**.
2. Open the folder in which you want to store the page template.
3. In the **Create Object** drop-down list, click **Page Template**.
   The Page Template Editor opens, displaying the Main Settings page.

4. Select a column layout for the page. Click **Select Page Layout**, then, in the Select Page Layout dialog box, select the layout you want and click **Finish**.

   Note:

   - If you intend to add a content canvas portlet (a portlet that straddles more than one column), you must select a layout that includes a dark gray section.
   - Narrow portlets can display in either narrow or wide columns, but wide portlets can display only in wide columns.

5. Select portlets and position them on the page:

   - If you selected a layout that includes a section for a content canvas portlet, click **Add Content Canvas**, then, in the Content Canvas Portlets dialog box, select a content canvas portlet and click **OK**.
   - To add a portlet to the page, click **Add Portlets**, then, in the Add Portlets dialog box, select the portlets you want to add and click **Finish**.
   - To see what a portlet looks like, click **Preview** under the portlet.
   - To remove a portlet, click **Remove** under the portlet.
   - To reposition a portlet, drag the portlet to the desired position (by clicking the portlet, holding down the mouse button, and moving the mouse), then release the mouse button.

   Note: When users create pages from this template or create communities from a community template that includes this page template, they can choose to inherit the template. Inheriting the template has two effects on resulting pages, users cannot reposition or remove any portlets that are included as part of the template, and any changes to the template are mirrored in the resulting pages.

6. Optionally, specify a default name for pages created from this template:
   a) Click the **Default Page Name** page.
   b) In the **Default Page Name** box, type a name.
   c) If your portal administrator did not set a mandatory object language, in the **Primary Language** drop-down list, select the language for the name you entered.

      If your portal administrator did set a mandatory object language, you see the mandatory language instead of a drop-down list. You cannot change this setting. The name you entered must be in the mandatory language.

      If a localized name is not available in a user's selected language, the user will see the name in the specified primary language.

d) If you want to add default names for other languages, select **Supports Localized Names**, then, in the **Localized Names and Descriptions** section, add or edit the localized names:

- To add an entry for a language, click  **New Localized Name**, then, in the Name and Description dialog box, enter the localized name and/or description, select the appropriate language, and click **Finish**.
- To edit an existing entry, click the entry you want to change, then, in the Name and Description dialog box, edit the entry as necessary, and click **Finish**.
- To remove existing entries, select the entries you want to remove and click  .

  To select or clear all entries, select or clear the check box to the left of **Name**.

Creating a Community Page

You can create different pages in a community to categorize information for your community audience. For example, you might have a project community that includes a page for each department that is involved in the project.

To create a community page you must have the following rights and privileges:

- If you are creating the page from the Administrative Objects Directory or from the Community Editor, you must have the Access Administration activity right and at least Select access to the page template on which the page will be based
- At least Edit access to the community
- At least Select access to any portlets you want to add to the page

There are several ways to create a community page:

- *Creating a Community Page with One Click* on page 255
- *Creating a Community Page From the Community Editor* on page 256
- *Creating a Community Page From the Administrative Objects Directory* on page 257

Creating a Community Page with One Click

You can create a new community page with one click.

To create a community page you must have the following rights and privileges:

- At least Edit access to the community
- At least Select access to any portlets you want to add to the page

1. Display the community to which you want to add a page.
2. Click **Create Page**.
   The new page is created.
3. To add portlets to the page, click **Edit Page**.

Creating a Community Page From the Community Editor

You can create a community page while creating or editing a community.

To create a community page you must have the following rights and privileges:

- Access Administration activity right
- At least Edit access to the community
- At least Select access to the page template on which the page will be based
- At least Select access to any portlets you want to add to the page

1. If the Community Editor is not already open, open it now.
2. Click **New Page**.
   The New Page dialog box opens.
3. Select the page template on which you want to base this page, and click **OK**.
4. Specify whether you want to inherit the template, and click **OK**.

   Note: If you inherit the template, you cannot reposition or remove portlets that are included as part of the template, and any changes made to the template are mirrored in the page you create.

   The Page Editor opens.
5. Select a column layout for the page. Click **Select Page Layout**, then, in the Select Page Layout dialog box, select the layout you want and click **Finish**.

   Note:

   - If you intend to add a content canvas portlet (a portlet that straddles more than one column), you must select a layout that includes a dark gray section.

- Narrow portlets can display in either narrow or wide columns, but wide portlets can display only in wide columns.

6. Select portlets and position them on the page:
   - If you selected a layout that includes a section for a content canvas portlet, click  **Add Content Canvas**, then, in the Content Canvas Portlets dialog box, select a content canvas portlet and click **OK**.
   - To add a portlet to the page, click  **Add Portlets**, then, in the Add Portlets dialog box, select the portlets you want to add and click **Finish**.
   - To create a portlet and add it to the page, click  **Create Portlets**. On the Choose Portlet Template page, select a Portlet Template, click **Next >>**, and complete the Portlet Editor.
   - To see what a portlet looks like, click **Preview** under the portlet.
   - To remove a portlet, click **Remove** under the portlet.
   - To reposition a portlet, drag the portlet to the desired position (by clicking the portlet, holding down the mouse button, and moving the mouse), then release the mouse button.

Creating a Community Page From the Administrative Objects Directory

You can create a community page from the Administrative Objects Directory.

To create a community page you must have the following rights and privileges:

- Access Administration activity right
- At least Edit access to the community
- At least Select access to the page template on which the page will be based
- At least Select access to any portlets you want to add to the page

1. Click **Administration**.
2. Open the community folder to which you want to add the page.
3. In the **Create Object** drop-down list, click **Page**.
   The New Page dialog box opens.
4. Select the page template on which you want to base this page, and click **OK**.
5. Specify whether you want to inherit the template, and click **OK**.

Note: If you inherit the template, you cannot reposition or remove portlets that are included as part of the template, and any changes made to the template are mirrored in the page you create.

The Page Editor opens.

6. Select a column layout for the page. Click **Select Page Layout**, then, in the Select Page Layout dialog box, select the layout you want and click **Finish**.

   Note:

   - If you intend to add a content canvas portlet (a portlet that straddles more than one column), you must select a layout that includes a dark gray section.
   - Narrow portlets can display in either narrow or wide columns, but wide portlets can display only in wide columns.

7. Select portlets and position them on the page:

   - If you selected a layout that includes a section for a content canvas portlet, click **Add Content Canvas**, then, in the Content Canvas Portlets dialog box, select a content canvas portlet and click **OK**.
   - To add a portlet to the page, click **Add Portlets**, then, in the Add Portlets dialog box, select the portlets you want to add and click **Finish**.
   - To create a portlet and add it to the page, click **Create Portlets**. On the Choose Portlet Template page, select a Portlet Template, click **Next >>**, and complete the Portlet Editor.
   - To see what a portlet looks like, click **Preview** under the portlet.
   - To remove a portlet, click **Remove** under the portlet.
   - To reposition a portlet, drag the portlet to the desired position (by clicking the portlet, holding down the mouse button, and moving the mouse), then release the mouse button.

Editing a Page in the Flyout Editor

You can rename a page, add portlets, recommend portlets, and reposition portlets while viewing the page.

Click **Edit Page**. The Flyout Editor appears, enabling you to perform the following actions:

Note: The Flyout Editor appears only if you are viewing an adaptive page layout (not a legacy page layout).

- To rename the page, in the **Change Page Name** box, type the new name.
- To add a portlet to the page, under the portlet name, click **Add to Page**.
  A placeholder for the portlet is added to the page below the Flyout Editor.
- To remove a portlet, under the portlet name, click **Remove**, or click in the portlet's title bar.
- To see what a portlet looks like, under the portlet name, click **Preview**.

  From the Preview Portlet page you can perform the following actions:

  - To add the portlet to your page and close the preview, click **Add this portlet**.
  - To view a description of the portlet, click **View Description**. When you are finished, click **Close**.
  - To return to the list of portlets without adding the portlet to your page, click **Close**.
- To view a list of the portlets in a portlet bundle, under the bundle name, click **Open**.
- To add all the portlets from a bundle, under the bundle name, click **Add**.
  A placeholder for each portlet is added to the page below the Flyout Editor.

  Note: You can remove a portlet added as part of a bundle by clicking in the portlet's title bar.
- To recommend a portlet to other users:

  a) Under the portlet name, click **Invite**.
  b) In the invitation dialog box, copy the text, and click **Close**.
  c) In your e-mail application, paste the text into an e-mail message and send it.

  When other portal users click the URL in your e-mail, they are taken to the portlet preview and given the option to add the portlet to one of their My Pages. Users that do not have permission to see the portlet receive an error message.
- To search for portlets and portlet bundles, in the **Search for Portlets** box, type the text you want to search for and click **Search**.

  For searching tips, see *Using Text Search Rules* on page 364.

  To remove your search criteria, click **Search** again.
- To change the sort order of portlets, in the **Sort By** drop-down list, select an option: Item Name Ascending, Item Name Descending, Date Modified Ascending, Date Modified Descending.
- To page through the list of portlets, click **<< Previous**, **Next >>**, or a particular page number.
- To browse through administrative folders, click **Browse All Folders**.

  Note: This link displays folders that might not contain portlets or portlet bundles.

To view the portlets and portlet bundles in a folder, click the folder name.

- To reposition a portlet, in the area under the Edit Page section, drag the portlet to the desired position (by clicking the portlet, holding down the mouse button, and moving the mouse), then release the mouse button.

  Each column on the page is represented by a gray box. To change the column structure, click **Go to Advanced Editor**, then click **Select Page Layout**.

- To close the Flyout Editor, click **Close**, or, in the Edit Page section, click or **Close Editor**.

Deleting a Page from a Community

You can delete any page that is not inherited from the community template.

To delete a community page, you must have at least Edit access to the community.

1. If the Community Editor is not already open, open it now.
2. Select the page you want to delete and click .

   Note: You can only delete pages that say **No** in the **From Community Template** column. If you chose to inherit changes from the community template, pages that are part of that template say **Yes** in the **From Community Template** column and cannot be deleted.

   The page is deleted from the community folder.

Managing 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.

Customizing Search Service Behavior

Search Result Types

You can limit your search to particular types of objects.

Result Type	Description
Documents	Returns documents from the portal Knowledge Directory.
Knowledge Directory Folders	Returns folders from the portal Knowledge Directory.
Users	Returns portal users.
Communities	Returns communities.
Community Pages	Returns pages in a community

Result Type	Description
Portlets	Returns portlets.
Collaboration Items	Returns documents, discussions, and task lists from Collaboration Server.
Publisher Items	Returns documents from Publisher.

Search Results Sorting Options

You can specify how your search results are sorted.

Option	Description
Relevance	Displays your results according to how closely they match your search query. **Note:** Best bets are only shown in search results when sorting by relevance.
Last Modified Date	Displays your results in the order in which they were most recently edited. The result that was modified most recently will display first.
Folder	Groups your results by the folders in which they are stored and displays a list of the folders that contain search results.
Object Type	Groups your results by type of object (such as documents, users, communities, or portlets) and displays a list of the types of objects returned by your search.

About Best Bets and Top Best Bets

Best bets associate specific search phrases (specified by the portal administrator) with a set of search results. When end users enter a search query that matches a best bet search phrase, the best bet results appear as the first results in the relevance-ranked result list. Additionally, users can choose to go directly to the highest ranking best bet, the top best bet, instead of seeing the normal search results. For example, the top best bet for the term "HR" might be the Human Resources community. If users use the top best bet feature, they go directly to the Human Resources community instead of seeing all search results for the term "HR."

Note:

- Best bets are not case-sensitive.

- 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 or object selection search.
- Users go to the top best bet (for example, a community) only if they have at least Read access to it. If they do not, they see the list of search results to which they do have access.
- The phrase “Best Bet” appears next to each best bet result to inform the user that the result has been judged especially relevant to the query.

Creating Best Bets

Best bets associate specific search phrases (specified by the portal administrator) with a set of search results. When end users enter a search query that matches a best bet search phrase, the best bet results appear as the first results in the relevance-ranked result list. Additionally, users can choose to go directly to the highest ranking best bet, the top best bet, instead of seeing the normal search results. For example, the top best bet for the term "HR" might be the Human Resources community. If users use the top best bet feature, they go directly to the Human Resources community instead of seeing all search results for the term "HR."

To access the Search Results Manager you must have the following rights:

- Access Administration activity right
- Access Utilities activity right
- Access Search Results Manager activity right

Note:

- Best bets are case-insensitive.
- 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.

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.

Users who search for the phrase you specified now see the best bets you created. Users can go directly to the top best bet by using the top best bet operator (>) in their search or by clicking  rather than **Search**.

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/*.

How Banner Field Settings Affect Search Results

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 and ranks those results based on the specified weighting settings.

Note: Banner field settings apply to all searches: search from the portal banner, advanced search, object selection search, and any other portal search interfaces.

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.

Controlling Search Results with Banner Fields and Weighting

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 and ranks the results based on the specified weighting settings. The default banner field properties are Name, Description, and Full-Text Content, with a high weight applied to Name. However, you can add other properties (such as Keyword, Department, or Author) and change the weighting to further refine the search results.

To access the Search Results Manager you must have the following rights:

- Access Administration activity right
- Access Utilities activity right
- Access Search Results Manager activity right

Note: Since banner field settings 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.

1. Click **Administration**.
2. From the **Select Utility** drop-down list, select **Search Results Manager**.
3. Click the **Banner Fields** page.
   The Banner Fields page displays the properties that the portal searches. The following information is displayed for each banner field.

Column	Description
Property	The property on the search banner.
Percent Weight	The proportion of the weight assigned to the property field.
Weight	The relevance ranking of the property field. Type in a weight for each property field. If you want to attach more weight to a particular property field, increase the weight number.

4. Add, edit, or delete banner fields to improve search results.
   - To add a new property field and set its weight:
     1. Click  **Add Field**.
        New fields appear.
     2. From the **Property** drop-down list, select the property field you want to add.
     3. Set the weight by typing a number in the text box in the **Weight** column.
   - To delete a property field, select it and click  .
   - To remove any customizations you have made, click  **Restore Defaults**.
   - To change the weight of a property field, type a number in the text box in the **Weight** column.

     The values in the **Percent Weight** column are automatically updated when you change the value in the **Weight** column.

About Spell Correction for Searches

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. Users can disable spell correction on a per-search basis by using the <WORD> operator.

Enabling and Disabling Spell Correction for Searches

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.

To access the Search Results Manager you must have the following rights:

- Access Administration activity right
- Access Utilities activity right
- Access Search Results Manager activity right

Note: Spell correction is enabled by default.

1. Click **Administration**.
2. From the **Select Utility** drop-down list, select **Search Results Manager**.
3. Click the **Thesaurus and Spell Correction** page.
4. Enable or disable spell correction.
   - To enable spell correction, select the **Apply Spell Correction** check box.
   - To disable spell correction, clear the **Apply Spell Correction** check box.

About the 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 “web applications” might not return documents that discuss portlets or web services. However, by creating a thesaurus entry

for “web applications,” 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, “web applications” is not really a synonym for web services, so a document that actually contains web applications should rank higher than one that contains web services.

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

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

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

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.

About the Thesaurus File

The thesaurus is a comma-delimited file, in which each line 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]
```

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.

Creating and Implementing the Synonym List for the Thesaurus

The thesaurus is a comma-delimited file, in which each line represents a single thesaurus entry. After you create the file, you load it into the Search Service.

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

   For details on how to format entries, see *About the Thesaurus File* on page 267.

   Note: Thesaurus entries must be in lower-case.

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 is located in the `bin\native` directory of the Search Service installation (for example, `C:\bea\alui\ptsearchserver\6.5\bin\native\customize.exe`). The utility must be run from a command prompt, taking command-line arguments for the thesaurus file and the path to the Search Service installation:

   ```
   customize -r thesaurus_file SEARCH_HOME
   ```

   Replace *thesaurus file* with the path to the thesaurus file and replace*SEARCH_HOME* with the root directory of the Search Service installation.

   For example, if your thesaurus file is located in `\temp` and your Search Service was installed in the default location, type:

   ```
   customize -r \temp\thesaurus.cdf C:\bea\alui\ptsearchserver\6.5
   ```

   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.

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

If you have not already done so, you must enable the thesaurus in the Search Results Manager.

Enabling the 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.

To access the Search Results Manager you must have the following rights:

- Access Administration activity right
- Access Utilities activity right
- Access Search Results Manager activity right

1. Click **Administration**.
2. From the **Select Utility** drop-down list, select **Search Results Manager**.
3. Click the **Thesaurus and Spell Correction** page.
4. Select **Use the Thesaurus**.

If you have not already done so, you must create the synonym list in the database.

Reverting to the Default Thesaurus Mappings

The customize utility has a command-line mode for reverting to the set of mappings files that shipped with the Search Service (removing any thesaurus customizations).

1. Stop the Search Service.
2. At a command prompt, run the customize utility.

   The customize utility is located in the `bin\native` directory of the Search Service installation (for example, `C:\bea\alui\ptsearchserver\6.5\bin\native\customize.exe`). The utility must be run from a command prompt, taking a command-line argument for the path to the Search Service installation:

   ```
   customize -default SEARCH_HOME
   ```

   Replace *SEARCH_HOME* with the root directory of the Search Service installation.

   For example, if your Search Service was installed in the default location, type:

   ```
   customize -default C:\bea\alui\ptsearchserver\6.5
   ```

   The files in *SEARCH_HOME*`\common` are removed and replaced with the original thesaurus file contents.

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

If you no longer want to use thesaurus expansion, disable the thesaurus in the Search Results Manager.

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 Objects Directory. You can customize this drop-down list to include additional categories relevant for your users. For example, if you use a property in your portal documents named Region, 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.

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. See *Mapping Source Document Attributes to Portal Properties Using the Global Document Property Map* on page 160 and *Associating Properties with Portal Objects Using the Global Object Property Map* on page 164.

You must also 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 Custom Search Results Sorting

You can enable sorting by custom properties by editing the `portalconfig.xml` file.

1. Get the object ID of the properties you want to configure for custom sorting.
   a) Navigate to the property in the Administrative Objects Directory.
   b) Right-click the link to the property and then choose **Properties**.
   This will yield a link that looks something like this:

   http://portal.company.com/portal/server.pt?open=3&objID=200&parentname=ObjMgr&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 you would complete the CategoryField entry as follows:

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

2. Open `portalconfig.xml` in a text editor and find the `<Search>` section.
3. Add the following entries to the `<Search>` section:

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

   Replace *CategoryName* with the name you want to appear in the **Sort By** drop-down list (for example, Region).

   Replace *ObjectID* with the integer that identifies the property object.

4. To add more entries, repeat Step 3, adding analogous tags named `CategoryName_2`, `CategoryField_2`, `CategoryName_3`, `CategoryField_3`, and so forth.

   Note: The Category tags must be numbered consecutively without skipping. For example, if there is a `CategoryName_3` tag, it must be preceded by tags for `CategoryName_1` and `CategoryName_2`.

About 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.

About Checkpoints

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.

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.

About 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.

About Search 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\6.5\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`.

About the Command Line Admin Utility

The Command Line Admin Utility lets you to perform the same functions you can perform in the Search Cluster Manager as well as change the run level of the cluster and purge and reset 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\6.5\bin\native\cadmin.exe`). Invoking the command with no arguments displays a summary of the available options:

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

This Command Line Admin Utility lets you perform the following actions:

- *Requesting Search Cluster Status for a Particular Node* on page 276
- *Requesting Search Cluster Status* on page 275
- *Changing the Run Level of the Cluster* on page 277
- *Purging the Search Collection* on page 280
- *Initiating a Cluster Checkpoint* on page 277
- *Reloading from a Checkpoint* on page 277
- *Changing Cluster Topology* on page 278
- *Aborting a Checkpoint or Reconfiguration Operation* on page 280

Requesting Search Cluster Status

You can use the Command Line Admin Utility to view the status of your search cluster.

The Command Line Admin Utility is located in `bin\native` folder in the Search Service installation folder (for example, `C:\bea\alui\ptsearchserver\6.5\bin\native\cadmin.exe`).

- Run the status command to display the status of the cluster.

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

By default, the status command displays a terse, one-line summary of the current state of the cluster:
2005-04-22 13:54:13 checkpoint_xxx 0/1/198 0/1/230 impaired

- Run the status command with the verbose flag to display 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
- Run the status command with the period flag 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 Search Cluster Status for a Particular Node

You can use the Command Line Admin Utility to request information about specific nodes within the cluster.

The Command Line Admin Utility is located in `bin\native` folder in the Search Service installation folder (for example, `C:\bea\alui\ptsearchserver\6.5\bin\native\cadmin.exe`).

- Run the nodestatus command to display the status of a particular node.

  ```
  % cadmin nodestatus p0n0 p1n0
  ```

  This displays the same type of information that is displayed as part of the verbose cluster status request:
  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
- Run the nodestatus command with the period flag to repeatedly emit status requests at a specified interval.

  ```
  % 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 use the Command Line Admin Utility to 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.

The Command Line Admin Utility is located in `bin\native` folder in the Search Service installation folder (for example, `C:\bea\alui\ptsearchserver\6.5\bin\native\cadmin.exe`).

- 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
  ```

Initiating a Cluster Checkpoint

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.

The Command Line Admin Utility is located in `bin\native` folder in the Search Service installation folder (for example, `C:\bea\alui\ptsearchserver\6.5\bin\native\cadmin.exe`).

Reloading from a Checkpoint

You can use the Command Line Admin Utility to reload your cluster from a saved checkpoint.

The Command Line Admin Utility is located in `bin\native` folder in the Search Service installation folder (for example, `C:\bea\alui\ptsearchserver\6.5\bin\native\cadmin.exe`).

Run the restore command.

```
% cadmin restore
```

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

Example Output from Restoring from a Checkpoint

```
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 can use the Command Line Admin Utility to add or remove nodes from the search cluster or repartition the search cluster.

The Command Line Admin Utility is located in `bin\native` folder in the Search Service installation folder (for example, `C:\bea\alui\ptsearchserver\6.5\bin\native\cadmin.exe`).

1. Run the topology command.

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

2. 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.

Example Output from Adding and Removing 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
```

Example Output from Repartitioning 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.

- To exit from the command line utility, type `CTRL+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.

Purging and Rebuilding the Search Collection

You can purge and rebuild the contents of the search collection. You might want to do this in a dire situation where the contents of the cluster are corrupted beyond repair and good checkpoints are not available for recovery.

1. Put all cluster nodes in standby mode.
2. Purge the collection.
3. Rebuild the collection.
4. If your portal deployment includes AquaLogic Interaction Collaboration, rebuild the Collaboration index.

Purging 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 (and rebuilding) 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.

As a safeguard against purging the collection by accident, all cluster nodes must be in standby mode.

The Command Line Admin Utility is located in `bin\native` folder in the Search Service installation folder (for example, `C:\bea\alui\ptsearchserver\6.5\bin\native\cadmin.exe`).

By default, the checkpoints and index queue are left in place. This allows you to rebuild the local index on a node where the archive appears to be corrupted. You can add a flag to the command to remove the checkpoints.

- To purge the search collection, but keep checkpoints:

  ```
  % cadmin purge
  ```

  As a safeguard against purging the collection by accident, 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
  ```

Sometimes the purge command does not work properly, where it fails to purge the index files. During this scenario the search file structure gets tainted and the Search Service will not start up.

If this occurs you receive an error similar to this:
Failed Unexpected exception while sending PURGE to searchserver01, Error during parsing: Failed -- purge failed on streetfighter01 This node should be shutdown, purged, and restarted manually.

At this point you will not be able to shut down, purge, or restart the Search Service. You must re-run the AquaLogic Integration installer. Select only the Search Service option and choose overwrite.

If you are purging the collection to correct a problem (for example if your collection was corrupted or you had to reinstall the Search Service), your next step is to rebuild the collection.

Rebuilding the Search Collection

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.

To rebuild the search collection you must have the following rights:

- Access Administration activity right
- Access Search Results Manager activity right

The Search Service Manager lets you specify when and how often the Search Update Agent repairs your search index. Rather than synchronizing particular objects, the repair 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.

1. Configure the Search Service to repair itself.
   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** again.
2. Wait one minute for the setting to update.
3. Run one of the Search Update jobs in verbose mode.
   a) Open the **Intrinsic Operations** folder.
   b) Open one of the Search Update jobs.
   The Job Editor opens.
   c) Change the **Logging Level** to **Verbose** and click **Finish**.

      Note: Make note of the logging mode before you change it, so that you can change it back after the repair is complete.

   d) Select the job you just edited and click **Run Once**.

      By running the job this way, you avoid having to go back into the job and revert to the previous schedule settings.

4. Ensure that the job is running in repair mode.
   a) Open the job you just created; it should be called something like Search Update 1 — Run Once.
      The Job Editor opens.
   b) Click the **Job History** page.
   c) Click the job name.
      The job log opens.
   d) Ensure that the job is running in repair mode.

      The second line of the job log should be similar to this:
      Mar 1, 2008 9:10:02 AM- PTIndexer.ctor : Indexing will extract at most 1000000 encoded bytes of text from each document.
      About half-way down the first page of the log you should see a message that should be similar to this:
      Mar 1, 2008 9:10:02 AM- Search Update Agent is repairing the directories...

5. Reinstall the Search Service and select **Overwrite the existing search index**. For details on installing the Search Service, refer to the *Installation Guide for AquaLogic Interaction*.

Rebuilding the Collaboration Search Collection

Rebuilding reconciles data between the AquaLogic Interaction Collaboration database and Search Service index. Since this is a lengthy and computationally expensive process, use the rebuild operation only when absolutely necessary.

To access the Collaboration Administration Utility you must have the following rights:

- Access Administration activity right
- Access Utilities activity right
- Manage Collaboration activity right

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **Collaboration Administration**.
3. Click the **Search Service** page.
4. Click **Rebuild Search Collection**.

About the Search Update Job

The Search Update job performs the following actions on the search index: updates the index, releases expired locks on users and objects, and repairs the search index according to the Search Service Manager repair settings.

The Search Update job is located in the Intrinsic Operations administrative folder. The default frequency of the Search Update job is one hour, which is suitable for most portal deployments. If your search index is very large, the Search Update Agent might not be able to finish in one hour, so you should edit the frequency of the job.

How the Search Index is Updated

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).

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.

About Providing Search Access to External Repositories with Federated Searches

Federated searches allow you to establish search relationships with other sources (including other portals, web sites, or custom databases). 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.

There are incoming and outgoing federated searches:

- An incoming federated search allows other AquaLogic Interaction portals to search your portal.
- An outgoing federated search enables users of your portal to search other AquaLogic Interaction portals or other external repositories.

Search Web Services

Search web services allow you to specify general settings for your remote search repository, leaving the security settings to be set in the associated outgoing federated searches. This allows you to segregate access to your search repository through multiple outgoing federated searches.

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 IDK that enables you to write your own search web services in Java or .NET. For details, refer to the *BEA AquaLogic User Interaction Development Center*.

Note: The settings for outgoing federated search objects will often be specific to the search web services that implement the searches. In these cases, the web services themselves provide the configuration options as Service Configuration Interface (SCI) pages.

Portal-to-Portal Searches

One AquaLogic Interaction portal can request and/or serve content to another AquaLogic Interaction portal. 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?
- What portal identification name 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's outgoing federated search and the serving portal's incoming federated search.
- What content from the serving portal will be available to the requesting portal?

  If both portals share a common external database of users, such as an LDAP server or Active Directory domain, and those users have been imported into both portals, 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 do not share a database of user information, users from the requesting portal must impersonate users from the serving portal. Because impersonation is specified on a group basis (that is a group from the requesting portal is set to impersonate a user from the serving portal), you should create a different serving portal user for each requesting portal group that needs access to different content in the serving portal.

  Note: You should create new serving portal users specifically for the purpose of impersonation, then communicate the user names and what they can access to the administrator of the requesting portal.

  The serving portal can also allow unauthenticated users to search the portal as the Guest user.

After you and the other administrators involved have determined how this relationship will work, you are ready to establish your incoming or outgoing federated searches.

For an example of how requesting portal users can impersonate serving portal users to gain access to secured content, see *Example of Impersonating Serving Portal Users* on page 292.

To learn how multiple portals accessing the same user repository can share content, see *Building a Composite Portal with Federated Searches* on page 286.

Building a Composite Portal with Federated Searches

Multiple portals accessing the same user repository can share content. All portals involved in the relationship must import the same users and groups from a single user repository.

There are two scenarios for building a composite portal:

- Multiple content portals each possess links to a large number of documents. A single user can visit the separate content portals, always with the same user name and password, and always receive access to the correct content.

  In this scenario, each portal acts as both a serving and a requesting portal.
- One portal is set up as a master portal rather than a content portal. Through this master portal, users can access content from the various content portals.

  In this scenario, the master portal acts as the requesting portal, and the content portals act as the serving portals.

- On each serving portal, the administrator creates an incoming federated search that includes the authentication sources that the serving portals share with the requesting portals.

  All users making requests to a serving portal need to be imported into the portal through one of these common authentication sources.
- On each requesting portal, the administrator creates an outgoing federated search for each content portal, selecting **No** for **Send portal authentication**.

  Users will make the requests using their own user accounts.

Creating a Search Web Service

Search web services allow you to specify general settings for your remote search repository, leaving the security settings to be set in the associated outgoing federated searches. This allows you to segregate access to your search repository through multiple outgoing federated searches.

Before you create a search web service, you must:

- Install the search provider on the computer that hosts the portal or on another computer
- Create a remote server pointing to the computer that hosts the search provider (optional, but recommended)

To create a search web service you must have the following rights and privileges:

- Access Administration activity right
- Create Web Service Infrastructure activity right
- At least Edit access to the parent folder (the folder that will store the search web service)
- At least Select access to the remote server that the search web service will use

1. Click **Administration**.
2. Open the folder in which you want to store the search web service.

3. In the **Create Object** drop-down list, click **Web Service — Search**.
   The Search Web Service Editor opens.
4. On the **Main Settings** page, complete the following task:
   -
5. Click the **HTTP Configuration** page and complete the following task:
   -
6. Click the **Advanced URL Settings** page and complete the following task:
   -
7. Click the **Advanced Settings** page and complete the following task:
   -
8. Click the **Authentication Settings** page and complete the following task:
   -
9. Click the **Preferences** page and complete the following task:
   -
10. Click the **User Information** page and complete the following task:
    -
11. Click the **Debug Settings** page and complete the following task:
    -
12. Click the **Properties and Names** page and complete the following tasks:
    - *Naming and Describing an Object* on page 217

      You can instead enter a name and description when you save this search web service.
    - *Localizing the Name and Description for an Object* on page 342 (optional)
    - *Managing Object Properties* on page 219 (optional)

The default security for this search web service is based on the security of the parent folder. You can change the security when you save this search web service (on the **Security** tab page in the

Save As dialog box), or by editing this search web service (on the **Security** page of the Search Web Service Editor).

Portal administrators with at least Select access to this search web service can create outgoing federated searches based on the web service.

Allowing Other Portals to Search Your Portal

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

Before you create an incoming federated search, you must:

- Agree upon a portal identification name and password with the administrator of the requesting portal.
- If the users from the requesting portal do not exist in your portal, create one or more portal users that can be impersonated by users of the requesting portal.

To create an outgoing federated search you must have the following rights and privileges:

- Access Administration activity right
- Create Federated Searches activity right
- At least Edit access to the parent folder (the folder that will store the federated search)
- If the users from the requesting portal do not exist in your portal, at least Select access to the authentication sources or groups that include the impersonated users

1. Click **Administration**.
2. Open the folder in which you want to store the federated search.
3. In the **Create Object** drop-down list, click **Federated Search - Incoming**.
4. In the **Portal identification name** box, type the agreed upon name.
5. In the **Portal identification password** box, type the agreed upon password.
6. In the **Password confirmation** box, type the password again.
7. In the **Served links are valid for** box, type the minimum number of minutes or which these results should be cached.

   After a requesting portal issues a search of your portal, the links returned by the search are stored for at least as long as you specify here. After this period has elapsed, the user on the requesting portal might need to re-issue the search.
8. To allow unauthenticated users to search the portal as a guest, click the **Allow unauthenticated users to search as the Guest user** box.

9. If the users from the requesting portal do not exist in your portal, select the authentication sources or groups that include the impersonated users.

   Incoming search requests include the name of a local portal user (that is, a user from the serving portal) to impersonate during the search. The request is honored only if the impersonated user is a member of one of the authentication sources or one of the groups you specify.

   - To add an authentication source, click **Add Authentication Source**, in the Choose Authentication Sources dialog box, select an authentication source, and click **OK**.
   - To add a group, click **Add Group**, in the Choose Groups dialog box, select a group, and click **OK**.
   - To delete an authentication source or a group, select it and click .

     To select or clear all of the authentication source or group boxes, select or clear the box to the left of **Authentication Sources** or **Groups**.
   - To toggle the order in which the authentication sources or groups are sorted, click **Authentication Sources** or **Groups**.

Providing Search Access to External Repositories with Outgoing Federated Searches

An outgoing federated search enables users of your portal to search other AquaLogic Interaction portals or other external repositories.

Before you create an outgoing federated search, you must:

- Create a search web service.
- Agree upon a portal identification name and password with the administrator of the serving portal.
- If your portal users do not exist in the serving portal, work with the serving portal user to determine the serving portal users that can be impersonated and what they can access.

To create an outgoing federated search you must have the following rights and privileges:

- Access Administration activity right
- Create Federated Searches activity right
- At least Edit access to the parent folder (the folder that will store the federated search)

- If your portal users do not exist on the serving portal, at least Select access to the groups that need to impersonate serving portal users

1. Click **Administration**.
2. Open the folder in which you want to store the federated search.
3. In the **Create Object** drop-down list, click **Federated Search - Outgoing**.
   The Choose Web Service dialog box opens.
4. Select the web service that provides the basic settings for your outgoing federated search and click **OK**.
   The Outgoing Federated Search Editor opens, displaying the Portal to Portal Settings page.
5. If you are not searching another AquaLogic Interaction portal, leave **No** selected.

   If you are searching another AquaLogic Interaction portal:

   a) Next to **Send portal authentication**, select **Yes**.
   b) In the **Portal identification name** box, type the agreed upon name.
   c) In the **Portal identification password** box, type the agreed upon password.
   d) In the **Password confirmation** box, type the password again.
   e) If your portal users do not exist on the serving portal, under **User Name Aliasing**, map groups from your portal to users from the serving portal that they can impersonate:

      Note: When a requesting user tries to search a serving portal, the requesting portal examines the list of mapped groups from the top down; the first group in the list to which the requesting user belongs is used to determine what serving portal user the requesting user will impersonate. Therefore, groups with high levels of security should be mapped first (at the top of the list), so that requesting users are granted the highest level of security available to them.

      1. Click  **Add Group**, in the Select Group dialog box, select the groups you want to add and click **OK**.
      2. To the far right of the group, click  .
      3. In the **Use this user name alias** column box, type the name of the serving portal user whom you want this group of requesting users to impersonate.
      4. Click  to save the mapping.

      To delete a group, select it and click  .

      To select or clear all of the group boxes, select or clear the box to the left of **Members of this group**.

After the administrator of the serving portal has set up the incoming federated search, your users can use federated search to search content from the other portal.

Example of Impersonating Serving Portal Users

This example shows how a search relationship might be set up between two separate portals.

The fictional company *Servicor* wants to share content with its fictional partner *Requesticon*. In this case, Servicor's portal is the serving portal, and Requesticon's portal is the requesting portal.

Configuring the Serving Portal

First, the administrator of the Servicor portal creates two portal users: *Requesticon Engineer* and *Requesticon Executive*. Both of these users are added to the portal group named *Requesticon Visitors*.

These users are then individually granted access to appropriate content on the Servicor portal. Requesticon Engineer is granted Read access to the *Engineering*, *QA*, and *Product Management* folders of the Servicor Knowledge Directory. Requesticon Executive is granted Read access to the *Servicor Market* and *Investor Relations* folders.

The administrator of the Servicor portal then sets up an incoming federated search. On the Main Settings page of the Incoming Federated Search Editor, the Servicor portal administrator includes the *AquaLogic Interaction Authentication Source* and the group *Requesticon Visitors*. The AquaLogic Interaction Authentication Source is included because the Requesticon Engineer and the Requesticon Executive users were both created in the portal; had they been imported through another authentication source, then *that* authentication source would need to be included instead. The Requesticon Visitors group is included here to prevent users of the requesting portal from attempting to impersonate any user other than Requesticon Engineer or Requesticon Executive.

With the serving portal configured this way, only requests issued by Requesticon Engineer and Requesticon Executive are answered, and only appropriate content is visible.

Configuring the Requesting Portal

On the Main Settings page of the Outgoing Federated Search Editor, the administrator of the Requesticon portal selects **Yes** for **Send portal authentication**. Then, under User Name Aliasing, the Requesticon portal administrator maps the group *Executives* to the Servicor user named *Requesticon Executive* and the group *Engineers* to the Servicor user named *Requesticon Engineer*. This way, all users that are members of the Engineers group impersonate Requesticon Engineer when issuing requests, and all users that are members of the Executives group issue requests as Requesticon Executive.

Note: The Requesticon Engineer and Requesticon Executive exist only in the Servicor portal, not in the Requesticon portal; these users were created specifically for impersonation by Requesticon users.

When a requesting user tries to search a serving portal, the requesting portal examines the list of mapped groups from the top down; the first group in the list to which the requesting user belongs is used to determine what serving portal user the requesting user will impersonate. Therefore, groups with high levels of security should be mapped at the top of the list. The requesting portal administrator made sure to add the Executives group *before* the Engineers group so that if any user on the requesting portal is a member of both the Executives group and the Engineers group, then that user will impersonate the Requesticon Executive user. Being an executive, this user is likely to be granted access to more content.

CHAPTER 10

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.

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)

About Portal Agents

The portal comes with several operations that can only be accessed through the jobs with which they are associated. These special operations are referred to as agents.

The following agent jobs are stored, by default, in the **Intrinsic Operations** folder:

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

  The Audit Log Management Agent job also archives and deletes audit files according to the schedule set in the Audit Manager utility.
- The Bulk Subscriptions Agent job subscribes users in bulk to the communities and portlets you specify in the Bulk Add editor.
- The Document Refresh Agent job performs background maintenance on your Knowledge Directory, such as refreshing document links and properties, and deleting expired documents.
- The Dynamic Membership Update Agent job updates dynamic portal group memberships.
- The Search Update Agent job 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 job also repairs the search index according to the frequency set in the Search Service Manager utility.
- The Weekly Housekeeping Agent job 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.

About Running Scripts Through the Portal

An external operation enables you to run shell scripts (for example, .sh or .bat files) through the portal and schedule these actions through portal jobs. For example, you might want to create a script that queries documents, pings portal servers, e-mails snapshot query results to users, or runs some other custom job, then create an external operation that points to the script, and use a job to run the script on a specified schedule.

External Operations Created Upon Installation

When you install the portal, there are two working example external operations that are created in the **Intrinsic Operations** portal folder:

- **Bulk Subscriber**: This external operation subscribes users to communities and groups when you use bulk add.
- **Snapshot Query Mailer**: This is a sample external operation that e-mails the results of snapshot queries to users.

  For more information on this external operation, see *E-mailing the Results of a Snapshot Query* on page 202.

Creating External Operations to Run Scripts Through the Portal

An external operation enables you to run shell scripts (for example, .sh or .bat files) through the portal and schedule these actions through portal jobs. For example, you might want to create a script that queries documents, pings portal servers, e-mails snapshot query results to users, or runs some other custom job, then create an external operation that points to the script, and use a job to run the script on a specified schedule.

To create an external operation you must have the following rights and privileges:

- Access Administration activity right
- Create External Operations activity right
- At least Edit access to the parent folder (the folder that will store the external operation)
- At least Edit access to the job that will run this external operation

Important:

- Because the standard error output from the command or script is captured to the job log, avoid the use of new shells, redirects, and pipes.
- Passing arguments to `cmd` or `start` in shell programs might disable the time-out mechanism.
- When you are extending scripts in the External Operation Editor, carefully consider all potential effects of the scripts. Make sure that your script does not introduce a security risk.

1. Click **Administration**.
2. Open the folder in which you want to store the external operation.
3. In the **Create Object** drop-down list, click **External Operation**.
4. In the **Operating System Command** box, type the relative path and file name of the script enclosed in quotes (").

Important: All external operation scripts must reside in the scripts directory of each of the Automation Services that will run them. The scripts directory is located on the computer that hosts the Automation Service, in the AquaLogic User Interaction installation directory (for example, `C:\bea\alui\ptportal\scripts`). The Automation Service will not run any scripts that are not in this directory.

The following tokens in the command line will be substituted:

- names of environment variables surrounded with percent signs (%)
- <user_id>
- <security_token>
- <job_id>
- <operation_id>
- <last_job_runtime>

Expanded tokens that contain spaces or special characters which are not surrounded with quotes (") are enclosed in quotes automatically.

5. In the **Time-out in seconds** box, type the number of seconds after which, if this operation is still running, you want the job to stop.

   If you do not want to set a time-out, leave this setting at **0** (infinite).

To run this operation, you must associate it with a job and schedule the job to run.

Registering Automation Services

Before you can run jobs, you need to register any computers hosting Automation Services and register job folders with those Automation Services. The primary Automation Service is registered when you install the Automation Service and execute the related database scripts described in the *Installation Guide for AquaLogic Interaction* .

To access the Automation Service Utility you must be a member of the Administrators Group.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **Automation Service**.
3. Click  **Add Automation Service**.
   The Register Automation Service dialog box opens.

4. Type the name of the computer that hosts the Automation Service.
   Use the host name only (for example, automationserver1), not the fully qualified domain name.
5. Type the network address that identifies the machine.
6. Click **Finish**.

The new Automation Service appears in the list. To the right of each Automation Service, you can see if the server is online or offline and when the job folders associated with the server were last updated.

You must assign job folders to this Automation Service before you can run jobs with it.

Registering Job Folders to Run Jobs

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 access the Automation Service Utility you must be a member of the Administrators Group.

Note: You must register each folder separately. An Automation Service does not monitor child folders of registered folders.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **Automation Service**.
3. Click the name of the Automation Service that you want to run the jobs.
   The Register Folders Editor opens.
4. Click  **Add Folder**.
   The Add Job Folder dialog box opens.

The job folder appears in the list. Under each registered job folder, you can see the jobs stored in that folder and the next time each job is scheduled to run.

- To edit a job, click its name.
- To remove a job folder, select the folder and click  .

To select or clear all of the folder check boxes, select or clear the box to the left of **Folders**.

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 Guide for AquaLogic Interaction* .

Creating 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.

1. Click **Administration**.
2. Open an administrative folder.
3. In the **Create Object** drop-down list, click **Job**.
4.

Schedule your job to run.

Associating an Object with a Job

On the Set Job page, you can associate an object with a new or existing job. You run jobs to import users with authentication sources, import content with content crawlers, run external operations, and import user information with profile sources.

Before you can run jobs, you must:

- 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 300.
- Register the Automation Service with the portal, as described in *Registering Automation Services* on page 298.
- Assign administrative folders to the registered Automation Services, as described in *Registering Job Folders to Run Jobs* on page 299.

1. Open the object's editor by creating a new object or editing an existing object.
2. Click the **Set Job** page.
3. Associate the object with one or more jobs:
   - To run this object with an existing job, click **Add Job**; then, in the Choose Jobs dialog box, select the jobs you want to add this object to and click **OK**.
   - To create a new job to run this object, click **Create Job**, then, in the Job Editor, schedule your job and click **Finish**.
   - To remove a job, select the job and click .

     To select or clear all of the job check boxes, select or clear the check box to the left of **Job Name**.
   - To edit a job, click the job name.

     If you added this object to an existing job, you might want to verify that the job is scheduled to run.
   - To change the order in which the jobs are sorted, click **Job Name**.

Viewing Job Status and Job Logs

You can view a history for a job as well as the logs from each job on the Job History page of the Automation Service Utility.

To access the Automation Service Utility you must be a member of the Administrators group.

1. Click **Administration**.

2. In the **Select Utility** drop-down list, click **Automation Service**.
3. Click **Job History** page.
4. View the history of jobs that have run and the logs for individual jobs.
   - To view the detailed job log for a job, click the name of the job.

     The log appears in a dialog box.
   - To search the log, in the job log dialog box, enter a search term (using wildcards such as *) and click **Search Log**.
   - To limit the histories displayed on this page, in the **View** and **to** boxes, type the earliest and latest dates for which you want to view job histories and click >>.

     To choose the date from a calendar, click  .
   - To remove the filter clear the **View** and **to** boxes and click >>.
   - To refresh the job history data, click  .
   - To download a text file version of a detailed job log, click  to the far-right of the job name; when asked to open or save the file, click **Save**, specify a location in which to save the file, and click **Save** again.

Job History Information

The Job History page of the Automation Service Utility provides information about in-process and completed jobs.

Column	Description
Job Name	Displays the name of the job. Click the job name to view the detailed job log.
Server	Displays the name of the Automation Service that ran the job.
Next Run	Displays the next date and time the job is scheduled to run.
Start	Displays the starting date and time for the last time the job ran.
Finish	Displays the ending date and time for the last time the job ran.
Status	Displays what happened when the job ran last: • **Succeeded** indicates that the job was able to complete.

Column	Description
	• **Failed** (in red text) indicates that the job experienced errors and was not able to complete. • **In Process** indicates that the job is running now. • **Interrupted** indicates that the job was terminated unexpectedly. • **Suspended** indicates that the job stopped before completing its work and will resume its work at the next scheduled runtime.
Download	Click the button to download a text file version of the detailed job log.

Deleting Job Histories

On the Job History page of the Automation Service Utility, you can delete job histories from the database.

To access the Automation Service Utility you must be a member of the Administrators group.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **Automation Service**.
3. Click **Job History** page.
4. To delete job histories from the database, in the **Delete to** box, type the latest date for which you want to delete histories and click >>.

   To choose the date from a calendar, click  .

   Any histories with this date or an earlier date are deleted from the database.

Aborting In-Process Jobs

On the Job History page of the Automation Service Utility, you can stop a job that is processing.

To access the Automation Service Utility you must be a member of the Administrators group.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **Automation Service**.

3. Click **Job History** page.
4. To cancel a job, select the job and click  **Abort**.

   Note: You cannot cancel jobs that have already completed. The check boxes next to completed jobs are unavailable (grayed out).

CHAPTER 11

Migrating Portal Objects

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

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.

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.

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.

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 the Access Utilities activity right 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

Migration Feature	Description
	to overwrite previously migrated objects, you have the option of creating a new instance of the same object, with a new UUID.

Requesting That an Object Be Migrated

You can request that an object be added to a migration package to be exported to another portal.

Note: You must have at least Edit access to the object for which you want to request migration.

1. Search for the object or click **Administration** and navigate to the object.
2. Select the object and click  .
3. In the Script Prompt dialog box, describe why you want this object migrated and click **OK**.

To view the status of your migration request, open the object's editor and click the **Migration History and Status** page. Under **Migration Status**, you see whether your request is waiting for approval, has been approved, or has been rejected, as well as your comments and any comments from the portal administrator.

Approving Objects for Migration

When users want an object to be migrated, they submit a migration request. A portal administrator can then approve the request, and the object is added to the migration package.

There are two ways to approve objects for migration: in the object's editor or using the Approve Objects for Migration Utility.

- To approve a single object requested for migration, use the object's editor.
  See *Approving an Object for Migration* on page 307.
- To approve all objects requested for migration, use the Approve Objects for Migration Utility.
  See *Approving Objects for Migration Through the Administrative Utility* on page 308 .

Approving an Object for Migration

When users want an object to be migrated, they submit a migration request. A portal administrator can then approve the request, and the object is added to the migration package.

1. Open the object's editor by creating a new object or editing an existing object.
2. Click the **Migration History and Status** page.
   Under **Migration Status**, you see whether this object has been requested for migration, and, if so, whether it is waiting for approval, has been approved, or has been rejected.
3. If you are a member of the Administrators group, and you want to add this object to the migration package to be migrated to another portal, select **Approve this object for migration**.

   Note: Users who are not members of the Administrators group do not see this option.

After approving objects for migration, you can use the Migration - Export Utility to create a migration package.

Approving Objects for Migration Through the Administrative Utility

When users want an object to be migrated, they submit a migration request. A portal administrator can then approve the request, and the object is added to the migration package.

To use the Approve Objects for Migration Utility you must be a member of the Administrators Group.

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, you must be a member of the Administrators group.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **Migration - Export**.
3. On the **Portal Resources** page, complete the following task:
   - *Selecting Objects to Export in a Migration Package* on page 309
4. Click the **Package Settings** page and complete the following task:
   - *Specifying a Name, Description, and Contact for a Migration Package* on page 309
5. Click the **Add Existing Package Resources** page and complete the following task:
   - *Adding Resources from Another Migration Package* on page 310
6. 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.

You can now use the migration package to import the migrated resources into another portal.

Specifying a Name, Description, and Contact for a Migration Package

You must specify a name, description, and contact person for a migration package.

To create a migration package, you must have at least Edit access to the objects you want to add to the package.

1. If the Migration — Export Utility is not already open, open it now and display the **Package Settings** page.
2. In the **Package name** box, type the name for the package file that will be created when you click **Finish** (this file is given a .pte extension).
3. In the **Package description** box, type a description that clarifies the purpose of this export package to other portal administrators.
4. In the **Publisher name** box, type the person to contact with any questions about this export package.

Selecting Objects to Export in a Migration Package

You can select objects to export on the Portal Resources page of the Migration — Export Utility.

To create a migration package, you must have at least Edit access to the objects you want to add to the package.

1. If the Migration — Export Utility is not already open, open it now and display the **Portal Resources** page.
2. Under **Resources**, select the objects you want to add and what you want to export.
3. To select individual objects, select the type of object from the **Select Resources** drop-down list, then, in the dialog box, select the objects you want to add and click **OK**.
4. To add all objects that have been approved for migration, click  **Add All Approved**.
5. If the object is a folder and you want to export the folder's contents, select **Export Contents**.
6. To export the object's dependencies, select **Export Dependencies**.

Dependencies are any other objects that are required by the object you are exporting.

7. To remove an object, select the object and click ✕.

   To select or clear all of the object check boxes in a column, select or clear the check box above the column.

8. To toggle the sort order of the objects, click **Resources**.
9. Under **Export Settings**, select the check box if you also want to export parent folders of objects you marked for Export Dependencies.

   If you do not select this option, only references to those parent folders will be exported.

Adding Resources from Another Migration Package

You can add objects from an existing migration package to the package you are creating on the Add Existing Package Resources page of the Migration — Export Utility.

To create a migration package, you must have at least Edit access to the objects you want to add to the package.

Note: If the portal settings on an object in the existing migration package have changed since the package was created, the current portal settings on the object will be exported.

1. If the Migration — Export Utility is not already open, open it now and display the **Add Existing Package Resources** page.
2. Under **Migration Package**, specify the package from which you want to add objects.
   - To browse to a package on your computer or in your network, click **Browse**, select the migration package file (a .pte file), and click **Open**.
   - To get the package from a web address, select **Web Address** and type the URL to the package.

     If you need to enter login credential to access the web address, type the information in the **Username** and **Password** boxes.

3. Click **Load Package**.
   The package name, description, and publisher appear, and the objects from this package are displayed on the **Portal Resources** page.

   Note: Only objects that are currently in the portal will be displayed on the **Portal Resources** page.

Select resources from this package on the **Portal Resources** page.

Creating a Migration Package Using the Command Line Tool

You can create a migration package to export portal objects from one portal to be imported into another portal.

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. See *Creating a Migration Package in the Portal* on page 308.

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.

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

You can import objects from another portal using the Migration — Import Utility..

Before you import a migration package (.pte file), place the package to a location that is accessible over your network.

Note: If you are also importing Collaboration or Publisher objects, you must place the .zip file in the same location as the .pte fle.

To import a migration package you must be a member of the Administrators group.

1. Click **Administration**.
2. In the **Select Utility** drop-down list, click **Migration - Import**.
3. On the **Package Settings** page, complete the following task:
   - *Specifying the Location and Import Settings for the Migration Package* on page 312
4. Click the **Unresolved Dependencies** page and complete the following task:
   - *Resolving Import Dependencies* on page 314

   Note: This page appears only if there is an unresolved dependency.
5. Click the **Portal Resources** page and complete the following task:
   - *Selecting Objects to Import from a Migration Package* on page 313
6. 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.

Specifying the Location and Import Settings for the Migration Package

To import objects you specify the location of the migration package and what you want you want to import.

To import a migration package, you must have at least Edit access to the objects you want to add to the package.

1. Under **Migration Package**, specify the package from which you want to add objects.
   - To browse to a package on your computer or in your network, click **Browse**, select the migration package file (a .pte file), and click **Open**.
   - To get the package from a web address, select **Web Address** and type the URL to the package.

     If you need to enter login credential to access the web address, type the information in the **Username** and **Password** boxes.
2. Click **Load Package**.

The package name, description, and publisher appear, and the objects from this package are displayed on the **Portal Resources** page.

3. Under Import Settings, select options for import.

Option	Description
Import ACLs	Select this to import the Access Control Lists (user and group security settings) for all the objects you are importing.
Overwrite Remote Servers	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.
Remember Dependency Settings	Select this if, on subsequent imports, you want the objects you are now importing to retain the new dependencies that you select in the importing portal. (You select new dependencies on the Unresolved Dependencies page.)
Always Create New Object Instances (Create Duplicates of Existing Objects)	Select this if you want to create new object instances instead of overwriting objects that may already exist on the importing portal.

Selecting Objects to Import from a Migration Package

You can select objects to import on the Portal Resources page of the Migration — Import Utility.

To import a migration package, you must have at least Edit access to the objects you want to add to the package.

This page displays the objects contained in the migration package you loaded. All objects will be imported when you click **Finish**, unless you remove them.

1. If the Migration — Import Utility is not already open, open it now and display the **Portal Resources** page.
2. Remove objects you do not want to import by selecting them and clicking  .

If there are objects to import from Collaboration or Publisher, you can choose those objects on the **Collaboration Resources** or **Publisher Resources** pages.

Resolving Import Dependencies

If any of the objects you are importing depend on resources that are not included in the package, those missing resources are listed on the Unresolved Dependencies page of the Migration — Import Utility. For example, in your migration package you may have a portlet that depends on a web service, which is not included in the package. You can resolve those dependencies by pointing to existing objects in your portal.

To create a migration package, you must have at least Edit access to the objects you want to add to the package.

Note: This page appears only if there is an unresolved dependency.

1. If the Migration — Import Utility is not already open, open it now and display the **Unresolved Dependencies** page.
   Any missing resources are displayed in the **Dependency** column. The associated object in the migration package is listed under each missing dependency.
2. Click  beside the dependency.
3. Select a replacement object from this portal and click **OK**.
   The replacement object is displayed in the **Replacement** column.

Importing Objects Using a Command Line Tool

You can import objects from another portal with 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

Parameter	Description
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 the Portal

You can back up your system without taking it offline.

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.

Restoring the Portal

1. Stop the web service on all machines hosting the portal application.
2. Stop the BEA ALI Automation Service on all computers hosting 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.

Rebuilding the Search Collection

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.

To rebuild the search collection you must have the following rights:

- Access Administration activity right
- Access Search Results Manager activity right

The Search Service Manager lets you specify when and how often the Search Update Agent repairs your search index. Rather than synchronizing particular objects, the repair 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.

1. Configure the Search Service to repair itself.
   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** again.
2. Wait one minute for the setting to update.

3. Run one of the Search Update jobs in verbose mode.
   a) Open the **Intrinsic Operations** folder.
   b) Open one of the Search Update jobs.
      The Job Editor opens.
   c) Change the **Logging Level** to **Verbose** and click **Finish**.

      Note: Make note of the logging mode before you change it, so that you can change it back after the repair is complete.

   d) Select the job you just edited and click **Run Once**.

      By running the job this way, you avoid having to go back into the job and revert to the previous schedule settings.

4. Ensure that the job is running in repair mode.
   a) Open the job you just created; it should be called something like Search Update 1 — Run Once.
      The Job Editor opens.
   b) Click the **Job History** page.
   c) Click the job name.
      The job log opens.
   d) Ensure that the job is running in repair mode.

      The second line of the job log should be similar to this:
      Mar 1, 2008 9:10:02 AM- PTIndexer.ctor : Indexing will extract at most 1000000 encoded bytes of text from each document.
      About half-way down the first page of the log you should see a message that should be similar to this:
      Mar 1, 2008 9:10:02 AM- Search Update Agent is repairing the directories...

5. Reinstall the Search Service and select **Overwrite the existing search index**. For details on installing the Search Service, refer to the *Installation Guide for AquaLogic Interaction*.

Configuring Advanced Properties in the Portal Configuration Files

This appendix describes how to modify the portal configuration files.

AquaLogic Interaction Configuration Manager

AquaLogic Interaction Configuration Manager enables you to manage the configuration settings of AquaLogic User Interaction products through a user interface rather than having to edit .xml files. This section describes the settings in the Portal Service section of the Configuration Manager.

Setting	Description
Analytics Communication	
Enable	Select this option if you want the portal to pass information to Analytics.
Crawler Settings	
Web Crawler file timeout (seconds)	Enter the number of seconds you want a web crawler to try to get to a web page before it times out.
SOAP file timeout (seconds)	Enter the number of seconds you want a remote crawler to try to get to a document before it times out.
Gateway	

Setting	Description
Gateway Enabled	Select this option if you want to allow portal content to be gatewayed.
Gateway temporary directory	Enter the full path to the directory where temporary files will be stored.
Gateway max upload (in bytes)	Enter the maximum file size allowed for uploads.
Gateway min upload (in bytes)	Enter the minimum file size allowed for uploads.
Gateway min streamable (in bytes)	Enter the minimum size of streamable content.
Gateway temporary file pool size	Enter a value that is greater than the number of ASP/JAVA worker threads for the portal.
Logging	
Server	Enter the application name that will uniquely identify log messages sent from this application. ALI Logging Utilities will use this string to determine the location from which log messages originate. The application name can be any string that meets the following restrictions: it must not be empty, it must not exceed 128 in length, and it may only contain non-white-space visible ASCII characters and the space character. Most ALI products follow the naming convention [product-name].[machine-name].[user-name].
Local only	Select this option to limit broadcast of this application's logging messages to only the computer on which this application is installed.
Main Portal Settings	
Web home directory	Enter the directory for the Portal's JAR files.
Image Server base URL	Enter the base URL for the Portal's Image Server.
Image Server secure base URL	Enter the base URL for the Portal Image Server running HTTPS.
Image Server connection URL	Enter the base URL that is used when the Portal Server connects to the Image Server to retrieve JSRegistry information. In many configurations this URL is the same as the Image Server Base URL.

Setting	Description
Image Server connection URL timeout (seconds)	Enter the timeout for the Image Server Connection URL, in seconds. -1 means do not check during startup. Note that if this is invalid, your portal will not work.
Administrative Portal base URL	Enter the base URL to the Administrative Portal. It is required that absolute URL be used in Security Mode 3.
Temporary file home directory	Enter the temporary files directory to be used by the portal. This should not be Web accessible.
Portal Database	
Vendor	Enter the database vendor.
Host	Enter the name of the computer that hosts the database.
User name	Enter the name of the database user.
Password	Enter the password for the database user.
Port	Enter the port number on which the database services requests.
Minimum pooled database connections	Enter the minimum number of pooled connections.
Maximum pooled database connections	Enter the maximum number of pooled connections.
Database Name (Microsoft SQL Server only)	Enter the name of the database.
Portal International Settings	
Locale	Enter the default locale for users. Setting it to the string "UseBrowser" causes the portal to derive the locale from the browser's language settings and store it as the user's locale.
Time Zone (numeric value)	Enter the default time zone for the user. If it is -1 then the time zone of the computer on which Portal is deployed is used.
Mandatory Object Language	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

Setting	Description
	should be a locale string (for example, it should match the name of a folder under the msgs directory.)
Portal System Properties	
Server name	Enter the name of the portal server or virtual load-balanced server (for example, the main load balanced server name: portal.mycompany.com).
Machine name	Enter the machine name (for example, the physical name of the individual portal server machine behind the load-balancer: portalserver1243).
Performance comments	This setting specifies whether the performance comments in the HTML source are enabled. 0 means the comments and stacktraces are disabled. 1 means the comments and stacktraces are enabled. 2 means the comments are enabled but stacktraces are disabled. 3 means the comments are disabled but stacktraces are enabled.
Debugging mode	This setting specifies whether debugging features are enabled. This is disabled by default. You can enable this to debug portal startup issues, especially if you have made customizations to XML files. Make sure to set to disable debugging when you are done troubleshooting. 1 means debugging is enabled. 0 means debugging is disabled.
Doctype specification	1: none, 2: HTML 3.2, 3: HTML 4.0 Transitional, 4: HTML 4.0 Frameset, 5: HTML 4.0 Strict. Note that neither the portal nor any other ALI product has been verified to support either the transitional or strict document type specification. By default the portal does not include a specific doctype declaration in its HTML because doing so would limit the display of portlets in which the doctypes were invalid or inconsistent with the portals declared doctype. If you are confident that all your portlets adhere to a particular doctype you should set that document type here.
Adaptive layout mode	Enter a numeric value indicating whether or not page and/or portlet layout modes are enabled. Both are enabled by default. 0 means layout mode is disabled. 1 means page and portlet

Setting	Description
	layout modes are enabled. 2 means page layout mode is only enabled. 3 means portlet layout mode is only enabled.
Virtual directory path	Enter the virtual directory path. It is typically /portal/.
HTTP entry point	Enter the portal main Servlet mapping name. This has to be the same has the mapping name for HTTPInterpreter in web.xml. It is typically server.pt.
HTTP Port	Enter the port number for the Portal running HTTP.
HTTP Secure Port	Enter the port number for the Portal running HTTPS.
SSO virtual directory path	Enter the SSO Virtual directory path. It is typically /portal/.
SSO servlet name	Enter the SSO Servlet mapping name. This has to be the same has the mapping name for SSOLoginPage in web.xml. It is typically SSOServlet.

Customizing the Tokens in Friendly URLs

By default, the portal areas are represented by the following tokens in friendly URLs: mypage, community, user, directory, document. However, the portal administrator can change these tokens to fit the needs of the company.

1. Open the following file in a text editor:
   Install_Dir`\bea\`*alui*`\settings\portal\FriendlyURLs.xml`.
   For example: `C:\bea\alui\settings\portal\FriendlyURLs.xml` or `/opt/bea/alui/settings/portal/FriendlyURLs.xml`
2. Edit the `<key>` values as desired.
   For example, you might change the token for community to “site” as in the following code:

```
<FriendlyURLMapping>
 <key>site</key>
 <classId>512</classId>
</FriendlyURLMapping>
```

About 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 *Install_Dir*/ptsearchserver/6.5/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

The following parameters appear in ignite.ini by default.

Parameter	Description
RFINDEX	Directory used to store Search Service index files. By default, the installer puts these files in the *Install_Dir*/ptsearchserver/6.5/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

Parameter	Description
	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

Parameter	Description
	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`.

Parameter	Description
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.

Parameter	Description
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

Parameter	Description
	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

Parameter	Description
	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).

Parameter	Description
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 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.

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.

Key Performance Counters

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.

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 336.

Category	Instances	Counters
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

Category	Instances	Counters
OpenHTTPLowLevelNetworkCounter 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
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 files.

In the *Install_Dir*\bin directory, enter the following command:

- On Windows:

```
opencounterslogger.bat ServerName ContextName
-d LogDirectory -l LogOutputStyle -s MaxLogSize -r
PollingInterval -f FilterExpression
```

- On Unix:

```
opencounterslogger.sh 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 configuration.xml file (in *Install_Dir*\settings). 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 configuration.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 (for example, 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. Refer to *Key Performance Counters* on page 332 for 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 336.

  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 category 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.

1. In the *Install_Dir*\bin directory, enter the following command:

   - On Windows:

     `opencountersconsole.bat` *`ServerName ContextName`*

   - On Unix:

     `opencountersconsole.sh` *`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 `configuration.xml` file (in Install_Dir\settings). 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 `configuration.xml` file. This value is case-sensitive.

   For 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.

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]>

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.

     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.

     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 a 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.

1. Click **Start ➤ Run** .
2. Type `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 `configuration.xml` file (in *`Install_Dir`*`\settings`).

The Perfmon adaptor adds a few percentage points of overhead to overall system performance, so you might want to disable it after viewing the counter data. In *`Install_Dir`*`\settings\configuration.xml`, set `opencounters:perfmon-enabled` to `false`.

Localizing Your Portal

This chapter describes how you can internationalize and localize your portal.

Localizing Object Names and Descriptions

You can localize object names and descriptions, so that users see the names and descriptions in their chosen language. For example, if you have an object called "Engineering," you could add "Ingénierie" as the localized entry for French. A user viewing the French user interface would see the Ingénierie as the object name, as well as any other names and descriptions localized into French.

There are two ways to localize object names and descriptions: in the object's editor or using the Localization Manager.

- To localize the name and description for a single object, use the object's editor.
  See *Localizing the Name and Description for an Object* on page 342.
- To localize all object names and descriptions, use the Localization Manager.
  See *Localizing All Object Names and Descriptions* on page 342.

Localizing the Name and Description for an Object

You can localize object names and descriptions, so that users see the names and descriptions in their chosen language. For example, if you have an object called "Engineering," you could add "Ingénierie" as the localized entry for French. A user viewing the French user interface would see the Ingénierie as the object name, as well as any other names and descriptions localized into French.

Note:

- You can localize names and descriptions into only the languages supported by the portal.
- You cannot localize names and descriptions for users.
- You can localize the names and descriptions for all objects at the same time using the Localization Manager.

1. Open the object's editor by creating a new object or editing an existing object.
2. Select **Supports Localized Names**.
   The **Localized Names and Descriptions** section appears.
3. Add or edit the localized names and descriptions:
   - To add an entry for a language, click  **New Localized Name**, then, in the Name and Description dialog box, enter the localized name and/or description, select the appropriate language, and click **Finish**.
   - To edit an existing entry, click the entry you want to change, then, in the Name and Description dialog box, edit the entry as necessary, and click **Finish**.
   - To remove existing entries, select the entries you want to remove and click  .

     To select or clear all entries, select or clear the check box to the left of **Name**.

Localizing All Object Names and Descriptions

You can localize object names and descriptions, so that users see the names and descriptions in their chosen language. For example, if you have an object called "Engineering," you could add "Ingénierie" as the localized entry for French. A user viewing the French user interface would see the Ingénierie as the object name, as well as any other names and descriptions localized into French.

To access the Localization Manager you must be a member of the Administrators Group.

Although you can supply localized names and descriptions 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. You can then edit the .xml file and upload it back into the portal.

Note:

- You must configure each object to allow localization before downloading the .xml file.
- You can localize names and descriptions into only the languages supported by the portal.
- You cannot localize names and descriptions for users.
- You can localize each object individually using the object's editor.

1. Click **Administration**.
2. In the Select Utility drop-down list, click **Localization Manager**.
3. Export the names and descriptions for all portal objects that support localization by clicking **Download**, and saving the XML file to your computer.
4. Open the file with a text editor, add or edit localized values for the objects, and save your changes.
   For example, to edit the French term for Everyone, you might make the following change:
   ```
   <target language="fr">Tous</target>
   ```
5. Return to the Localization Manager, and click **Browse**.
6. Navigate to the exported file and click **Open**.
7. Upload and apply your changes to the portal by clicking **Upload**.

Localization Manager XML

Each localized object is represented by an entry in the Localization Manager XML file.

```
<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.

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

About Search Service Internationalization Support

The portal provides support for 61 languages. The portal uses Unicode characters 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 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.

Some languages supported by the portal include support for word stemming and compound decomposition. This additional information is used to enhance results of the full-text index. For a list of supported languages, including which have enhanced support, see *Search Service Language Support* on page 363.

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.

APPENDIX D

Deploying Single Sign-On

This appendix describes how to deploy Single Sign-On (SSO) capabilities in the portal environment.

Common SSO Questions

This section contains links to common SSO questions and the answers.

Why Doesn't SSO Work for a Particular User?

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. For more information, see *Configuring ALI Logging Utilities* in the *ALUI Development Documentation*.

Why Isn't the SSO Cookie Forwarded to Remote Servers or Portlets?

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`.

Does the Portal with SSO Support Guest User Sessions?

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 **Log In** 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 **Log In**. 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.

How Can I Change Login Credentials From an SSO Session?

If you need to log in as Administrator or other portal user from within an SSO session, you can perform the following steps:

1. Click **Log Off** in the portal banner.
   This logs you out of the portal and takes you to the portal login page, as if SSO were disabled.
2. From this page you can log in as a non-SSO user or you can browse the portal as guest.
3. When you want to log back in as an SSO user, click **Log In** in the portal banner.
   You are automatically logged in to the portal in an SSO session.

Why Can't I Access the Portal Through SSOLogin.aspx or the SSOServlet?

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.`

Why Do Users Get JavaScript Errors and Portal Menus Fail to Load if I Configure the SSO Authentication Server to Protect the Image Service Virtual Directory?

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.

How Can I Debug My SSO Deployment?

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 ALI Logging Utilities* in the *ALUI Development Documentation*.

How Do I Configure Reverse Proxy with My SSO Deployment Using Oblix Netpoint Access Server (versions 6.1.1 or 6.5) with an Apache WebGate?

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 Apache 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.

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.
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.

How Do I Configure Reverse Proxy with My SSO Deployment Using Apache HTTP Server?

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`:

   ```
   ProxyRequests Off
   ProxyPass /portal
   http://your_portal_server.domain.com:7001/portal
   ProxyPassReverse /portal
   http://your_portal_server.domain.com:7001/portal
   ```

   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). You must specify the fully qualified domain name here and for all other times you type in the server names. For more information on Apache reverse proxy, see *http://httpd.apache.org/docs/mod/mod_proxy.html*.

4. Start or reboot the Apache HTTP server.

How Do I Configure Reverse Proxy with My SSO Deployment Using a Java Application Server?

1. Open *Install_Dir*/ptportal/6.5/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="*"/>
   ```

3. Replace proxy_server.domain.com with the fully qualified domain name for the Apache HTTP server.
4. Configure the <SSOVirtualDirectoryPath> element so that it is similar to the following example:

   ```
   <SSOVirtualDirectoryPath
   value="http://proxy_server.domain.com/portal/"/>
   ```

5. Replace proxy_server.domain.com with the fully qualified domain name for the Apache HTTP server.
6. Reboot the application server.

Default Behavior of Search Service

This appendix describes the default behavior of the portal searches.

About the Different 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.

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

Portal Area	Search Type	Description
		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.
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.

Elements of Search Syntax

There are several syntax elements that work together in search.

About Operator Modes

The Search Service parses queries to determine which 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.

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.

Search String Operators

Operator	Description	Example Search Text	Example Search Results
<AND> Alternative: AND, & (ampersand)	Connects two terms that must both be included in each item returned.	holiday <AND> schedule	Holiday Schedule
<OR> Alternative: OR, ACCRUE, ANY, | (vertical bar), , (comma)	Connects two terms where at least one must be included in each item returned.	holiday <OR> vacation	Holiday Schedule, Christmas Holiday Party, Scheduling Vacation
<NOT> Alternative: NOT, AND NOT	Term must not appear in items returned.	holiday <NOT> vacation	Holiday Schedule, Christmas Holiday Party
<NEAR/N> Alternative: NEAR	Terms must appear within N words of each other, regardless of order, in items returned.	early <NEAR/10> retirement	Plan early for your retirement
<ORDER>	Both terms must appear in items returned, and the first term must precede the second term.	song <ORDER> bird	song bird (not bird song)

Operator	Description	Example Search Text	Example Search Results
<WORD>	Turns off stemming, alternate case, and spell correction.		
<PHRASE> Alternative: Surround terms in “ (double quotes)	Both terms must appear sequentially, in a phrase in items returned.		
<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.		
* (asterisk)	The wildcard specifies that the result must match 0 or more characters at the beginning or end of a word.	sub*	subdirectory, subject, subjective
> (right angle bracket)	The top best bet operator brings the user directly to the top best bet for a term, such as a community, document, or portlet.	>HR	You are navigated to the HR Community.

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:

Query	Equivalent to...
The young and the restless	“the young” <AND> “the restless”
File not found	file <AND> <NOT> found
Error -217439239	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:

1. NEAR, ORDER, PHRASE, SENTENCE, PARAGRAPH
2. NOT
3. AND
4. 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 in searches.

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.

Note:

- 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.

Case Sensitivity

All searches are case-insensitive, except when the <WORD> operator is used.

Table 3: Case Sensitivity Examples

Query	Matches
`BEA`	Items containing BEA, bea, or any other case variant.

Query	Matches
`"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.)

Note:

- Terms generated by wildcard expansion are not stemmed.
- Stemming is not applied to terms within a quoted phrase.

Wildcards

The wildcard operator (*) is used to search for partial matches (prefixes, suffixes, and substrings) of indexed terms.

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.

Note:

- Terms generated by wildcard expansion are not stemmed.
- Wildcards cannot be used within quoted phrases.

Table 4: 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).

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.

Note:

- Stemming is not applied to terms within a quoted phrase.
- Wildcards cannot be used within quoted phrases.

Thesaurus Expansion

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.

Thesaurus expansion has the following characteristics:

- It is applied to each term in a basic search.
- It is applied in all three search modes (Internet Style, Query Operators, and Bag of Words).
- It 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, thesaurus expansion is always applied to all individual search terms.

How Language Settings Apply to Search

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.

Search Service Language Support

The portal provides support for 61 languages.

Of the languages supported by the portal, 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)	German	Norwegian (Bokmal)
Chinese (Traditional)	Greek	Polish
Czech	Hungarian	Portuguese
Danish	Italian	Russian
Dutch	Japanese	Spanish
English	Korean	Swedish
Finnish	Norwegian (Bokmal)	Turkish
French		

The following languages are supported at a reduced level.

Afrikaans	Gallegan	Marathi
Albanian	Hebrew	Persian
Arabic	Hindi	Romanian
Basque	Icelandic	Serbian
Belarusian	Indonesian	Serbian-Croatian
Bengali	Irish	Slovak
Bulgarian	Kalaallisut	Slovenian
Catalan	Konkani	Swahili
Cornish	Latvian	Tamil

Croatian	Lithuanian	Telugu
Esperanto	Macedonian	Thai
Estonian	Maltese	Ukranian
Faeroese	Manx	Vietnamese

Using Text Search Rules

When you search for text, you generally can just type the text you are looking for (the search string). However, there are a few rules you should be aware of:

Note: Search strings are case-insensitive; that is, uppercase A is the same as lowercase a.

- To find objects or documents containing all terms in your search string, separate your terms with spaces.

  This is the same as using AND.
- To find objects or documents containing one or more of the terms in your search string, separate your terms with commas.

  This is the same as using OR.
- To search for an exact phrase, type quotation marks (") around the phrase.
- To specify that a term must be included in each result, type a plus (+) in front of the term.
- To exclude a term from the results, type a minus (-) in front of the term.

  Note: Do not include a space after the plus or minus.

  Note: Do not use the plus or minus in the same search with other search string operators.

Search 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.

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.
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.

Query	Expected Behavior
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.
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.

How Search Results Are Ranked

Search results are ranked according to relevance, by default. There are several factors that determine relevance.

How Term Frequency Factors in Relevance

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.

About 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.

How Phrases and Proximity Factor in Relevance

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:

1. All words joined together as a single phrase
2. Stem variants of all words and all quoted phrases <ORDER><NEAR> each other
3. 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.

About Basic Search Behavior

Basic search adds some special features in order to increase the chances that a search will return relevant results.

Basic search has several characteristics:

- 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 and would therefore retrieve documents containing dog and and.
- Term proximity can boost the relevancy ranking in basic search.
- Automatic spelling correction is applied only in basic search.

About 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.