BEA AquaLogic™ User Interaction

Deployment Guide

Document Version: G6 3.0
Document Revised: November 17 2006

Copyright

Restricted Rights Legend

Trademarks and Service Marks

Contents

4. Defining Administrative Roles

5. Customizing the User Interface

6. Provisioning Host Computers

7. Implementing Network Security

F. Component-Host Worksheets

CHAPTER 1

Welcome

This book describes how to plan an AquaLogic User Interaction deployment.

How to Use This Book

Audience

This guide is intended for IT and management responsible for planning and executing the deployment of the AquaLogic User Interaction solution.

Organization

This guide includes the following chapters:

- Chapter 2, "Components of AquaLogic User Interaction." This chapter summarizes the functional components and features of the AquaLogic User Interaction solution.
- Chapter 3, "Planning Portal Structure and Content." This chapter describes portal structure and content at a high level.
- Chapter 4, "Defining Administrative Roles." This chapter describes administrative roles at a high level.
- Chapter 5, "Customizing the User Interface." This chapter summarizes AquaLogic User Interaction user interface customization techniques.

- Chapter 6, "Provisioning Host Computers." This chapter summarizes host requirements for deployment components.
- Chapter 7, "Implementing Network Security." This chapter describes security options you can implement for your deployment.
- Chapter 8, "Using Migration Features to Stage Your Deployment." This chapter summarizes migration capabilities for AquaLogic User Interaction deployments.
- Chapter 9, "Localizing Your Deployment." This chapter describes AquaLogic User Interaction localization features.
- Chapter 10, "Developing a Production Maintenance Plan." This chapter provides an overview of portal maintenance tasks and tools.

Typographical Conventions

This book uses the following typographical conventions.

Table 1-1 Typographical Conventions

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

BEA Documentation and Resources

This section describes other documentation and resources provided by BEA.

Table 1-2 BEA Documentation and Resources

Resource	Description
Installation and Upgrade Guide	This guide describes the prerequisites (such as required software) and procedures for installing or upgrading ProductName. It is available on edocs.bea.com and on the application CD.
Installation Guide	This guide describes the prerequisites (such as required software) and procedures for installing ProductName. It is available on edocs.bea.com and on the application CD.
Installation Worksheet	This worksheet allows you to record prerequisite information necessary for installing ProductName. It is available on edocs.bea.com and on the application CD.
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 and on the application CD.
Administrator Guide	This guide describes how to manage and maintain ProductNameShort. It is available on edocs.bea.com and on the application CD.
User Guide	This guide describes how to create and configure portlets with the ProductNameShort. It is available on edocs.bea.com and on the application CD.
Online Help	The online help is written for all levels of ProductNameShort users. It describes the user interface for ProductNameShort and gives detailed instructions for completing tasks in ProductNameShort. 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.

Table 1-2 BEA Documentation and 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.
dev2dev.bea.com	Download developer tools and documentation, get help with your development project, and interact with other developers via BEA's dev2dev newsgroups.
AquaLogic User Interaction Support Center	The AquaLogic User Interaction Support Center is a comprehensive repository for technical information on AquaLogic User Interaction products. From the Support Center, you can access products and documentation, search knowledge base articles, read the latest news and information, participate in a support community, get training, and find tools to meet most of your AquaLogic User Interaction-related needs. The Support Center encompasses the following communities: **Technical Support Center** Submit and track support incidents and feature requests, search the knowledge base, access documentation, and download service packs and hotfixes. **User Group** Visit the User Group section to collaborate with peers and view upcoming meetings. **Product Center** Download products, read release notes, access recent product documentation, and view interoperability information. **Developer Center** Download developer tools and documentation, get help with your development project, and interact with other developers via BEA's dev2dev newsgroups. **Education Services** Find information about available training courses, purchase training credits, and register for upcoming classes. If you do not see the Support Center when you log in to http://support.plumtree.com, contact ALUIsupport@bea.com for the appropriate access privileges.

Table 1-2 BEA Documentation and Resources

Resource	Description
Technical Support	If you cannot resolve an issue using the above resources, BEA Technical Support is happy to assist. Our staff is available 24 hours a day, 7 days a week to handle all your technical support needs. E-mail: ALUIsupport@bea.com Phone Numbers: U.S.A. +1 866.262.PLUM (7586) or +1 415.263.1696 Europe +44 1494 559127 France +33 1.46.91.86.79 Australia/NZ +61 2.9923.4030 Asia Pacific +61 2.9931.7822 Singapore +1 800.1811.202

CHAPTER 2

Components of AquaLogic User Interaction

This chapter summarizes the functional components and features of the AquaLogic User Interaction (ALUI) solution.

For a conceptual overview of ALUI, contact a sales representative or visit the Product Center for the latest white papers.

For detailed product specifications and compatibility requirements, contact a sales representative or visit the Product Center for the latest product data sheets.

The purpose of this chapter is to help you identify the components you plan to include in your ALUI deployment so that you can provision host computers and IT resources and plan administrative- and development-related assignments for your deployment.

At the end of this chapter, you should be able to complete Appendix A, "AquaLogic Interaction Products Worksheet."

The following table provides a brief summary of the functionality provided by the components of the ALUI solution.

	Application Suite	AquaLogic Interaction	Sold Separately	Description
User Interface	X	X		AquaLogic Interaction (ALI) portal and Image Service components provide the basic building blocks of the end-user experience, including UI templates for ruled-based Experience definitions, personalized pages, communities, and the Knowledge Directory. The Administrative Portal provides a centralized management UI for all deployment components.
			X	AquaLogic Interaction Development Kit (IDK) includes the UI Customization Kit and the Pluggable Navigation Kit to enable you to rapidly deploy customizations to the out-of-the-box UI look, feel, and functionality.
Web Services	X	X		AquaLogic Interaction is an application that acts as the gateway between user requests and composite applications.
	X	X		ALI API Service enables rapid integration with AquaLogic User Interaction applications, as well as remote applications.
User Management	X	X		AquaLogic Interaction enables rule-based Experience definitions, allowing you to manage content access and activity rights for users and groups.
			X	Identity Service - Active Directory enables rapid integration with Active Directory authentication services. Identity Service - LDAP enables rapid integration with LDAP authentication services.

	Application Suite	AquaLogic Interaction	Sold Separately	Description
Content Management	X	X		Document Repository Service enables content from file system locations, URLs, Collaboration projects, or Publisher targets to be made available through the portal.
	X	X		Content Upload Service enables portal users who might not have access to data staging locations (for example, extranet users) to upload files to the Knowledge Directory.
	X		X	Publisher provides a UI and templates for creating content and its metadata and managing this content with workflow, scheduled publishing, and expiration controls.
			X	Content Service - Windows Files enables document discovery for remote data sources on Windows file system. Content Service - Documentum enables document discovery for remote data sources on Documentum Docbases. Content Service - Lotus Notes enables document discovery for remote data sources on Lotus Domino Servers. Content Service - Microsoft Exchange enables document discovery for remote data sources on Microsoft Exchange Servers.
Security	X	X		AquaLogic User Interaction architecture allows most Web applications to be hosted behind a firewall, with only ALI directly serving user requests. ALI supports SSO. ALI provides object-level security so you can manage user access to content with ACL and user activity by designating activity rights.

	Application Suite	AquaLogic Interaction	Sold Separately	Description
Collaboration	X		X	Collaboration provides portlets to facilitate collaborative workspaces, including document source control, threaded discussions, and announcements.
			X	Process is a set of business process modeling (BPM) tools that provides the ability to design, activate, and deploy business processes into a live environment. The platform lets users quickly combine dissimilar applications into integrated business processes. Users can then modify these processes in real time, enabling business to react dynamically to changing market conditions.
Search	X	X		Search indexes all of the resources accessible through the portal pages, communities and Web applications deployed across the enterprise. These resources include: Content indexed from file systems, Web sites, and document databases Project documents and Web pages created and stored by Collaboration and Publisher Applications, portlets, communities, and users
Scheduled Operations	X	X		Automation Service enables the ALUI administrator to schedule and run administrative and maintenance jobs, such as user and group synchronization or data source crawls.
Capacity Planning	X	X		ALUI integrates with performance monitors common to operating systems, Web application servers, and back-end data source servers.
			X	Analytics provides portlets and portlet templates for tracking user activity and content usage.

	Application Suite	AquaLogic Interaction	Sold Separately	Description
Business Applications	X	X		ALI provides intrinsic portlets.
	X		X	Studio enables rapid development of portlets, such as telephone lists, work order processes, calendars and surveys, without any coding. These plug-in applications feature a user interface, application logic and a database, and operate using Web services technologies.
			X	Portlet Suite - Exchange provides portlets to support user access to mail, calendar, and address books stored on a remote Exchange Server. Portlet Suite - IMAP provides portlets to support user access to mail, calendar, and address books stored on a remote IMAP Server. Portlet Suite - Lotus Notes provides portlets to support user access to mail, calendar, and address books stored on a remote Lotus Notes Domino Server. Portlet Suite - Documentum provides portlets to support user access to documents stored on a remote Docbase. Portlet Framework - Microsoft Excel provides a framework for developing portlets that enable collaboration on Excel projects.
			X	Integration Services - PeopleSoft enables rapid development of portlets to support user access to PeopleSoft data sources. Integration Services - SAP enables rapid development of portlets to support user access to SAP data sources. Integration Services - Siebel enables rapid development of portlets to support user access to Siebel data sources.
			X	AquaLogic Interaction Development Kit (IDK) is a set of APIs, documentation and sample code that work in both Java and .NET-based development environments, allowing you to develop with the IDE of your choice.

CHAPTER 3

Planning Portal Structure and Content

This chapter describes portal structure and content at a high level.

The purpose of this chapter is to help you develop a plan to assign administrative responsibility for managing portal structure and content.

At the end of this chapter, you should be able to scope the effort involved in deploying the initial administrative objects and content. You should also be ready to assign administrative responsibility for the initial deployment.

Before you proceed, print Appendix B, "Portal Content Responsibilities Worksheet." You can use the worksheet to record your decisions and assignments.

This chapter includes the following topics to help you scope and assign tasks for your deployment effort

.

Task	Topic
Decide whether you want to implement a single portal with multiple experience definitions or to create a federation of multiple portals.	"Single Portals and Federated Portals" on page 3-2
Determine different types of audiences for which you want to support custom views so that you can implement corresponding, appropriate Experience Definitions.	"Experience Definitions" on page 3-4

Task	Topic
Determine the content types and content sources you want to make accessible through the Knowledge Directory.	"Knowledge Directory" on page 3-5
Determine the types of communities you want to include in your initial deployment.	"Communities" on page 3-7
Determine the types of portlets you want to include in your initial deployment.	"Portlets" on page 3-11
Decide how you plan to incorporate collaboration features into your deployment.	"Collaboration" on page 3-13
Determine your preferences for customizing Search behavior.	"Search" on page 3-16

The Administrator Guide provides detailed procedures for implementing these objects.

Single Portals and Federated Portals

First, decide whether you plan to deploy one portal with multiple Experience Definitions or multiple, federated portals. Figure 3-1 and Figure 3-2 illustrate the difference in deployments of a single portal, in which a single instance of ALI manages various service requests from users; and a federated portal, in which multiple ALI instances manage various service requests from separate portals.

Figure 3-1 Single Portal with Multiple Experience Definitions

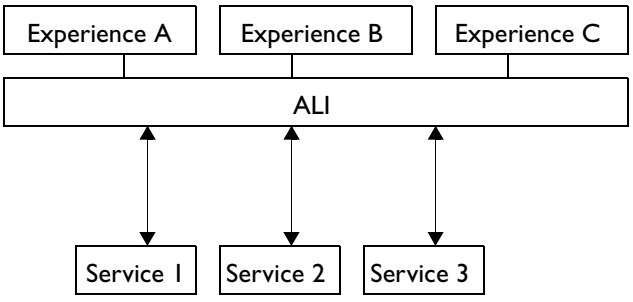

Figure 3-2 Federated Portal

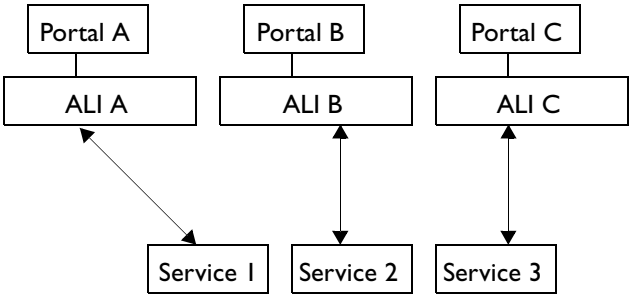

In the past, an enterprise might have used a federated model for any of the following reasons:

- Different technology: many applications offer portal interfaces
- Different sponsors: many business units want autonomy because they are comfortable with different platforms or devoted to separate brands
- Different audiences: audiences with deep functional needs or with different security—consumers, employees, partners
- Different project schedules: business unit frustration

Consider the many benefits of having only one portal:

- Users have different experiences but are managed in one place
- IT can more easily scale the portal
- It is easy to distribute enterprise-wide communications
- It is easy to instigate enterprise-wide business processes
- Common content management system

If you are starting with multiple portals, there are several ways to take advantage of what you have already done:

- Leverage the ALI Web services architecture and common portal standards (JSR-168, WSRP) to integrate services from other portals. The ALUI can integrate applications written in any modern programming language.

- Leverage your existing authentication service or single sign-on to synchronize users and user privileges across multiple portals or sites.
- Leverage Federated Search services to connect to search engines for multiple content portals.
- Leverage customizable navigation features to link to other portals.

For more information on the business cases in which single portals or federated portals are appropriate, visit the Product Center.

In Appendix B, "Portal Content Responsibilities Worksheet," record your decision.

Experience Definitions

Next, determine how many different Experience Definitions you plan to deploy.

Experience Definitions allow you to display different branding and features to different audiences.

Experience Definitions control the overall look, feel, and access features for a group of users—the header and footer; the navigation scheme; the default page displayed at login; access to My Pages, communities, and the Knowledge Directory; and any mandatory links displayed in the navigation.

For example, you might create an Experience Definition for a particular customer that includes the customer's logo and company colors and that includes access to My Communities and the Knowledge Directory, but not to My Pages or Administration.

Experience Definitions are applied according to rules you configure with the Experience Rules Manager. For example, you can create rules that sort users from a particular customer IP address into an Experience Definition you have configured for such customers.

For information on configuring Experience Definitions and Experience Rules, see the Administrator Guide.

Figure 3-3 illustrates the correspondence between various audiences and Experience Definitions. For example, if you have three divisions in your company that should have somewhat different access, you might create Experience Definitions A, B, and C. Similarly, if you have four customers, you might want them to view customized support sites F, G, H, and I. You might create definition J for guests from trusted domains and definition K for guests from non-trusted domains.

Figure 3-3 Supporting Various Audiences with Experience Definitions

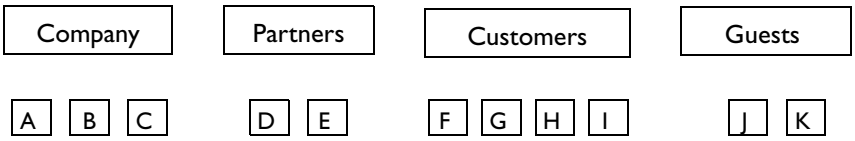

In Appendix B, "Portal Content Responsibilities Worksheet," identify the Experience Definitions you want to implement and assign a leader to be responsible for their implementation.

BEA provides the following resources to prepare the leader for this assignment.

Resource	Description
Consulting Services	Get advice from the experts.
Education Services	Take classes from the experts.
Knowledge Base	1. Register and log into the Support Center. 2. Search the Knowledge Base for "Experience Definition" topics.
Administrator Guide and Online Help	Learn how to implement the Experience Definition and Experience Rule portal objects.

Knowledge Directory

Next, determine the types of content you plan to make available through the portal Knowledge Directory.

Content might be located in multiple systems: file repositories, ERP systems, databases and so on. ALUI pulls content into the portal through various mechanisms, including Content Services, searches, portlets, Publisher, and Collaboration. In the portal, content is discoverable through the Knowledge Directory and through Search indexes.

All common file types can be full-text indexed, but ALI has extra capabilities to extract metadata from the following types of file repositories.

Repository	Deployment Considerations
Windows Files	The most common source of information is files located on a Microsoft Windows network. A Content Service (CS) attempts to import content from each document it finds. The CS starts crawling at one file folder of any UNC-compatible directory (\\<*computer_name*>\<*folder_name*>\<*subfolder_name*>) and continue crawling into subfolders. Things to consider about Windows network system files include: • Is there information in a secure folder that will require authentication to access? What domain and user name can be used to gain access? • Do you want to import Windows security information with the document or rely on ALI security?
Web	Content Services start crawling at a single Web page and continue to follow links to connected pages. You can control the content that enters your Knowledge Directory by using filters and special CS options. The important questions to ask about Web information are: • Is there a proxy server? What is the configuration of the proxy server? • What is the target site security?
Lotus Notes	If you use Lotus Notes, you can pull application records and documents directly into the Knowledge Directory. Considerations include: • Which databases do you want to start with? • Do you want to import security?
Microsoft Exchange	E-mail is everywhere and is a critical source of information that most people need to do their jobs. However, in most organizations, e-mail information stays with the user. With ALI, however, you can crawl information from Microsoft Exchange or Lotus Notes and have it be searchable from the portal. Considerations to think about include: • Are there public folders set up on the mail system? • What type of access will be required?

Repository	Deployment Considerations
Documentum	If you store documents in Documentum, then you can pull them directory into the Knowledge Directory. Considerations to think about include: • Which docbases do you want to start with? • Do you want to import security?
Custom	If you have documents or records in other systems, you can use the IDK to develop a Java or .NET Content Service to pull the documents from these systems into the Knowledge Directory. For information on developing custom applications with the IDK, visit the http://dev2dev.bea.com.

In Appendix B, "Portal Content Responsibilities Worksheet," identify the content types you want to support and assign a leader to be responsible for implementation of each.

BEA provides the following resources to prepare the leader for this assignment.

Resource	Description
Consulting Services	Get advice from the experts.
Education Services	Take classes from the experts. To learn how to integrate content, enroll in the series of classes developed for Content Managers.
Knowledge Base	Register and log into the Support Center. Search the Knowledge Base for content-related topics, such as "Knowledge Directory", "Content Services", "Exchange", "Windows Files", and the like.
Administrator Guide and Online Help	Learn how to implement a Knowledge Directory taxonomy, content types, content filters, and Content Services objects.

Communities

Next, consider the numbers and types of Communities you plan to include initially in your deployment.

Communities are pages shared between the members of a group to collaborate and communicate on a particular project or on departmental goals.

Communities provide the following benefits to users:

- Consistent user experience and navigation

- Systematic knowledge capture and sharing
- Location of and interaction with experts
- Shortened learning curve and time-to-contribution

Here are a few ideas for how to use communities:

- Business Unit Resource Center (Line of Business Communities)

Audience	• Business unit or department • Customers of that business unit or department
What to put in it	• Community documents, links, calendar • Metrics • Expert finder • Q&A
Success indicators	• Strong departmental or group identity • Existing intranet as content source • Motivated community owner
Pitfalls to avoid	• Static page that people visit and forget
Suggested AquaLogic Interaction tools	• Portlet Framework - Microsoft Excel • Publisher

- Interactive Workspace (Collaborative Communities)

Audience	• Ad hoc or established project workgroups
What to put in it	• Project task list • Document management and archive • Project calendar • Threaded discussions • Project metrics
Success indicators	• Members spread out • Project has specific objectives and milestones • Project has outgrown e-mail and file-shares

Pitfalls to avoid	• Dustbin of history: old projects, communities that do not go away • Ghost town: two or three people are probably not enough
Suggested AquaLogic Interaction tools	• Collaboration • Studio • Portlet Framework - Microsoft Excel

- Customer or Partner Management Site (Sales and Service-Oriented Communities)

Audience	• Customers or partners
What to put in it	• Key customer or partner resources: documents, calendar • Self-service access to CRM or PRM system • Feedback mechanism • Customer-to-customer or partner-to-partner: facilitate community
Success indicators	• Portal-only access for critical information • Responsiveness to customer/partner feedback
Pitfalls to avoid	• No human input
Suggested AquaLogic Interaction tools	• Collaboration • Studio • Portlet Framework - Microsoft Excel

- Dashboards (Analytic Communities)

Audience	• Management
What to put in it	• Performance metrics • Financial documents
Success indicators	• Support to enforce consistent data formatting • Portal-only access for critical information • Culture of accountability based on metrics

Pitfalls to avoid	• Make sure the dashboards have fresh data • Make sure security works appropriately
Suggested AquaLogic Interaction tools	• Portlet Framework - Microsoft Excel • Integration Services for SAP and PeopleSoft

- Business Process Applications (Process Communities)

Audience	• Users involved in process
What to put in it	• Published content • Data from multiple systems • Workflow • Metrics
Success indicators	• Simple navigation, consistent branding a priority • Unified search criteria • Looking to utilize reusable components, common foundation
Pitfalls to avoid	• If you do not have a process, the software will not do it for you
Suggested AquaLogic Interaction tools	• Publisher • Collaboration • Process • Search • Studio • Portlet Framework - Microsoft Excel

On Appendix B, "Portal Content Responsibilities Worksheet," identify the types of communities you want to include initially in your deployment and assign a leader to be responsible for implementation of each.

BEA provides the following resources to prepare the leader for this assignment.

Resource	Description
Consulting Services	Get advice from the experts.
Education Services	Take classes from the experts. To learn how to develop community content, enroll in the series of classes developed for Community Managers.
Knowledge Base	1. Register and log into the Support Center. 2. Search the Knowledge Base for community-related topics, such as "Community", "Subcommunities", and the like.
Administrator Guide	Learn how to implement and manage Community objects and related portal objects, such as Community Templates and Subcommunities.

Portlets

Next, consider the numbers and types of portlets, as well as tools you might use to develop additional portlets.

Portlets are applications that are embedded in a portal and can be interactive or merely informational. They are able to communicate preferences with the portal and communicate with other portlets.

There are several components involved in a portlet. A portlet must be based on a web service. The web service controls the bulk of the portlet settings, for example, the URL and cache settings. The portlet controls the name, width, type, and administrative preferences (if available). You can also create portlet templates allowing you to create multiple instances of the same portlet, with each instance looking different or displaying different information based on the administrative preferences set in that instance.

BEA provides pre-packaged portlets to support integration with commonly requested applications or services. For a complete list of pre-packaged portlets, see Chapter 2, "Components of AquaLogic User Interaction." For a description of these products, visit the Product Center.

You can also create your own portlets using Studio or Publisher. For a description of these products, visit the Product Center.

You can also develop portlets using the IDK. For information, visit the ALUI Developer Center http://dev2dev.bea.com/aluserinteraction/.

There are many reasons you might want to build your own portlets:

- You can simplify the user experience by avoiding separate logins for each service, avoiding the complexity of exposing users to a complete application, or you can customize and condense application experiences.
- You can Web-enable proprietary systems to avoid Web site sprawl, avoid security and user interface costs, and provide Internet access through the ALI gateway.
- You can draw users to a broad experience like quality and safety checks or benefits enrollment.

Consider the following questions when deciding which portlets to build first:

- What do people want?
- What is a difficult or time-consuming to do?
- What do people use regularly?
- What is important to people's jobs?
- What drives the business?
- Which systems are proprietary?

Here are some ideas for portlets:

• E-mail	• Invoice lookup	• WYSIWYG publishing
• Pager	• Time and expense	• Time zone calculators
• Calendar	• Work order status	• Process-based portlets
• ERP access	• Employee directory	• Project artifact dashboard
• Yellow pages	• Executive dashboard	• Business initiatives calendar
• Sales analysis	• Simplified navigation	• Conference room scheduling
• Sales support		

In Appendix B, "Portal Content Responsibilities Worksheet," identify the types of portlets you want to initially include in your deployment and assign a leader to be responsible for implementation of each.

BEA provides the following resources to prepare the leader for this assignment.

Resource	Description
Consulting Services	Get advice from the experts.
Education Services	Take classes from the experts. To learn how to develop portlets, enroll in the series of classes developed for Content Managers or Portal Developers.
Knowledge Base	1. Register and log into the Support Center. 2. Search the Knowledge Base for portlet-related topics, such as “portlet”, “+mail +portlet”, and the like.
Administrator Guide and Online Help	Learn how to implement and manage portlet objects and related portal objects, such as Portlet Templates and Web Services.
Development Documentation	You can also develop portlets using the IDK. For information, visit http://dev2dev.bea.com/aluserinteraction/.

Collaboration

Next, consider requirements for incorporating collaboration features in your site.

This section describes the following topics for Content Managers to consider:

- “Incorporating AquaLogic Interaction Collaboration” on page 3-13
- “Incorporating AquaLogic Interaction Process” on page 3-15

Incorporating AquaLogic Interaction Collaboration

This section describes how you can deploy Collaboration to support applications with a focus on producing documents, transferring knowledge to a large external audience, or simply organizing a small team around a common goal. It includes the following examples:

- “Example: Sales Process Automation” on page 3-14
- “Example: Customer Support” on page 3-14
- “Example: Finance Application” on page 3-15
- “Example: E-learning Application” on page 3-15

Example: Sales Process Automation

Suppose each account manager at a given organization manages several prospective clients. A sales process automation application powered by AquaLogic User Interaction would enable account managers to create Web properties for each prospective client in their portfolio. Each of these Web applications would have a Collaboration project as part of the application template. The project itself would be based on a Collaboration project template that was region specific. For example, there might be two project templates, one for United States sales and one for International. The template would then create a project with all of the appropriate legal contracts, product schedules, marketing documentation, and support contracts automatically preloaded.

Account managers would simply have to enter a new prospect's account name and e-mail address into the portal. The Web property would be automatically created and Collaboration would send a notification directly to the client. The account managers would also subscribe to the Collaboration project as a whole so that they would be notified when clients begin to interact. This online property would then be available to the prospective client from anywhere in the world. Clients could share documents, upload an RFP, assign various tasks to their account managers, post discussion questions on topics of concern, or schedule meeting events through the Collaboration project available in this application.

Not only would this application increase contact and communication between sales representatives and prospective clients, but it would serve to minimize unnecessary work-load on the sales representatives. Managing prospective clients becomes much easier with Collaboration notifications. Each time a client poses a question or uploads a new document, the sales representatives are instantly notified. In addition, because Collaboration is fully integrated into the portal, sales representatives can easily invite other employees with valid expertise into the Collaboration project.

Example: Customer Support

Most organizations today have some sort of support Web site. Collaboration can be used in this support context to improve response rate, decrease incident resolution time, and give customers an inside look into how their issues are being solved.

Assume that a customer can log in to an external support Web site and file incident reports for problems they are experiencing with a product. For each new incident that is filed, a Collaboration project is created and the relevant support engineer is notified. The engineer and the customer can then enter a discussion through the Discussions portion of the Collaboration project. Collaboration will keep track of all the messages that go back and forth and in the end will provide a clear record of the incident resolution. If this particular incident becomes common,

the support engineer might forward the discussion information to the entire group so that everyone is immediately aware of the resolution.

The Collaboration project thus becomes the hub for every incident that Customer Support deals with. Each incident can be populated with the appropriate tasks and documents that are necessary to solve the customer's issue. The customer then has a real-time outlook into exactly what is being done to resolve the issue.

Example: Finance Application

At the end of each quarter, the accounting departments of most organizations run through a number of tasks and produce several documents detailing the organization's financial performance. Many customers have created a Finance Dashboard that incorporates Collaboration to automate the process of closing their books.

Normally the heart of these applications is a Collaboration project template which outlines all of the necessary tasks, documents, and dates that must be met to close the books. Each quarter a new project is created from this template and all of the necessary data is loaded up and all of the relevant employees are automatically assigned their tasks for the quarter. In addition, notifications are set up so that employees know when their particular tasks should start, when they are due, and which of their tasks are overdue. Employees come to the Collaboration project to examine their tasks, post questions or concerns, and check-in relevant documentation. The head of the finance department then has a consistent outlook onto the close process and is able to communicate results more quickly and accurately to the executive team.

Once the close process is complete, the project is archived and stored as a record. In the event of an audit, the finance department can easily pull up all of the actions that were performed each quarter by simply restoring old Collaboration projects.

Example: E-learning Application

A number of customers in the education field use Collaboration as an online classroom. Normally each class has an associated Collaboration project where the instructor posts a syllabus, assignments, and answers student questions. A similar application can be applied to e-learning in the corporate environment.

Incorporating AquaLogic Interaction Process

Process builds upon AquaLogic User Interaction to allow people to rapidly build and deploy applications which automate business processes in their enterprise.

Process is a set of business tools that provides the ability to design, activate, and deploy business processes into a live environment. The platform lets users quickly combine dissimilar applications into integrated business processes. Users can then modify these processes in realtime, enabling business to react dynamically to changing market conditions.

Process is used to create and manage departmental and enterprise business processes and workflows. Common examples of these processes include purchase order requisitions, performance reviews, travel authorizations, and work order requests. All businesses have tens or hundreds of processes that can potentially be automated with Process. These processes generally involve human participants viewing and entering data through a Web browser and a flow of data through a dynamic sequence of steps encompassing both humans and systems.

For more information on Process, visit http://edocs.bea.com/alui/index.html.

Search

Next, consider requirements for incorporating search features in your site.

Search allows users to quickly and efficiently find a wide variety of information from sources across the enterprise, both inside and outside the Portal and related Publisher and Collaboration products. Search works with the ALI Directories and Web Services infrastructure to help employees do their jobs. Salespeople can find contract resources needed to close deals; marketing executives can find in-progress design documents for new products; customer service representatives can find resources stored in a variety of CRM, file, and Web repositories.

There are a number of possible sources of searchable content, and it is important to understand the options for providing that content to end-users:

- **Knowledge Directory**: The core of the ALUI knowledge management infrastructure is the Knowledge Directory—a hierarchy of folders that contain links to files of various formats, stored in different types of repositories. Files can be crawled into the Directory or manually submitted, and can be filtered into the folder hierarchy (also known as a taxonomy) in order to provide an entry point to high-quality, organized content. In addition to the out-of-the-box functionality, virtually any repository can be made searchable through the creation of Content Services. All items in the Knowledge Directory can be searchable.
- **Collaboration**: The project workspaces provided by Collaboration contain documents, threaded discussions, announcements, and task lists contributed and managed by distributed teams. All items in Collaboration can be searchable.

- **Publisher**: The form-based data entry and file management provided by Publisher allows specialized content submitted by users to be published and surfaced in the portal through portlets. All published content items associated with portlets can be searchable.
- **Portal Administrative Objects**: Users, web services, portlets, Content Services—all the objects that make up the administrative infrastructure of a portal are searchable. End-users can search for users (to view profile and expertise information), communities (to visit or join), and portlets (to add to a My Page). Administrators (who need to create and manipulate all types of objects) can search for a wider variety of items and have more advanced options in their search results.
- **Non-portal Searchable Content**: Legacy search engines and repositories with pre-existing search or query functionality can often contain valuable sources of content that for various reasons cannot be crawled into the portal or managed through Collaboration or Publisher. With search web services, any repository that can respond to queries can be extended with a web services adapter so that it can be searched from the portal. Results from a number of disparate search providers (both inside the enterprise and on the internet) can be aggregated in this way.

The Search administrator is responsible for creating and scheduling the initial search index jobs as well as update jobs. The Search administrator is also responsible for customizing search "Best Bets" and the search thesaurus.

On Appendix B, "Portal Content Responsibilities Worksheet," assign a leader to be responsible for implementing Search.

BEA provides the following resources to prepare the leader for this assignment.

Resource	Description
Consulting Services	Get advice from the experts.
Education Services	Take classes from the experts. To learn how content is discoverable in the portal, enroll in the series of classes developed for Content Managers.
Knowledge Base	1. Register and log into the Support Center. 2. Search the Knowledge Base for search-related topics, such as "Search", "Best Bets", "thesaurus", and the like.
Administrator Guide	Learn how to set up and maintain the Search index, as well as how to configure Search features, such as Best Bets and the thesaurus.

Developing and Integrating Custom Applications

ALI offers integration packages for commonly used applications. If ALI does not offer an integration package for an application you use, you can use the AquaLogic Interaction Development Kit (IDK) to develop integration packages for existing applications, packaged or custom, or to develop new applications to meet the needs of a dynamic enterprise.

This section provides the following examples of applications you can develop with the IDK:

- "Example: Access and Personalization" on page 3-18
- "Example: Searching for Data and Documents" on page 3-19
- "Example: Customer Branded Support Site" on page 3-19

For complete information on the IDK and more examples of custom-developed applications, see http://dev2dev.bea.com/aluserinteraction/.

Example: Access and Personalization

Suppose your customer accounts are set up in a CRM system such as Siebel. As new accounts are added to Siebel, you need to provide customers access to the portal and certain communities on the portal.

You can set up an authentication web service to enable the portal to periodically connect to the Siebel system, query for customer users, and provision accounts for them on the portal. In addition, this authentication web service can also be configured to accept the username and password at login and authenticate the user against Siebel. This eliminates the need for creating portal accounts manually. Also, since access is based on the back-end system, the moment the customer account is suspended on Siebel, access on the portal is also automatically denied.

You can also set up a profile web service in conjunction with the authentication web service to tap into each customer's profile information on Siebel and augment the portal user profile. This profile information can be used to add the customer user to the appropriate groups that will drive personalization. This profile information might be the basis for providing access to relevant communities; for example, a company that sells wines could provide their distributor with access to the appropriate community depending on their geographical address, a portlet that lists all the special events could be personalized based on the user profile information that tracks geography.

Example: Searching for Data and Documents

Suppose that you want to provide a way for customers to look at the status of their orders and these orders are stored in a Lawson ERP system (it could be in any ERP system, packaged or homegrown). One solution is to develop a custom search web service that, given keywords such as start and end dates, pulls in a set of sales orders and associated status information. The search web service would use the customer user credentials and information for identification purposes and return only those sales orders corresponding to that customer. This search could be conducted from the Federated Search page or, alternatively, set up as a portlet inside a community.

Suppose that you want to be able to provide the customer with access to all the contracts that they have signed with your company and that these contracts are stored in an Windows file server. You can use the Content Service -Windows Files to crawl these documents into the Knowledge Directory, and the customer can browse the Knowledge Directory to find the contracts. Alternatively, you can create a content snapshot portlet that provides them with links to the Knowledge Directory that correspond to the contracts that pertain to them.

What if instead of just finished contracts, you need to keep various versions of contracts that are being negotiated between your sales representative and the customer. Use a Collaboration project to store these documents. Collaboration projects are appropriate for work in progress. When the document is finished, it can be stored in a document repository.

Example: Customer Branded Support Site

You can choose to give each customer a personalized experience replete with custom branding. You can create a Experience Definition that includes the customer’s own branding. As soon as the customer logs in to the portal, they are directed to their branded Experience Definition. You can build custom navigation using the IDK Pluggable Navigation, replacing the portal's regular navigation. The look of a Experience Definition can be completely different from that of the parent portal.

Suppose that whenever a user from a particular customer site logs on, you want to be notified. Using the IDK Portal Event Interface (PEI), you can write custom code that gets triggered upon the customer logging in and notifies you via e-mail.

CHAPTER 4

Defining Administrative Roles

This chapter describes administrative roles at a high level.

The purpose of this chapter is to help you develop a plan to assign administrative responsibility for managing portal objects.

At the end of this chapter, you should be able to scope the effort involved in deploying the initial administrative objects. You should also be ready to assign administrative responsibility for the initial deployment.

Before you proceed, print Appendix C, "Administrative Roles Worksheet." You can use the worksheet to record your assignments and note the activity rights and access privileges for groups of users.

This chapter includes the following topics to help you scope and assign tasks for your deployment effort.

Task	Topic
1. Learn about access privileges and activity rights.	"About Access Privileges and Activity Rights" on page 4-2
2. Draft a role-based group hierarchy.	"Creating a Group Hierarchy" on page 4-3
3. Note Activity Rights required for particular roles.	"Assigning Activity Rights" on page 4-4

Task	Topic
4. Create an administrative object hierarchy.	"Defining an Administrative Object Hierarchy" on page 4-5
5. Determine which objects you want to manage through the migration approval process.	"Managing Quality through Object Migration" on page 4-6

ALI provides the following resources to prepare the leader for this assignment.

Resource	Description
Consulting Services	Get advice from the experts.
Education Services	Take classes from the experts. To learn how to develop a group hierarchy, enroll in the series of classes developed for ALI administrators.
Knowledge Base	1. Register and log into the Support Center. 2. Search the Knowledge Base for role- and group-related topics, such as "Roles", "Groups", "Activity Rights", "ACL", and the like.
Administrator Guide and Online Help	Learn how to implement roles by configuring properties for groups and users.

About Access Privileges and Activity Rights

You can manage what users read, select, and modify by assigning *access privileges* to administrative or content objects (folders) and *activity rights* to user objects (users and groups).

You delegate administrative responsibility by assigning roles to administrative groups. Generally, roles that include activity rights are held by different people in each department. For example, each department probably has its own community administrator. Keeping the access privileges separate from the activity rights allows you to define activity roles once and then apply them throughout your company, instead of having to define the same set of activity rights for each department. For example, if you want to create a content administrator role, you can create a group with the following activity rights: Access Administration, Create Content Service, Create Job, Create Filter, Create Document Type, and Access Utilities.

You can usually assign access privileges to your existing user groups, because access is generally granted along departmental lines, and, more than likely, your user groups are also along departmental lines. For example, you might give the Marketing group Select access on the

Marketing folder and all Marketing objects. Depending on the size of your department, you might also need to create roles for Admin and/or Edit access; if there are only a few people who need Admin or Edit access to an area of the portal, you might want to just change the access for those individual users. You do not need to give creation activity rights (such as Create Portlets) to every role; you only need Edit access on an object to edit the object. Therefore, you create roles that can manage existing objects but cannot create anything.

Creating a Group Hierarchy

Put the most powerful groups at the bottom (the deepest level) of a group hierarchy. Because groups inherit the rights of the parent groups, the groups with most rights are the groups furthest down the group hierarchy. For example, the IT Managers group should be a child group of the IT group, not the other way around. This is especially true of groups with activity rights directly assigned to them. They should be as low in the hierarchy as possible.

In Figure 4-1, the first-level group includes all users in the subdomain. You can configure corresponding access privileges and rights for such users. In the first domain, Community Manager 1 is a sub-group of the Domain Users 1 group and its users have all the rights of the Domain Users group, plus additional privileges required for managing communities. The Content Manager 1.1 and Content Manager 1.2 groups have all the rights of the Domain Users 1 group, plus the rights of the Community Manager 1 group, plus additional privileges for managing content. The lowest group in the hierarchy is the Administrator, who has all rights.

Figure 4-1 Representation of Group Hierarchy: Users, Community Managers, Content Managers, Admins

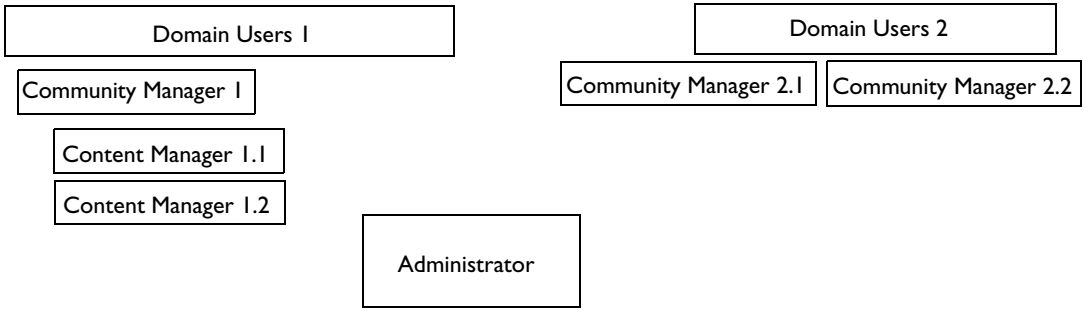

Assigning Activity Rights

Here are a few suggestions for common roles used to assign sets of activity rights.

Role	Suggested Activity Rights
Content/Document Administrator	• Access Administration – to access the administration hierarchy • Edit Knowledge Directory – to create new document folders • Create Content Services – to create new Content Services • Create Data Sources – to access secured documents • Create Document Types – to force metadata onto documents • Create Filters – to automatically manage folders • Create Jobs – to run jobs • Access Utilities – to approve documents • Access Smart Sort – to re-sort entire folders of already categorized documents
Community Creator	• Access Administration • Create Communities – to create communities • Create Community Infrastructure – to create community and page templates
Portlet Creator	• Access Administration • Create Portlets – to create portlets • Create Web Service Infrastructure – to create the remote server and web service to create truly custom portlets
Group/User Creator	• Access Administration • Create Admin Folders – to make new admin folders to store users • Create Experience Definitions – to modify the user experience of users • Access Utilities – to create default profiles to apply initial layouts to users • Create Authentication Sources – to create authentication sources • Create Jobs • Create Profile Sources – to apply user information to synchronized users • Create Groups – to create groups • Create Users – to create users • Delegate Rights – to delegate rights to users (create activity groups)

Defining an Administrative Object Hierarchy

The Administrative Object Directory is similar to the Knowledge Directory; it is a hierarchical folder structure (up to 10 levels deep), but it stores administrative objects rather than files.

The folders can store any type of administrative object (for example, Content Services, portlets, or users). Within each folder, objects are automatically grouped by object type to ease management. Each folder is secured, and objects created in that folder default to the ACL of the parent folder.

Consider the following tips when creating your administrative object hierarchy:

- Start with the end-user hierarchy. Although you can create a hierarchy based on the organizational or management structure, this is often not the best organization for end-users (users who browse the portal rather than manage parts of the portal). End-users can see the administrative hierarchy in a few places in the portal. For example, by default, the Add Portlets and Join Communities pages search the administrative hierarchy for available portlets and communities and try to display a list of objects without showing their parent folders. However, users can choose to browse the hierarchy.

  The best thing to do is to start by creating the hierarchy for only communities and portlets (including portlet bundles) and hiding the administrative objects created during installation. For example, you might want to move all objects meant for administrators to a particular folder and restrict access to the folder so that end-users will not see it if they browse the hierarchy.

  The organization of the objects meant for administrators can be based on administrative structure or by object type or by topic.

- Hide folders that are intended for administrators (as mentioned in the previous bullet).

- Organize objects by topic rather than by object type (objects are automatically grouped by type within each folder).

- Preset folder ACLs. If your group hierarchy is fairly stable, preset the ACL on Admin folders with both groups and their subgroups. This takes some planning but it will be of great use to portal managers. Since it is easier to remove than add groups and users to the ACL of objects, pre-add as many groups to the ACL of objects folders. When an object is created in that folder, it defaults to the ACL of the folder. If the object is to be restricted to a subset of the folder, the owner of the object can then remove groups from the object. For example, assume all the members of the IT Managers group are members of the IT group. It would be helpful to add both the IT Managers and the general IT group to the ACL of the IT folder. For objects that are meant for just IT Managers, simply remove the general IT group. If only the general IT group was initially added to the parent folder, you would

have to remove the general IT group from the object and search for the IT Managers group to add it to the object. If your are constantly adding groups, this might become cumbersome.

- Do not always rely on folder ACLs to control access to objects. Remember, technically a user only needs Read access to a portlet if that portlet is already on that user's My Page or community. It is better to set security on the object than the folder.
- Always manage by groups. It is always easier to manage by groups than by users, especially because you can put groups in groups. For example, assume IT manages 100 objects. user U is the IT manager. If only user U was added to the ACL of objects then it would be very difficult to also let user X manage IT objects since user X would have to be added individually to all IT objects. If the IT Manager group is added to the ACL of the IT objects, then any user added to that group would inherit the rights of that group.
- Start from the top. When you are manipulating folder security for a lot of folders, you should always start from the top. Remember that propagation goes from the top down. Therefore if you change the ACL of a subfolder and then change the ACL of the parent folder and select propagation, the changes you made to the subfolder will be lost. Since propagation is handled by jobs that are automatically created in the Intrinsic Operations folder, you will not see the changes immediately, but the jobs will run and the proper security will be set.

Managing Quality through Object Migration

You can set up a staging system for development, testing, and production deployments and use object migration to move objects from one deployment to another. This allows the ALUI administrator to strictly control who has the ability to create new objects, to test object security and process load, and to make objects available to users only when they are working as planned. For information on object migration, see Chapter 8, "Using Migration Features to Stage Your Deployment."

CHAPTER 5

Customizing the User Interface

This chapter summarizes ALUI user interface customization techniques, most of which do not require special programming skills.

The purpose of this chapter is to help you scope the effort of implementing the user interface for your deployment so that you can assign responsibility to lead.

This chapter includes the following topics:

About Experience Definitions

Experience Definitions determine many aspects of the user interface for broad groups of users. Experience Definitions control your start page when you log in, the features available to you, what your navigation looks like, and what mandatory links are shown in your navigation.

You can create a number of Experience Definitions for different audiences, including unauthenticated or guest users.

For information on configuring Experience Definitions, refer to the Administrator Guide.

Navigation

Portal navigation is customizable, in fact, the portal ships with eight different navigation options out of the box. Navigation controls everything outside the center of the page, not including the header and footer. To change the navigation presented to a user, edit the Navigation Options in the Experience Definition editor.

Navigations are pluggable; that is, you can develop new navigations using programming languages like C++, Visual Basic .NET, Java, or by simply defining them in XML.

Although navigations are associated with Experience Definitions, each navigation can be very dynamic, displaying a completely different look for each page type, each user, or any other settings you like. For example, the Support Center navigation shows completely different HTML when you are on a Support Center community, even though it is all done within a single pluggable navigation view. The Support Center navigation is downloadable from the Developer Center.

Style Sheets and Portlets

If you are happy with the layout of your existing portal and simply want to change things such as fonts, colors, logos, and images, you can override all those settings by changing your style sheet. Since we support localization of our style sheets into many languages, the easiest way to modify a style sheet for a multi-language portal is to use the Style Sheet Mill, which takes values from template files and uses them to generate style sheets in multiple languages using localized text from our translation files.

To find out more about the Style Sheet Mill, see http://dev2dev.bea.com/aluserinteraction/.

Branding

You can apply different headers and footers to different Experience Definitions and communities. The headers and footer set for the Experience Definition are applied to the entire interface, except for communities that have their own custom headers and footers. Headers and footers are just special portlets that can be used

AquaLogic Interaction Publisher provides three branding portlet templates that enable you to customize the look and feel of Experience Definitions and communities:

- Header Portlet: enables you to create customized headers. A header portlet appears at the top of the page, in the portal's banner area.

- Footer Portlet: enables you to create customized footers. A footer portlet appears at the bottom of the page.
- Content Canvas: enables you to create a branding portlet that spans the entire space just below the banner and above the page's other portlets.

You can create and configure branding portlets from these portlet templates using the Portlet Editor and Publisher's Configure Portlet Wizard. An ALUI license enables you to download and install Publisher to create and customize these branding portlets for your communities and Experience Definitions.

You apply header and footer portlets at the community template, community, or Experience Definition level, and you apply content canvas portlets at the page level. Header and footer portlets applied at the community template level override header and footer portlets applied at the community level. Header and footer portlets applied at the Experience Definition level also override those applied at the community level if the Force Community to use Header and Footer from Experience Definition option is selected for the community template.

You can also customize the way a header or footer appears in a particular community by using the Community Preferences page from the Community Editor.

For more information, see the *Administrator Guide for AquaLogic Interaction Publisher*.

Pluggable Event Interfaces (PEIs)

Sometimes you want to add new functionality rather than modify existing functionality. ALI has a large number of event categories you can hook into, each with several different event types. For example, you might change the behavior after a user logs in; for users who had not yet filled in their user profiles, you could have them redirected to the user profile form. To accomplish this, you would need to implement not only a PEI, but a custom activity space, model, view, and controller for any special landing pages you wanted to write from scratch. You could also use Dynamic Discovery to override a view class for an existing page.

For information on PEIs, see http://dev2dev.bea.com/aluserinteraction/.

Custom Activity Spaces

You might want to precisely control the exact look of the center of the page, as well. For example, you might want to control how portlets are rendered on the page. The file MyPortalContentView renders portlets into columns based on your page layout style. You might want to redesign that page center so the portlets are arranged in rows instead of columns. For this, you would need to

override the default view with your own, using dynamic discovery, also outlined in http://dev2dev.bea.com/aluserinteraction/.

A more forward-compatible approach involves extending ActivitySpaces, by creating new ActivitySpaces and views that extend existing ones and directing PEIs and other links to those new spaces. Then, as ALI improves existing activity space components, your code will benefit.

CHAPTER 6

Provisioning Host Computers

This chapter summarizes host requirements for deployment components.

The purpose of this chapter is to help you provision host computers for components you plan to deploy in a production environment.

Appendix F, "Component-Host Worksheets," provides example worksheets that characterize host requirements for typical deployment scenarios.

Use the worksheets provided in Appendix F, "Component-Host Worksheets," to record decisions for your deployments.

This chapter includes the following topics:

BEA provides the following resources to help you to make these decisions.

Resource	Description
Consulting Services	Get advice from the experts.
Knowledge Base	1. Register and log into the Support Center. 2. Search the Knowledge Base for capacity-planning topics, such as “Load Balancing”, “Failover”, and the like.

Component Host Requirements

The following table provides guidelines for provisioning host computers for ALI components.

Component	Host Requirements
Portal Service	**Minimum** Refer to Appendix D, "Evaluating Hardware for the Portal Component," for guidance in determining adequate hardware. **Recommended** • Dual processor, 1 Ghz or greater • 2 GB RAM • Separate host or share with Administrative portal and/or Image Service. **Scaling Guide** For large deployments, install multiple Portal components and configure load balancing and failover. **Portal Load Balancing** The portal can be used with any load balancing system that supports sticky IPs, such as Cisco LocalDirector, F5 Big-IP, and Windows NLB load balancing systems. Session states are maintained on the ALUI Web servers themselves. Therefore, if a Web server is taken out of the Web farm, sessions on that server are lost. If users have not set their Web browser to Remember My Password, they will have to log back in to the portal. It is possible for the portal to become unresponsive while the Web site is still operational. In that case, the load balancer should assume that the portal is still operational and continue to send requests. The load balancer should perform content verification to ensure that the portal is actually available. Since users use the Portal component in different ways, the load balancer should send requests to the computer with the most available resources instead of simply performing a round-robin distribution of requests. For maximum fault tolerance, BEA recommends that load balancers be clustered, so if one load balancer fails, another will continue to distribute requests. Consult manufacturer guidelines on clustering load balancers. **Security Guide** Separate the Portal component from other system components to increase security. When you separate the Portal component and other components, persistent data (search and database) and back-end tasks (Automation Service) are not on the same computer. If you run .NET portals in extranet environments, install the Portal component on its own computer and place that computer in a DMZ; install all other components (except the Image Service) behind the internal firewall.

Component	Host Requirements
Administrative Portal	**Minimum** • Can function also as a Portal component in a Web farm. • Can be installed on the same host as Portal component and/or Image Service. • If not functioning also as a Portal component, can be on the same host as Automation Service. **Recommended** Dedicate a CPU. Some administrative actions are CPU-intensive. **Scaling Guide** Install one Administrative Portal for your deployment. **Security Guide** If you prefer, you can install the Administrative Portal on a separate host that is located in a physical environment that only the ALI administrator can access.
Image Service	**Minimum** 256 MB RAM Can be installed on the same host computer as the Portal component. **Recommended** 512 MB RAM More processing power is required if you use SSL or compression. **Scaling Guide** You can install one or many Image Service instances. Typically, you install one instance and specify this location when you install other ALUI components. **Security Guide** The Image Service contains static content that is typically not sensitive. Therefore, it is not imperative that you install the Image Service behind a firewall

Component	Host Requirements
Document Repository Service	**Minimum** 256 MB RAM **Recommended** • 512 MB RAM • Fault tolerant disk for doc storage. **Scaling Guide** Multiple instances of the Document Repository Service can be load balanced and failed over using IP load balancing such as NLB or a hardware load balancer. This will also provide partial failover for the Document Repository Service. However, the host for the Document Repository Service requires a single writable file system backing store. This backing store cannot be load balanced, but it can be failed over with one of the following: • A shared local disk, failed over via MSCS • An external shared network drive, implemented using either NAS or MSCS **Security Guide** Install the Document Repository Service behind a firewall and restrict access so that only computers that host ALUI components can access the Document Repository Service host. End-users do not need to access to the Document Repository Service host. In Windows deployments, the Document Repository Service runs as a Windows service. In UNIX or Linux deployments, the Document Repository Service runs as a daemon or console process.

Component	Host Requirements
Automation Service	**Minimum** Must be on a separate host from Portal component; otherwise, you must schedule all jobs to run during off-peak hours. **Recommended** • Dual processor, 1 Ghz or greater • 1 GB RAM Separate host or share with Administrative Portal and/or Image Service. **Scaling Guide** If you anticipate intensive use of Identity Services and Content Services jobs, install multiple Automation Services and configure load balancing. However, because Search performs document indexing and cannot be horizontally scaled, adding multiple Automation Services for the sole purpose of crawling content does not greatly improve system performance. **Automation Service Load Balancing** Automation Services do not require any special technology to provide load balancing or failover. Installing multiple instances of the Automation Services in a portal system will provide load balancing, as jobs can be designated to run on any set of available servers. In case of server failure mid-job, the job will not complete on another server. However, jobs are typically scheduled to recur, and the next instance of any standard ALI job will complete the processing. Automation Services can be load balanced by registering job folders to multiple Automation Services. The Automation Services poll the database and pick up the next available job. Should one Automation Service fail, another Automation Service will run the necessary jobs. **Security Guide** Install the Automation Service behind a firewall and restrict access so that only computers that host ALI components can access the Automation Service host.

Component	Host Requirements
Search	**Minimum** Small (up to 250,000 documents) Dual CPU, 2 GB RAM Medium (up to 500,000 documents) Dual CPU, 4GB RAM Larger 64-bit Solaris or AIX host, Dual CPU 1.2 Ghz or greater; 4-8 GB RAM **Recommended** 64-bit Solaris or AIX host, Dual CPU 1.2 Ghz or greater; 4-8 GB RAM **Scaling Guide** CPU requirements are directly proportional to the number of users the component can support. Indexing speed is proportional to the speed of an individual CPU. RAM supports internal caching done by Search. RAM requirements are proportional to the size and number of documents indexed. **Search Load Balancing** You can improve performance by installing multiple Search instances and dedicating one instance for indexing jobs and the remaining instances for serving queries. The Search indexing instance cannot be load balanced and does not support failover. If you experience performance problems with the Search indexing instance, enhance host capacity. You can implement three levels of load capacity, with increasing deployment complexity: 1. Single server performing both indexing and serving queries. 2. One server performing indexing, with another taking the query load. 3. One server performing indexing, with two or more taking the query load. The query servers can be proxied through a third-party load balancer.

Component	Host Requirements
Search (cont.)	**Failover** You can implement three levels of failover, with increasing deployment complexity: 1. Single server. With a single server performing both indexing and querying, an additional server can be configured for query failover. The failover server will NOT provide load balancing in this configuration. 2. Two servers. With two servers splitting indexing and querying, an additional server can be configured for query failover. The failover server will NOT provide load balancing in this configuration. 3. Externally managed query server pool. The third party load balancer will provide failover to other servers in the pool. **Security Guide** Install Search behind a firewall and restrict access so that only computers that host ALI components can access the Search host. In Windows deployments, Search runs as a Windows service. In UNIX and Linux deployments, Search runs as a daemon. Search connects to other components through TCP.
Analytics	**Minimum** • Dual processor, 1 Ghz • 1 GB RAM **Recommended** Install on a separate host from the Portal component. **Scaling Guide** Install one Analytics. **Security Guide** Enable Unicast UDP on port 31314 for communication between Analytics and the Portal component. End-user access to Analytics is gatewayed by the Portal component, so the Analytics host computer can reside behind a DMZ firewall.

Component	Host Requirements
Collaboration	**Minimum** • Dual processor, 1 Ghz • 1 GB RAM Can reside on same host computer as other components that generate portlets, such as the Publisher, Studio, and Analytics. **Recommended** Install Collaboration on a separate host computer from other components to preclude contention for the JVM. **Scaling Guide** For large deployments, install multiple Collaborations and configure load balancing and failover. For details, refer to Collaboration documentation. **Security Guide** End-user access to the Collaboration is gatewayed by the Portal component, so the Collaboration host computer can reside behind a DMZ firewall. Collaboration connects to the Portal component through portlets via HTTP, to the ALI API Service via HTTP/SOAP, to the Collaboration database and portal database through JDBC, and to Search through TCP.

Component	Host Requirements
Publisher	**Minimum** • Dual processor, 1.2Ghz • 1 GB RAM Can reside on same host computer as other components that generate portlets, such as the Collaboration, Studio, and Analytics. **Recommended** Install Publisher on a separate host computer from other components to preclude contention for the JVM. **Scaling Guide** Install one Publisher. If capacity is an issue, install Publisher on a separate host with premium hardware. **Security Guide** End-user access to the Publisher is gatewayed by the Portal component, so the Publisher host computer can reside behind a DMZ firewall. Publisher connects to the Portal component through portlets via HTTP, to the ALI API Service via HTTP/SOAP, to the Publisher database and portal database through JDBC, and to Search through TCP. Publisher publishes HTML pages and image files to a Web server (called the "publishing target") via file copy or FTP. The publishing target can have the same host as the Publisher or a separate host.

Component	Host Requirements
Studio	**Minimum** • Dual processor, 1Ghz • 1 GB RAM Can reside on same host computer as other components that generate portlets, such as the Collaboration, Publisher, and Analytics. **Recommended** Install Studio on a separate host computer from other components to preclude contention for the JVM. **Scaling Guide** Install one Studio. If capacity is an issue, install Studio on a separate host with premium hardware. **Security Guide** End-user access to the Studio is gatewayed by the Portal component, so the Studio host computer can reside behind a DMZ firewall. Studio connects to the Portal component through portlets via HTTP, to the ALI API Service via HTTP/SOAP, and to the Studio database and portal database through JDBC.
ALI API Service	**Minimum** Can be on the same host as a Portal component. **Recommended** Install on the same host as a Portal component, unless you want to keep SOAP API behind a firewall. If subject to heavy use, consider one or more separate hosts. **Scaling Guide** Install one ALI API Service. **Security Guide** If you do not want to expose the SOAP API through the extranet, install the ALI API Service on a separate host from a Portal component and locate the ALI API Service host behind a firewall.

Component	Host Requirements
Database Server	**Minimum** • 1 CPU, 1Ghz • 1 GB RAM **Recommended** • 2-8 CPU • 4 GB RAM Install on separate host computer. **Scaling Guide** Database Server Load Balancing The database server can be scaled using any database-compatible clustering technology. Currently, this means that scaling can only be provided by a larger machine. If necessary, each portal database can be placed on a separate computer and scaled separately. If running on Windows, failover of databases can be provided with Microsoft Cluster Services, and geographic load balancing and failover can be provided using SQL Server replication. However, this method is technically and administratively challenging and is not recommended unless availability requirements cannot be met otherwise. Oracle databases can be deployed for high availability. ALI supports both client-side connection and server-side connection failover with Oracle RAC. For more details, see the Knowledge Base article DA_288256 *"How to configure Plumtree products to use Oracle RAC."* **Security Guide** Install the database server behind a firewall and restrict access so that only computers that host ALI components can access the database server host. End users do not need access to the database server host.

Component	Host Requirements
Remote Server - Identity Services (IDS)	**Minimum** • Dual processor, 1Ghz • 1 GB memory • 2 GB disk space **Recommended** Install on a separate host from the Portal component. To maximize performance, install in a network location that is in close proximity to back-end components. **Scaling Guide** Install additional Automation Services, as necessary, to accommodate a large number of IDS jobs. **Security Guide** End-user access to IDS portlets is gatewayed by the Portal component, so the IDS host computer can reside behind a DMZ firewall.

Component	Host Requirements
Remote Server - Content Services (CS)	**Minimum** Install on a separate host from the Portal component. **Recommended** To maximize performance, install in a network location that is in close proximity to back-end data sources. **Scaling Guide** Install additional Automation Services, as necessary, to accommodate a large number of CS jobs. **Security Guide** End-user access to CS portlets is gatewayed by the Portal component, so the CS host computer can reside behind a DMZ firewall.
Remote Server - Portlets	**Minimum** Can share a host with other portlets and Web services. **Recommended** Install on a separate host from the Portal component. To maximize performance, install in a network location that is in close proximity to back-end components. **Scaling Guide** In general, caching enables static portlets with minimal personalization to scale very well to any number of users. Dynamic portlets with more personalization cannot be as effectively cached and so require more processing power. If necessary, you can improve performance by installing dynamic portlets on hosts with premium hardware. **Remote Server Load Balancing** Remote servers can be load balanced using Parallel Portal Engine load balancing. Refer to "Load Balancing" on page 6-18 for instructions on configuring this feature. Remote servers can also be load balanced in a similar way to Portal components using the same kind of load balancing hardware. **Security Guide** End-user access to portlets is gatewayed by the Portal component, so the remote server host computer for portlets can reside behind a DMZ firewall.

Optimization Strategies

The following table characterizes optimization strategies you might consider when you provision computer resources for your site.

Goal	Approach
Low initial hardware cost	Organizations optimizing for low initial hardware cost seek to buy the least expensive machines necessary to make the software work reliably. Given a choice between repurposing two existing 1x700 MHz Pentium III servers and spending $7,500 on one 2x2.4GHz Pentium IV Server, they would choose the former.
Low hardware maintenance cost	Organizations optimizing for low hardware maintenance costs seek to reduce the number of machines needed to host the software. Because each additional computer incurs a minimum fixed cost in terms of administrative overhead, power consumption, space, and operating system license, these organizations would rather combine multiple ALUI components on a single, more powerful computer than distribute those components over multiple, less expensive machines.
High availability	Organizations optimizing for high availability are willing to spend extra money and effort to ensure that the portal and other ALUI components are available reliably to their users at all times. Such organizations typically purchase more computers and load balance them where possible, creating redundant configurations.
Low software maintenance cost	Organizations optimizing for low software maintenance cost assume that at some point in the life of the system, some part of the software will malfunction, and they seek both to lessen the chance that malfunctions will occur and lessen their impact when they do occur. Such organizations would typically purchase more individual computers to ensure that system components do not interfere with one another, and to reduce the risk that taking a computer out of the system to install new software will impact multiple system functions.
Scalability	Organizations optimizing for scalability assume that their deployments will be required to handle a large number of users. Such organizations would typically purchase extra hardware, and more expensive hardware, in order to create excess capacity in the system.

Goal	Approach
Performance	Organizations optimizing for performance seek to make their systems operate as fast as possible, especially in their ability to render pages quickly for end-users. Like organizations seeking to lower software maintenance costs, these organizations would distribute system components across a larger number of computers to ensure that each component has unrestricted access to the computing power it needs to perform its tasks the moment those tasks are called for.
Network Security	Organizations optimizing for network security seek to ensure that end-users touch only machines hosting the smallest amount of code and data. Such organizations also typically install firewalls between layers of their deployment, to ensure that if an intruder compromises one layer, the potential damage is limited. Such organizations tend to purchase more computers in order to isolate the Portal component, which end-users touch directly, from other components.

Load Balancing

The Parallel Portal Engine Load Balancer (PPE-LB) is a built in feature that allows you to load balance your Remote Servers to make better use of the Parallel Portal Engine (PPE). PPE-LB is a solution for middle-tier HTTP messaging (between the Portal component and the Remote Servers). It provides robust failover services for high availability and eliminates the need for a third party load balancing solution in front of Portlets. PPE-LB is designed to be as easy to configure as round robin DNS and readily solves proxy and SSL problems that are typically encountered with load balancing devices in middle-tier messaging.

On the DNS server, configure the Remote Server cluster name (for example, gs.portal.company.com) to resolve to multiple IP addresses. This is similar to setting up DNS round robin, except that PPE load balancing will failover, provide stickiness, and act as a load balancer. Each Remote Server in a cluster must have a unique IP address and must have the same software installed.

Note: Editing the hosts file on a Windows machine is not equivalent to configuring the DNS server. Windows caches and returns only the first IP address, instead of returning multiple IP addresses the way a DNS server does. If you are not able to configure the DNS server, contact Customer Support for registry settings you can add to provide equivalent functionality.

The entry in the DNS server should look something like this using BIND on a Unix DNS server:

```
remoteserver  60     IN     A      10.10.10.1
remoteserver  60     IN     A      10.10.10.2
remoteserver  60     IN     A      10.10.10.3
```

If the domain is *company.com*, then the *remoteserver.company.com* host name should be resolved to this list of IP addresses by the DNS server.

Portlet Support

Most Portlets should work correctly with PPE load balancing, but some Portlets may do in-memory caching that assumes the underlying database will not be modified by another application. Consult the Portlet documentation or Portlet developer to determine if specific Portlets can be load balanced.

PPE Load Balancing and SSL

If your Remote Servers use Secure Sockets Layer (SSL), BEA recommends creating a single SSL certificate by name and adding it to each machine in a Remote Server cluster.

Verifying That PPE Load Balancing is Configured Correctly

You can verify the DNS server configuration by using a tool called *nslookup*. For example, try using nslookup on www.microsoft.com:

1. Open a command line prompt.
2. Run nslookup. At the command prompt enter:

   ```
   nslookup www.microsoft.com
   ```

   This command will return something similar to the following lines:

   ```
   Server:  plumdc1.plumtree.com
   Address:  10.1.88.4

   Non-authoritative answer:
   Name:    www.microsoft.akadns.net
   Addresses:  207.46.197.100, 207.46.197.102, 207.46.230.218
   Aliases:  www.microsoft.com
   ```

   Notice that www.microsoft.com is using round robin DNS and three different IP addresses.

The PPE updates itself from the DNS server. The PPE algorithm refreshes the list of IP addresses in a Remote Server cluster more frequently as more load is placed on it; it is not based on a timed update. It starts load balancing without requiring you to restart the server.

PPE Configuration Settings

PPE is implemented in the OpenHTTP standard. OpenHTTP settings are configured through the Portal component **/settings/common/serverconfig.xml** file. The default configuration includes many of built-in and internal settings for your deployment. You can configure the following additional settings.

Setting	Description
ForceHttp10	Sends HTTP/1.0 requests instead of HTTP/1.1. The sockets are closed after sending a single request.
TraceBodyAndHeaders	Fo debugging only. Traces the values of headers and some parts of the body of the requests/responses to PTSpy. Turned off by default because headers might contain passwords in cleartext.
HttpCacheSizeMb	Defines maximum size of the cached data. Cache uses LRU algorithm to decide which old entry should be kicked out in order to accommodate newer data.
ConnectionCacheTimeoutSec	Defines the time that socket remains unused in the cache before being closed by OpenHTTP.
MinimumDNSThreads	Specifies the minimum number of threads that are used to perform DNS lookups.
MaximumDNSThreads	Specifies the maximum number of threads that are used to perform DNS lookups.
ProxyURL	Specifies the URL for a proxy host.
ProxyUser	Specifies an authentication user name for the proxy connection.
ProxyPassword	Specifies an authentication password for the proxy connection.
ProxyBypass	Contains a list of hosts accessed directly instead of through the proxy.
ProxyBypassLocal	Boolean flag specifies that all hosts that are in the same domain should not be accessed through the proxy. If a hostname does not have any “.” (dots) in its name, it is considered local, in the same DNS domain.

Before you configure other OpenHTTP settings, contact BEA Professional Services.

For more information on configuring OpenHTTP, see the Support Center Knowledge Base.

External Service Load Balancing

The portal is dependent on many other servers and services to function. Each of these services must provide some failover. At a minimum, these services include:

- Authentication sources such as the customer/partner directory and the employee domain
- File and Web servers housing documents and other content
- External applications to which portlets provide access

Scaling Using Federated Portals

One way of scaling a portal is to use multiple networked (federated) portals rather than one very large portal. This is especially useful if you require more than 25-50 GB of indexed content in your Knowledge Directory. It also makes sense if disparate departments need to share some data but use mostly different portlets, communities, and content. Sometimes the politics or organization of a company's business lends itself to a federated portal solution. Different groups can administer and control their own content separately using smaller systems that require less planning and maintenance. This can also be accomplished in a large portal system to some degree by having different departments control remote servers that serve secure content to the portal. In a federated portal system, information is shared via federated search and, possibly, shared portlets. For these systems, identify the scaling needs of each portal in the network and decide how the portals should be connected. It is very important that the various groups agree on how content is to be shared and how much load they can expect other portals to place on their portal. Size each system as you would a single large portal, but take into account potentially higher load on shared remote servers and federated search pages.

CHAPTER 7

Implementing Network Security

This chapter describes security options you can implement for your deployment.

The purpose of this chapter is to help you understand the security requirements so that you can assign responsibility for developing a plan for security in your deployment.

This chapter includes the following topics:

BEA provides the following additional resources to prepare the security leader for this assignment.

Resource	Description
Consulting Services	Get advice from the experts.
Knowledge Base	1. Register and log into the Support Center. 2. Search the Knowledge Base for security-related topics, such as "SSL", "SSO", "Security Modes", "firewall", and the like.

Security Architecture

Figure 7-1 provides a high-level representation of the recommended network security for deployment of ALI components.

Figure 7-1 Representation of Network

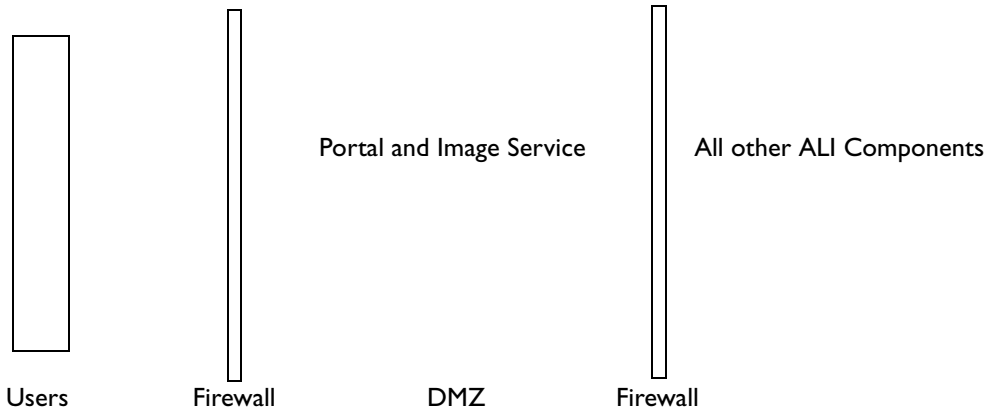

Firewalls and Security

A firewall can be a valuable component in an overall security strategy. However, firewalls alone do not create security. Firewalls typically provide the first line of defense, intelligently routing requests and filtering out those that do not meet requirements configured into the device (or software). Depending on the sophistication of the firewall product, more or less intelligence can be built into the decision tree affecting whether a packet should pass through the firewall. Having a firewall in place can provide a false sense of security, however. Consider the following real world scenario:

Suppose a Web server operating in a DMZ behind a firewall restricts traffic to port 443 (HTTPS) requests. A second firewall insulates the internal network from the computers in the DMZ. A hacker from the Internet sends a buffer overrun attack to the IP address of the Web server. The data looks like a regular HTTPS request, it goes to port 443, and is passed through to the Web server as a TCP stream. The Web server tries to decrypt the stream in the normal fashion. The data inside the request exploits a weakness in the Web server that allows it to overrun the memory stack of one of its threads. The thread executes some code sent by the hacker. The code gains control of the Web server, opens a new socket and sends a similar malicious request to the next server in the HTTP(S) chain. The request goes unnoticed by the second firewall (it still uses

HTTPS TCP port 443). The target Web server is controlled in the same fashion. If the second server is a member of your primary domain, the hacker has a good access point to your network. If the buffer overruns were done carefully and security audits for successes (versus failures) are not implemented on the Web servers in the chain, it is unlikely the attack would even be noticed.

This is not to imply that firewalls are useless. On the contrary, firewalls can dramatically limit the nature and even the source of potential attacks. Firewalls must be supplemented by good internal security policies and measures. For example, by restricting the rights and privileges of the user in whose process space the Web server runs, risks can be minimized or eliminated. Furthermore, with careful configuration of network trusts and operating system security audits and alerts, an attack such as that described above could be very difficult to implement.

Implementing AquaLogic Interaction in a DMZ

A DMZ (sometimes referred to as a demilitarized zone or perimeter network) is a computer or small network inserted as a neutral zone between a private network (intranet) and the outside public network (Internet or extranet). A DMZ uses some combination of firewalls and gateways to prevent outside users from having direct access to a server that holds company data.

The remainder of this section focuses on the positioning of ALI components with respect to firewalls and perimeter networks (DMZs). BEA does not advocate the use of any specific configuration. This section presents several possible topologies that incorporate firewalls, since they are a common element of many company infrastructures.

The most important security measure is the “hardening” (establishing maximum possible security) of the computers involved in the portal configuration, especially those that receive direct user requests. These Web server computers are sometimes referred to as bastion hosts, since they are typically the most vulnerable to attack. Establishing the appropriate privileges for the ALUI user is a critical component of this activity (as is installing all of the appropriate security patches for the operating system and application server). However, once steps are taken to secure bastion hosts and other computers, you can provide extra layers of security to protect against unknown vulnerabilities in the operating system or Web server software.

Web Services and Internal Network Security

Residing in the DMZ, the Portal component requires the most scrutiny in designing for security. Except for the Database and Search, all requests from the Portal component into the internal network are made through web services protocols using TCP/IP and HTTP 1.1. This web services architecture provides the following security advantages:

- HTTP 1.1 is a well-known protocol for tools to monitor and audit.

- Each web service runs over a single, configurable port number, which is easy to protect with a firewall.
- The Portal component implements the full range of HTTP security, including SSL/TLS, certificates, and basic authentication when making requests.
- Single sign-on (SSO) products that are designed to protect HTTP traffic can be used to protect web services residing in the internal network. The portal is designed to forward SSO tokens as needed.
- Full connections to systems of record are not provided to the DMZ. Each web service is designed to provide remote calls to specific functionality. This is analogous to providing only a set of stored procedures instead of full DBA-level database access to a client-server application.
- On Windows networks, the portal can authenticate against multiple Windows domains with no trust relationship by using multiple Windows authentication web services.
- Administrators can ban all traffic except known HTTP connections from the DMZ into the internal network. Protocols considered “unsafe” for use in a DMZ are limited to use on the internal network.

ALI provides many out-of-the-box integration web services to connect to Windows Active Directory, LDAP Servers, Documentum, Microsoft Exchange, Lotus Notes, SAP, PeopleSoft, and Siebel, just to name a few. Organizations are encouraged to develop custom web services to integrate with internal systems. From a security standpoint, these are all the same and all would be deployed on the internal network. Organizations can use firewalls to further compartmentalize internal networks as needed.

Risk Mitigation Scenarios

Following are some simplified examples of perimeter network topologies. In all cases, the target audience for the portal is both internal and external, thus some form of perimeter network is implemented. VPN topology is deliberately omitted, although it is a very common means of accessing internal portal content from outside the firewall. For the purposes of this discussion, VPN is considered equivalent to internal network access.

Scenario 1: The Reverse Proxy + IIS

One of the simplest implementations of a perimeter network involves a reverse proxy. Much as proxy servers route traffic from the internal network to the Internet, reverse proxies route traffic in the opposite direction. A reverse proxy can be hardware (for example, F5 Big IP Application Switch) or a software component (for example, MS Proxy Server), and can incorporate other

functionality, such as packet inspection and intelligent routing. The reverse proxy generally acts as the firewall between the outside world and the internal network, but can also be used in concert with multiple firewalls. Reverse proxies are somewhat controversial among network administrators, but are commonly used by organizations who view reverse proxy servers as one component in an overall solution.

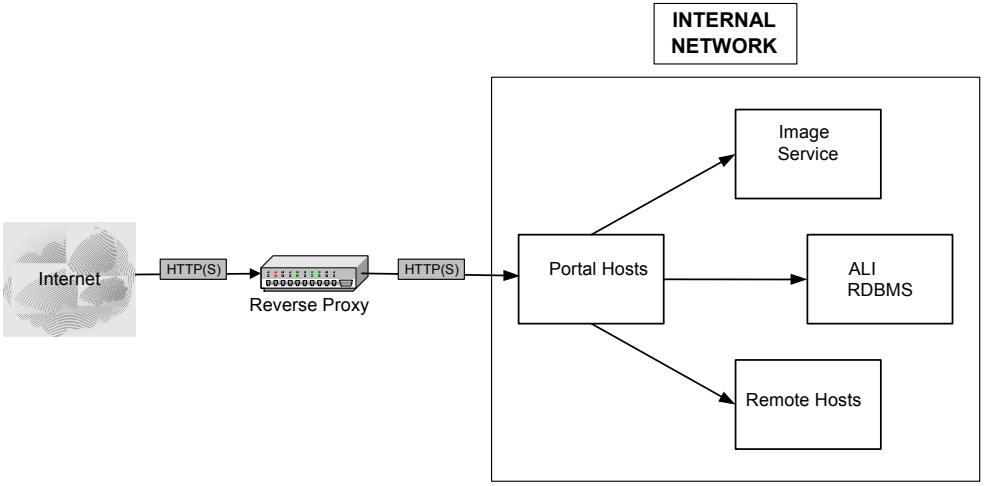

Scenario 2: Multiple Network Cards

A second very common practice is to create multiple networks in a single computer through the use of two or more Network Interface Cards (NICs). With this scenario, you might also incorporate a firewall in addition to multiple NICs, with the firewall and one NIC serving as the perimeter network. The firewall blocks all traffic except HTTP requests, which are received by the Portal component on the first NIC. The Portal component uses the second NIC to communicate with other hosts residing on the internal network. Multiple NICs on the other hosts with yet another firewall can serve to separate them from the true internal network computers. In this case, adding another layer of indirection offers considerable benefit, without any associated performance degradation. The extra network card doubles the network bandwidth of the Portal

component, improving scalability, increasing network security, and eliminating the need for server-to-server encryption, since there is no access to the dedicated subnet.

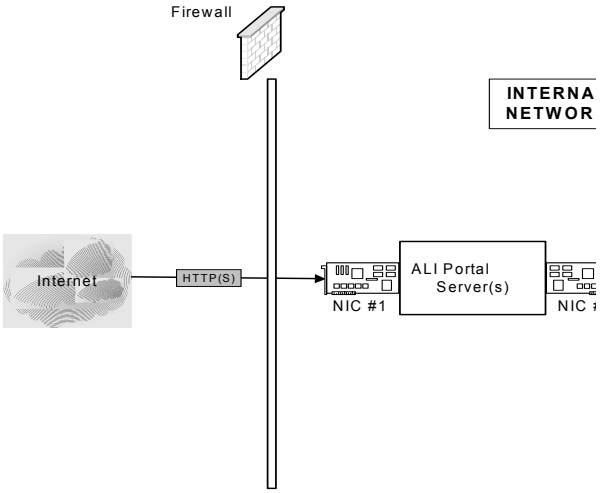

Scenario 3: Limited Functionality for External Users

Given the ALUI architecture, it is possible to deploy multiple ALI Web servers that implement different functionality. For example, one Web server might implement administrative functionality (for example, Content Service creation), while another might not. Administrative users are redirected to a separate URL to take advantage of the administrative portal pages. Similarly, external users can be granted more limited use of the portal than internal users. (Users that access the portal through a VPN are considered internal users.) Relegating all externally accessible resources to the perimeter network can eliminate virtually all traffic through the firewall for external users. At some sites, the only resource accessible through the firewall is the database server using vendor-specific database proxy software. For example, you might not want

to allow internal documents to be searched or viewed by external users. Likewise, all external users are ALUI users, not LDAP users, since no HTTP or LDAP connection is enabled.

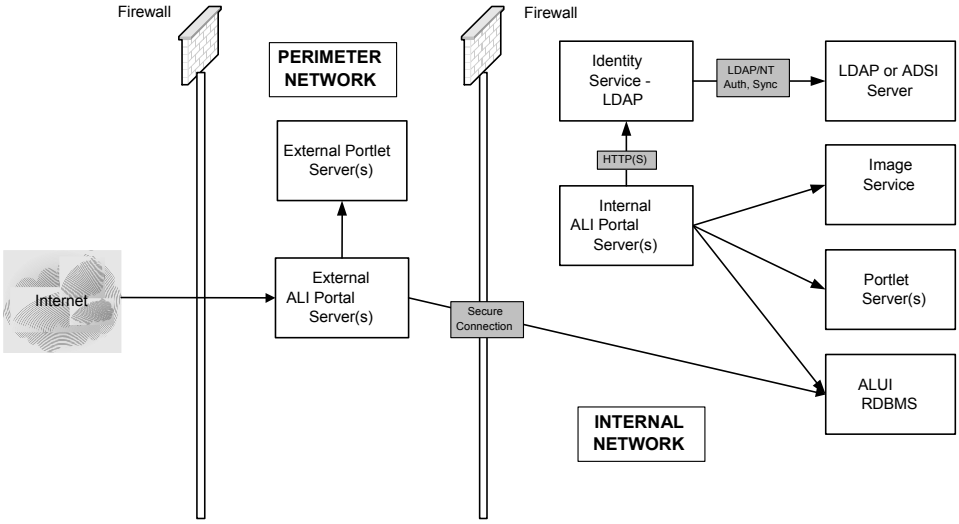

Security Modes

After you install ALI components, you can configure portal communication to use any of the following security modes.

Security Mode	Description
0	Portal pages remain in whatever security mode, http or https, that the user initially uses to access the portal. For example, if a user accesses the portal via http, all the portal pages will remain http; if a user accesses the portal via https, all the portal pages will remain https. This is the default setting. Use this mode only when the deployment is used internally, behind a firewall, and without an SSL accelerator. For example, you might want to use this mode for testing or development deployments.
1	Certain portal pages are always secured via SSL and other pages are not. For example, the login page may always be secured but a directory browsing page may not. The page types that are secured are configurable. This mode is not generally recommended.
2	All portal pages are always secured via SSL, that is, pages are accessed via https. Use this mode if you are not using an SSL accelerator. In this mode, the Web server should provide an SSL endpoint. We recommend against configuring the SSL endpoint directly on the Tomcat application server. Although application servers can handle Web requests directly, for scalability and security reasons, we recommended that you not permit your users to connect to the application server directly. Instead of securing the application server itself, you secure the front-end Web server and the channel between the Web server and the application server. Therefore, you must set up SSL and install an SSL certificate on the Web server.
3	The portal uses an SSL accelerator. This is the most common set-up for production deployments. Use this mode if you are using an SSL accelerator. As with Security Mode 2, users are not connecting to the application server directly, so you need to secure the front-end application server and the channel between the accelerator and the application sever. Therefore, you must set up SSL and install an SSL certificate on the SSL accelerator.

For detailed information on configuring these settings, see the Administrator Guide. For additional information on SSL, see “Deploying SSL” on page 7-9.

Deploying SSL

Secure Sockets Layer (SSL) is a protocol developed by Netscape for transmitting private documents on the Internet. SSL works by using a public key to encrypt data that is transferred over the SSL connection. All the major browsers support SSL, and many Web sites use the protocol to obtain confidential user information, such as credit card numbers. By convention, URLs that require an SSL connection start with https rather than http.

SSL prevents potential eavesdroppers from intercepting information (such as passwords) sent to Web sites and information (such as secure documents) the Web site sends back. Encrypting and decrypting communications requires processing and increases the time required for pages to load and display. A Web site can use SSL on some pages and not others, so most sites find a compromise between performance and security by using SSL only in strategic locations. For example, most e-commerce sites allow you to browse selections and add items to your shopping cart over HTTP, then they switch to HTTPS when you check out to protect personal information, such as your name and credit card number. Although there is some overhead in establishing an SSL session, the overhead for encrypting and decrypting during an established session is usually insignificant.

SSL can also be used to encrypt traffic between Web servers and Automation Services and portlet, Content Service, Search, or Authentication components. The Parallel Portal Engine supports SSL encryption of the parallel requests made to these remote servers, allowing safe transmission of portlet preferences and other data delivered from Portal components to other ALI components. To implement SSL between Portal and other hosts, register the remote server (the remote server object is used for all four types of web services) with an HTTPS URL. On Unix and Linux, the PPE implements SSL using OpenSSL libraries. The SSL/TLS strength and algorithm used is determined by the negotiation between the connecting parties. On Windows, the portal engine uses Microsoft SSL libraries.

About Encryption

Generally speaking, encryption is the process of converting plain text into code in order to prevent any but the intended recipient from seeing that data. There are many types of data encryption, and they are a key element of network security. For portals, encryption is important in three primary contexts. These are:

- Data transmitted between a user's Web browser and the Portal component; this is the most important context, since data might pass over the Internet.

- Data transmitted between the Portal component and other Web services that are components of your solution. In this case, a physically secure network is the ideal solution. Secure tunnels or virtual private networks can also be used. Encryption should only be used if physical security cannot be implemented; that is, different organizations control the servers, or they are physically separated.
- Data persisted by the portal.

For the first two cases, SSL provides a convenient, standard means of encrypting any data transmitted over HTTP. ALI uses an alternate mechanism for encrypting persistent sensitive data.

Encryption of Persistent Data

ALUI provides a means of encrypting sensitive information that is persisted to any of a variety of repositories (for example, the ALUI database, the Windows registry, configuration files).

Data stored by ALUI components are most commonly encrypted using a 128-Bit AES algorithm. The algorithm is fixed in the software and cannot be altered. This encryption represents the final line of defense; however, additional measures are strongly recommended to secure this data, such as:

- Both the repository and the communication channel can be secured independently using vendor-specific techniques (for example, Oracle Advanced Networking Option and data encryption)
- The database server should live in a cold room, physically isolated from other computers, on an isolated subnet where it can only talk to other ALUI servers.

About Public Key Infrastructure (PKI)

Public Key Cryptography

PKI systems use two mathematically related keys known as the public key and the private key. The public key can be distributed publicly without compromising the security of a system, as long as the private key remains secret. Anyone can use the public key to encrypt a message such that only someone with access to the private key can read it. Of more importance to this discussion, someone with access to the private key can digitally sign a message. Anyone can read a digitally signed message and can use the public key to verify that it was signed with the private key, proving the identity of the author.

Attack 1: The Imposter and the Bank

Suppose a bank receives an e-mail message directing the transfer of $10,000 from a particular account to a numbered Swiss bank account. The message is, purportedly, from a bank customer and is digitally signed. The bank officers attempt to verify the digital signature, but, as they have never received e-mail from that particular customer before, they do not have the customer's public key in the system. The e-mail contains a link to a public key server, so the officials follow the link, look up the key, and use it to verify the signature. Everything checks out, so they transfer the money to the account of the impostor who has just robbed the bank customer's account.

How? The impostor set up a key server and entered his own public key under the bank customer's name, then signed the directive using his own private key. The signature is perfectly valid since the public and private keys correspond. The bank was fooled into thinking the impostor's public key was that of the customer, making the directive seem genuine.

Defense: X.509 Digital Certificates

The X.509 digital certificate standard was designed for a single purpose: to certify the owner of a public key. A certificate contains a public key and the name and contact information of its owner, all signed by a trusted certificate authority. The certificate authority has taken steps to ensure the bank customer's identity before issuing the certificate, and, verifying the authority's digital signature ensures that the certificate is not a forgery.

Returning to the example above, the impostor will be unable to obtain a certificate for his public key in the name of the bank customer from a reputable authority. If he tries to use his own certificate, the bank officials will note that the owner of the key is not the purported signer of the e-mail and block the transaction.

Attack 2: Using a Digital Certificate to Prove Identity

Suppose the bank receives another e-mail directing a transfer, again purporting to come from one of its customers. This message is not digitally signed, but attached is a copy of the customer's X.509 digital certificate. The certificate contains the customer's name and address and is properly signed by a trusted certificate authority. Should the bank officers approve the transfer?

No, not if they understand PKI. A public key is, by definition, public information. For people to verify a signature, the owner of the signature posts it to public key servers and distributes it to anyone who asks. The message could have come from anyone who has access to a computer.

Conclusion: Signed and Certified

To be confident of identity in a PKI system, two pieces of information are needed:

- A valid certificate, proving ownership of the public key
- A valid signature, proving knowledge of the corresponding private key

SSL employs both of these pieces of information to provide secure client authentication as an optional part of the SSL handshake. If requested by the server, an SSL client provides a copy of its client certificate and digitally signs a digest of the handshake. An impostor can easily supply a copy of a stolen certificate, but cannot forge the signature without knowing the corresponding private key.

About Delegation and Portals

Delegation is a technical term for a process in a computer security system that allows a system to act on behalf of a particular user, particularly when accessing other systems. As an example, consider an e-mail portlet in a portal. When the portal system attempts to access the e-mail system, the latter will respond with a challenge for credentials. At this point, the portal can prompt the user for credentials or use credentials stored earlier. What you want to be able to do, however, is delegate to the portal system the authority to access the e-mail system on your behalf. This grant of authority should last only as long as you are connected to the Portal component and should not require storing credentials permanently. Kerberos is an excellent example of a system that permits this sort of delegation.

PKI Does Not Permit Delegation

In any PKI system, the private key is jealously guarded, since access to it grants the privilege of signing messages. Since signing messages is the only way to prove identity within PKI, the only way to delegate authority is to share the private key. This is not controlled, secure delegation, but rather an unlimited, permanent grant of authority.

Application to Portals

Returning to the portal example cited previously, consider the case where the portal has been configured to authenticate using the client certificate option of SSL. When a user connects a Web browser, it sends a signed message accompanied by the user's certificate, and these together prove the user's identity to the portal. Can the portal now pass the user's certificate on to the e-mail system to retrieve mail? No, for the same reason the bank officials will reject attack 2. If the e-mail system accepted the certificate as proof of identity, anyone could easily access an e-mail account using the publicly available certificate.

So, to access e-mail on a user's behalf, the portal system needs the private key, which it does not have after the SSL handshake. Browsers do not supply any automatic way to transfer this key (nor

should they), so the user needs to go through some configuration process of extracting the private key from the key store on the local machine and uploading it to the portal for storage. The user would either repeat this process on every login or allow the portal to store the certificate permanently. In the latter case, the portal becomes a holding point for every certificate in the system, a sort of master key room, which can potentially grant a hacker access to every system as any user.

Using PKI in Your Portal

Delegation is the major problem, but digital certificates have other drawbacks as a portal single-sign-on solution:

- Digital certificates are hard for users to set up. While the Web browser might come with support for certificates, the user still needs to install a personal certificate.
- What about users accessing from multiple computers, including, for example, an airport kiosk? Either administrators must prohibit this, severely restricting functionality, or users must remember to uninstall their certificates or risk leaving their credentials installed on public machines.
- Certificates are generally valid for many days, or even years, after they are issued. This creates the problem of certificate revocation, as an organization will want to cancel a certificate before the expiration date if an employee leaves or a certificate is stolen or compromised. Doing this requires issuing and maintaining certificate revocation lists, which are not yet standardized and create administrative and performance issues.

Complete Solution: PKI with an SSO Server

The BEA OEM version of Oblix NetPoint SSO software supports the use of digital certificates to securely and transparently authenticate against the SSO server, which then issues a delegable credential (in the form of an HTTP cookie). Using out-of-the-box configuration options, the portal can accept this credential for user authentication and forward it to selected portlet web services. Further, the Oblix NetPoint product provides tools to simplify the process of issuing and managing certificates and can be configured to accept other forms of authentication (for example, passwords) for users on the road without access to their desktop certificate. To implement this option, consult the Oblix documentation to enable and deploy PKI within that system.

ALUI also supports integration with SSO products from other vendors.

Stand-Alone Solution without Delegation

In the absence of a cookie-based SSO solution, there are two approaches to configuring the portal to accept client certificates for authentication. Both require using SSL on the login page (security modes 1, 2, or 3). When using certificates tied to Windows domain accounts, you can configure IIS to accept client certificates and then use the out-of-the-box Windows SSO configuration on your portal. This is the easiest approach to take since it leverages the built-in features of Windows, and is the recommended approach whenever possible. To accept certificates from users unrecognized by IIS, you will need to implement a custom SSO solution, which entails writing a custom SSO vendor class in Java or C#.

Summary

You can install server digital certificates on any ALUI component (Administrative Portal, Portal, Remote Servers, and so on) and enable SSL. ALI supports SSL between browser and server as well as between a Portal component and Remote Servers.

You can use a digital certificate and SSL to communicate with your LDAP server.

You can use client digital certificates with SSL to authenticate users to the portal. This can be done with Windows Integrated SSO, with a custom SSO solution, or in conjunction with a supported third-party SSO product (for example, Netegrity, Oblix).

The portal *cannot* "passthrough" digital certificates to Remote Servers. This is impossible because SSL does not permit delegation.

You can do SSO to portlet Web services and use digital certificates to log in to the portal, but only if you use a third-party SSO product that supports both cookie-based SSO and digital certificates (this includes Oblix WebGate and Netegrity SiteMinder). In this case, users use the digital certificate to log in to the SSO server and obtain the SSO cookie, and ALI accepts the SSO cookie and forwards it to portlet Web services.

Setting Up SSL

There are several steps involved in setting up SSL for your deployment. This section provides a brief overview of the steps you need to complete.

1. Set up SSL on the Web servers or SSL accelerators that run the Portal component and Image Service. Refer to your Web server or application server documentation for instructions on setting up SSL and creating, signing, and installing an SSL certificate.

2. Configure the Portal component by editing the configuration file—j_config.xml for Java deployments, n_config.xml for .NET deployments. The configuration file is located in your portal installation directory, for example, \settings\config\j_config.xml.

   a. Make sure HTTPSecurePort and HTTPort are set to the ports you want to use.

   b. Change ApplicationURL0 from `*` to

      `http://computer_name:port_number/portal/server.pt.`

      Note: You do not need to include the *`port_number`* for .NET deployments.

   c. Change SecureApplicationURL0 from `*` to

      `https://computer_name:port_number/portal/server.pt.`

      Note: You do not need to include the *`port_number`* for .NET deployments.

   d. If you have more than one URL mapping entry, you might need to change those entries as well. Refer to the comments in the configuration file for more information on URL mapping.

   e. Change SecurityMode from `0` to `1`, `2`, or `3`.

   f. Change ImageServerSecureBaseURL from `http` to `https`, and change the Image Service port to the correct one.

3. If you set the IMAGESERVERCONNECTIONURL in the portal configuration file to an Image Service running in SSL (not recommended) you must import onto the Portal component the certificate of the CA that signed the certificate used by the Image Service. For details, see "Importing CA Certificates into the Keystore" on page 7-16.

   Note: If you have any portlets or Remote Servers that use JSControls or Adaptive Portlets you must also import the CA certificate into those runtimes. (The JSControls libraries are embedded in server and IDK products and are initialized by HTTP-downloading an XML configuration file from the Image Service.)

4. If you have a Remote Server (a server running remote Web Services such as portlets, authentication sources, profile sources, or Content Services) running in SSL, you need to import onto the Portal component the certificate of the CA that signed the certificate used by the Remote Server.

5. If you are running Publisher, Workflow, Collaboration, or Studio, refer to the following sections to configure them to use a secure Image Service and Portal:

   - "Setting Up Publisher to Use a Secure Image Service or Portal" on page 7-18

- "Setting Up Workflow to Use a Secure Image Service or Portal" on page 7-19
- "Setting Up Collaboration to Use a Secure Image Service" on page 7-19
- "Setting Up Studio to Use a Secure Image Service or Portal" on page 7-20

Importing CA Certificates into the Keystore

Importing CA Certificates into the cacerts Keystore (for Java Portals)

For each machine that makes requests to a secured server running in SSL, you must import into the cacerts keystore the certificate of the CA that signed the certificate used by the secured server:

1. On the machine that makes requests to a secured server, open a command prompt.
2. Copy the CA certificate to this machine.

   To obtain the CA certificate, navigate to the CA and save the .der encoded certificate file as a .cer file; you might want to use *imgsvr.cer* for an Image Service or *portal.cer* for a Portal, or you might want to use the server hostname.
3. Import the certificate:

   ```
   keytool -v -import -trustcacerts -alias CA_alias -file
   CA_certificate_path -keystore cacerts_keystore_path
   ```

   Replace the variables with the following information:

   - *CA_alias* - the alias for your CA, for example, *verisign* or the server hostname
   - *CA_certificate_path* - the path and filename to the CA certificate you copied to the Portal component
   - *cacerts_keystore_path* - the path to your cacerts keystore, located at "jre\lib\security\cacerts" in the home of the JVM that runs your Java application server, for example:
     - For Tomcat, `jakarta-tomcat-4.1.30-LE-j2sdk1.4.1_05-win32\j2sdk1.4.1_05\jre\lib\security\cacerts`
     - For WebLogic, `bea\weblogic700\server\lib\cacerts`
     - For WebSphere, `java\jre\lib\security\cacerts`
4. Enter the password to the cacerts keystore.

Importing CA Certificates into MMC (for .NET Portals)

For each machine that makes requests to a secured server running in SSL, you must import into the MMC the certificate of the CA that signed the certificate used by the secured server:

1. On the machine that makes requests to a secured server, open a command prompt.
2. Copy the CA certificate to this machine.

   To obtain the CA certificate, navigate to the CA and save the .der encoded certificate file as a .cer file; you might want to use *imgsvr.cer* for an Image Service or *portal.cer* for a Portal, or you might want to use the server hostname.
3. Open MMC:

   ```
   C:\>mmc
   ```
4. Click **Console** | **Add/Remove Snap-in**.
5. Click **Add**.
6. Click **Certificates**.
7. Click **Computer Account** and then click **Next**.
8. Click **local computer** and then click **Finish**.
9. Click **Close** to close the Add Standalone Snap-in dialog box.
10. Click **OK** to close the Add/Remove Snap-in dialog box.
11. In the MMC tree, expand to **Console Root** | **Certificates** | **Trusted Root Certificate Authorities** | **Certificates**.
12. Right-click **Certificates** and select **All Tasks** | **Import**.
13. Click **Next**.
14. Select your certificate.
15. Click **Next**.
16. Choose to place all certificates in the following store: **Trusted Root Certification Authorities**.
17. Click **Next** and then click **Finish**.
18. Restart IIS.

Setting Up Publisher to Use a Secure Image Service or Portal

1. If you are using a secure Image Service:

   a. In a text editor, open **content.properties** (located in your Publisher installation directory, for example, C:\Program Files\plumtree\ptcs\6.0\settings\config\content.properties).

   b. Change Image Service entries:

   – If you are using Security Modes 1 or 2, find and replace *all* occurrences of `http://machine_name/imageserver` with `https://machine_name/imageserver`, where *machine_name* is the name of the machine hosting Publisher.

   – If you are using Security Mode 3, change the Image Service entries as follows (note that some variables use http and some use https):

   - `CommunityImagePublishBrowserLocation=https://machine_name/imageserver/plumtree/portal/templates`
   - `CommunityImagePreviewBrowserLocation=https://machine_name/imageserver/plumtree/portal/templates/preview`
   - `CommunityStyleSheetListURL=http://machine_name/imageserver/plumtree/common/public/css/community-themes.txt`
   - `JSComponents.AlternateImageServerUrl=http://machine_name/imageserver`

   If the Image Service is on Java, be sure to change the port to the correct one (for example, `https://machine_name:ssl_port_number/imageserver`).

2. If you are using Security Modes 1 or 2, import the certificate of the CA that signed the Image Service and/or Portal certificate into Publisher. For details, see "Importing CA Certificates into the Keystore" on page 7-16.

3. If Publisher runs on WebSphere against a .NET portal in Security Mode 1, complete the following steps:

   a. Open the WebSphere Admin console.

   b. Navigate to the Default Server.

   c. Click the **JVM Settings** tab.

   d. Under System Properties, click **Add** (this should add a line).

e. For **Name** type “java.protocol.handler.pkgs” and for **Value** type “com.ibm.net.ssl.www.protocol”.

4. Restart Publisher.

Setting Up Workflow to Use a Secure Image Service or Portal

1. If you changed the AlternateImageServerURL in the content.properties file, perform the following steps so that Publisher can communicate the alternate Image Service URL to Workflow:

   a. Restart Publisher. This writes the URL to Workflow.

   b. After Publisher has restarted, restart Workflow. This forces the Workflow Web application to re-query Publisher for the alternate Image Service URL.

2. Import the certificate of the CA that signed the Image Service and/or Portal certificate into Workflow. For details, see For details, see “Importing CA Certificates into the Keystore” on page 7-16.

Setting Up Collaboration to Use a Secure Image Service

Collaboration does not require any changes to function in security modes 1 or 2, as it uses the portal's Image Service settings. A certificate is not required.

If you are using Security Mode 3, import the certificate of the CA that signed the Image Service and/or Portal certificate into Collaboration. For details, see “Importing CA Certificates into the Keystore” on page 7-16.

However, if the host/port of the normal Image Service URL used by browsing users is not accessible from Collaboration (for example, the Image Service is on a different machine than Collaboration), you must change the jscontrols component that Collaboration uses. The symptom of this problem is error messages displayed in the Calendar portlets. To avoid the errors:

1. In a text editor, open **config.xml** (located in your Collaboration installation directory, for example, C:\Program Files\plumtree\ptcollab\5.0\settings\config\config.xml).

2. In the following line set the URL to the value in the portal configuration file (j_config.xml or n_config.xml).

```
<jscontrols>
   <imageServerConnectionURL>[URL]</imageServerConnectionURL>
```

Setting Up Studio to Use a Secure Image Service or Portal

Import the certificate of the CA that signed the Image Service and/or Portal certificate into Studio. For details, see "Importing CA Certificates into the Keystore" on page 7-16.

Troubleshooting

Component	Guidelines
KeyTool	• If the keytool command is not recognized, it might be because Java is not in your path. Change your directory to the Tomcat installation directory, for example, C:\Program Files\jakarta-tomcat-4.1.18-LE-j2sdk1.4.1_02\j2sdk1.4.1_02\jre\bin. • If, when running the keytool command, you get an "alias already exists" error, change your command's "-alias" argument to use a different alias.
Publisher	The following errors indicate that Publisher has not been set up with SSL: • Workflow Tracker portlet displays "`Unable to display this page`" • The Community Templating Style Sheets, Community Branding Image Publishing Target, and Community Branding Image Preview Target display "`Write Channel Closed, possible SSL handshaking or trust failure`." • Attempts to create a content item or branding portlet display an http status 500 error: `org.apache.jasper.JasperException: jscomponent file for jscontrols not found or failed to load. Exceptions encountered: com.plumtree.openfoundation.io.XPIOException: java.security.cert.CertificateException: Couldn't find trusted certificate – com.plumtree.openfoundation.io.XPIOException: Unexpected end of file from server` • Navigating to the image directory of the Community Directory Section in Publisher Explorer displays a blank page when trying to view the images.
Workflow	If Workflow portlets are not displayed when using an alternate Image Service URL, refer to Knowledge Base article 230242. If the My Activities portlet displays "`An unknown error has occurred`", the Workflow Server is not working. Import the certificate of the CA that signed the Image Service and/or Portal certificate into the Workflow Server JRE

Single Sign-On Options

Single sign-on (SSO) has many different meanings in different contexts. It can mean protecting your Web server with a product from an SSO vendor, such as Oblix. It can mean preventing your users from having to enter credentials more than once, or at all. It can mean identity management or a way to store credentials for many systems to simplify user experience and administrative management. Usually it means some combination of these notions.

The key to understanding your SSO requirements is realizing that the portal is based on a loosely coupled architecture; different tiers of components communicate with each other, primarily over HTTP. The end-user connects to the portal tier over HTTP; the portal connects to numerous other systems over HTTP, including many systems that act on behalf of the end-user. The key is to make this simple for the end-user, simple enough for the administrator, and secure enough for the security team.

Delegating to Remote Authentication or SSO

When users point their browsers at the portal, the default experience is for users to be presented a login screen. This login screen allows users to authenticate against any authentication source, which might be a remote system such as LDAP, the portal database, or an SSO provider.

When delegating authentication to an LDAP authentication source, the portal can be configured to keep the user credentials in a safe location within memory for later use. The sequence of events for LDAP is as follows:

1. User goes to portal HTTP address; enters credentials.
2. Portal stores credentials in safe section of memory.
3. Portal sends request to LDAP authentication source.
4. LDAP authentication source returns OK.
5. User is granted access to their profile in the portal.

Delegating authentication to an SSO source can circumvent the ALUI login screen and engage the end-user in the SSO login mechanism (could be a login screen, a key card, or some other mechanism). Common SSO sources include Oblix, Netegrity, and Windows Integrated Authentication (WIA). The sequence of events for Oblix would go something like this:

1. User goes to portal HTTP address.
2. Oblix intercepts this HTTP request, realizing user is missing a cookie. If user already has cookie, skip to Step 6.
3. Oblix redirects to Oblix server.
4. Oblix server authenticates; sets browser cookie.
5. Oblix redirects to original HTTP address.

6. Oblix intercepts this HTTP request, recognizes valid cookie and instructs ALUI to grant user access to their profile.

In both cases the authentication was delegated to an external source. In both cases this was likely ultimately delegated to LDAP. In the first case, however, the portal has the user's password for later use (if configured to do so). In the second case, the SSO vendor might have employed any of several authentication mechanisms that it supports, whether login screen or keycard.

Logging in to the Portal with Auto-Authentication

When using Windows Integrated Authentication, the user must be logged in to a machine on a Windows network. The browser on this machine is smart enough to pick up the user's identity, so the browser can negotiate with the portal to establish the users credentials. The sequence of events for WIA is as follows:

1. User logs onto a Windows network; opens a browser.
2. User goes to portal HTTP address.
3. Portal challenges the browser with an WIA challenge (401 Unauthorized).
4. The browser asserts an encoded piece of information (Negotiate).
5. The portal challenges the browser with a piece of information (Challenge).
6. The browser asserts another encoded piece of information (Response).
7. WIA accepts the user's HTTP request; if the credential were incorrect, the user would be challenged with a login screen in Step 6.
8. The portal accepts the user's identity brokered by WIA and grants user access to their profile.

Notice that the portal was never able to capture the user's password, and the user needed to be logged in to a Windows network for the authentication to succeed. Additionally, a multi-pass handshake occurred between the browser and the portal. Companies often request that the portal be able to repeat this WIA authentication between the portal and the remote server. The portal cannot do this, because there is no forwardable token; Although WIA supports not only NTLM but also Kerberos5, which theoretically supports delegation, no supported browsers implement delegation. So, not only is the portal unable to broker this multi-pass handshake, WIA will fail across any HTTP proxy.

Other important considerations when using WIA are that the user must be an Active Directory or Windows user, Internet Explorer or Netscape 7.1+ must be used, and proxies between the browser and portal are not allowed.

Brokering Credentials to the Remote Tier

After the user has logged in to the portal, the portal wants to serve the user exciting applications. These applications might be pulling from custom systems as well as enterprise systems such as SAP. Let us discuss how the portal can connect to SAP.

ALI provides a Integration Services- SAP that allows business users to provide SAP functionality in their applications. This extension is a remote component that calls into various SAP APIs. These APIs require a username and password. By default, the extension gives user credentials to the SAP system.

ALUI can pass user credentials in any of several ways:

- Preferences: Users can set credentials as preferences, and these preferences can be sent in the HTTP request. This is useful if SAP runs off of different credentials than the main network and if no mapping exists between SAP credentials and the main user credential (such as LDAP). Preferences are stored in the portal database (encrypted), and controlled by the end-user.
- UserInfo: The administrator can map properties from LDAP or another system into user profiles, and segments of the user profile can be sent as UserInfo. If LDAP contains entries for SAP credentials for each user, then the user will never need to enter their SAP credentials, but the LDAP system will need to be the system of record.
- Lockbox: As part of the Credential Vault, credential pairs can be put in a Lockbox, which ensures that they are never transmitted in clear text. They can be sent to remote services over SSL.
- PassThrough: The login credentials the user provided at login can be sent to the remote tier as a Basic Authentication header. This is very useful if both the portal and SAP are based on a user's LDAP credentials. In this case, all communication between the portal and the remote tier should be over a secure channel (such as SSL) to protect the user password. For this method to work, the password had to be captured on login, which was not the case with the NTLM or Oblix examples described previously.
- SSO: An SSO token can be forwarded to the remote tier. This will not, however, let the SAP extension log in to SAP unless the SAP API accepts the SSO token as a valid user login. The SAP API that ALI uses (JCO) is not enabled to accept tokens of third-party SSO systems, such as Oblix. Systems such as Oblix can be used in conjunction with the SAP extension, but Preferences, UserInfo or PassThrough will need to be used to log in to SAP. It may be necessary, however to forward the SSO token if the remote server has been protected by Oblix; in this case, without the SSO token, the portal's request will be rejected. Also, the SSO token can be used with an SSO vendor API to reconvert to name

and password, but this is highly dependent on the SSO vendor and the particular company configuration of the SSO software.

Summary

The following table summarizes some features of different SSO options.

Feature	No SSO	Oblix	Netegrity	WIA	Custom
Can use windows identity	X	✓	✓	✓	C
Auto-login to all systems	C	C	C	C	C
Remember password	✓	✓	✓	✓	C
Forward password to remote tier	✓	C	C	X	C
Forward SSO token to remote tier	X	✓	✓	X	C
Supports X.509 client certificates	X	✓	✓	✓	C
Supports two-factor authentication	X	✓	✓	X	C
Supports non-windows applications	✓	✓	✓	X	C

✓ = out of box

C = custom work required

X = not supported

CHAPTER 8

Using Migration Features to Stage Your Deployment

This chapter summarizes migration capabilities for ALUI deployments.

The purpose of this chapter is to make you aware that in the ALUI environment, objects can be staged initially in a testing environment before they "go live" in a production environment.

BEA recommends you use this feature to test quality and gain acceptance on a development stage before you migrate components to a production stage.

Figure 8-1 Components that Can Be Migrated Are Indicated by an Arrow

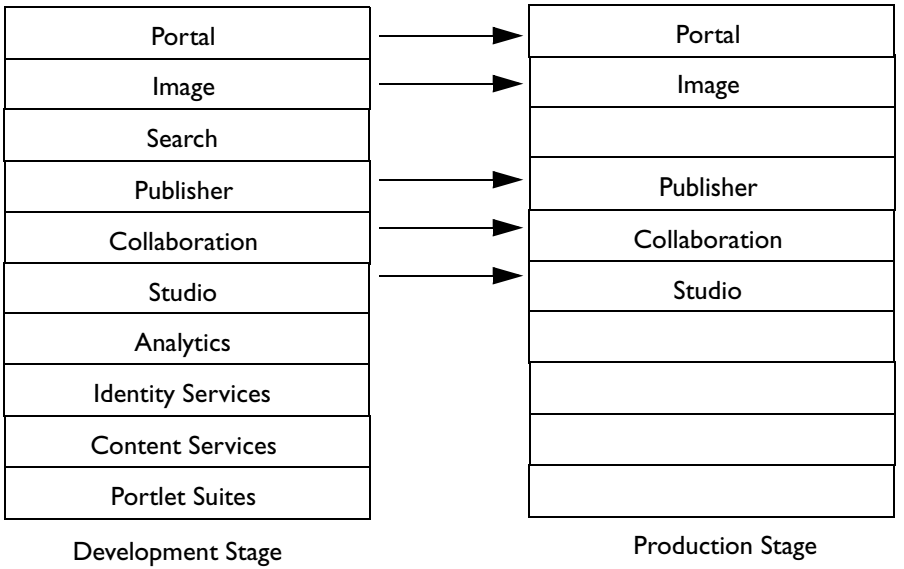

The following table summarizes migration capabilities.

Component	Migration Guidelines
Portal and Image Service	Follow the guidelines in the Administrator Guide.
Search	If the portal is being deployed live for the very first time, BEA recommends that only the portal database, Web server, and so on, *not* Search be ported directly from the staging environment to production. The production Search installation should start out as a freshly installed component with an empty index. A Search Repair should then be scheduled, using Search Manager utility in the portal. The next run of the Search Update Agent will ensure that all portal objects and documents get indexed by Search. If the portal is already live, and some new portal objects are tested out in a staging environment and then pushed out onto production, no extra action needs to be taken in order for the new objects to be indexed. The Search Update Agent will soon run (it is typically scheduled to run every 10 minutes) and index the new objects for searching.

Component	Migration Guidelines
• Identity Services • Content Services • Portlet Suite	BEA recommends you do not use the Migration Utility for these migrations. Instead, import the original **.pte** package and complete configuration as described in the installation guide for the BEA product.
Publisher	Follow the guidelines provided in the Publisher documentation.
Collaboration	Follow the guidelines provided in Collaboration documentation.
Studio	To export a Studio portlet 1. Export it to a **.pte** file. 2. Use the Migration Utility to import the .pte, which creates the portlet on the portal, as well as creates the underlying Studio application and database. Keep in mind the following: • Migration is relevant only when the source portal and destination portal have independent Studio installations. • You can migrate only portlets created with compatible versions. • You do not need to include dependencies when migrating Studio portlets. • Data in the source Studio portlet database will not be migrated. You can use the Studio data export/import functionality to move data from one Studio portlet to another. • If multiple Studio portlets share the same underlying database, BEA recommends that all of these portlets be migrated in a batch, rather than individually, in order to maintain the relationship between the portlets and their shared database.
Analytics	BEA recommends that the Analytics installation not be migrated but freshly installed in a production environment.

CHAPTER 9

Localizing Your Deployment

This chapter describes ALUI localization features.

The purpose of this chapter is to make you aware of localization features so that you can assign a leader for any localization effort for your portal.

This chapter includes the following topics:

About Localization

ALUI products are fully Unicode enabled.

The ALI user interface uses the UTF-8 encoding of Unicode when delivering HTTP content to the browser.

The ALI user interface is localized into the following languages:

- Dutch
- English
- French
- German

- Italian
- Portuguese
- Spanish
- Japanese
- Korean
- Simplified Chinese
- Traditional Chinese

Each portal user can choose their preferred language by changing their locale under **My Account | Edit Locale Settings**. For example, if a portal user changes their locale setting to any of the German locales (Austria, Germany, Luxembourg, or Switzerland), the user interface language will change to German.

Localizing Names and Descriptions of Objects Stored in the Database

You can localize names and descriptions of objects stored in the portal database. For example, if you have created a portlet with the name "Travel Portlet", you can give that portlet an associated German name of "Dienstreise Portlet". The German name will display for users who have chosen German as their user interface language.

Names and descriptions can be added or modified using the administrative user interface for each object. To add a localized name or description, open the object in the administrative editor and, on the Properties and Names page, specify the localized name and description and choose the appropriate language. For more information, refer to the online help.

The Localization Manager enables you to export and import localized names and descriptions in bulk. The Localization Manager can be used to translate a large number of names and descriptions at once. For example, if you want to translate all the portlet names and descriptions into French, German, and Italian, you can download an XML file containing the names and descriptions of all the objects in the ALUI system, edit the list, and then upload it back into the portal, effectively replacing all localized names and descriptions for all the objects in the portal database. (If an object is not set to support localized names, it is not included in the names and descriptions that are downloaded.)

Here is a small sample of a downloaded Names and Descriptions .xml file:

```
<localizationtable>
  <languages count='9'>
    <language>de</language>
    <language>en</language>
    <language>es</language>
    <language>fr</language>
    <language>it</language>
    <language>ja</language>
    <language>ko</language>
    <language>pt</language>
    <language>zh</language>
  </languages>
  <segments count='554'>
    <segment stringid='0' itemid='1' classid='2'>
      <source language='en'>Administrators Group</source>
      <target language='de'>Administratorengruppe</target>
      <target language='en'></target>
      <target language='es'>Grupo Administradores</target>
      <target language='fr'>Groupe d'administrateurs</target>
      <target language='it'>Gruppo Amministratori</target>
      <target language='ja'>[redacted]</target>
      <target language='ko'>[redacted]</target>
      <target language='pt'>Grupo de administradores</target>
      <target language='zh'>[redacted]</target>
    </segment>
    <segment stringid='1' itemid='1' classid='2'>
      <source language='en'>Plumtree Administrators Group</source>
```

```
        <target language='de'>Plumtree-Administratorengruppe</target>
        <target language='en'></target>
        <target language='es'>Grupo Administradores de Plumtree</target>
        <target language='fr'>Groupe d'administrateurs Plumtree</target>
        <target language='it'>Gruppo Amministratori Plumtree</target>
        <target language='ja'>                    </target>
        <target language='ko'>Plumtree        </target>
        <target language='pt'>Grupo de administradores Plumtree</target>
        <target language='zh'>Plumtree        </target>
      </segment>
…
    </segments>
</localizationtable>
```

The Localization Manager uses XML so that the translations for all names and descriptions in all languages can be kept in one file. The “languages” element lists all the user interface languages in the portal. The “segments” element indicates the number of names and descriptions in the portal. Finally each “segment” element contains one name or description in the portal. The source element contains the source text for translation and the translated text is stored in the “target” element for each target language. (Languages are identified using standard ISO 639-1 two letter language identifiers.)

Adding a User Interface Language

There are two types of languages that affect product deployment: user interface languages and search languages. Eleven user interface languages and can be easily extended to support additional languages. 62 search languages are hard-coded and not extensible.

The user interface languages are automatically detected when the portal is started by detecting the language folders that exist under the root folder, for example, \settings\i18n.

A user interface language can be easily added by creating a new language folder (using the ISO-639-1 language code). For example, if you wanted to add Czech as a user interface language, you would copy the “en” language folder under the i18n root folder and rename the folder “cz”. After you do this and restart the portal, you will notice that you can select Czech as your language

in the **My Account** | **Edit Locale Settings** page. If you do this, however, you will notice immediately that you are missing the style sheets for Czech. The additional elements required for a complete portal localization include:

- Style sheets
- Online help
- Javascript language files

Adding Language Style Sheets

If you are adding a user interface language to the portal, you need to add the corresponding style sheets for that language. The CSSMILL was designed to make adding languages and generating all the language style sheets relatively easy. The folder **\ptimages\tools\cssmill\prop-text** is where all the language files are kept. Each language file in the **prop-text** folder has language-specific values for font style, font size, text style, and so on. This design makes it easy to change the default font for each language. For example, if you want the default font for the Japanese user interface to be Tahoma, then you can add Tahoma to the **ja** language file in the **prop-text** folder. Besides adding a language file, you must also edit the **build.xml** file to generate the new language style sheets.

For example, suppose you wanted to add Czech as a portal user interface language. Here are the precise steps to follow to add the Czech style sheets:

1. Navigate to the **\ptimages\tools\cssmill\prop-text** folder on the Image Service. Copy one of the existing files to the same folder and rename it using the language conventions in ISO-639-1 and ISO-3166. For example for Czech, we would rename the file to be “cz”.
2. Open the new file in a text editor and make any necessary modifications for the new language. For example, if you want to add a new default language, you could change the line

   ```
   font.largest=20px verdana,arial,helvetica,”sans-serif”
   ```

   to

   ```
   font.largest=21px Tahoma,”MS PGothic”,Verdana,”sans-serif”
   ```

   Be sure to add the new font for each font attribute in the language file.
3. Navigate to the **\ptimages\tools\cssmill\prop-color** folder on the Image Service. Edit every one of the existing color properties files and add the translation for the name of the color for the new language. For example, edit the file **color.1.properties**, copy the last colorscheme.name entry. Change the name according to the new language ID chosen in Step 1. In this example, we would copy and edit the line

```
colorscheme.name.zh=\\u6DE1\\u7D2B
```

to

```
colorscheme.name.cz=Lavendelblauw
```

4. Modify the Ant build script (**build.xml**) to add the new language to the style sheet collection by following the steps below. (This is the only way the script knows to create versions of the new style sheet for each of the languages supported by the portal.)

   a. Navigate to the **\cssmill** directory and open the **build.xml** file in a text editor.

   b. Add an entry for the new language within the **make_main_css** target: Copy the last `<antcall target="make_main_language_css">` entry and paste it at the end of the list. Modify the `<param name="LANGUAGE" value="pt"/>` tag by changing the value ("pt") to the language id used in the name of the new language file created in Step 1 above. In this example, the new language ID for Czech is "cz".

   c. Add an entry for the new language within the **make_508_css** target: Copy the last `<antcall target="make_508_language_css">` entry and paste it at the end of the list. Modify the `<param name="LANGUAGE" value="pt"/>` tag by changing the value ("pt") to the language id used in the name of the new language file created in Step 1 above. In this example, the new language id for Czech is "cz".

   d. Add an entry for the new language within the **make_comm_color_css** target: Copy the last `<antcall target="make_comm_lang_color_css">` entry and paste it at the end of the list. Modify the `<param name="LANGUAGE" value="pt"/>` tag by changing the value ("**pt**") to the language id used in the name of the new language file created in Step 1 above. In this example, the new language id for Czech is "cz".

   e. Add an entry for the new language within the **append_index_for_color** target: Copy the last `<concat destfile="${INDEX}" append="true">mainstyle${COLOR}-pt.css=${colorscheme.name.pt}</concat>` entry and paste it at the end of the list. Change the language ID in the new line to the new language id. In this example, change the language ID "pt" to the new language id for Czech "cz". The new line would look like this:

   f. <concat destfile="${INDEX}" append="true">mainstyle${COLOR}-nl.css=${colorscheme.name.nl}</concat>

   g. Save the file and close it.

5. After the build script modifications have been made, create the new style sheets by running the **make_all** batch file. (See the directions in the previous section.)

6. Verify that the new language style sheets were created based on the new language property file. Navigate to the **cssmill\css** directory and make sure that there are 20 new style files with the new language ID you chose in Step 1 (that is, mainstyle-nl.css). For further verification, open the **community-themes.txt** file (in the **\css** directory) and confirm that there is a new entry corresponding to the language ID used in the new language property file.
7. After confirming that your changes are correct, move the new style sheet files from the **\cssmill\css** folder to the **\ptimages\imageserver\plumtree\common\public\css** folder.
8. Restart the Application Server.

Adding an Online Help Language

The online help is localized into core languages: English, French, German, Italian, Portuguese, Spanish, Japanese, Korean, Simplified Chinese, and Traditional Chinese. The online help files are stored on the Image Service under the root folder, for example, \imageserver\plumtree\portal\private\help.

The online help is compiled using RoboHelp X5. Each language folder contains a separate help project for that language. All Western European language projects are compiled using the standard English version of RoboHelp. The Asian languages are compiled using the corresponding Asian language version.

Adding Javascript Language Files

Several user interface components are written in client-side Javascript. These components also contain user interface text messages. These string files must also be localized when adding a language to the portal.

The Javascript components are located on the Image Service. The Javascript component string files are located in the portal installation directory (for example, C:\Program Files\plumtree), in these folders:

- \ptimages\imageserver\plumtree\common\private\js\jscontrols\120449\strings
- \ptimages\imageserver\plumtree\common\private\js\jsdatepicker\118523\strings
- \ptimages\imageserver\plumtree\common\private\js\jsutil\118981\strings

The ALI convention for Javascript string files is somewhat different than for XML files. Instead of placing each language file in a separate folder, we have given each language file a suffix consisting of a dash and the language code. For example, to create a language file for Czech, you would copy the English file and replace the “-en” suffix with “-cz”.

Language Support in the Knowledge Directory

The search index is stored in Unicode (UTF-8) and supports 62 languages in total. The Search engine supports advanced stemming and tokenization for 23 languages and basic tokenization for an additional 39 languages. The languages supported for advanced stemming and tokenization are:

• Chinese (Simplified)	• Chinese (Traditional)	• Czech
• Danish	• Dutch	• English
• French	• Finnish	• German
• Greek	• Hungarian	• Italian
• Japanese	• Korean	• Norwegian (Bokmål)
• Norwegian (Nynorsk)	• Polish	• Portuguese
• Romanian	• Russian	• Spanish
• Swedish	• Turkish	•

CHAPTER 10

Developing a Production Maintenance Plan

This chapter provides an overview of portal maintenance tasks and tools.

The purpose of this chapter is to help you scope administrative responsibilities for the portal so that you can develop a maintenance plan.

This chapter includes the following topics:

Periodic Tasks

The following table provides a few suggestions for periodic tasks you should consider to maintain your production system.

Frequency	Task
Daily	Modify security of portlets, communities, and other objects in the portal. Modify permission roles for users. Publish new and existing applications/portlets to Remote Servers. Monitor Portal, Database, and Remote Server alerts for CPU, memory, and hard disk usage to ensure availability.
Weekly	Install releases to one or more software components.
Monthly	Add new hardware to the environment (for example, new Remote Server, new hard disk, and so on).
Ad Hoc	Install portal patches. Install server patches from Microsoft/Dell/Antivirus.

Monitoring ALUI Services

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 ALUI 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.

For detailed information on the Counter Monitoring System, see the Administrator Guide.

Monitoring Databases and Java Application Servers

Databases support Performance Monitor counters on Windows. WebLogic, Tomcat, and WebSphere do not. For information on performance monitoring for application servers, refer to the related application server documentation.

Monitoring Usage

AquaLogic Analytics is an advanced usage tracking and analytics tool designed exclusively for ALI. This portal add-on enables you to assess portal ROI and define future opportunities with usage trends in mind. Analytics delivers the following features out of the box:

- Usage Tracking Metrics: Tracks metrics for common portal functions, including community, portlet, collaboration project, and document hits, as well as search queries, logins, and more.
- Behavior tracking: Tracks usage patterns, such as number and duration of visits.
- User Profile Correlation: Correlate metrics with user profile information. In this way, usage tracking reports can be viewed and filtered by profile data, such as country, company and department.

AquaLogic Interaction Analytics includes the following reports that can be customized by setting filtering, grouping, and presentation options.

Report	Description	Features
Community Traffic	Displays traffic information for each community in the portal.	Traffic is displayed in three ways: • Hits: Count of page views within the community. • Visits: Count of visits to the community, each visit can consist of several hits. • Users: Count of unique users who have visited the community. Users can select to see the most active, least active, or a select list of communities.
Community Response Time	Displays average, maximum and minimum response time for each community within the portal.	The response time is calculated as the time between the portal receiving a community page request until the time an HTML response is sent to the client. Users can select to see the slowest response times, fastest response times or response times for a select list of communities.

Report	Description	Features
Portlet Usage	Shows usage statistics within gatewayed portlets.	The traffic is displayed in two ways: • Activity: Count of hits on an object (for example, a button or link) within a portlet. • Users: Count of unique users who have performed an activity within the portlet. Users can select to see the most active, least active, or a select list of portlets.
Portal Traffic	Shows an aggregate of all portal page views within the portal.	
Portal Users	Displays statistics regarding portal user accounts.	The following four figures are displayed to help explain user inception and activity. • Total: Total user accounts in the portal. • Added: Added (new) user accounts created in the portal during a given date range. • Active: Active users defined by activity during a given date range. • Inactive: Inactive users defined by inactivity during a given date range
Portal Logins	Shows an aggregate of all portal logins.	
Portal Duration	Displays the length of visits to the portal.	Visit durations are calculated as the time between login and logoff or the time between login and inactivity for a configurable length of time. This report shows both average and maximum visit duration.

Report	Description	Features
Search Keywords	Shows the top search keywords entered in searches within the portal.	See the top 5, 10, 25, 50 or 100 search keyword phrases entered within the portal.
Document Views	Shows statistics for document views in the portal.	These statistics can be displayed in two ways: • Top Documents: List of top documents viewed with view count.· • Folders: Count of all document views by folders in the knowledge directory.

Troubleshooting Tools

This section describes logging and troubleshooting tools. It includes the following topics:

ALI Logging Utilities

ALI Logging Utilities includes three *log message receivers* that allow for a wide variety of logging solutions. In the ALI OpenLog Framework, log message receivers act to display or store log messages generated by *log message senders*, such as the portal, Collaboration, or Publisher. ALI Logging Utilities include:

- ALI Logging Spy. ALI Logging Spy (previously called PTSpy) is the primary log message receiver for the OpenLog Framework. In addition to displaying log messages from the portal and other ALUI products and services, ALI Logging Spy provides features including fine-grained filtering, viewing of saved log files, highlighting of errors, and the searching and sorting of log messages.
- ALI Logger. Logger runs as an unattended background process that receives log messages from the OpenLog Framework and writes the messages to the file system. In addition to this primary use, the Logger can be configured to output in other ways, such as sending log messages to an e-mail system.

- ALI Console Logger. The Console Logger runs in a console window, writing log messages to the console standard output. The Console Logger has limited use; in most cases, it is preferable to use Logging Spy.

For information on using these utilities, see the *Installation Guide for AquaLogic Interaction Logging Utilities*.

View Source

HTML code creates Web pages. In turn, ALI Activity Spaces generate HTML code. Along with HTML from the View and Display pages, the underlying framework inserts some general information for each page. If there is an error on the page, the Error framework might insert additional debugging information. You can review the HTML source for any given Web page to gather this information. Often the HTML for a given error page contains detailed information about the error.

When to Use View Source

Use View Source to gather more information when you receive an error on a portal page or when you want some general information about the page. For example, use View Source if you receive the following error message on a portal page: “An unexpected error occurred when trying to start the Editor.” The message itself gives no clues to the source of the error, but when you view the HTML source code for the page, you might be able to determine the source of the error.

How to Use View Source

While viewing the Web page, in the browser menu, click **View** | **Source**. This displays the HTML for that page. If the browser menu is unavailable, sometimes it is possible to view source by right-clicking the page, and then clicking **View Source**. With this approach, be aware that if there are frames, only the source for the frame in which you right-clicked will display. When the source displays, you can search for specific pieces of information as described in the next section.

What Is Available in View Source

Each portal page contains several pieces of general information:

- To determine the server hosting the portal, search for “Hostname:”. The hostname of the server is commented in the source: “<!--Hostname: My Server-->”.
- To find information about the build of the portal, search for “Portal Version:”, “Clingiest:”, and “Build Date:”.

- To find information about general timing data points, search for “Total Request Time:”, “Control Time:”, “Page Construction Time:”, and “Page Display Time:”.

If there is an error on the page, View Source might provide extended information. There are three items that you can search for:

- To view the error, search for “alert Error Title”. You might have to repeat the search because several error related Tanglements might use that text.
- To view extended information, search for “Extended Error Message:”. The extended error is wrapped in an HTMLComment and thus does not show up on the page, “<!--Extended Error Message: Sample Extended Error message.-->”. The extended information, controlled by the developer and Activity Space, is frequently the same as the error message that displays in the user interface.
- You might also need to search for “unexpected error”. When the portal encounters an unexpected error, the stack trace for the error is often inserted into an HTMLComment. The following example informs the user where the error originates from. The user then has a starting point from which to perform further debugging:

```
<!--An unexpected error occurred when trying to start the Editor.:

com.plumtree.openfoundation.util.XPException: An unexpected error
occurred when trying to start the Editor.

at
com.plumtree.portalpages.admin.editors.group.GroupModel.DoTaskOnStartEd
itor(GroupModel.java:411)
```

APPENDIX A

AquaLogic Interaction Products Worksheet

In the following worksheet, note the products you have selected to deploy.

	Description
User Interface	ALI portal and Image Service components provide the basic building blocks of the end-user experience, including UI templates for ruled-based Managed Experiences, personalized pages, communities, and the Knowledge Directory. The Administrative Portal provides a centralized management UI for all deployment components.
	AquaLogic Interaction Development Kit (IDK) includes the UI Customization Kit and the Pluggable Navigation Kit to enable you to rapidly deploy customizations to the out-of-the-box UI look, feel, and functionality.
Web Services	ALI is a Web Service that acts as the gateway between user requests and composite applications.
	ALI API Service enables rapid integration with ALUI applications.
User Management	ALI allows you to manage group privileges hierarchically.
	• Identity Service - Active Directory enables rapid integration with Active Directory authentication services. • Identity Service - LDAP enables rapid integration with LDAP authentication services.

	Description
Content Management	AquaLogic Document Repository Service enables content from file system locations, URLs, Collaboration projects, or Publisher targets to be made available through the portal.
	AquaLogic Content Upload Service enables portal users who might not have access to data staging locations (for example, extranet users) to upload files to the Knowledge Directory.
	AquaLogic Publisher provides a UI and templates for creating content and its metadata and managing this content with workflow, scheduled publishing, and expiration controls.
	• AquaLogic Content Service - Windows Files enables document discovery for remote data sources on Windows file system. • AquaLogic Content Service - Documentum enables document discovery for remote data sources on Documentum Docbases. • AquaLogic Content Service - Lotus Notes enables document discovery for remote data sources on Lotus Domino Servers. • AquaLogic Content Service - Microsoft Exchange enables document discovery for remote data sources on Microsoft Exchange Servers.
Security	AquaLogic User Interaction architecture allows most Web applications to be hosted behind a firewall, with only the ALI portal directly serving user requests. ALI supports SSO. ALI provides object-level security so you can manage user access to content with ACL and user activity by designating activity rights.
Collaboration	AquaLogic Interaction Collaboration provides portlets to facilitate collaborative workspaces, including document source control, threaded discussions, and announcements.
	AquaLogic Interaction Process is a set of business process management (BPM) tools that provides the ability to design, activate, and deploy business processes into a live environment. The platform lets users quickly combine dissimilar applications into integrated business processes. Users can then modify these processes in real time, enabling their business to react dynamically to changing market conditions.

	Description
Search	AquaLogic Interaction Search indexes all of the resources accessible through the portal pages, communities and Web applications deployed across the enterprise. These resources include: • Content indexed from file systems, Web sites, and document databases • Project documents and Web pages created and stored by Collaboration and Publisher • Applications, portlets, communities, and users
Scheduled Operations	AquaLogic Interaction Automation Service enables the ALI administrator to schedule and run administrative and maintenance jobs, such as user and group synchronization or data source crawls.
Capacity Planning	AquaLogic User Interaction integrates with performance monitors common to operating systems, Web application servers, and back-end data source servers.
	AquaLogic Interaction Analytics provides portlets and portlet templates for tracking user activity and content usage.

	Description
Business Applications	AquaLogic Interaction Studio enables rapid development of portlets, such as telephone lists, work order processes, calendars and surveys, without any coding. These plug-in applications feature a user interface, application logic and a database, and operate using Web services technologies.
	• AquaLogic Portlet Suite - Exchange provides portlets to support user access to mail, calendar, and address books stored on a remote Exchange Server. • AquaLogic Portlet Suite - IMAP provides portlets to support user access to mail, calendar, and address books stored on a remote IMAP Server. • AquaLogic Portlet Suite - Lotus Notes provides portlets to support user access to mail, calendar, and address books stored on a remote Lotus Notes Domino Server. • AquaLogic Portlet Suite - Documentum provides portlets to support user access to documents stored on a remote Docbase. • AquaLogic Portlet Framework - Microsoft Excel provides a framework for developing portlets that enable collaboration on Excel projects.
	• Integration Services- PeopleSoft enables rapid development of portlets to support user access to PeopleSoft data sources. • Integration Services- SAP enables rapid development of portlets to support user access to SAP data sources. • Integration Services - Siebel enables rapid development of portlets to support user access to Siebel data sources.
	AquaLogic Interaction Development Kit (IDK) is a set of APIs, documentation and sample code that work in both Java and .NET-based development environments, allowing you to develop with the IDE of your choice. Developers building service-oriented applications use the IDK to ensure that services from different systems and application servers work together. To speed development, ALI provides a broad range of integration Web services that allow developers to integrate data and functionality from enterprise systems and provides shared application services for content management, search, collaboration and knowledge management.

Portal Content Responsibilities Worksheet

In the following worksheet, assign responsibility for deployment of content to your portal.

Content	Description	Assigned to:
Portal	Decide whether to implement: • a single portal • federated portals	
Experience Definitions	1. _____ 2. _____ 3. _____ 4. _____ 5. _____ 6. _____ 7. _____ 8. _____	

Content	Description	Assigned to:
Document Types	Windows Files	
	HTML	
	Microsoft Office	
	Lotus Notes	
	Microsoft Exchange	
	Adobe PDF	
	Documentum	
	Novell	
	Custom Type: _____	
	Custom Type: _____	
Search	Assign a manager for AquaLogic Search.	
Collaboration	Assign a manager for AquaLogic Collaboration.	
	Assign a manager for AquaLogic Process.	

Content	Description	Assigned to:
Portlets	AquaLogic Portlet Suite - Exchange provides portlets to support user access to mail, calendar, and address books stored on a remote Exchange Server.	
	AquaLogic Portlet Suite - IMAP provides portlets to support user access to mail, calendar, and address books stored on a remote IMAP Server.	
	AquaLogic Portlet Suite - Lotus Notes provides portlets to support user access to mail, calendar, and address books stored on a remote Lotus Notes Domino Server.	
	AquaLogic Portlet Suite - Documentum provides portlets to support user access to documents stored on a remote Docbase.	
	AquaLogic Portlet Framework - Microsoft Excel provides a framework for developing portlets that enable collaboration on Excel projects.	
	Integration Services - PeopleSoft enables rapid development of portlets to support user access to PeopleSoft data sources.	
	Integration Services - SAP enables rapid development of portlets to support user access to SAP data sources.	
	Integration Services- Siebel enables rapid development of portlets to support user access to Siebel data sources.	
	Other: Using AquaLogic Studio	
	Other: Using AquaLogic Publisher	
	Other: Using AquaLogic IDK	

Content	Description	Assigned to:
Communities	1. _____	
	2. _____	
	3. _____	
	4. _____	
	5. _____	
	6. _____	
	7. _____	
	8. _____	

Administrative Roles Worksheet

In the following worksheet, note the Activity Rights and access privileges required for administrative roles and assign them.

Role	Rights Required	Assigned to:
ALI administrator	Activity Rights • All ACL • Select on /	
Content Manager I	Activity Rights 1. _____ 2. _____ 3. _____ 4. _____ 5. _____ ACL 1. _____ 2. _____ 3. _____ 4. _____ 5. _____	

Role	Rights Required	Assigned to:
Content Manager II	Activity Rights 1. _____ 2. _____ 3. _____ 4. _____ 5. _____ ACL 1. _____ 2. _____ 3. _____ 4. _____ 5. _____	
Content Manager III	Activity Rights 1. _____ 2. _____ 3. _____ 4. _____ 5. _____ ACL 1. _____ 2. _____ 3. _____ 4. _____ 5. _____	

Role	Rights Required	Assigned to:
Content Manager IV	Activity Rights 1. _____ 2. _____ 3. _____ 4. _____ 5. _____ ACL 1. _____ 2. _____ 3. _____ 4. _____ 5. _____	
Community Manager I	Activity Rights 1. _____ 2. _____ 3. _____ 4. _____ 5. _____ ACL 1. _____ 2. _____ 3. _____ 4. _____ 5. _____	

Role	Rights Required	Assigned to:
Community Manager II	Activity Rights 1. _____ 2. _____ 3. _____ 4. _____ 5. _____ ACL 1. _____ 2. _____ 3. _____ 4. _____ 5. _____	

Role	Rights Required	Assigned to:
Community Manager III	Activity Rights 1. _____ 2. _____ 3. _____ 4. _____ 5. _____ ACL 1. _____ 2. _____ 3. _____ 4. _____ 5. _____	
Community Manager IV	Activity Rights 1. _____ 2. _____ 3. _____ 4. _____ 5. _____ ACL 1. _____ 2. _____ 3. _____ 4. _____ 5. _____	

For information on mapping roles to administrative groups and users, see the Administrator Guide.

APPENDIX D

Evaluating Hardware for the Portal Component

This appendix describes how to evaluate hardware options for the ALUI portal component so that you can provision hardware with appropriate CPU, RAM, and disk capacity.

Complete the steps in the following worksheet to evaluate hardware options.

Step	Calculation
Step 1	Estimate peak load using the following calculation: `Pages/sec = ((Power user pages/hr * #power users + Normal user pages/hr * #normal users + Infrequent user pages/hr * #infrequent users pages/hr) / (3600 sec/hr) * fraction of users who could log on who are actually connected` **Note:** Base your calculations on historical data for existing Web sites that serve a similar function. Use the following conventions to identify users: • Power users. A power user is one who routinely adds or deletes portal content. • Normal users. A normal user is one who routinely reads content. • Infrequent users. An infrequent user is one who does not routinely use the portal. Record your estimated peak load here: _____ pages/sec
Step 2	• Review the benchmark charts on "Portal Performance on Various Hardware Hosts" on page D-3. • Choose a configuration that supports the peak load calculation from Step 1. In general, you want to provision a number of Portal components that support a total of 2 to 3 times the estimated peak load from Step 1. For example, if you estimate peak load to be 15 pages/sec, you want to provision either: – One (1) Portal component that can support 30-45 pages/sec – Two (2) or three (3) Portal components that each support 15 pages/sec. • Record the benchmark capacity here: _____ pages/sec • Follow the steps described in Steps 3-10 to adjust this benchmark capacity to a real-life estimate of expected use.
Step 3	If users use My Pages more than communities, revise the number upward by approximately 5%.
Step 4	If users use the Knowledge Directory more than 20% of the time, revise the number downward by approximately 10%.
Step 5	If the deployment runs under SSL (security mode 2) on the Portal component, without an SSL accelerator, subtract 25%.
Step 6	If the Portal component or server hosting the Portal component performs HTTP compression, subtract 10%.
Step 7	If the conditions in both Step 5 and Step 6 are true--that is, you use SSL without an accelerator and use HTTP compression in the portal component--add 18%.

Step	Calculation
Step 8	If this Portal component also serves the Administrative Portal, revise the number downward by 5%.
Step 9	If you use a virus scanner on the Portal component, subtract 0-10%, depending on the virus scanner settings.
Step 10	If you deploy the Portal component on a Tomcat Web application server, subtract 20%. • Tomcat does not perform or scale as well as WebLogic, WebSphere, or IIS with .NET. The capacity can be as much as 20% lower than with these alternatives. • Tomcat is also very sensitive to configuration. A poorly tuned Tomcat configuration could perform at less than 50% of the capacity of these other application servers. • Set the following Tomcat configuration variables in the Coyote HTTP 1.1 connector to the following settings for optimal performance: – connectionTimeout = 8000 – maxThreads = 100 – maxKeepAliveRequests = 1 (disable HTTP 1.1 Keep-Alive) – acceptCount = 500
Step 11	After you have made the adjustments in Steps 3-10, does the configuration you selected in Step 2 still meet your capacity requirements?

Portal Performance on Various Hardware Hosts

Portal performance demonstrates the following general trends:

- Performance varies significantly on different types of x86 server hardware.
- Performance is not dependent on the operating system, where platforms are otherwise similar.
- Performance is dependent on the JVM or CLR used and how these are tuned.
- In general, .NET and Java show similar performance, being nearly equal on most two-processor servers. However these vary somewhat in the sensitivity to processor frequency and system memory performance:
  - Java tends to be more sensitive to system memory performance.
  - .NET is more sensitive to processor cache size and processor frequency.

These differences run approximately within a plus or minus 20% performance range at the very extreme.

- Overall performance is highly dependent on memory subsystem performance, which tends to be the most important performance-related property of a server. Memory subsystem performance can be characterized by the total aggregate system bandwidth to memory as well as the latency of memory access. For Intel Xeon-based systems, this is correlated with the processor bus speed. Systems with a 800Mhz bus significantly outperform those with a 400Mhz bus. Pentium III Xeon-based systems are also limited by their memory subsystem and scale poorly with extra CPUs.

  The performance data in the table that follows is indicative of these general trends. This table provides benchmark data for the current version of the G6 Portal component. For each representative system, the load shown is the maximum sustainable load on the server with an average mix of page views on an uncustomized system. Various factors will influence the maximum sustainable load of individual deployments such as UI customizations, effective use of portlet caching, and different mixes of page types.

System	System Details	Pages/Second
Pentium III Xeon 1.4Ghz, 133 Mhz SDRAM (Dell PowerEdge 1650)	1 x 1.4Ghz Processor 512K L2	33
	2 x 1.4Ghz Processors 512K	44
Xeon 2.4Ghz to 3.06Ghz, 533Mhz system bus, HyperThreading enabled (Dell PowerEdge 1750) **Note:** Mhz is less important than total processor cache for performance. Disabling HyperThreading decreases performance 12% with two CPUs and 25% with one.	1 x 2.4Ghz Processor 512K L2	56
	2 x 2.4Ghz Processors 512K L2	71
	1 x 2.8Ghz Processor 512K L2	58
	2 x 2.8Ghz Processors 512K L2	73
	1 x 3.06Ghz Processor 512K L21M L3	67
	2 x 3.06Ghz Processor 512K L2 1M L3	82

System	System Details	Pages/Second
Xeon 3.2Ghz, 800Mhz system bus, HyperThreading enabled (Dell PowerEdge 1850) **Note:** The 800Mhz bus improves the performance of Xeon-based systems greatly. Disabling HyperThreading decreases performance by approximately 15% for dual-CPU systems and 25% for single CPU systems.	1 x 3.2Ghz Processor 1M L2	75
	2 x 3.2Ghz Processors 1M L2	100
Xeon MP 2.7Ghz, 512K L2 2M L3 caches, 400Mhz system bus, HyperThreading enabled (Dell PowerEdge 6650) **Note:** The 400Mhz bus of this system severely limits performance. This system is the only one with large differences in performance between .NET and Java platforms: .NET performs 25% better than the performance indicated here for the 4 processor node and performs equally well as Java for one processor.	1 x 2.7Ghz Processor 512K L2 2M L3	34
	2 x 2.7Ghz Processors 512K L2 2M L3	53
	4 x 2.7Ghz Processors 512K L2 2M L3	66
UltraSparc III 1.0Ghz (Sun Fire 280R) **Note:** Most Sun systems scale the memory bandwidth with each CPU, and will scale well with increased CPU count.	1 x 1Ghz CPU	27
	2 x 1Ghz CPU	50
	4 x 1Ghz CPU	90
Power 5 1.5Ghz (IBM iSeries 520) **Note:** Power 5 based IBM systems scale memory with each pair of CPUs, and performance will scale with CPU count as a result.	2 x 1.5Ghz CPU	70

APPENDIX E

Component-Host Templates

This appendix provides example worksheets that characterize host requirements for typical deployment scenarios.

Guidelines

Components Required for all Deployments

You must install the following components:

- Administrative Portal
- Portal
- Image Service
- Plumtree Search
- Automation Service
- ALUI Database

Optionally, install the following additional components to add specific functionality:

ALI API Service	Required for access to the ALI SOAP API.
Document Repository Service	Required to support document upload, Publisher, and Collaboration.

Publisher Publisher Database	Required to provide template-based branding for portal communities and the full set of Publisher features and portlets.
Collaboration Collaboration Database	Required for Collaboration features and portlets.
Studio, Studio Database	Required for Studio features and portlets.
Analytics Analytics Database	Required for Analytics features and portlets.

General Principles

- All machines should have multiple CPUs.
- More RAM generates better performance for all products.
- Disk speed is more important for the Database Server, Document Repository Service, and Search.

One-Host Portals

Usage

Demonstration, proof of concept, development, or test environment.

Configuration Worksheet

Host	Portal	Image Service	Automation Service	ALI API Service	Administrative Portal	Database Server	Document Repository Service	Search	Collaboration	Publisher	Studio	Analytics
Host1	X	X	X	X	X	X	X	X	X	X	X	X

Two-Host Portals

Usage

Demonstration, proof of concept, development, or testing environment.

Configuration Worksheet

Host	Portal	Image Service	Automation Service	ALI API Service	Administrative Portal	Database Server	Document Repository Service	Search	Collaboration	Publisher	Studio	Analytics
Host1	X	X	X		X							
Host2				X		X	X	X	X	X	X	X

Risks and Mitigation

A two-machine configuration might be suitable for user acceptance testing for 100-200 users, where system availability is not a priority. If you plan to deploy a two-host portal for a small number of users, consider the following risks and mitigating factors.

Risks	Suggestions
The Automation Service might draw CPU and memory from the Portal component during crawls, maintenance operations, and user synchronization.	• Schedule the Automation Service to run its jobs at night.
Collaboration, Publisher, Studio, Search, Analytics, and the database compete for memory, slowing performance.	• Do not index a large number of documents, host many Studio databases, or publish a lot of content. • Increase the amount of RAM on the second machine. • Set all components to use a smaller memory footprint. • If performance is still unacceptable, move the database to a separate server and follow the three-machine configuration option.

Three-Host Portals

Usage

Minimum host requirements for a production portal.

If you plan to deploy Publisher, Collaboration, Studio, or Analytics, BEA recommends a three-host portal only for demonstration, proof of concept, development, or testing environment.

Configuration Optimized for a Large Amount of Content

Host	Portal	Image Service	Automation Service	ALI API Service	Administrative Portal	Database Server	Document Repository Service	Search	Collaboration	Publisher	Studio	Analytics
Host1	X	X		X	X				X	X	X	
Host2			X					X				X
Host3						X	X					

Caution: If you deploy on Oracle databases, install the database servers on a host separate from all other portal components. For a three-host deployment, follow the template for a two-host deployment and install only the database servers on the third host.

Configuration Optimized for a Large Number of Users

Host	Portal	Image Service	Automation Service	ALI API Service	Administrative Portal	Database Server	Document Repository Service	Search	Collaboration	Publisher	Studio	Analytics
Host1	X	X			X							
Host2				X					X	X	X	X
Host3			X			X	X	X				

Four-Host Portals

Usage

Minimum host requirements for a production portal that includes Publisher, Collaboration, Studio, or Analytics.

Configuration Optimized Equally for Content and Users

Host	Portal	Image Service	Automation Service	ALI API Service	Administrative Portal	Database Server	Document Repository Service	Search	Collaboration	Publisher	Studio	Analytics
Host1	X	X			X							
Host2				X					X	X	X	X
Host3			X					X				
Host4						X	X					

Configuration Optimized for a Large Amount of Content

Host	Portal	Image Service	Automation Service	ALI API Service	Administrative Portal	Database Server	Document Repository Service	Search	Collaboration	Publisher	Studio	Analytics
Host1	X	X		X	X				X	X	X	X
Host2			X									
Host3								X				
Host4						X	X					

Configuration Optimized for a Large Number of Users

Host	Portal	Image Service	Automation Service	ALI API Service	Administrative Portal	Database Server	Document Repository Service	Search	Collaboration	Publisher	Studio	Analytics
Host1	X	X			X							
Host2	X	X		X								
Host3									X	X	X	X
Host4			X			X	X	X				

APPENDIX F

Component-Host Worksheets

This appendix provides blank component-host worksheets you can use to record host assignments and host specifications for portal components.

Development Environment

Host	CPU	RAM	Portal	Image Service	Automation Service	ALI API Service	Administrative Portal	Database Server	Document Repository Service	Search	Collaboration	Publisher	Studio	Analytics	Remote Server - IS	Remote Server - CS	Remote Server - SS	Remote Server - Portlets
Host1																		
Host2																		
Host3																		
Host4																		
Host5																		
Host6																		
Host7																		
Host8																		
Host9																		
Host10																		

Production Environment

Host	CPU	RAM	Portal	Image Service	Automation Service	ALI API Service	Administrative Portal	Database Server	Document Repository Service	Search	Collaboration	Publisher	Studio	Analytics	Remote Server - IS	Remote Server - CS	Remote Server - SS	Remote Server - Portlets
Host1																		
Host2																		
Host3																		
Host4																		
Host5																		
Host6																		
Host7																		
Host8																		
Host9																		
Host10																		