BEA AquaLogic® User Interaction

Deployment Maintenance Guide

Document Revised: December 19 2007

Contents

A. Java Virtual Machine Configuration

CHAPTER 1

Welcome

This book describes strategies for maintaining an AquaLogic User Interacton deployment.

For products and versions covered by this book, see the section Products Covered by the Deployment Guide in the *Deployment Overview* book.

How to Use This Book

Audience

This guide is written to provide guidance to people responsible for maintaining the AquaLogic User Interaction system. Access to resources with strong knowledge of the platform operating system, database, web and application servers, and any other third-party software is recommended.

Organization

This guide includes the following chapters:

- This chapter provides information on how to use this guide and describes general resources available to assist in the AquaLogic User Interaction deployment.
- Chapter 2, "Developing a Production Maintenance Plan" provides an overview of portal maintenance tasks and tools.

- Chapter 3, "Performance Tuning" details the process of tuning application servers and standalone AquaLogic User Interaction components to the needs of your AquaLogic User Interaction deployment.
- Appendix A, "Java Virtual Machine Configuration" describes how to adjust JVM memory parameters and turn garbage collection logging on and off.

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 Guides	These guides describe the prerequisites (such as required software) and procedures for installing or upgrading the various AquaLogic User Interaction componets. These guides are available under the appropriate product sections on edocs.bea.com.

Table 1-2 BEA Documentation and Resources

Resource	Description
Administrator Guides	These guides describe how to manage and maintain the various AquaLogic User Interaction components. These guides are available under the appropriate product sections on edocs.bea.com.
Release Notes	The release notes provide information about new features, issues addressed, and known issues in the release of various AquaLogic User Interaction products. They are available on edocs.bea.com.
Online Help	The online help is written for all levels of AquaLogic User Interaction users. It describes the user interface for AquaLogic User Interaction components and gives detailed instructions for completing tasks in AquaLogic User Interaction products. To access online help, click the help icon.
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.

Table 1-2 BEA Documentation and Resources

Resource	Description
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.
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 Australia/NZ +61 2.9923.4030 Korea +82 27676 888 Singapore +1 800.1811.202

CHAPTER 2

Developing a Production Maintenance Plan

This chapter provides an overview of AquaLogic User Interaction maintenance tasks and tools.

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

This chapter includes the following topics:

Periodic Tasks

The following table provides suggestions for periodic tasks that you should consider as part of your production system maintenance plan.

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 servers 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 servers, new hard disk, and so on).
Ad Hoc	Install AquaLogic Interaction patches. Install server patches from critical third-party software providers, such as operating system and anti-virus software.

Monitoring AquaLogic User Interaction 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 AquaLogic User Interaction application that resides on a remote host, provided the both the remote host and the counter monitoring system 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.
- Use the Windows Perfmon utility to view portal counter data, if you use a Windows system.

For detailed information on the Counter Monitoring System, see the *Administrator Guide for AquaLogic Interaction*.

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 Interaction Analytics is an advanced usage tracking and analytics tool designed exclusively for AquaLogic Interaction. 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: Correlates 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.

Analytics includes the following reports that you can customize by setting filtering, grouping, and presentation options.

Report	Description	Features
Community Traffic	Displays traffic information for each community in the portal.	Displays traffic 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.	Calculates response time 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.	Displays traffic 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.	Displays the following four figures to help explain user inception and activity. • Total user accounts in the portal. • Added (new) user accounts created in the portal during a given date range. • Active users defined by activity during a given date range. • 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.	Calculate visit durations 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.	Can display these statistics 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:

AquaLogic Interaction Logging Utilities

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

- AquaLogic Interaction Logging Spy. Previously called PTSpy, this utility is the primary log message receiver for the OpenLog Framework. In addition to displaying log messages from the portal and other AquaLogic User Interaction products and services, AquaLogic Interaction Logging Spy provides fine-grained filtering, viewing of saved log files, highlighting of errors, and the searching and sorting of log messages.
- AquaLogic Interaction 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 provide output in other ways, such as sending log messages to an e-mail system.

- AquaLogic Interaction 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, AquaLogic Interaction 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)
```

CHAPTER 3

Performance Tuning

This chapter details the process of tuning application servers and standalone AquaLogic User Interaction products to the needs of your AquaLogic User Interaction deployment.

The standalone AquaLogic User Interaction products are:

- AquaLogic Interaction Analytics
- AquaLogic Interaction Automation Server
- AquaLogic Interaction Collaboration
- AquaLogic Interaction Document Repository
- AquaLogic Interaction Notification
- AquaLogic Interaction PTUpload
- AquaLogic Interaction Publisher
- AquaLogic Interaction Studio
- AquaLogic Interaction Workflow

Tuning a Java Application Server or Standalone AquaLogic User Interaction Product

For Java application servers and standalone AquaLogic User Interaction products, tuning is a matter of adjusting various Java Virtual Machine (JVM) settings to optimize garbage collection.

Sun provides a comprehensive document on this subject, *Tuning Garbage Collection with the 1.4.2 Java[tm] Virtual Machine*, which you can find at http://java.sun.com/docs/hotspot/gc1.4.2/**.**

The following provides a brief background on the garbage collection process and a detailed, AquaLogic User Interaction focused process for tuning JVM garbage collection.

Garbage Collection Concepts

Garbage collection is the process the JVM undergoes to remove unused objects from memory. The following description of the garbage collection process is simplified for the purpose of this guide.

The JVM stores objects in two sections of the heap: the *young generation* and the *tenured generation***.** The young generation is where objects are first created and provides the quickest, least CPU intensive access to objects. When the young generation fills, older active objects are transfered to the larger tenured generation. Objects in the tenured generation are more CPU intensive to access than those in the young generation.

The JVM undertakes two types of garbage collection. The *minor collection* runs when the young generation fills. It clears garbage objects and copies surviving objects to the tenured generation. The *major collection* runs when the tenured generation fills. The minor collection is signficantly less CPU-intensive than the major collection.

Garbage Collection Logs

In order to analyze garbage collection impact on application server performance, a garbage collection log needs to be collected. The process is:

1. Enable garbage collection logging in the JVM. This is done in different places for each of the supported application servers and standalone AquaLogic User Interaction products. For details on enabling garbage collection logging, see Appendix A, “Java Virtual Machine Configuration.”
2. Restart the application service to start logging garbage collection memory usage.
3. Run the until the problem occurs. If the problem is continuous, collect approximately a 24 hours of data.

Note: Every time the application server is restarted, the garbage collection log is overwritten. It is important to turn off automatic restarting of services, especially if you are investigating an issue that yields a server crash.

Analyzing the Garbage Collection Log

Tagtraum industries (http://tagtraum.com) provides a free utility, *gcviewer*, for analyzing garbage collection logs generated by the JVM. Load the garbage collection log into gcviewer and determine which issue is occuring based on the descriptions in Table 3-1.

Table 3-1 Garbage Collection Performance Issues

Issue	Symptoms in Garbage Collection Log	Impact of the Issue
Insufficient total (heap) memory allocated	Memory usage trends upwards and reaches the top of the total memory allocated.	Reduces the performance or potentially crashes the AquaLogic User Interaction product.
Excessive total (heap) memory allocated	Memory usage peaks much lower than total memory allocated.	Can cause a slowdown across all applications on the server. The application server or AquaLogic User Interaction product is taking up too much of the system memory.
Insufficient young generation memory allocated	Sawtoothed memory usage.	Reduces the performance of the AquaLogic User Interaction product. This represents excessive minor garbage collector runs, which increases the number of objects in the tenured generation. Objects in the tenured generation are more resource intensive when called.

Resolving Garbage Collection Performance Issues

Resolving the issues described in Analyzing the Garbage Collection Log is a matter of adjusting the JVM memory settings and reanalyzing the garbage collection log. Table 3-2 shows what memory settings to adjust for each issue. For details on how to adjust these settings for each

supported application server and standalone AquaLogic User Interaction product, see Appendix A, "Java Virtual Machine Configuration."

Table 3-2 Garbage Collection Performance Issue Resolution

Issue	Resolution	JVM Memory Parameter
Insufficient total (heap) memory allocated	Increase total heap memory allocation until memory usage stays reasonably below total memory.	Increase `-Xmx`
Excessive total (heap) memory allocated	Decrease total heap memory allocation until memory usage is reasonably, but not excessively, below total memory.	Decrease `-Xmx`
Insufficient young generation memory allocated	Increase young generation memory allocation until the memory usage trend is horizontal.	Adjust `-XX:NewRatio`

APPENDIX A

Java Virtual Machine Configuration

This appendix describes how to adjust JVM memory parameters and turn garbage collection logging on and off. Instructions below cover applications supported by AquaLogic Interaction and those standalone AquaLogic User Interaction products.

The standalone AquaLogic User Interaction products are:

- AquaLogic Interaction Analytics
- AquaLogic Interaction Automation Server
- AquaLogic Interaction Collaboration
- AquaLogic Interaction Document Repository
- AquaLogic Interaction Notification
- AquaLogic Interaction PTUpload
- AquaLogic Interaction Publisher
- AquaLogic Interaction Studio
- AquaLogic Interaction Workflow

Java Memory Switches

The following are Java memory switches used to tune JVM garbage collection. Use these switches in conjunction with the instructions specific to your application server or AquaLogic User Interaction product.

- `-Xloggc:<path/filename>`

  This switch turns on garbage collection logging for the JVM. Replace **`<path/filename>`** with the location where the garbage collection log should be generated.

- `-Xms` and `-Xmx`

  These switches set the minimum (`-Xms`) and maximum (`-Xmx`) heap size for the JVM. The JVM adjusts heap size based on object usage and bounded by these two switches. Setting these switches to the same value increases predictability by removing the ability of the JVM to adjust the heap size.

  Caution: Fixing the heap size to a specific value requires special attention to memory tuning.

- `-XX:NewRatio`

  This switch sets the ration of the young generation to the tenured generation. For example

  `-XX:NewRatio=3`

  would mean that the tenured generation is 3x the size of the young generation, or, in other words, the young generation is one quarter of the heap and the tenured generation is three-quarters of the heap.

Application Servers

Tomcat 4.1.x

To update Java options for Tomcat 4.1.x on Windows:

1. Edit **setenv.bat** in the **<TOMCAT_HOME>\bin** directory.
2. Update the JVM memory switches in the line

   `set JAVA_OPTS=`

3. Save **setenv.bat**. Reinstall the Tomcat service.

   Note: The **setenv.bat** parameters are only loaded into the service settings during service creation.

Tomcat 5.x

To update Java options for Tomcat 5.x on Windows:

1. Run **<TOMCAT_HOME>\tomcat5w.exe**
2. Click the **Java** tab.
3. Update the Java memory switches in the **Java Options:** box.
4. Click **OK**. Restart the Tomcat service.

WebLogic 8.1

To update Java options for WebLogic 8.1:

1. Edit **installservice.cmd** in **<BEA_HOME>/user/projects/domains/<domain name>**.
2. Add JVM memory switches to the line

   ```
   set MEM_ARGS=
   ```

3. Reinstall the WebLogic service.

Note: The `MEM_ARGS` parameter can also be updated in the startup script for the WebLogic domain.

ALUI Standalone Products

Analytics

To update Java options for the Analytics JVM:

1. Edit **wrapper.conf** in **<PT_HOME>/ptanalytics/<ver>/settings/config**.
2. Add or modify parameters in the section **Additional -D Java Properties**.

   Note: Java parameter numbers must be continuous and incremental, and are set in both **wrapper_base.conf** and **wrapper.conf**. Check both files to ensure added parameters use the next number in sequence.

Restart the Analytics service.

Automation Service

To update Java options for the Automation Service JVM:

1. Edit **wrapper.conf** in **<PT_HOME/ptportal/<ver>/settings/config**.

2. Add or modify parameters in the section **Additional -D Java Properties**.

   Note: Java parameter numbers must be continuous and incremental, and are set in both **wrapper_base.conf** and **wrapper.conf**. Check both files to ensure added parameters use the next number in sequence.

Restart the Automation service.

Collaboration

To update Java options for the Collaboration JVM:

1. Edit **wrapper.conf** in **<PT_HOME>/ptcollab/<ver>/settings/config**.
2. Add or modify parameters in the section **Additional -D Java Properties**.

   Note: Java parameter numbers must be continuous and incremental, and are set in both **wrapper_base.conf** and **wrapper.conf**. Check both files to ensure added parameters use the next number in sequence.

3. Restart the Collaboration service.

Document Repository

To update Java options for the Document Repository JVM:

1. Edit **wrapper.conf** in **<PT_HOME>/ptdr/<ver>/settings/config**.
2. Add or modify parameters in the section **Additional -D Java Properties**.

   Note: Java parameter numbers must be continuous and incremental, and are set in both **wrapper_base.conf** and **wrapper.conf**. Check both files to ensure added parameters use the next number in sequence.

Restart the Document Repository service.

Notification

To update Java options for the Notification JVM:

1. Edit **wrapper.conf** in **<PT_HOME>/ptnotification/<ver>/settings/config**.
2. Add or modify parameters in the section **Additional -D Java Properties**.

Note: Java parameter numbers must be continuous and incremental, and are set in both **wrapper_base.conf** and **wrapper.conf**. Check both files to ensure added parameters use the next number in sequence.

Restart the Notification service.

PTUpload

To update Java options for the PTUpload JVM:

1. Edit **wrapper.conf** in **<PT_HOME>/ptupload/<ver>/settings/config**.
2. Add or modify parameters in the section **Additional -D Java Properties**.

   Note: Java parameter numbers must be continuous and incremental, and are set in both **wrapper_base.conf** and **wrapper.conf**. Check both files to ensure added parameters use the next number in sequence.

Restart the PTUpload service.

Publisher

To update Java options for the Publisher JVM:

1. Edit **service.conf** in **<PT_HOME>/ptcs/<ver>/settings/config**.
2. Add a new parameter or modify existing parameters in the section **Java Additional Parameters**.

   For example, locate

   ```
   # Java Additional Parameters
   wrapper.java.additional.1=-Dprogram.name=cswfserver
   wrapper.java.additional.2=-Djava.awt.headless=true
   wrapper.java.additional.3=-Dplumtree.container.home=./../../../../../pt
   cs/6.2/container
   wrapper.java.additional.4=-Dplumtree.container.logs=./../../../../../pt
   cs/6.2/logs
   wrapper.java.additional.5=-Dorg.jboss.net.protocol.file.decodeFilePaths
   =true
   ```

   and add the garbage collection logging parameter

   ```
   # Java Additional Parameters
   wrapper.java.additional.1=-Dprogram.name=cswfserver
   wrapper.java.additional.2=-Djava.awt.headless=true
   wrapper.java.additional.3=-Dplumtree.container.home=./../../../../../pt
   ```

```
cs/6.2/container
wrapper.java.additional.4=-Dplumtree.container.logs=./../../../../../pt
cs/6.2/logs
wrapper.java.additional.5=-Dorg.jboss.net.protocol.file.decodeFilePaths
=true
wrapper.java.additional.6=-Xloggc:c:\publishergclog
```

3. Restart the Publisher service.

Note: Publisher and Workflow run on the same JVM. Garbage collection logging and memory tunings are for both services.

Studio

To update Java options for the Studio JVM:

1. Edit **wrapper.conf** in **<PT_HOME>/ptstudio/<ver>/settings/config**.
2. Add or modify parameters in the section **Additional -D Java Properties**.

   Note: Java parameter numbers must be continuous and incremental, and are set in both **wrapper_base.conf** and **wrapper.conf**. Check both files to ensure added parameters use the next number in sequence.

Restart the Studio service.

Workflow

Workflow runs on the same JVM as Publisher and does not require seperate tuning.