



BEA AquaLogic Service Bus™

User Guide

Copyright

Copyright © 1995-2006 BEA Systems, Inc. All Rights Reserved.

Restricted Rights Legend

This software is protected by copyright, and may be protected by patent laws. No copying or other use of this software is permitted unless you have entered into a license agreement with BEA authorizing such use. This document is protected by copyright and may not be copied photocopied, reproduced, translated, or reduced to any electronic medium or machine readable form, in whole or in part, without prior consent, in writing, from BEA Systems, Inc.

Information in this document is subject to change without notice and does not represent a commitment on the part of BEA Systems. THE DOCUMENTATION IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. FURTHER, BEA SYSTEMS DOES NOT WARRANT, GUARANTEE, OR MAKE ANY REPRESENTATIONS REGARDING THE USE, OR THE RESULTS OF THE USE, OF THE DOCUMENT IN TERMS OF CORRECTNESS, ACCURACY, RELIABILITY, OR OTHERWISE.

Trademarks and Service Marks

Copyright © 1995-2006 BEA Systems, Inc. All Rights Reserved. BEA, BEA JRocket, BEA WebLogic Portal, BEA WebLogic Server, BEA WebLogic Workshop, Built on BEA, Jolt, JoltBeans, SteelThread, Top End, Tuxedo, and WebLogic are registered trademarks of BEA Systems, Inc. BEA AquaLogic, BEA AquaLogic Data Services Platform, BEA AquaLogic Enterprise Security, BEA AquaLogic Interaction, BEA AquaLogic Interaction Analytics, BEA AquaLogic Interaction Collaboration, BEA AquaLogic Interaction Content Services, BEA AquaLogic Interaction Data Services, BEA AquaLogic Interaction Integration Services, BEA AquaLogic Interaction Process, BEA AquaLogic Interaction Publisher, BEA AquaLogic Interaction Studio, BEA AquaLogic Service Bus, BEA AquaLogic Service Registry, BEA Builder, BEA Campaign Manager for WebLogic, BEA eLink, BEA Kodo, BEA Liquid Data for WebLogic, BEA Manager, BEA MessageQ, BEA SALT, BEA Service Architecture Leveraging Tuxedo, BEA WebLogic Commerce Server, BEA WebLogic Communications Platform, BEA WebLogic Enterprise, BEA WebLogic Enterprise Platform, BEA WebLogic Enterprise Security, BEA WebLogic Express, BEA WebLogic Integration, BEA WebLogic Java Adapter for Mainframe, BEA WebLogic JDriver, BEA WebLogic Log Central, BEA WebLogic Mobility Server, BEA WebLogic Network Gatekeeper, BEA WebLogic Personalization Server, BEA WebLogic Personal Messaging API, BEA WebLogic Platform, BEA WebLogic Portlets for Groupware Integration, BEA WebLogic Real Time, BEA WebLogic RFID Compliance Express, BEA WebLogic RFID Edge Server, BEA WebLogic RFID Enterprise Server, BEA WebLogic Server Process Edition, BEA WebLogic SIP Server, BEA WebLogic WorkGroup Edition, BEA Workshop for WebLogic Platform, BEA Workshop JSP, BEA Workshop JSP Editor, BEA Workshop Struts, BEA Workshop Studio, Dev2Dev, Liquid Computing, and Think Liquid are trademarks of BEA Systems, Inc. Accelerated Knowledge Transfer, AKT, BEA Mission Critical Support, BEA Mission Critical Support Continuum, and BEA SOA Self Assessment are service marks of BEA Systems, Inc.

All other names and marks are property of their respective owners.

Contents

Introduction to AquaLogic Service Bus

| | |
|---|-----|
| Overview of AquaLogic Service Bus | 1-1 |
| Document Scope and Audience | 1-2 |
| Document Organization | 1-2 |

Modeling Message Flow in AquaLogic Service Bus

| | |
|---|------|
| About AquaLogic Service Bus Message Flow | 2-2 |
| Building a Message Flow | 2-3 |
| Message Execution | 2-5 |
| Pipelines | 2-6 |
| Branching in Message Flows | 2-9 |
| Operational Branching | 2-9 |
| Conditional Branching | 2-9 |
| Performing Transformations | 2-11 |
| Transformations and Publish Actions | 2-12 |
| Transformations and Route Nodes | 2-12 |
| Configuring Single and Multiple Stages in Pipelines | 2-13 |
| Communication | 2-14 |
| Flow Control | 2-15 |
| Message Processing | 2-16 |
| Reporting | 2-16 |
| Using Multiple Stages | 2-17 |

| | |
|---|------|
| Handling Errors | 2-17 |
| Generating the Error Message, Reporting, and Replying | 2-18 |
| Example of Action Configuration in Error Handlers | 2-19 |
| Selecting a Service Type | 2-21 |
| Using a WSDL to Define a Service | 2-23 |
| SOAP Document Wrapped Web Services | 2-24 |
| SOAP Document Style Web Services | 2-24 |
| SOAP RPC Web Services | 2-26 |
| Binding a Service to a WSDL Port Instead of to a Binding | 2-30 |
| Using Any SOAP or Any XML Service Types | 2-31 |
| Using the Messaging Service Type | 2-31 |
| Viewing Resource Details | 2-31 |
| Using Dynamic Routing | 2-32 |
| Sample XML File | 2-33 |
| Creating an XQuery Resource From the Sample XML | 2-34 |
| Creating and Configuring the Proxy Service to Implement Dynamic Routing | 2-35 |
| Understanding Message Context | 2-36 |
| Message Context Components | 2-37 |
| Guidelines for Viewing and Altering Message Context | 2-38 |
| Copying JMS Properties From Inbound to Outbound | 2-39 |
| Working with Variable Structures | 2-40 |
| Using the Inline XQuery Expression Editor | 2-40 |
| Using Variable Structures | 2-41 |
| Creating Variable Structure Mappings | 2-42 |
| Sample WSDL | 2-43 |
| Creating the Resources You Need for the Examples | 2-44 |
| Example 1: Selecting a Predefined Variable Structure | 2-47 |
| Example 2: Creating a Variable Structure That Maps a Variable to a Type | 2-48 |

| | |
|--|------|
| Example 3: Creating a Variable Structure that Maps a Variable to an Element | 2-50 |
| Example 4: Creating a Variable Structure That Maps a Variable to a Child Element | 2-51 |
| Quality of Service | 2-55 |
| Delivery Guarantees | 2-55 |
| Overriding the Default Element Attribute | 2-57 |
| Delivery Guarantee Rules | 2-58 |
| Threading Model | 2-60 |
| Splitting Proxy Services | 2-60 |
| Outbound Message Retries | 2-61 |
| Content Types, JMS Type, and Encoding | 2-61 |
| Throttling Pattern | 2-62 |
| WS-I Compliance | 2-62 |
| WS-I Compliance Checks | 2-64 |

Monitoring

| | |
|------------------------------------|-----|
| Monitoring Scenarios | 3-1 |
| Operational Health | 3-2 |
| Monitoring Alerts | 3-3 |
| Monitoring Statistics | 3-3 |
| Verifying Service Level Agreements | 3-4 |
| Pipeline Alert Action | 3-4 |
| Alert Destination | 3-5 |
| AquaLogic Service Bus Console | 3-5 |
| E-mail Alert Destination | 3-5 |
| SNMP Traps | 3-6 |
| JMS | 3-8 |
| Reporting | 3-8 |

| | |
|--|------|
| About Monitoring | 3-8 |
| Aggregation Interval | 3-9 |
| Monitoring Architecture | 3-10 |
| Monitoring Services..... | 3-11 |
| Refresh Rate of Monitored Information..... | 3-11 |
| Dashboard | 3-12 |
| Service Summary | 3-14 |
| About the Service Summary | 3-14 |
| Service Monitoring Summary | 3-15 |
| Service Monitoring Details | 3-17 |
| Server Summary | 3-21 |
| About the Server Summary | 3-22 |
| Log Summary..... | 3-22 |
| Server Summary..... | 3-25 |
| Server Details..... | 3-27 |
| Alert Summary | 3-29 |
| About the Alert Summary | 3-29 |
| Pipeline Alerts..... | 3-29 |
| Service Level Agreement Alerts (SLA) | 3-30 |
| System Alerts History | 3-32 |
| System Alert Details | 3-35 |
| View Alert Rule Details | 3-36 |
| Alert Rules..... | 3-38 |
| About Alert Rules | 3-38 |
| Some Uses for Alerts | 3-39 |
| Understanding Alert Rules..... | 3-39 |
| Statistics Associated With Different Resources | 3-42 |
| SERVICE | 3-42 |

| | |
|-------------------------------------|------|
| FLOW_COMPONENT | 3-43 |
| WEBSERVICE_OPERATION..... | 3-43 |
| Auditing..... | 3-44 |
| Configuration Change Auditing | 3-44 |
| Runtime Auditing of Messages | 3-44 |
| Security Auditing | 3-44 |

Using the Test Console

| | |
|--|------|
| Features | 4-2 |
| Prerequisites..... | 4-2 |
| Testing Proxy Services | 4-3 |
| Direct Calls..... | 4-3 |
| Indirect Calls | 4-4 |
| HTTP Requests..... | 4-4 |
| Testing Business Services | 4-5 |
| Transport Security | 4-5 |
| Recommended Approaches to Testing Proxy and Business Services..... | 4-6 |
| Tracing Proxy Services Using the Test Console..... | 4-7 |
| Example: Testing and Tracing a Proxy Service | 4-8 |
| Testing Resources | 4-12 |
| MFL | 4-12 |
| XSLT | 4-14 |
| XQuery..... | 4-14 |
| Performing In-line XQuery Testing..... | 4-16 |
| Testing Services With Web Service Security | 4-16 |
| Test Console Transport Settings | 4-21 |

Reporting

| | |
|--|------|
| Reporting Scenarios | 5-2 |
| Message Tracking | 5-2 |
| Search for a Particular Message. | 5-2 |
| Logging for Regulatory Auditing | 5-2 |
| Alert Reporting Provider | 5-3 |
| Reporting Framework | 5-3 |
| JMS Reporting Provider | 5-5 |
| About the JMS Reporting Provider | 5-5 |
| How to Enable Message Reporting | 5-6 |
| Using the Reporting Module | 5-8 |
| Summary of Messages | 5-9 |
| View Message Details | 5-10 |
| Purging Messages. | 5-13 |
| Configuring a Database for the JMS Reporting Provider Store. | 5-14 |
| Configuring a Database in a Development Environment | 5-14 |
| Configuring a Database for Production | 5-15 |
| Removing, Stopping, or Untargeting a Reporting Provider | 5-15 |
| Stopping a Reporting Provider when the Server is Running | 5-16 |
| Untargeting a Reporting Provider when the Server is Running. | 5-18 |
| Untargeting the JMS Reporting Provider—Server Not Running | 5-19 |

Tracing

| | |
|----------------------------|-----|
| To Enable Tracing. | 6-1 |
|----------------------------|-----|

UDDI

| | |
|---|-----|
| Overview of BEA AquaLogic Service Bus and UDDI. | 7-1 |
| Basic Concepts of the UDDI Specification | 7-3 |

| | |
|---|------|
| Benefits of Using a UDDI Registry with AquaLogic Service Bus | 7-3 |
| Introduction to UDDI Entities | 7-4 |
| Prerequisites | 7-5 |
| Certification | 7-5 |
| Features | 7-5 |
| What is the BEA AquaLogic Service Registry? | 7-6 |
| Sample Business Scenario for AquaLogic Service Bus and UDDI | 7-6 |
| Basic Proxy Service Communication with a UDDI Registry | 7-7 |
| Cross-Domain Deployment in AquaLogic Service Bus | 7-7 |
| Using AquaLogic Service Bus and UDDI | 7-8 |
| UDDI Workflow | 7-8 |
| Configuring a Registry | 7-9 |
| Publishing a Proxy Service to a UDDI Registry | 7-10 |
| Using Auto-Publish | 7-11 |
| Importing a Service from a Registry | 7-11 |
| Using Auto-Import | 7-13 |
| Mapping AquaLogic Service Bus Proxy Services to UDDI Entities | 7-15 |
| UDDI Mapping Details for an AquaLogic Service Bus Proxy Service | 7-17 |
| Transport Attributes | 7-20 |
| Service Type Attributes | 7-22 |
| Canonical tModels Supporting AquaLogic Service Bus Services | 7-23 |
| Example | 7-25 |

Transports

| | |
|---|-----|
| E-mail | 8-1 |
| Configuring Proxy Services using E-mail Transport Protocol | 8-2 |
| Configuring Business Services using E-mail Transport Protocol | 8-3 |
| EJB | 8-4 |

| | |
|--|------|
| File | 8-4 |
| Configuring Proxy Services using File Transport Protocol | 8-5 |
| Configuring Business Services using File Transport Protocol | 8-6 |
| FTP | 8-6 |
| Configuring Proxy Services using FTP Transport Protocol | 8-7 |
| Configuring Business Services using FTP Transport Protocol | 8-8 |
| HTTP | 8-9 |
| Configuring Proxy Services using HTTP Transport Protocol | 8-9 |
| Configuring Business Services using HTTP Transport Protocol | 8-10 |
| HTTP(S) | 8-11 |
| Configuring Proxy Services using HTTP(S) Transport Protocol | 8-11 |
| Configuring Business Services using HTTP(S) Transport Protocol | 8-12 |
| JMS | 8-13 |
| Configuring Proxy Services using JMS Transport Protocol | 8-13 |
| Configuring Business Services using JMS Transport Protocol | 8-16 |
| Local | 8-18 |
| Tuxedo | 8-18 |

EJB Transport

| | |
|--|-----|
| Introduction | 9-1 |
| Invoking EJBs from AquaLogic Service Bus | 9-3 |
| Register a JNDI Provider Resource | 9-3 |
| Adding a JNDI Provider | 9-4 |
| Register an EJB Client JAR Resource | 9-4 |
| Adding a Client JAR | 9-4 |
| Create a Service Account (Optional) | 9-4 |
| Locate an EJB in the JNDI Tree | 9-5 |
| Create an EJB Business Service | 9-5 |

| | |
|--|------|
| General Configuration | 9-5 |
| EJB Transport-Specific Configuration | 9-7 |
| EJB Business Service Interface Configuration | 9-9 |
| Invoking EJB Business Services | 9-11 |
| Exposing EJBs as Web Services | 9-11 |
| Advanced Topics | 9-11 |
| Transaction Processing, Retries, and Errors Handling | 9-12 |
| Transactions | 9-12 |
| Retries and Failover | 9-13 |
| Error Handling | 9-14 |
| Supported Types and Converter Class | 9-14 |
| Converter Classes | 9-15 |
| Troubleshooting | 9-15 |

Local Transport

| | |
|--|------|
| Introduction | 10-1 |
| Features and Characteristics of Local Transport Proxy Services | 10-2 |
| Usage of Local Transport Proxy Services | 10-3 |
| Limitations | 10-4 |

Extensibility Using Java Callouts and POJOs

| | |
|----------------------------|------|
| Usage Guidelines | 11-1 |
| Best Practices | 11-2 |

Tuning AquaLogic Service Bus

Debugging AquaLogic Service Bus

XQuery Implementation

| | |
|---|-----|
| Supported Function Extensions from AquaLogic Data Services Platform | C-1 |
|---|-----|

Function Extensions from AquaLogic Service Bus C-2

Introduction to AquaLogic Service Bus

BEA AquaLogic Service Bus is part of the BEA AquaLogic™ family of Service Infrastructure Products. AquaLogic Service Bus manages the routing and transformation of messages in an enterprise system. Combining these functions with its monitoring and administration capability, AquaLogic Service Bus provides a unified software product for implementing and deploying your Service-Oriented Architecture (SOA).

The following sections provide an overview of AquaLogic Service Bus and of this document:

- [“Overview of AquaLogic Service Bus”](#) on page 1-1
- [“Document Scope and Audience”](#) on page 1-2
- [“Document Organization”](#) on page 1-2

Overview of AquaLogic Service Bus

AquaLogic Service Bus is a configuration-based, policy-driven Enterprise Service Bus (ESB). From the AquaLogic Service Bus Console, you can monitor your services, servers, and operational tasks. You configure proxy and business services, set up security, manage resources, and capture data for tracking or regulatory auditing. The AquaLogic Service Bus Console enables you to respond rapidly and effectively to changes in your service-oriented environment.

AquaLogic Service Bus relies on WebLogic Server run-time facilities. It leverages WebLogic Server capabilities to deliver functionality that is highly available, scalable, and reliable.

Document Scope and Audience

This guide provides detailed information on using and configuring AquaLogic Service Bus. It is intended for those responsible for messaging and SOA, such as enterprise architects, operations specialists, security architects and developers, application architects and developers, server and application administrators, and support engineers.

While sometimes providing procedural information, this guide does not provide detailed information on how to configure resources using the AquaLogic Service Bus Console. For more information on using the AquaLogic Service Bus Console, see [Using the AquaLogic Service Bus Console](#).

Document Organization

This document includes the following topics:

- [Modeling Message Flow in AquaLogic Service Bus](#)—Guidelines for modeling message flows in AquaLogic Service Bus. A message flow defines the implementation of a proxy service, which is the AquaLogic Service Bus definition of an intermediary Web services that is hosted locally on AquaLogic Service Bus. In AquaLogic Service Bus, service clients exchange messages with an intermediary proxy service rather than directly with a business service.
- [Monitoring](#): Monitoring and collecting run-time information for systems operations and business auditing purposes. This chapter describes how you can monitor the health of the system, including the state of the services, servers, and Service Level Agreement (SLA) violations.
- [Using the Test Console](#): Using the test console to test proxy services and business services created and used AquaLogic Service Bus. You can also use the test console to test the resources created and used in AquaLogic Service Bus.
- [Reporting](#): Capturing message data for tracking messages or regulatory auditing. This section also contains information about setting up your own reporting provider; using the JMS reporting provider included with AquaLogic Service Bus; using the Reporting module in AquaLogic Service Bus Console; and configuring a reporting provider for data on alerts. Alerts contain information about SLA violations.
- [Tracing](#): Tracing messages without shutting down the server. This feature is useful in both a development and production environment. This feature allows you to troubleshoot and diagnose a message flow in one or more proxy services.

- [UDDI](#): Using Universal Description, Discovery and Integration (UDDI) registries with AquaLogic Service Bus. The UDDI protocol is one of the major building blocks required for successful Web services. UDDI provides a standard interoperable platform that enables enterprises and applications to find and use Web services over the Internet.
- [Transports](#): Transport protocols available in AquaLogic Service Bus.
- [EJB Transport](#): EJB Transport features and business services.
- [Local Transport](#): Local transport features and use cases.
- [Extensibility Using Java Callouts and POJOs](#): Guidelines for using the Java callout action with POJOs.
- [Tuning AquaLogic Service Bus](#): Optimizing the AquaLogic Service Bus performance in a production environment.
- [Debugging AquaLogic Service Bus](#): Enabling debugging for different modules in AquaLogic Service Bus.
- [XQuery Implementation](#): Valid extensions of BEA AquaLogic Data Services Platform functions and AquaLogic Service Bus functions.

Introduction to AquaLogic Service Bus

Modeling Message Flow in AquaLogic Service Bus

In BEA AquaLogic Service Bus, Message Flow defines the implementation of a proxy service. This section presents guidelines to follow when you model message flows. You configure AquaLogic Service Bus in the AquaLogic Service Bus Console, which is described in [Using the AquaLogic Service Bus Console](#).

The following sections provide information about message flow:

- [“About AquaLogic Service Bus Message Flow”](#) on page 2-2
- [“Pipelines”](#) on page 2-6
- [“Branching in Message Flows”](#) on page 2-9
- [“Performing Transformations”](#) on page 2-11
- [“Configuring Single and Multiple Stages in Pipelines”](#) on page 2-13
- [“Handling Errors”](#) on page 2-17
- [“Selecting a Service Type”](#) on page 2-21
- [“Using a WSDL to Define a Service”](#) on page 2-23
- [“Viewing Resource Details”](#) on page 2-31
- [“Using Dynamic Routing”](#) on page 2-32
- [“Understanding Message Context”](#) on page 2-36
- [“Working with Variable Structures”](#) on page 2-40

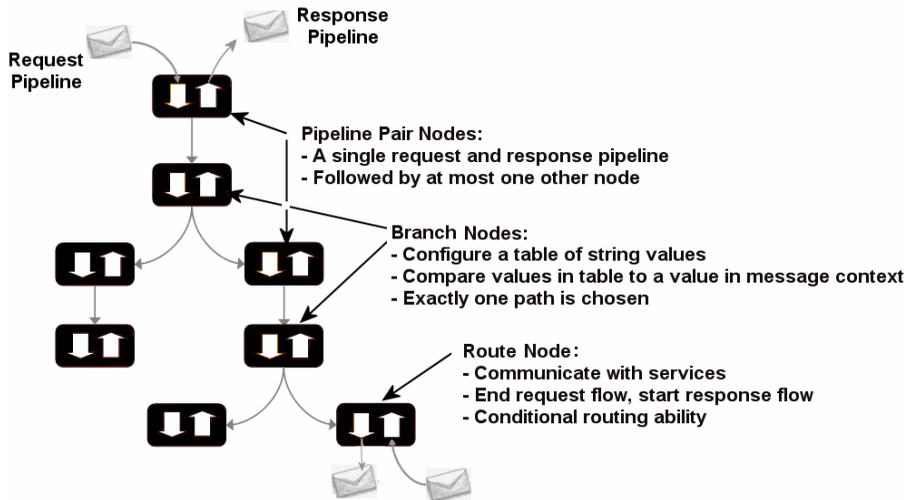
- “Quality of Service” on page 2-55
- “Content Types, JMS Type, and Encoding” on page 2-61
- “Throttling Pattern” on page 2-62
- “WS-I Compliance” on page 2-62

About AquaLogic Service Bus Message Flow

A message flow consists of the pipelines, branch nodes, and route nodes that together define the implementation of an AquaLogic Service Bus proxy service. A proxy service is an AquaLogic Service Bus definition of an intermediary Web Service that is hosted locally on AquaLogic Service Bus. Using the AquaLogic Service Bus Console, you can configure the logic for the manipulation of messages in proxy service message flow definitions. This logic includes such activities as transformation, publishing, and reporting—the logic is configured in individual actions within the message flow.

The following figure shows a high level view of the components of the message flow definition.

Figure 2-1 Components of Message Flow



This topic includes the following sections:

- [“Building a Message Flow” on page 2-3](#)
- [“Message Execution” on page 2-5](#)

Building a Message Flow

Any component can be at the root of a message flow. (For a description of the components, see [Table 2-1, “Message Flow Components,” on page 2-4](#)). One of the simplest of message flow designs is to have only a route node representing the entire flow. No restrictions exist on what two components you can chain together to create a message flow. For example, two pipeline pair nodes can be linked together without a branch node in between. In the case of branch nodes, each branch node can start with a different element. One branch can terminate with a route node, another can be followed by a pipeline pair, and yet another may have no descendant. In the latter case, a branch with no descendants means that at run time, when this branch is executed, response processing begins immediately. However, in general a message flow is likely to be designed in one of the following forms:

- In the case of non-operational services (services that are not based on WSDLs with operations), the flow likely consists of a single pipeline pair at the root followed by a route node.
- In the case of operational services, the flow likely consists of a single pipeline pair at the root, followed by a branch node based on an operation, with each branch consisting of a pipeline pair followed by a route node.

A message flow is constructed by linking together instances of the top-level components described in the following table. Subsequent sections in this topic describe the node types in more detail.

Table 2-1 Message Flow Components

| Node Type | Summary |
|--|---|
| <p>Pipeline Pair See “Pipelines” on page 2-6.</p> | <p>A pipeline pair combines a single request and a single response pipeline into one top-level element. A pipeline pair node can have only one direct descendant in the message flow. During request processing, only the request pipeline is executed when AquaLogic Service Bus processes a pipeline pair node. The execution path is reversed when AquaLogic Service Bus processes the response pipeline.</p> <p>For an example of a simple pipeline pair node, see Figure 2-3.</p> <p>To learn how to configure a pipeline pair node, see “Adding a Pipeline Pair Node” in Proxy Services: Message Flow in <i>Using the AquaLogic Service Bus Console</i>.</p> |
| <p>Branch See “Branching in Message Flows” on page 2-9.</p> | <p>A branch node allows processing to proceed along exactly one of several possible paths. Branching is driven by an XPath-based switch table. Each branch in the table specifies a condition (for example, <500) that is evaluated in order down the message flow against a single XPath expression (for example, <code>./ns:PurchaseOrder/ns:totalCost</code> on \$body). Whichever condition is satisfied first determines which branch is followed. If no branch condition is satisfied, then the default branch is followed. A branch node may have several descendants in the message flow: one for each branch, including the default branch.</p> <p>Note: It is highly recommended that you define a default branch whenever your message flow involves conditional branching.</p> <p>To learn how to add a branch node, see “Adding a Conditional Branch Node” in Proxy Services: Message Flow in <i>Using the AquaLogic Service Bus Console</i>.</p> <p>For information about working with the message context variables to design conditions, see Message Context in <i>Using the AquaLogic Service Bus Console</i>.</p> |

Table 2-1 Message Flow Components

| Node Type | Summary |
|--------------|---|
| Route | <p>A route node is used to perform request/response communication with another service. It represents the boundary between request and response processing for the proxy service. When the route node dispatches a request message, the request processing is considered complete. When the route node receives a response message, the response processing begins. The route node supports conditional routing as well as request and response transformations.</p> <p>Because a route node represents the boundary between request and response processing, it cannot have any descendants in the message flow.</p> <p>To learn how to add a route node, see Adding a Route Node in Proxy Services: Message Flow in the <i>Using the AquaLogic Service Bus Console</i>.</p> |

To create a message flow, see “Viewing and Changing Message Flow” in [Proxy Services: Message Flow](#) in *Using the AquaLogic Service Bus Console*.

Message Execution

The following table gives brief description of the components in a typical message flow

Table 2-2 Path Of a Message during a Message Flow

| Message Flow Node | What Happens During Message Processing? |
|----------------------------|--|
| Request Processing | Request processing begins at the root of the message flow. |
| Pipeline Pair | Executes the request pipeline only. |
| Branch | Evaluates the branch table and proceeds down the relevant branch. |
| Route | <p>Performs the route along with any request transformations.</p> <p>Note: In the message flow, regardless of whether routing takes place or not, the route node represents the change-over from processing a request to processing a response. At the route node, the direction of the message flow is reversed. If a request path does not have a route node, the response processing is initiated in the reverse direction without waiting for any response.</p> |
| Response Processing | Skips any branch nodes and continues with the node that preceded the branch. |

Table 2-2 Path Of a Message during a Message Flow

| Message Flow Node | What Happens During Message Processing? |
|--------------------------|--|
| Route | Executes any response transformations. See “Route” on page 2-5 for Request Processing. |
| Branch | Skips any branch nodes and continues with the node that preceded the branch. |
| Pipeline Pair | Executes the response pipeline. |
| Root of the Message Flow | Sends the response back to the client. |

Pipelines

The principal component in a proxy service implementation is the *pipeline*. A pipeline is a named sequence of stages representing a non-branching one-way processing path.

Pipelines belong to one of the following categories:

- Request—Request pipelines process the request path of the message flow.
- Response—Response pipelines process the response path of the message flow.
- Error—Error pipelines handle errors for stages and nodes in a message flow, and also at the level of the message flow (service).

To create the request and response paths, you pair request and response pipelines and organize them into a single node called a *pipeline pair node*.

[“Message Flow Definition for a Proxy Service” on page 2-7](#) shows an example of a simple message flow. It defines a proxy service named `loanGateway3`.

Figure 2-2 Message Flow Definition for a Proxy Service

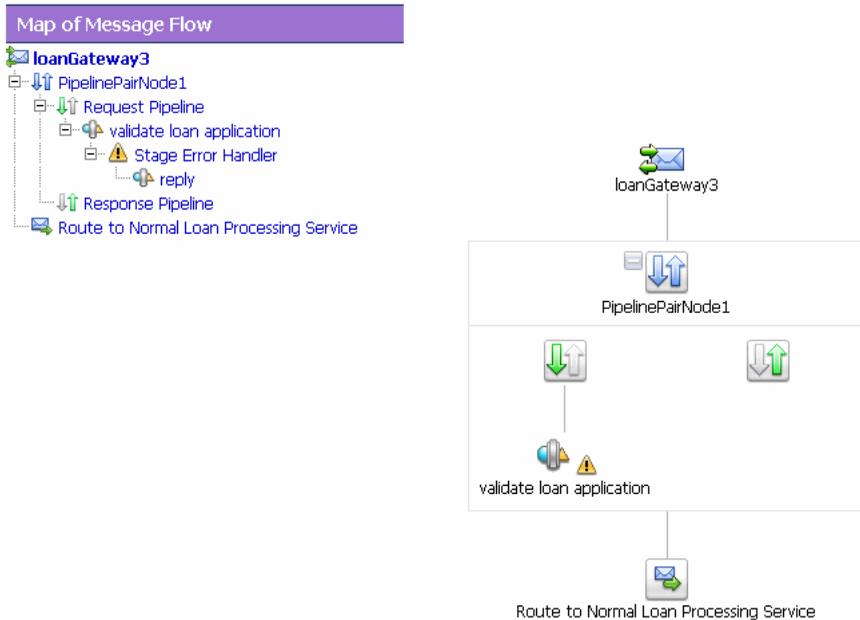


The message flow in the preceding figure shows:

- A start node is the root of the tree structure for the `loanGateway3` proxy service.
- A pipeline pair node (`PipelinePairNode1`), which includes request and response pipelines. The request pipeline includes one stage (`validate loan application`). The  icon associated with the `validate loan application` stage indicates that an error handler is defined for this stage. For more information about error handlers, which are also implemented as message flows, see [“Handling Errors” on page 2-17](#).
- A Route node (`Route to Normal Loan Processing Service`)

In addition to the view of the message flow shown in the preceding figure, the AquaLogic Service Bus Console displays the corresponding tree view map of the message flow to help you navigate components of a message flow at design time.

Figure 2-3 Message Flow Definition for a Proxy Service



To view or edit the components of the message flow, click the component in the **Map of Message Flow** view. To edit or view a component from the tree view map, click the component and select the appropriate action from the list.

This flow structure provides a clear overview of the message flow behavior at design time, making both routes and branch conditions explicit parts of the overall design, rather than locating them out of view inside a pipeline stage or route node. A branch node allows you to conditionally execute these pipeline pairs, and route nodes at the ends of the branches perform the request and response dispatching. For more information about branch nodes, see [“Branching in Message Flows”](#) on page 2-9.

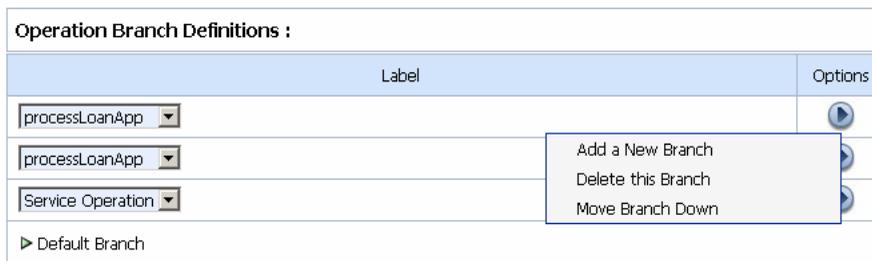
Branching in Message Flows

Two kinds of branching are supported in message flows: *operational* and *conditional* branching. The following sections explain when to use operational branching and when to use conditional branching.

Operational Branching

When message flows define Web Services Description Language (WSDL)-based proxy services, operation-specific processing is required. Instead of configuring a branching node based on operations manually, AquaLogic Service Bus provides a minimal configuration branching node that automatically branches based on operations. In other words, when you create an operational branch node in a message flow, you can quickly build your branching logic based on the operations defined in the WSDL because the AquaLogic Service Bus Console presents those operations in the branch node configuration page ([Figure 2-4](#)).

Figure 2-4 Definition for an Operation Branch



You must use operational branching in situations when a proxy service is based on a WSDL with multiple operations, you can consider using an operational branch node to handle messages separately for each operation. To learn how to configure operational branch nodes, see “Adding an Operational Branch Node” and “Viewing and Changing Operational Branch Details” in [Proxy Services: Message Flow](#) in *Using the AquaLogic Service Bus Console*.

Conditional Branching

If the proxy service is not based on a WSDL and receives multiple document types as input, consider using a conditional branch node.

Conditional branching is driven by a lookup table with each branch tagged with a simple, but unique, string value. A variable in the message context is designated as the lookup variable for

that node, and at run time, its value is used to determine which branch to follow. If no branch matches the value of the lookup variable, then the default branch is followed. You should design the proxy service in such a way that the value of the lookup variable is set before reaching the branch node.

Note: It is highly recommended that you define a default branch whenever your message flow involves conditional branching.

For example, consider a case when a proxy service is of type **Any SOAP** or **Any XML**, and you need to determine the type of the message is so that you can perform conditional branching. In this case you can design a stage action to identify the message type and then design a conditional branching node in the flow to separate processing based on the message type you receive. When you design the conditional branch node in a message flow, you build the branching logic based on evaluation of the value of the variable populated in the preceding stage.

For more information on conditional branch nodes, see “Adding a Conditional Branch Node” in [Proxy Services: Message Flow](#) in *Using the AquaLogic Service Bus Console*.

You can also use conditional branching to expose the routing alternatives at the top level flow view. For example, if you invoke service A or service B based on a condition, instead of configuring conditional branching by using a routing table within the route node, you can expose this branching in the message flow itself and use simple route nodes as the subflows for each of the branches.

[Figure 2-5](#) shows a simple message flow with a top-level branch node (`BranchNode1`) and two subordinate route nodes. At run time, one branch is executed, causing messages to be routed to either service A or service B.

Figure 2-5 Branch Node

For more information on configuring a conditional branch in a route node, see “Adding Route Node Actions” in [Proxy Services: Message Flow](#) in *Using the AquaLogic Service Bus Console*.

Consider your business scenario before deciding whether you configure branching in the message flow or in a stage or route node. When making your decision, remember that configuring branches in the message flow can be awkward in the design interface if a large number of branches extend from the branch node.

For more information, see “Overview of Message Flow” in [Proxy Services: Message Flow](#) in *Using the AquaLogic Service Bus Console*.

Performing Transformations

This section presents guidelines to follow when you design transformations. Transformation maps describe the mapping between two data types. AquaLogic Service Bus supports data mapping that uses XQuery and the eXtensible Stylesheet Language Transformation (XSLT) standards. XSLT maps describe XML-to-XML mappings, whereas XQuery maps can describe XML-to-XML, XML to non-XML, and non-XML to XML mappings. For more information, see [XQuery Transformations](#) and [XSL Transformations](#) in *Using the AquaLogic Service Bus Console*. For information on using the BEA XQuery Mapper to create XQueries, see [Transforming Data Using the XQuery Mapper](#) in *Transforming Data Using the XQuery Mapper*.

The point in a message flow at which you specify a transformation depends on whether:

- The message format relies on target services—that is, the message format must be in a format acceptable by the route destination. This applies when the transformation is performed in a route node or in one of the publish actions.

Publish actions identify a target service for a message and configure how the message is packaged and sent to that service. AquaLogic Service Bus provides Publish Table actions also. A Publish Table action consists of a set of routes wrapped in a switch-style condition table. It is a shorthand construct that allows different routes to be selected, based upon the results of a single XQuery expression.

- You perform the transformation on the response or request message regardless of the route destination. In this case, you can configure the transformations in the request or response pipeline stages.

Transformations and Publish Actions

When transformations are designed in publish actions, the transformations have a local copy of the `$outbound` variable and message-related variables (`$header`, `$body`, and `$attachments`). Any changes you make to an outbound message in a publish action affect only the published message. In other words, the changes you make in the publish action are rolled back before the message flow proceeds to any actions that follow the publish action in your message flow. For more information, see [Proxy Services: Actions](#) and [Message Context](#) in *Using the AquaLogic Service Bus Console*.

For example, consider a message flow that deals with a large purchase order, and you have to send the summary of the purchase order, through e-mail, to the manager. The summary of the of the purchase order is created in the SOAP body of the incoming message when you include a publish action in the request pipeline. In the publish action, the purchase order data is transformed into a summary of the purchase order—for example, all the attachments in `$attachments` can be deleted because they are not required in the summary of the purchase order.

Transformations and Route Nodes

In a situation in which you need to route messages to one of two possible destinations, based on a WS-addressing header, content-based routing and the second destination requires the newer version of the document in the SOAP body. In this situation, you can configure the route node to conditionally route to one of the two destinations. You can configure a transformation in the route node to transform the document for the second destination.

You can also set the control elements in the outbound context variable (`$outbound`) to influence the behavior of the system for the outbound message (for example, you can set the Quality of Service). See “Inbound and Outbound Variables” and “Constructing Messages to Dispatch” in [Message Context](#) in *Using the AquaLogic Service Bus Console* for information about the sub-elements of the inbound and outbound variables and how the content of messages is constructed using the values of the variables in the message context.

For more information about:

- Quality of Service: See “Quality of Service” on page 2-55.
- Configuring pipelines: See “Pipelines” in “Overview of Message Flow” in [Proxy Services: Message Flow](#) in *Using the AquaLogic Service Bus Console*.
- Actions: See “Adding an Action” in [Proxy Services: Actions](#) in *Using the AquaLogic Service Bus Console*.
- Route nodes: See “Adding a Route Node” in [Proxy Services: Message Flow](#) in *Using the AquaLogic Service Bus Console*.

Configuring Single and Multiple Stages in Pipelines

In AquaLogic Service Bus message flows, stages are the containers for actions that define the logic of the message flow. In most cases it is sufficient to use a single stage in a pipeline. However, some situations require the use of multiple stages. Section “[Using Multiple Stages](#)” on page 2-17 explains the usage of multiple stages in a pipeline. For information about configuring a stage, see “Adding a Stage” in [Proxy Services: Actions](#) in *Using the AquaLogic Service Bus Console*.

The BEA AquaLogic Service Bus provides a wide range of actions with which you can configure a stage in message flows. The actions are divided into following categories:

- “Communication” on page 2-14
- “Flow Control” on page 2-15
- “Message Processing” on page 2-16
- “Reporting” on page 2-16

Note:

- Only communication and flow control actions are available in a stage under a route node.

- The number of actions that are available to you in each category differ between the stage under a message flow pipeline are different from that in a stage under the route node.
- In a stage under a pipeline error handler all the categories of actions similar to that of the message flow pipeline are available.

Communication

The actions in this category control the message flow in the pipeline. You use them to specify the target URL for a message flow, a mode of packaging for a message flow, and a mode to configure a synchronous callout to an AquaLogic Service Bus registered proxy service or a business service. Communication actions in a stage in a message flow pipeline include:

- Dynamic Publish
- Publish Overview
- Publish Table
- Routing Options
- Service Callout
- Transport Headers

For more information on communication actions, see [Proxy Services: Action](#) in *Using the AquaLogic Service Bus Console*. The communication actions available to you in a route node are:

- Dynamic Routing
- Routing
- Routing Table

Note: For more information on adding action to a stage in a route node, see [Proxy Services: Message Flow-Adding Route Node Actions](#) in *Using the AquaLogic Service Bus Console*.

The communication actions available in an error handler stage are:

- Dynamic Publish
- Publish Table
- Routing Options
- Service Callout

- Transport Headers

Flow Control

The actions in this category control the message flow in the pipeline. You use them to implement conditional routing, conditional looping, and error handling within a stage in a message flow. Also you can use them to notify the invoker of success or to skip rest of the actions in the stage. Flow actions in a stage in a pipeline include:

- For Each
- If... Then...
- Raise Error
- Reply
- Skip

For more information on actions in this category, see [Proxy Services: Actions](#) in *Using the AquaLogic Service Bus Console*.

The flow control action available in a route node is If... Then...

Note: For more information on adding an action to a stage in a route node, see [Proxy Services: Message Flow-Adding Route Node Actions](#) in *Using the AquaLogic Service Bus Console*.

The flow control actions available in an error handler stage are

- For Each
- If... Then...
- Raise Error
- Reply
- Resume
- Skip

Note: For more information on adding an action to a stage in a route node, see [Proxy Services: Error Handlers](#) in *Using the AquaLogic Service Bus Console*.

Message Processing

The actions in this category process the message flow according to your requirements. You can use the actions under this category to modify the XPath expressions, invoke Java methods for processing, transform the message format, and set transport headers. Message Processing actions in a stage in a message flow pipeline include:

- Assign
- Delete
- Insert
- Java Callout
- MFL Transform
- Rename
- Replace
- Validate

For more information on message processing actions, see [Proxy Services: Actions](#) in *Using the AquaLogic Service Bus Console*.

Message processing actions available in a stage in a pipeline error handler includes:

- Assign
- Delete
- Insert
- Java Callout
- MFL Transform
- Rename
- Replace
- Validate

Reporting

You use the actions in this category to log or report errors and generate alerts if required in a message flow within a stage. Reporting actions in a stage in a message flow pipeline include:

- Alert

- Log
- Report

Reporting actions in a stage under a pipeline error handler includes:

- Alert
- Log
- Report

For more information on the reporting actions, see [Proxy Services: Actions](#) in *Using the AquaLogic Service Bus Console*.

Using Multiple Stages

Having multiple stages in a message flow enables you to define error handlers at a modular level. Each stage in a message flow can have a separate error handling pipeline. You can use two types of actions to control runtime execution of the actions in a stage:

- **Resume:** This is typically used in the error handlers to resume the next action in the message flow pipeline.

Note: The message flow processing resumes at the next stage in the pipeline.

- **Skip:** On encountering this action the processing of the current stage is skipped and the processing continues with the next stage in the message flow.

For more information, see “Adding a Stage” and “Viewing and Changing Stage Configuration Details” in [Proxy Services: Message Flow](#) in *Using the AquaLogic Service Bus Console*.

Handling Errors

The process described in the next paragraph constitutes an error handling pipeline for the error handling stage. In addition, an error pipeline can be defined for a pipeline (request or response) or for an entire proxy service.

The error handler at the stage level is invoked for handling an error; If the stage-level error handler is not able to handle a given type of error, the pipeline error handler is invoked. If the pipeline-level error handler also fails to handle the error the service level error handler is

invoked. If the service level error handler also fails, the error is handled by the system. The following table summarizes the scope of the error handlers at various levels in the message flow.

Table 2-3 Scope of Error Handlers

| Level | Scope |
|----------|--|
| Stage | Handles all the errors within a stage. |
| Pipeline | Handles all the errors in a pipeline, along with any unhandled errors from any stage in a pipeline. |
| Service | Handles all the errors in a proxy service, along with any unhandled errors in any pipeline in a service. Note: All WS-Security errors are handled at this level. |
| System | Handles all the errors that are not handled anywhere else in a pipeline. |

Note: There are exceptions to the scope of error handlers. For example, an exception thrown by a non-XML transformation at the Stage level is only caught by the Service level error handler. Suppose a transformation occurs that transforms XML to MFL for an outgoing proxy service response message, it always occurs in the binding layer. Therefore, for example, if a non-XML output is missing a mandatory field at the stage level, only a service level error handler can catch this error.

For more information on error messages and error handling, see “Error Messages and Handling” in [Proxy Services: Error Handlers](#) in *Using the AquaLogic Service Bus Console*.

You can handle errors by configuring a test that checks if an assertion is true and use the reply action configured false. You can repeat this test at various levels. Also you can have an error without an error handler at a lower level and handle it through an error handler at a higher level in message flow. In general, it is easier to handle specific errors at a stage level of the message flow and use error handlers at the higher level for more general default processing of errors that are not handled at the lower levels. It is good practice to explicitly handle anticipated errors in the pipelines and allow the service-level handler to handle unanticipated errors.

Note: You can only handle WS-Security related errors at the service level.

Generating the Error Message, Reporting, and Replying

A predefined context variable (the `fault` variable) is used to hold information about any error that occurs during message processing. When an error occurs, this variable is populated with

information before the appropriate error handler is invoked. The `fault` variable is defined only in error handler pipelines and is not set in request and response pipelines, or in route or branch nodes. For additional information about `$fault`, see “Predefined Context Variables” in [Message Context](#) in *Using the AquaLogic Service Bus Console*.

In the event of errors for request/response type inbound messages, it is often necessary to send a message back to the originator outlining the reason why an error occurred. You can accomplish this by using a *Reply with Failure* action after configuring the message context variables with the response you want to send. For example, when an HTTP message fails, *Reply with Failure* generates the `HTTP 500` status. When a JMS message fails, *Reply with Failure* sets the `JMS_BEA_Error` property to true. The AquaLogic Service Bus error actions are discussed in “Error Messages and Handling” in [Proxy Services: Error Handlers](#) in *Using the AquaLogic Service Bus Console*.

An error handling pipeline is invoked if a service invoked by a proxy service returns a SOAP fault or transport error. Any received SOAP fault is stored in `$body`, so if a *Reply with Failure* is executed without modifying `$body`, the original SOAP fault is returned to the client that invoked the service. If a reply action is not configured, the system error handler generates a new SOAP fault message. The proxy service recognizes that a SOAP fault is returned because a HTTP error status is set, or the JMS property `SERVER_Error` is set to true.

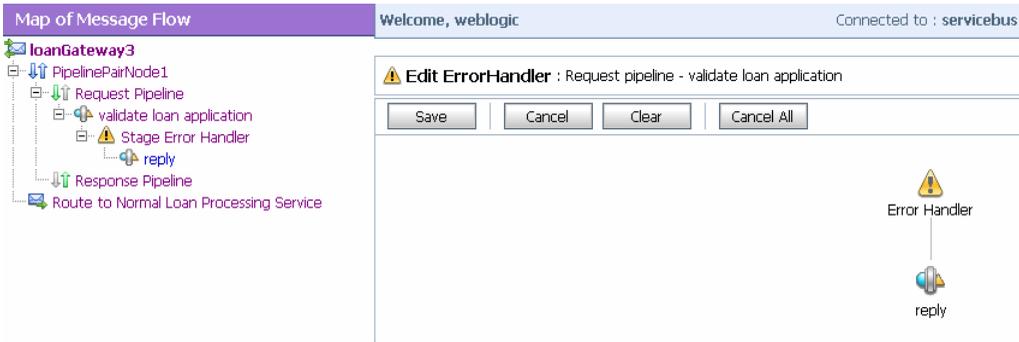
Some use cases require error reporting. You can use the `report` action in these situations. For example, consider a scenario in which the request pipeline reports a message for tracking purposes, but the service invoked by the route node fails after the reporting action. In this case, the reporting system logged the message, but there is no guarantee that the message was processed successfully, only that the message was successfully received.

You can use the AquaLogic Service Bus Console to track the message to obtain an accurate picture of the message flow. This allows you to view the original reported message indicating the message was submitted for processing, and also the subsequent reported error indicating that the message was not processed correctly. To learn how to configure a `Report` action and use the data reported at run time, see [Proxy Services: Actions](#) in *Using the AquaLogic Service Bus Console*.

Example of Action Configuration in Error Handlers

This example shows how you can configure the `Report` and `Reply` actions in error handlers. The message flow shown in [Figure 2-2](#) includes an error handler on the `validate loan application` stage. The error handler in this case is a simple message flow with a single stage configured—it is represented in the AquaLogic Service Bus Console as shown in the following figure.

Figure 2-6 Error Handler Message Flow



The stage is, in turn, configured with actions (Replace, Report, and Reply) as shown in the following figure.

Figure 2-7 Actions in Stage Error Handler

Replace `./exam:processL...` in variable `body` with `fault/ctx:reas...`

- Replace entire node
- Replace node contents

Report `$body` with search keys:

| Key Name | Key Value | Options |
|------------------------|---|---------|
| <code>errorCode</code> | <code>./ctx:errorCode</code> in variable <code>fault</code> | |

Reply

- With Success
- With Failure

The actions control the behavior of the stage in the pipeline error handler as follows:

- **Replace**—The contents of a specified element of the body variable are replaced with the contents of the `fault` context variable. The body variable element is specified by an XPath expression. The contents are replaced with the value returned by an XQuery expression—in this case `$fault/ctx:reason/text()`

- **Report**— Messages from the reporting action are written to the AquaLogic Service Bus Reporting Data Stream if the error handler configured with this action is invoked. The JMS Reporting Provider reports the messages on the AquaLogic Service Bus Dashboard. AquaLogic Service Bus provides the capability to deliver message data to one or more reporting providers. Message data is captured from the body of the message and from any other variables associated with the message, such as header or inbound variables. You can use the message delivered to the reporting provider for functions such as tracking messages or regulatory auditing.

When an error occurs, the contents of the fault context variable are reported. The key name is `errorCode`, and the key value is extracted from the fault variable using the following XPath expression: `./ctx:errorCode`. Key/value pairs are the key identifiers that identify these messages in the Dashboard at run time.

To configure a Report action and use the data reported at run time, see [Proxy Services: Actions](#) in *Using the AquaLogic Service Bus Console*.

- **Reply**— At run time, an immediate reply is sent to the invoker of the `loanGateway3` proxy service (see [Figure 2-2](#)) indicating that the message had a fault. The reply is `With Failure`.

For configuration information, see “Error Messages and Handling” in [Proxy Services: Error Handlers](#) in *Using the AquaLogic Service Bus Console*.

Selecting a Service Type

AquaLogic Service Bus supports a variety of service types that range from conventional Web services (using XML or SOAP bindings in WSDLs) to non-XML or generic services. This section provides guidelines on selecting a service type.

AquaLogic Service Bus service types for a proxy service include:

- **SOAP Services**—SOAP services receive and respond with SOAP messages. SOAP messages are constructed by wrapping the contents of the header and body variables inside a `<soap:Envelope>` element.
- **XML Services (Non SOAP)**—The messages to XML-based services are XML, but can be of any type allowed by the proxy service configuration.
- **Messaging Services**—Messaging services are those that can receive messages of one data type and respond with messages of a different data type. The supported data types include XML, MFL, text, and untyped binary.

- **WSDL Web Service**—In AquaLogic Service Bus you define proxy services based on WSDL. Although it is not mandatory, BEA recommends that you use a WSDL to define a proxy service. For more information on WSDL based services, see [“Using a WSDL to Define a Service” on page 2-23](#).

Note: All service types can send and receive attachments using MIME.

For more information on selecting a service type, see [Adding a Proxy Services](#) in *Using the AquaLogic Service Bus Console*

The following table shows the service types and the transports, which AquaLogic Service Bus supports.

Table 2-4 Supported Service Types and Transports

| Service Type | Transport Protocols |
|---------------------|----------------------------|
| SOAP or XML WSDL | HTTP |
| | HTTP(S) |
| | JMS |
| | Local |
| SOAP (no WSDL) | HTTP |
| | HTTP(S) |
| | JMS |
| | Local |

Table 2-4 Supported Service Types and Transports

| Service Type | Transport Protocols |
|---|---------------------|
| XML (no WSDL) | e-mail |
| | File |
| | FTP |
| | HTTP |
| | HTTP(S) |
| | JMS |
| | Local Tuxedo |
| Messaging Type (Binary, Text, MFL, XML) | e-mail |
| | File |
| | FTP |
| | HTTP |
| | HTTP(S) |
| | JMS |
| | Local Tuxedo |

Note: HTTP Get is supported by the XML (no WSDL) service type and Messaging Service.

The business service, which is of the Transport Typed supports only EJB type.

BEA recommends that you use the local transport for communication between two proxy services. For more information on local transport, see [Chapter 10, “Local Transport.”](#)

Using a WSDL to Define a Service

If a service has a well defined Web Services Description Language (WSDL) interface, it is recommended, although not required, that you use the WSDL to define the service. For more information on WSDL resources in AquaLogic Service Bus, see [WSDLs](#) in *Using the AquaLogic Service Bus Console*.

There are three types of WSDLs you can define. They are:

- [“SOAP Document Wrapped Web Services” on page 2-24](#)

- “SOAP Document Style Web Services” on page 2-24
- “SOAP RPC Web Services” on page 2-26

SOAP Document Wrapped Web Services

A document wrapped Web Service is described in a WSDL as a Document Style Service. However, it follows some additional conventions. Standard document-oriented Web Service operations take only one parameter or message part, typically an XML document. This means that the methods that implement the operations must also have only one parameter. Document-wrapped Web Service operations, however, can take any number of parameters, although the parameter values will be wrapped into one complex data type in a SOAP message. This wrapped complex data type will be described in the WSDL as the single document for the operation.

For more information on SOAP Document Wrapped Web Services see [Adding a Business Service](#) in *Using the AquaLogic Service Bus Console*.

SOAP Document Style Web Services

You can configure proxy services as SOAP style proxy services and configure business services as SOAP style business services.

The following listing provides an example of a WSDL for a sample document style Web service.

Listing 2-1 WSDL for a Sample Document Style Web Service

```
<definitions name="Lookup"
targetNamespace="http://example.com/lookup/service/defs"
xmlns:tns="http://example.com/lookup/service/defs"
xmlns:xs="http://www.w3.org/2001/XMLSchema"
xmlns:docs="http://example.com/lookup/docs"
xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
xmlns="http://schemas.xmlsoap.org/wsdl/">
  <types>
    <xs:schema targetNamespace="http://example.com/lookup/docs"
elementFormDefault="qualified">
      <xs:element name="PurchaseOrg" type="xs:string"/>
      <xs:element name="LegacyBoolean" type="xs:boolean"/>
    </xs:schema>
  </types>
  <message name="lookupReq">
    <part name="request" element="docs:purchaseorg"/>
  </message>
</definitions>
```

```

</message>
<message name="lookupResp">
  <part name="result" element="docs:legacyboolean"/>
</message>
<portType name="LookupPortType">
  <operation name="lookup">
    <input message="tns:lookupReq"/>
    <output message="tns:lookupResp"/>
  </operation>
</portType>
<binding name="LookupBinding" type="tns:lookupPortType">
  <soap:binding style="document"
transport="http://schemas.xmlsoap.org/soap/http"/>
  <operation name="lookup">
    <soap:operation/>
    <input>
      <soap:body use="literal" />
    </input>
    <output>
      <soap:body use="literal"/>
    </output>
  </operation>
</binding>
</definitions>

```

The service has an operation (equivalent to a method in a Java class) called `lookup`. The binding indicates that this is a SOAP document style Web service.

When the WSDL shown in the preceding listing is used for a request, the value of the body variable (`$body`) that the document style proxy service obtains is displayed in the following listing.

Note: Namespace declarations have been removed from the XML in the listings that follow for the sake of clarity.

Listing 2-2 Body Variable Value

```

<soap-env:body>
  <req:purchaseorg>BEA Systems</req:purchaseorg>
</soap-env:body>

```

In Listing 2-2, `soap-env` is the predefined SOAP name space and `req` is the namespace of the `PurchaseOrg` element (`<http://example.com/lookup/docs>`).

If the business service to which the proxy service is routing uses the above WSDL, the value for the body variable (`$body`) given above is the value of the body variable (`$body`) from the proxy service.

The value of the body variable (`$body`) for the response from the invoked business service that the proxy service receives is displayed in the following listing.

Note: Namespace declarations have been removed from the XML in the listings that follow for the sake of clarity.

Listing 2-3 Body Variable Value

```
<soap-env:body>
  <req:legacyboolean>true</req:legacyboolean>
</soap-env:body>
```

This is also the value of the body variable (`$body`) for the response returned by the proxy service using this WSDL.

There are many tools available (including BEA WebLogic Workshop tools) that take the WSDL of a proxy service (obtained by adding the `?WSDL` suffix to the URL of the proxy service in the browser) and generate a Java class with the appropriate request and response parameters to invoke the operations of the service. This Java class can be used to invoke the proxy service that uses this WSDL.

SOAP RPC Web Services

You can configure proxy services as RPC style proxy services and configure business services as RPC style business services.

The following listing provides an example of a WSDL for a sample RPC style Web service.

Listing 2-4 WSDL for a Sample RPC Style Web Service

```

<definitions name="Lookup"
targetNamespace="http://example.com/lookup/service/defs"
xmlns:tns="http://example.com/lookup/service/defs"
xmlns:xs="http://www.w3.org/2001/XMLSchema"
xmlns:docs="http://example.com/lookup/docs"
xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
xmlns="http://schemas.xmlsoap.org/wsdl/">
  <types>
    <xs:schema targetNamespace="http://example.com/lookup/docs"
elementFormDefault="qualified">
      <xs:complexType name="RequestDoc">
        <xs:sequence>
          <xs:element name="PurchaseOrg" type="xs:string"/>
        </xs:sequence>
      </xs:complexType>
      <xs:complexType name="ResponseDoc">
        <xs:sequence>
          <xs:element name="LegacyBoolean" type="xs:boolean"/>
        </xs:sequence>
      </xs:complexType>
    </xs:schema>
  </types>
  <message name="lookupReq">
    <part name="request" type="docs: RequestDoc"/>
  </message>
  <message name="lookupResp">
    <part name="result" type="docs: ResponseDoc"/>
  </message>
  <portType name="LookupPortType">
    <operation name="lookup">
      <input message="tns:lookupReq"/>
      <output message="tns:lookupResp"/>
    </operation>
  </portType>
  <binding name="LookupBinding" type="tns:lookupPortType">
    <soap:binding style="rpc"
transport="http://schemas.xmlsoap.org/soap/http"/>
    <operation name="lookup">
      <soap:operation/>
      <input>
        <soap:body use="literal"
namespace="http://example.com/lookup/service"/>
      </input>
      <output>
        <soap:body use="literal"
namespace="http://example.com/lookup/service"/>

```

```
</output>
</operation>
</binding>
</definitions>
```

The service described in the preceding listing includes an operation (equivalent to a method in a Java class) called `lookup`. The binding indicates that this is a SOAP RPC Web service. In other words, the Web service's operation receives a set of request parameters and returns a set of response parameters. The `lookup` operation has a parameter called `request` and a return parameter called `result`. The namespace of the operation in the binding is:

```
http://example.com/lookup/service
```

When the WSDL shown in [Listing 2-2](#) is used for a request, the value of the body variable (`$body`) that the SOAP RPC proxy service obtains is displayed in the following listing.

Note: Namespace declarations have been removed from the XML in the listings that follow for the sake of clarity.

Listing 2-5 Body Variable Value

```
<soap-env:body>
  <ns:lookup>
    <request>
      <req:purchaseorg>BEA Systems</req:purchaseorg>
    </request>
  </ns:lookup>
</soap-env:body>
```

Where `soap-env` is the predefined SOAP name space, `ns` is the operation namespace (`<http://example.com/lookup/service>`) and, `req` is the namespace of the `PurchaseOrg` element (`<http://example.com/lookup/docs>`).

If the business service to which the proxy service routes the messages uses the WSDL shown in [Listing 2-2](#), the value for the body variable (`$body`), shown in [Listing 2-3](#), is the value of the body variable (`$body`) from the proxy service.

When this WSDL is used for a request, the value of the body variable (`$body`) for the response from the invoked business service that the proxy service receives is displayed in the following listing.

Listing 2-6 Body Variable Value

```
<soap-env:body>
  <ns:lookupResponse>
    <result>
      <req:legacyboolean>true</req:legacyboolean>
    </result>
  </ns:lookupResponse>
</soap-env:body>
```

This is also the value of the body variable (`$body`) for the response returned by the proxy service using this WSDL.

There are many tools available (including BEA WebLogic Workshop tools) that take the WSDL of a proxy service (obtained by adding the `?WSDL` suffix to the URL of the proxy in the browser) and generate a Java class with the appropriate request and response parameters to invoke the operations of that service. You can use such Java classes to invoke the proxy services that use this WSDL.

The benefits of using a WSDL include the following:

- The system can provide metrics for each operation in a WSDL.
- Operational branching is possible in the pipeline. For more information, see [“Branching in Message Flows” on page 2-9](#).
- The `SOAPAction` header is automatically populated for services invoked by a proxy service.
- A WSDL is required for services using WS-Security. WS-Policies are attached to WSDLs. See [WS-Policies](#) in *Using the AquaLogic Service Bus Console*.
- The system supports the `<url>?WSDL` syntax, which allows you to dynamically obtain the WSDL of a HTTP proxy service. This is useful for a number of SOAP client generation tools, including BEA WebLogic Workshop.

- In the XQuery and XPath editors and condition builders, it is easy to manipulate the body content variable (`$body`) because the editor provides a default mapping of `$body` to the request message in the WSDL of a proxy service. See [Message Context](#) in *Using the AquaLogic Service Bus Console*.

Note: The run-time contents of `$body` for a specific action can be different from the default mapping displayed in the editor. This is because AquaLogic Service Bus is not a programming language in which typed variables are declared and used. Instead, variables are untyped and are created dynamically at run time when a value is assigned. In addition, the type of the variable is the type that is implied by its contents at any point in the message flow. To enable you to easily create XQuery and XPath expressions, the design time editor allows you to map the type for a given variable by mapping the variable to the type in the editor. To learn about using the XQuery and XPath editor to create expressions, see [“Working with Variable Structures” on page 2-40](#).

Binding a Service to a WSDL Port Instead of to a Binding

If you use a WSDL service type, it is useful to bind the service to a WSDL port instead of to a binding because:

- If the service is bound to port X in the template WSDL, then port X is also defined in the generated WSDL. Any other ports defined in the template WSDL are not included in the generated WSDL. Furthermore, if you base the proxy service on a WSDL port, the generated WSDL uses that port name and preserves any WS-Policies associated with that port.

(The template WSDL is the WSDL for the service upon which you based your proxy service; the generated WSDL is the WSDL created for the new proxy service.)

- If the service is bound to binding Y in the template WSDL, the generated WSDL defines one service and port (`<service-name>QSService` and `<port-name>QSPort`). None of the ports defined in the template WSDL are included in the generated WSDL.

You can get the WSDL for an HTTP or HTTP(S)-based proxy service by entering the following URL in your browser’s address field:

```
http://host:port/sbresource?PROXY/project/proxyname
```

In the WSDL returned by the `http://host:port/sbresource?PROXY/project/proxyname` URL or the WSDL, which is obtained from the URL for the proxy service, the port name is preserved if the proxy service is bound to a port on the WSDL and the URL accurately reflects the URL of the proxy service. This can be important to some tools, which generate a client. The URL in the WSDL port that is bound to the service is not used when you define a service, except

to populate the URL in the WSDL port as the default URL for a business service. You can overwrite the transport type and transport URL in the transport configuration UI for the service definition.

Any WS-Security policies at the port level apply. See “Overview of Proxy Services” in [Proxy Services](#) in *Using the AquaLogic Service Bus Console*.

Using Any SOAP or Any XML Service Types

If you want to expose one port to clients for a variety of enterprise applications, use **Any SOAP** or **Any XML** service types.

Using the Messaging Service Type

If one of the request or response messages is non-XML, you must use the messaging service type.

AquaLogic Service Bus does not automatically perform “misunderstand” SOAP header checking. However, you can use XQuery conditional expressions and validate actions to explicitly perform this type of check. For more information on the validate action, see “Validate” in [Proxy Services: Actions](#) in *Using the AquaLogic Service Bus Console*. For more information on conditional XQuery expressions, see “Using the XQuery Condition Editor” in [Proxy Services: Editors](#) in *Using the AquaLogic Service Bus Console*.

You can use AquaLogic Service Bus to configure a validate action and use XQuery conditional expressions to perform validation checks explicitly in the message flow.

For more information on service types, see “Overview of Proxy Services” in [Proxy Services](#) in *Using the AquaLogic Service Bus Console*.

Viewing Resource Details

AquaLogic Service Bus provides a resource servlet that is used to expose the resources registered in AquaLogic Service Bus. The resources registered with AquaLogic Service Bus include:

The format of the URLs used to expose the resources is as follows:

- WSDL (a WSDL registered as a resource in AquaLogic Service Bus)
- Schema
- MFL
- WS-Policy

- WSDL (a derived WSDL with resolved policies and port information for a proxy service—this derived WSDL is available if the proxy service was created using a WSDL).

You can use the following URL formats to expose the resource details:

- `http://host:port/sbresource?WSDL/project/...wsdlname`
- `http://host:port/sbresource?POLICY/project/...policyname`
- `http://host:port/sbresource?MFL/project/...mflname`
- `http://host:port/sbresource?SCHEMA/project/...schemaname`
- `http://host:port/sbresource?PROXY/project/...proxyname`

Note: The URLs used to expose the resources in AquaLogic Service Bus must be encoded in UTF-8 in order to escape special characters.

Using Dynamic Routing

When you do not know the service you need to invoke from the proxy service you are creating, you can use dynamic routing.

For any given proxy service, you can use one of the following techniques to dynamically route messages:

- In a message flow pipeline, design an XQuery expression to dynamically set the fully qualified service name in AquaLogic Service Bus and use the dynamic route or dynamic publish actions.

Note: Dynamic Routing can be achieved in a route node, whereas dynamic publishing can be achieved in a stage in a request pipeline or a response pipeline.

With this technique, the proxy service dynamically uses the service account of the endpoint business service to send user names and passwords in its outbound requests. For example, if a proxy service is routing a request to Business Service A, then the proxy service uses the service account from Business Service A to send user names and passwords in its outbound request. See [“Implementing Dynamic Routing” on page 2-33](#).

- Configure a proxy service to route or publish messages to a business service. Then, in the request actions section for the route action or publish action, add a Routing Options action that dynamically specifies the URI of a service.

With this technique, to send user names and passwords in its outbound requests, the proxy service uses the service account of the statically defined business service, regardless of the URI to which the request is actually sent.

For information on how to use this technique, see [“Implementing Dynamic Routing” on page 2-33](#).

Note: This technique is used when the overview of the interface is fixed. The overview of the interface includes message types, port types, and binding, and excludes the concrete interface. The concrete interface is the transport URL at which the service is located.

Implementing Dynamic Routing

You can use dynamic routing to determine the destination during the runtime of a proxy service. To achieve this you can use a routing table in an XML file to create an XQuery resource.

Note: Instead of using the XQuery resource, you can also directly use the XML file from which the resource is created.

An XML file or the Xquery resource can be maintained easily. At runtime you provide the entry in the routing table that will determine the routing or publishing destination of the proxy service. The XML file or the XQuery resource contains a routing table, which maps a logical identifier to (such as the name of a company) to the physical identifier (the fully qualified name of the service in AquaLogic Service Bus). The logical identifier, which is extracted from the message, maps on to the physical identifier, which is the name of the service you want to invoke.

Note: To use the dynamic route action, you need the fully qualified name of the service in AquaLogic Service Bus.

In a pipeline the logical identifier is obtained with an XPath into the message. You assign the XML table in the XQuery resource to a variable. You implement a query against the variable in the routing table to extract the physical identifier based on the corresponding logical identifier. Using this variable you will be able to invoke the required service. The following sections describe how to implement dynamic routing.

- [“Sample XML File” on page 2-33](#).
- [“Creating an XQuery Resource From the Sample XML” on page 2-34](#).
- [“Creating and Configuring the Proxy Service to Implement Dynamic Routing” on page 2-35](#)

Sample XML File

You can create an XQuery resource from the following XML file. Save this as sampleXquery.xml.

Listing 2-7 Sample XML File

```
<routing>
  <row>
    <logical>BEA Systems</logical>
    <physical>default/goldservice</physical>
  </row>
  <row>
    <logical>ABC Corp</logical>
    <physical>default/silverservice</physical>
  </row>
</routing>
```

Creating an XQuery Resource From the Sample XML

1. In an active session, select **Project Explorer** from the left navigation panel. The **Project View** page is displayed.
2. Select the project to which you want to add the XQuery resource.
3. In the Project view page select the XQuery resource from the **Select Resource Type** drop-down list. The **Create XQuery** page is displayed.
4. In the **Resource Name** field, enter the name of the resource. This is a mandatory.
5. In the **Resource Description** field provide the a description for the resource. This is optional.
6. In the XQuery field provide the path to the XML you are using as an XQuery resource. Click on the **Browse** to locate the file. Optionally you can copy and paste the XML in the XQuery field. This is mandatory.
7. Save the XQuery resource.
8. Activate the session.

Creating and Configuring the Proxy Service to Implement Dynamic Routing

1. In an active session select **Project Explorer** from the left navigation panel. The **Project View** page is displayed.
2. Select the project to which you want to add the proxy service.
3. In the **Project View** page, select the **Proxy Service** resource from the **Select Resource Type** drop-down list. The **General Configuration** page is displayed.
4. In the **Service Name** field of the **General Configuration** page enter the name of the proxy service. This is mandatory.
5. Select the type of service by clicking on the button adjacent to various types of services available under **Service Type**. For more information on selecting the service type, see [Proxy Services: Actions](#).
6. Click **Finish**. On the **Summary** page, click **Save** to save the proxy service.
7. On the **Project View** page, click the Edit Message Flow icon against the newly created proxy service in the **Resource** table. The **Edit Message Flow** page is displayed.
8. Click on the message flow to add a pipeline pair to the message flow.
9. Click on request pipeline icon select **Add Stage** from the menu.
10. Click on the stage1 icon to and select **Edit Stage** from the menu. The **Edit Stage Configuration** page appears.
11. Click **Add Action** icon. Choose **Add an Action** item from the menu.
12. Choose the **Assign** action from **Message Processing**.
13. Click on **Expression**. The **XQuery Expression Editor** is displayed.
14. Click on **XQuery Resources**. The browser displays the page where you can import the XQuery resource. Click on the **Browse** to locate the XQuery resource.
15. Click on **Validate** to validate the imported XQuery resource.
16. Save the imported XQuery resource on successful validation.
17. In the **Edit Stage Configuration** page enter the name of the variable in the field. By this you assign the XQuery resource to this variable.
This variable now contains the externalized routing table.
18. Click on the Assign action icon to add another assign action.

Note: To do this repeat [step 11](#) to [step 13](#)

19. Enter the following Xquery:

```
<ctx: route>
<ctx:
service>{$routingtable/row[logical/text()=$logicalidentifier]/physical/
text()}</ctx: service>
</ctx: route>
```

In the above code, replace `$logicalidentifier` by the actual XPath to extract the logical identifier from the message (example from `$body`).

20. Click on **Validate** to validate the Xquery.

21. Save the Xquery on successful validation.

22. In the **Edit Stage Configuration** page, enter the name of the variable (for example, `routeresult`) in the field.

By this you extract the XML used by the dynamic route action into this variable.

23. Click on the message flow to add a route node to the end of the message flow.

24. Click on the route node icon and select **Edit** from the menu.

25. Click the **Add Action** icon. Choose **Add an Action** item from the menu.

26. Choose the **Dynamic Route** action.

27. Click on **Expression**. The XQuery Expression Editor is displayed.

28. Enter the variable from [step 22](#) (for example, `$routeresult`)

Understanding Message Context

The message context is a set of variables that hold message context and information about messages as they are routed through the AquaLogic Service Bus. Together, the `header`, `body`, and `attachments` variables, (referenced as `$header`, `$body` and `$attachments` in XQuery statements) represent the message as it flows through AquaLogic Service Bus. The canonical form of the message is SOAP. Even if the service type is not SOAP, the message appears as SOAP in the AquaLogic Service Bus message context.

Message Context Components

In a Message Context `$header` contains a SOAP Header element, `$body` contains a SOAP Body element, and `$attachments` contains a wrapper element called `attachments` with one child attachment element per attachment. The attachment element has a `body` element with the actual attachment.

When a message is received by a proxy service, the message contents are used to initialize the header, body, and attachments variables. For SOAP services, the Header and Body elements are taken directly from the envelope of the received SOAP message and assigned to `$header` and `$body` respectively. For non-SOAP services, the entire content of the message is typically wrapped in a Body element and assigned to `$body`, and an empty Header element is assigned to `$header`.

Binary and MFL messages are initialized differently. For MFL messages, the equivalent XML document is inserted into the Body element that is assigned to `$body`. For binary messages, the message data is stored internally and a piece of reference XML is inserted into the Body element that is assigned to `$body`. The reference XML looks like `<binary-content ref="..." />`, where `"..."` contains a unique identifier assigned by the proxy service.

The message context is defined by an XML Schema. You must use XQuery expressions to manipulate the context variables in the message flow that defines a proxy service.

The predefined context variables provided by AquaLogic Service Bus can be grouped into the following types:

- Message-related variables
- Inbound and outbound variables
- Operation variable
- Fault variable

For information about the predefined context variables, see “Predefined Context Variables” in [Message Context](#) in *Using the AquaLogic Service Bus Console*.

The `$body` contains message payload variable. When a message is dispatched from AquaLogic Service Bus you can decide the variables, whose you want to include in the outgoing message. That determination is dependent upon whether the target endpoint is expecting a SOAP or a non-SOAP message:

- For a binary, any text or XML message content inside the Body element in `$body` is sent.

- For MFL messages, the Body element in `$body` contains the XML equivalent of the MFL document.
- For text messages, the Body element in `$body` contains the text. For text attachments, the body element in `$attachments` contains the text. If the contents are XML instead of simple text, the XML is sent as a text message.
- For XML messages, the Body element in `$body` contains the XML. For XML attachments, the body element in `$attachments` contains the XML.
- SOAP messages are constructed by wrapping the contents of the header and body variables inside a `<soap:Envelope>` element. If the body variable contains a piece of reference XML, it is sent. That is the referenced content is not substituted in the message.

For non-SOAP services, if the Body element of `$body` contains a binary-content element, then the referenced content stored internally is sent ‘as is’, regardless of the target service type.

For more information, see [Message Context](#) in *Using the AquaLogic Service Bus Console*.

The types for the message context variables are defined by the message context schema (`MessageContext.xsd`). When working with the message context variables in the BEA XQuery Mapper, you need to reference `MessageContext.xsd` and the transport-specific schemas, which are available in a JAR file at the following location in your AquaLogic Service Bus installation:

```
<BEA_HOME>\weblogic92\servicebus\lib\sb-schemas.jar
```

where `BEA_HOME` represents the directory in which you installed AquaLogic Service Bus.

To learn about the message context schema and the transport specific schemas, see “Message Context Schema” in [Message Context](#) in *Using the AquaLogic Service Bus Console*.

Guidelines for Viewing and Altering Message Context

Consider the following guidelines when you want to inspect or alter the message context:

- In an XQuery expression, the root element in a variable is not present in the path in a reference to an element in that variable. For example, the following XQuery expression obtains the `Content-Description` of the first attachment in a message:

```
$attachments/ctx:attachment[1]/ctx:content-Description
```

To obtain the second attachment

```
$attachments/ctx:attachment[2]/ctx:body/*
```

- A context variable can be empty or it can contain a single XML element or a string value. However, an XQuery expression often returns a sequence. When you use an XQuery

expression to assign a value to a variable, only the first element in the sequence returned by the expression is stored as the variable value. For example, if you want to assign the value of a WS-Addressing Message ID from a SOAP header (assuming there is one in the header) to a variable named `idvar`, the assign action specification is:

```
assign data($header/wsa:messageID to variable idvar
```

Note: In this case, if two WS-Addressing MessageID headers exist, the `idvar` variable will be assigned the value of the first one.

- The variables `$header`, `$body`, and `$attachments` are never empty. However, `$header` can contain an empty SOAP Header element, `$body` can contain an empty SOAP Body element, and `$attachments` can contain an empty attachment element.
- In cases in which you use a transformation resource (XSLT or XQuery), the transformation resource is defined to transform the document in the SOAP body of a message. To make this transformation case easy and efficient, the input parameter to the transformation can be an XQuery expression. For example, you can use the following XQuery expression to feed the business document in the Body element of a message (`$body`) as input to a transformation:

```
$body/* [1]
```

The result of the transformation can be put back in `$body` with a `Replace` action. That is replace the content of `$body`, which is the content of the Body element. For more information, see [XQuery Transformations](#) and [XSL Transformations](#) in *Using the AquaLogic Service Bus Console*.

- In addition to inserting or replacing a single element, you can also insert or replace a selected sequence of elements using an insert or replace action. You can configure an XQuery expression to return a sequence of elements. For example, you can use insert and replace actions to copy a set of transport headers from `$inbound` to `$outbound`. For information on adding an action, see “Adding an Action” in [Proxy Services: Actions](#) in *Using the AquaLogic Service Bus Console*. For an example, see “[Copying JMS Properties From Inbound to Outbound](#)” on page 2-39.

Copying JMS Properties From Inbound to Outbound

It is assumed that the interfaces of the proxy services and of the invoked business service may be different. Therefore, AquaLogic Service Bus does not propagate any information (such as the transport headers and JMS properties) from the inbound variable to the outbound variable.

The transport headers for the proxy service’s request and response messages are in `$inbound` and the transport headers for the invoked business service’s request and response are in `$outbound`.

For example, the following XQuery expression can be used in a case where the user-defined JMS properties for a one-way message (an invocation with no response) need to be copied from inbound message to outbound message:

Use the Transport Headers action to set

```
$inbound/ctx:transport/ctx:request/tp:headers/tp:user-header
```

as the first child of:

```
./ctx:transport/ctx:request/tp:headers
```

in the outbound variable.

To learn how to configure the Transport Header action in the AquaLogic Service Bus Console, see “Transport Headers” in [Proxy Services: Actions](#) in *Using the AquaLogic Service Bus Console*.

Working with Variable Structures

The following sections describe

- [“Using the Inline XQuery Expression Editor”](#) on page 2-40
- [“Using Variable Structures”](#) on page 2-41
- [“Creating Variable Structure Mappings”](#) on page 2-42

Using the Inline XQuery Expression Editor

AquaLogic Service Bus allows you to import XQueries that have been created with an external tool such as the BEA XQuery Mapper. You can use these XQueries anywhere in the proxy service message flow by binding the XQuery resource input to an Inline XQuery, and binding the XQuery resource output to an action that uses the result as the input; for example, the Assign, Replace, or Insert actions.

However, you can enter the XQuery inline as part of the action definition instead of entering the XQuery as a resource. You can also use Inline XQueries for the condition in an **If...Then...** action.

The Inline XQuery Expression Editor to enter simple XQueries that consist of the following:

- Fragments of XML with embedded XQueries.
- Simple variable paths along the child axis.

Note: For more complex XQueries, it is recommended that you use the XQuery Mapper, especially if you are not familiar with XQuery.

Inline XQueries can be used effectively to:

- Create variable structures by using the Inline XQuery Expression Editor. See [“Using Variable Structures” on page 2-41](#).
- Extract or access a business document or RPC parameter from the SOAP envelope elements in `$header` or `$body`.
- Extract or access an attachment document in `$attachments`.
- Set up the parameters of a Service Callout action by extracting it from the SOAP envelope.
- Insert the result parameter of a Service Callout action into the SOAP envelope.
- Extract a sequence from the SOAP envelope to drive a `for loop`.
- Update an item in the sequence in a `for loop` with an Update action.

Note: You can also use the Inline XQuery Expression Editor to create variable structures. For more information, see [“Using Variable Structures” on page 2-41](#)

Using Variable Structures

You can use the Inline XQuery Expression Editor to create variable structures, with which you define the structure of a given variable for design purposes. For example, it is easier to browse the XPath variable in the console rather than viewing the XML Schema of the XPath variable.

Note: It is not necessary to create variable structures for your runtime to work. Variable structures define the structure of the variable or the variable path but do not create the variable. Variables are created at runtime as the target of the `Assign` action in the stage.

In a typical programming language, the scope of variables is static. Their name and type are explicitly declared. The variable can be accessed anywhere within the static scope.

In AquaLogic Service Bus, there are some predefined variables, but you can also dynamically create variables and assign value to them using the `Assign` action or using the loop variable in the `for-loop`. When a value is assigned to a variable, the variable can be accessed anywhere in the proxy service message flow. The variable type is not declared but the type is essentially the underlying type of the value it contains at any point in time.

Note: The scope of the `for-loop` variable is limited and cannot be accessed outside the stage.

When you use the Inline XQuery Expression Editor, the XQuery has zero or more inputs and one output. Because you can display the structure of the inputs and the structure of the output visually in the Expression Editor itself, you do not need to open the XML Schema or WSDL resources to see their structure when you create the Inline XQuery. The graphical structure display also enables you to drag and drop simple variable paths along the child axis without predicates, into the composed XQuery.

Each variable structure mapping entry has a label and maps a variable or variable path to one or more structures. The scope of these mappings is the stage or route node. Because variables are not statically typed, a variable can have different structures at different points (or at the same point) in the stage or route node. Therefore, you can map a variable or a variable path to multiple structures, each with a different label. To view the structure, select the corresponding label with a drop-down list.

Note: You can also create variable structure mappings in the Inline XPath Expression Editor. However, although the variable or a variable path is mapped to a structure, the XPath generated when you select from the structure are XPaths relative to the variable. An example of a relative XPath is `./ctx:attachment/ctx:body`.

Creating Variable Structure Mappings

The following sections describe how to create several types of variable structure mappings:

- [“Sample WSDL” on page 2-43](#)
- [“Creating the Resources You Need for the Examples” on page 2-44](#)
- [“Example 1: Selecting a Predefined Variable Structure” on page 2-47](#)
- [“Example 2: Creating a Variable Structure That Maps a Variable to a Type” on page 2-48](#)
- [“Example 3: Creating a Variable Structure that Maps a Variable to an Element” on page 2-50](#)
- [“Example 4: Creating a Variable Structure That Maps a Variable to a Child Element” on page 2-51](#)
- [“Example 5: Creating a Variable Structure that Maps a Variable to a Business Service” on page 2-52](#)
- [“Example 6: Creating a Variable Structure That Maps a Child Element to Another Child Element” on page 2-53](#)

Sample WSDL

This sample WSDL is used in most of the examples in this section. You need to save this WSDL as a resource in your configuration. For more information, see [Creating the Resources You Need for the Examples](#).

Listing 2-8 Sample WSDL

```
<definitions
  name="samplewsdl"
  targetNamespace="http://example.org"
  xmlns="http://schemas.xmlsoap.org/wsdl/"
  xmlns:s0="http://www.bea.com"
  xmlns:s1="http://example.org"
  xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/">
<types>
  <xs:schema
    attributeFormDefault="unqualified"
    elementFormDefault="qualified"
    targetNamespace="http://www.bea.com"
    xmlns:xs="http://www.w3.org/2001/XMLSchema">
    <xs:element name="PO" type="s0:POType"/>
    <xs:complexType name="POType">
      <xs:all>
        <xs:element name="id" type="xs:string"/>
        <xs:element name="name" type="xs:string"/>
      </xs:all>
    </xs:complexType>
    <xs:element name="Invoice" type="s0:InvoiceType"/>
    <xs:complexType name="InvoiceType">
      <xs:all>
        <xs:element name="id" type="xs:string"/>
        <xs:element name="name" type="xs:string"/>
      </xs:all>
    </xs:complexType>
  </xs:schema>
</types>
<message name="POTypeMsg">
```

```
<part name="PO" type="s0:POType"/>
</message>
<message name="InvoiceTypeMsg">
  <part name="InvReturn" type="s0:InvoiceType"/>
</message>

<portType name="POPortType">
  <operation name="GetInvoiceType">
    <input message="s1:POTypeMsg"/>
    <output message="s1:InvoiceTypeMsg"/>
  </operation>
</portType>
<binding name="POBinding" type="s1:POPortType">
<soap:binding style="rpc"
transport="http://schemas.xmlsoap.org/soap/http"/>
  <operation name="GetInvoiceType">
    <soap:operation soapAction="http://example.com/GetInvoiceType"/>
    <input>
      <soap:body use="literal"/>
    </input>
    <output>
      <soap:body use="literal"/>
    </output>
  </operation>
</binding>
</definitions>
```

Creating the Resources You Need for the Examples

To make use of the examples that follow, you save the sample WSDL as a resource in your configuration and create the sample business service and proxy service that use the sample WSDL.

The tasks in this procedure include:

- [“Save the WSDL as a Resource” on page 2-45](#)
- [“Create a Proxy Service That Uses the Sample WSDL” on page 2-45](#)

- [“Build a Message Flow for the Sample Proxy Service”](#) on page 2-46
- [“Create a Business Service That Uses the Sample WSDL”](#) on page 2-46

Save the WSDL as a Resource

1. In the left navigation pane in the AquaLogic Service Bus Console, under **Change Center**, click **Create** to create a new session for making changes to the current configuration.
2. In the left navigation pane, click on **Project Explorer**.
3. In the **Project View** page, click the project to which you want to add the WSDL.
4. In the **Project View** page, in the **Create Resource** field, select **WSDL** under **Interface**.
5. In the **Create a New WSDL Resource** page in the **Resource Name** field, enter `SampleWSDL`. This is a required field.
6. In the **WSDL** field, copy and paste the text from the sample WSDL into this field.
Note: This is a required field.
7. Click **Save**. The new WSDL `SampleWSDL` is included in the list of resources and saved in the current session. You must now create a proxy service that uses this WSDL, see [“Create a Proxy Service That Uses the Sample WSDL”](#) on page 2-45.

Create a Proxy Service That Uses the Sample WSDL

1. In the left navigation pane, click **Project Explorer**.
2. In the **Project View** page, select the project to which you want to add the proxy service.
3. In the **Project View** page, in the **Create Resource** field, select **Proxy Service** under **Service**.
4. In the **Edit a Proxy Service - General Configuration** page, in the **Service Name** field, enter `ProxywithSampleWSDL`. This is a required field.
5. In the **Service Type** field, which defines the types and packaging of the messages exchanged by the service:
 - a. Select **WSDL Web Service** from under **Create a New Service**.
 - b. Click **Browse**. The **WSDL Browser** is displayed.
 - c. Select `SampleWSDL`, then select `POBinding` in the **Select WSDL Definitions** pane.
 - d. Click **Submit**.

6. Use the default values for all other fields on the **General Configuration** page, then click **Next**.
7. Use the default values for all fields on the **Transport Configuration** pages, then click **Next**.
8. In the **Operation Selection Configuration** page, make sure **SOAP Body Type** is selected in the **Selection Algorithm** field, then click **Next**.
9. Review the configuration data that you have entered for this proxy service, then click **Save**. The new proxy service `ProxywithSampleWSDL` is included in the list of resources and saved in the current session. To build message flow for this proxy service, see [“Build a Message Flow for the Sample Proxy Service” on page 2-46](#).

Build a Message Flow for the Sample Proxy Service

1. In the **Project View** page, in the **Actions** column, click the **Edit Message Flow** icon for the `ProxywithSampleWSDL` proxy service.
2. In the **Edit Message Flow** page, click the `ProxywithSampleWSDL` icon, then click **Add Pipeline Pair**. **PipelinePairNode1** is displayed, which includes request and response pipelines.
3. Click the request pipeline, then click **Add Stage**. The stage **stage1** is displayed.
4. Click **Save**. The basic message flow is created for the `ProxywithSampleWSDL` proxy service.

Create a Business Service That Uses the Sample WSDL

1. In the left navigation pane, click on **Project Explorer**. The **Project View** page is displayed.
2. Select the project to which you want to add the business service.
3. From the **Project View** page, in the **Create Resource** field, select **Business Service** from under **Service**. The **Edit a Business Service - General Configuration** page is displayed.
4. In the **Service Name** field, enter `BusinesswithSampleWSDL`. This is a required field.
5. In the **Service Type** field, which defines the types and packaging of the messages exchanged by the service, do the following:
 - a. Select **WSDL Web Service** from under **Create a New Service**.
 - b. Click **Browse**. The **WSDL Browser** is displayed.
 - c. Select `SampleWSDL`, then select **POBinding** in the **Select WSDL Definitions** pane.

- d. Click **Submit**.
6. Use the default values for all other fields on the **General Configuration** page. Click **Next**.
7. Use the default values for all fields on the **Transport Configuration** and **SOAP Binding Configuration** page. Click **Next**.
8. Review the configuration data that you have entered for this business service, and then click **Save**. The new business service `BusinesswithSampleWSDL` is included in the list of resources and is saved in the current session.
9. From the left navigation pane, click **Activate** under **Change Center**. The session ends and the configuration is deployed to run time. You are now ready to use the examples—continue in [“Example 1: Selecting a Predefined Variable Structure”](#) on page 2-47.

Example 1: Selecting a Predefined Variable Structure

In this example, you select a predefined variable structure using the proxy service `ProxyWithSampleWSDL`, which has a service type **WSDL Web Service** that uses the binding **POBinding** from `SampleWSDL`.

The proxy service message flow needs to know the structure of the message in order to manipulate it. To achieve this, AquaLogic Service Bus automatically provides a predefined structure that maps the `body` variable to the SOAP body structure as defined by the WSDL of the proxy service for all the messages in the interface. This predefined structure mapping is labeled `body`.

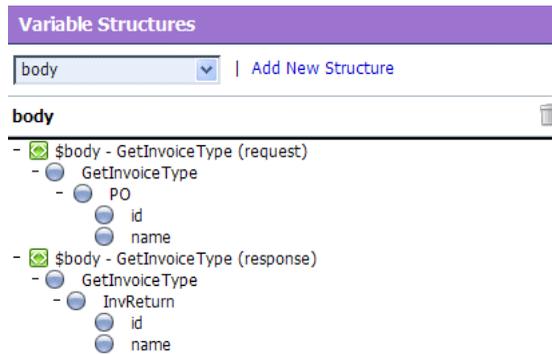
Note: This predefined structure is also supported for messaging services with a typed interface.

To select a predefined variable structure:

In the Variable Structures panel on the XQuery Expression Editor page, select `body` from the drop-down list of built-in structures.

The variable structure `body` is displayed as follows:

Figure 2-8 Variable Structures—body



Example 2: Creating a Variable Structure That Maps a Variable to a Type

Suppose the proxy service `ProxyWithSampleWSDL` invokes a service callout to the business service `BusinessWithSampleWSDL`, which also has a service type **WSDL Web Service** that uses the binding **POBinding** from `SampleWSDL`. The operation `GetInvoiceType` is invoked.

In this example, the message flow needs to know the structure of the response parameter in order to manipulate it. To achieve this, you can create a new variable structure that maps the response parameter variable to the type `InvoiceType`.

To map a variable to a type:

1. In the Variable Structures panel, click **Add New Structure**. Additional fields are displayed as follows:

Figure 2-9 Variable Structures—Add a New Structure

Variable Structures

► XML Type | Service Interface | Simple Type

Structure Label:

Structure Path:

Type:

Schema Element
 Set as child

Schema Type

MFL:
 Set as child

Namespace Definitions

XQuery Functions

► **Variable Structures**

2. Select the **XML Type**.
3. In the **Structure Label** field, enter `InvoiceType` as the display name for the variable structure you want to create. This display name enables you to give a meaningful name to the structure so you can recognize it at design time but it has no impact at run time.
4. In the **Structure Path** field, enter `$InvoiceType` as the path of the variable at run time.
5. To select the type **InvoiceType**, do the following:
 - a. Under the **Type** field, select the appropriate radio button, then select **WSDL Type** from the drop-down list.
 - b. Click **Browse**. The **WSDL Browser** is displayed.
 - c. In the **WSDL Browser**, select **SampleWSDL**, then select **InvoiceType** under **Types** in the **Select WSDL Definitions** pane.

- d. Click **Submit**. **InvoiceType** is displayed under your selection **WSDL Type**.
6. Click **Add**. The new variable structure **InvoiceType** is included under **XML Type** in the drop-down list of variable structures.

The variable structure **InvoiceType** is displayed as follows:

Figure 2-10 Variable Structures—InvoiceType



Example 3: Creating a Variable Structure that Maps a Variable to an Element

Suppose a temporary variable has the element **Invoice** described in the `SampleWSDL` WSDL. In this example, the `ProxyWithSampleWSDL` message flow needs to access this variable. To achieve this, you can create a new variable structure that maps the variable to the element **Invoice**.

To map a variable to an element:

1. In the Variable Structures panel, click **Add New Structure**.
2. Make sure you select the **XML Type**.
3. In the **Structure Label** field, enter `Invoice` as the meaningful display name for the variable structure you want to create.
4. In the **Structure Path** field, enter `$Invoice` as the path of the variable structure at run time.
5. To select the element **Invoice**, do the following:
 - a. For the **Type** field, make sure you select the appropriate radio button. Then select **WSDL Element** from the drop-down list.
 - b. Click **Browse**.
 - c. In the **WSDL Browser**, select `SampleWSDL`, then select **Invoice** under **Elements** in the **Select WSDL Definitions** pane.
 - d. Click **Submit**. **Invoice** is displayed under your selection **WSDL Element**.

- Click **Add**. The new variable structure **Invoice** is included under **XML Type** in the drop-down list of variable structures.

The variable structure **Invoice** is displayed as follows:

Figure 2-11 Variable Structures—Invoice



Example 4: Creating a Variable Structure That Maps a Variable to a Child Element

The `ProxyWithSampleWSDL` proxy service routes to the document style `Any` SOAP business service that returns the Purchase Order in the SOAP body. In this example, the `ProxyWithSampleWSDL` proxy service message flow must then manipulate the response. To achieve this, you can create a new structure that maps the **body** variable to the **PO** element, and specify the **PO** element as a child element of the variable. You need to specify it as a child element because the **body** variable contains the SOAP **Body** element and the **PO** element is a child of the **Body** element.

To map a variable to a child element:

- In the Variable Structures panel, click **Add New Structure**.
- Make sure you select the **XML Type**.
- In the **Structure Label** field, enter `body to PO` as the meaningful display name for the variable structure you want to create.
- In the **Structure Path** field, enter `$body` as the path of the variable structure at run time.
- To select the **PO** element:
 - Under the **Type** field, make sure you select the appropriate radio button, and then select **WSDL Element** from the drop-down list.
 - Click **Browse**.

- c. In the **WSDL Browser**, select `SampleWSDL`, then select `PO` under **Elements** in the **Select WSDL Definitions** pane.
- d. Click **Submit**.
6. Select the **Set as child** checkbox to set the `PO` element as a child of the **body to PO** variable structure.
7. Click **Add**. The new variable structure **body to PO** is included under **XML Type** in the drop-down list of variable structures.

The variable structure **body to PO** is displayed as follows:

Figure 2-12 Variable Structures—body to PO



Example 5: Creating a Variable Structure that Maps a Variable to a Business Service

The `ProxyWithSampleWSDL` proxy service routes the message to the `BusinessWithSampleWSDL` business service, which also has a service type **WSDL Web Service** that uses the binding `POBinding` from `SampleWSDL`. In this example, the message flow must then manipulate the response. To achieve this, you can define a new structure that maps the **body** variable to the `BusinessWithSampleWSDL` business service. This results in a map of the **body** variable to the SOAP body for all the messages in the WSDL interface of the service.

Note: This mapping is also supported for messaging services with a typed interface.

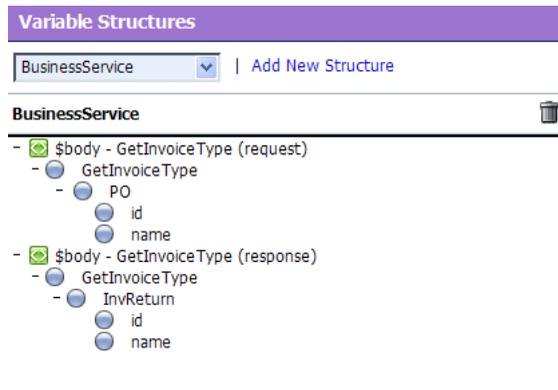
To map a variable to a business service:

1. In the Variable Structures panel, click **Add New Structure**.
2. Select **Service Interface**.
3. In the **Structure Label** field, enter `BusinessService` as the meaningful display name for the variable structure.

4. In the **Structure Path** field, `$body` is already set as the default. This is the path of the variable structure at run time.
5. To select the business service, do the following:
 - a. Under the **Service** field, click **Browse**. The **Service Browser** is displayed.
 - b. In the **Service Browser**, select the `BusinessWithSampleWSDL` business service, then click **Submit**. The business service is displayed under the **Service** field.
 - c. In the **Operation** field, select **All**.
6. Click **Add**. The new variable structure `BusinessService` is included under **Service Interface** in the drop-down list of variable structures.

The variable structure `BusinessService` is displayed as shown below:

Figure 2-13 Variable Structures—BusinessService



Example 6: Creating a Variable Structure That Maps a Child Element to Another Child Element

Modify the `SampleWSDL` so that the `ProxyWithSampleWSDL` proxy service receives a single attachment. The attachment is a Purchase Order. In this example, the proxy service message flow must then manipulate the Purchase Order. To achieve this, you can define a new structure that maps the **body** element in `$attachments` to the `PO` element, which is specified as a child element. The **body** element is specified as a variable path of the form:

```
$attachments/ctx:attachment/ctx:body
```

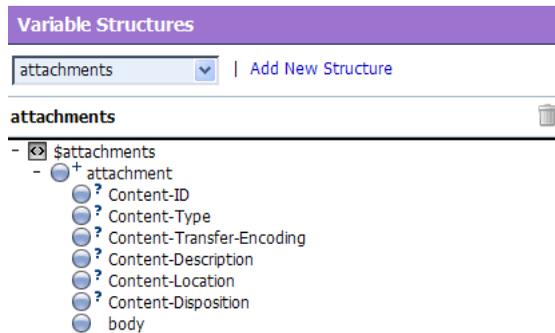
You can select and copy the **body** element from the predefined **attachments** structure, paste this element as the variable path to be mapped in the new mapping definition.

To map a child element to another child element:

1. In the Variable Structures panel, select **attachments** from the drop-down list of built-in structures.

The variable structure **attachments** is displayed as follows:

Figure 2-14 Variable Structures—attachments



2. Select the **body** child element in the attachments structure. The variable path of the body element is displayed in the Property Inspector on the right side of the page:

`$attachments/ctx:attachment/ctx:body`

3. Copy the variable path of the **body** element.
4. In the Variable Structures panel, click **Add New Structure**.
5. Select the **XML Type**.
6. In the **Structure Label** field, enter `po attachment` as the meaningful display name for this variable structure.
7. In the **Structure Path** field, paste the variable path of the body element:

`$attachments/ctx:attachment/ctx:body`

This is the path of the variable structure at run time.

8. To select the `po` element:

- a. Under the **Type** field, make sure the appropriate radio button is selected, then select **WSDL Element**.
 - b. Click **Browse**.
 - c. In the **WSDL Browser**, select `SampleWSDL`, then select **PO** under **Elements** in the **Select WSDL Definitions** pane.
 - d. Click **Submit**.
9. Select the **Set as child** checkbox to set the PO element as a child of the **body** element.
 10. Click **Add**. The new variable structure **PO attachment** is included under **XML Type** in the drop-down list of variable structures.
 11. If there are multiple attachments, add an index to the reference when you use fields from this structured variable in your XQueries. For example, if you drag the PO field to the XQuery field, but as PO will be the second attachment, change the inserted value from


```
$attachments/ctx:attachment/ctx:body/bea:PO/bea:id
```

 to


```
$attachments/ctx:attachment[2]/ctx:body/bea:PO/bea:id
```

Quality of Service

The following sections discuss quality of service features in AquaLogic Service Bus messaging:

- [“Delivery Guarantees” on page 2-55](#)
- [“Outbound Message Retries” on page 2-61](#)

Delivery Guarantees

BEA AquaLogic Service Bus supports reliable messaging. The value of the `qualityOfService` element in the outbound context variable provides AquaLogic Service Bus with a hint on the desired delivery behavior. When messages are routed to another service from a route node, the default `Quality of Service` element in `$outbound` is either `exactly-once` or `best-effort`.

The following delivery guarantee types are provided in AquaLogic Service Bus:

Table 2-5 Delivery Guarantee Types

| Delivery Reliability | Description |
|----------------------|--|
| Exactly once | <p data-bbox="428 430 1162 517"><i>Exactly once</i> means reliability is optimized. <i>Exactly once</i> delivery reliability is a hint, not a directive. When <i>exactly-once</i> is specified, <i>exactly-once</i> reliability is provided if possible.</p> <p data-bbox="428 531 1162 586">The default value of the <code>qualityOfService</code> element is <code>exactly-once</code> for a Route Node action for the following inbound transports:</p> <ul data-bbox="428 600 677 774" style="list-style-type: none"> • e-mail • FTP • File • JMS/XA • Transactional Tuxedo <p data-bbox="428 795 1162 822">Note: Do not retry the outbound transport when the QoS is exactly once</p> |

Table 2-5 Delivery Guarantee Types

| Delivery Reliability | Description |
|----------------------|--|
| At least once | <i>At least once</i> delivery semantics is attempted if <i>exactly once</i> is not possible but the <code>qualityOfService</code> element is <i>exactly-once</i> . |
| Best effort | <p><i>Best effort</i> means that performance or availability is optimized. It is performed if the <code>qualityOfService</code> element is <i>best-effort</i>. <i>Best effort</i> delivery is also performed if <i>exactly once</i> and <i>at least once</i> delivery semantics are not possible but the <code>qualityOfService</code> element is <i>exactly-once</i>.</p> <p>The default value of the <code>qualityOfService</code> element for a route node is <i>best-effort</i> for the following inbound transports:</p> <ul style="list-style-type: none"> • JMS/nonXA • HTTP • HTTP(S) • Non-Transactional Tuxedo <p>The default value of the <code>qualityOfService</code> element is always <i>best-effort</i> for the following:</p> <ul style="list-style-type: none"> • Service callout action — always <i>best-effort</i>, but can be changed if required. • Publish action — defaults to <i>best-effort</i>, modifiable <p>Note: When the value of the <code>qualityOfService</code> element is <i>best-effort</i> for a Publish action, all errors are ignored. However, when the value of the <code>qualityOfService</code> element is <i>best-effort</i> for a Route Node action or a Service callout action, any error will raise an exception.</p> |

Overriding the Default Element Attribute

You can override the default `qualityOfService` element attribute for the following:

- Route Node action
- Publish action
- Service Callout

To override the `qualityOfService` element attribute, you must use the Route Options action to route or publish, and also select the checkbox for a service callout action. See “Message Context Schema” in [Message Context](#) in *Using the AquaLogic Service Bus Console*.

Delivery Guarantee Rules

The delivery guarantee supported when a proxy service publishes a message or routes a request to a business service depends on the following conditions:

- The value of the `qualityOfService` element.
- The inbound transport (and connection factory, if applicable).
- The outbound transport (and connection factory, if applicable).

However, if the inbound proxy service is a Local Transport and is invoked by another proxy service, the inbound transport of the invoking proxy service is responsible for the delivery guarantee. That is because a proxy service that invokes another proxy service is optimized into a direct invocation if the transport of the invoked proxy service is a Local Transport. For more information on transport protocols, see “Adding a Proxy Service” and “Adding a Business Service” in [Proxy Services](#) in *Using the AquaLogic Service Bus Console*.

Note: No delivery guarantee is provided for responses from a proxy service.

The following rules govern delivery guarantees:

Table 2-6 Delivery Guarantee Rules

| Delivery Guarantee Provided | Rule |
|-----------------------------|---|
| <i>Exactly once</i> | The proxy service inbound transport is transactional and the value of the <code>qualityOfService</code> element is <code>exactly-once</code> to an outbound JMS/XA transport. |
| <i>At least once</i> | The proxy service inbound transport is file, FTP, or e-mail and the value of the <code>qualityOfService</code> element is <code>exactly-once</code> . |
| <i>At least once</i> | The proxy service inbound transport is transactional and the value of the <code>qualityOfService</code> element, where applicable, is <code>exactly-once</code> to an outbound transport that is not transactional. |
| No delivery guarantee | All other cases, including all response processing cases. |

Note: To support *at least once* and *exactly-once* delivery guarantees with JMS, you must exploit JMS transactions and configure a retry count and retry interval on the JMS queue to ensure that the message is redelivered in the event of a server crash or a failure that is not handled in an error handler with a *Reply* or *Resume* action. File, FTP, and e-mail transports also internally use a JMS/XA queue. The default retry count for a proxy

service with a JMS/XA transport is 1. For a list of the default JMS queues created by AquaLogic Service Bus, see *AquaLogic Service Bus Deployment Guide*.

The following are some more delivery guarantee rules:

- If the transport of the inbound proxy service is File, FTP, e-mail, Transactional Tuxedo, or JMS/XA, the request processing is performed in a transaction.
 - When the `qualityOfService` element is set to `exactly-once`, any Route node and Publish actions executed in the request flow to a transactional destination are performed in the same transaction.
 - When the `qualityOfService` element is set to `best-effort` for any action in a Route node, Service Callout or Publish actions are executed outside of the request flow transaction. Specifically, for JMS, Tuxedo, Transactional Tuxedo, or EJB transport, the request flow transaction is suspended and the Transactional Tuxedo work is done without a transaction or in a separate transaction that is immediately committed.
 - If an error occurs during request processing, but is caught by a user error handler that manages the error (by using the Resume or Reply action), the message is considered successfully processed and the transaction commits. A transaction is aborted if the system error handler receives the error—that is, if the error is not handled before reaching the system level. The transaction is also aborted if a server failure occurs during request pipeline processing.
- If a response is received by a proxy service that uses a JMS/XA transport to business service (and the proxy inbound is not Transactional Tuxedo), the response processing is performed in a single transaction.
 - When the `qualityOfService` element is set to `exactly-once`, all Route, Service Callout, and Publish actions are performed in the same transaction.
 - When the `qualityOfService` element is set to `best-effort`, all Publish actions and Service Callout actions are executed outside of the response flow transaction. Specifically, for JMS, EJB, or transactional Tuxedo types of transports, the response flow transaction is suspended and the service is invoked without a transaction or in a separate transaction that is immediately committed.
 - Proxy service responses executed in the response flow to a JMS/XA destination are always performed in the same transaction, regardless of the `qualityOfService` element setting.
- If the proxy service inbound transport is transactional Tuxedo, both the request processing and response processing are done in this transaction.

Note: You will encounter a run-time error when the inbound transport is transactional Tuxedo and the outbound is an asynchronous transport, for example, JMS/XA.

Threading Model

The BEA AquaLogic Service Bus threading model works as follows:

- The request and response flows in a proxy service execute in different threads.
- Service callouts are always blocking. An HTTP route or publish action is non-blocking (for request/response or one-way invocation), if the value of the `qualityOfService` element is `best-effort`.
- JMS Route actions or Publish actions are always non-blocking, but the response is lost if the server restarts after the request is sent because AquaLogic Service Bus has no persistent message processing state.

Note: In a request or response flow Publish action, responses are always discarded because Publish actions are inherently a one-way message send.

Splitting Proxy Services

You may want to split a proxy service in the following situations:

- When HTTP is the inbound and outbound transport for a proxy service, you may want to incorporate enhanced reliability into the middle of the message flow. To enable enhanced reliability in this way, split the proxy service into a front-end HTTP proxy service and a back-end JMS (one-way or request/response) proxy service with an HTTP outbound transport. In the event of a failure, the first proxy service must quickly place the message in the queue for the second proxy service, in order to avoid loss of messages.
- To disable the direct invocation optimization for a non-JMS transport when a proxy service, say `loanGateway1` invokes another proxy service, say `loanGateway2`. Route to the proxy service `loanGateway2` from the proxy service `loanGateway1` where the proxy service `loanGateway2` uses JMS transport.
- To have an HTTP proxy service publish to a JMS queue but have the Publish action rollback if there is a exception later on in the request processing, split the proxy service into a front-end HTTP proxy service and a back-end JMS proxy service. The Publish action specifies a `qualityOfService` element of `exactly-once` and uses an XA connection factory.

Outbound Message Retries

In addition to configuring inbound retries for messages using JMS, you can configure outbound retries and load balancing. Load balancing, failover, and retries work in conjunction to provide performance and high availability. For each message, the list of URLs you provide as failover URLs is automatically ordered based on the load balancing algorithm into a failover sequence. If the retry count is N, the entire sequence is retried N times before stopping. The system waits for the specified retry interval before commencing subsequent loops through the sequence. After completing the retry attempts, if there is still an error, the error handler pipeline for the route node is invoked. For more information on the error handler pipeline, see “Adding Pipeline Error Handling” in [Proxy Services](#) in *Using the AquaLogic Service Bus Console*.

Note: For HTTP and HTTP(S) transports, any HTTP status other than 200 or 202 is considered an error by AquaLogic Service Bus and must be retried. Because of this algorithm, it is possible that AquaLogic Service Bus retries errors like authentication failure that may never be rectified for that URL within the time period of interest. On the other hand, if AquaLogic Service Bus also fails over to a different URL for subsequent attempts to send a given message, the new URL may not give the error.

For `quality of service=exactly once` failover or retries will not be executed.

Content Types, JMS Type, and Encoding

To support interoperability with heterogeneous endpoints, AquaLogic Service Bus allows you to control the content type used, the JMS type used, and the encoding used.

AquaLogic Service Bus does not make assumptions about what the external client or service needs, and uses the information configured for this purpose in the service definition. AquaLogic Service Bus derives the content type for outbound messages from the service type and interface. Content type is a part of the e-mail and HTTP(S) protocols.

If the service type is:

- XML or SOAP with or without a WSDL, the content type is text/XML.
- Messaging and the interface is MFL or binary, the content type is binary/octet-stream.
- Messaging and the interface is text, the content type is text/plain.
- Messaging and the interface is XML, the content type is text/XML.

Additionally, there is a JMS type, which can be byte or text. You configure the JMS type to use when you define the service in AquaLogic Service Bus Console.

You can override the content type in the outbound context variable (`$outbound`) for proxy services invoking a service, and in the inbound context variable (`$inbound`) for a proxy service response. For more information on `$outbound` and `$inbound` context variables, see [Message Context](#) in *Using the AquaLogic Service Bus Console*.

Encoding is also explicitly configured in the service definition for all outbound messages. For more information on service definitions, see [Adding a Proxy Service](#) in and [Adding a Business Service](#) in *Using the AquaLogic Service Bus Console*.

Throttling Pattern

A throttling pattern is typically used with an HTTP Web service to restrict the degree of concurrency, that is to keep the number of outstanding requests without a response below a limit. Instead of accessing the business service directly, you access the business service through another proxy service. This proxy service typically uses the JMS one-way transport or JMS request response transport to communicate with the business service. You should define a work manager for the JMS request queue. For more information on defining a work manager, see [Work Manager](#). Configure the work manager to have the maximum number of threads. This restricts the number of requests that can be placed in the request queue. That is no requests can be placed in the queue if the number of incoming request exceeds the maximum number of threads configured in the work manager.

Note: Set the `qualityOfService` of the business service `$outbound` to `Exactly Once`. You can use the **Routing Options** action to set the required `qualityOfService`. For more information, see [Routing Options](#) *Using the AquaLogic Service Bus Console*.

When a throttling pattern is implemented in a cluster, the total number of requests across all the Managed Servers should be equal to
maximum number of threads on the work manager / number of managed servers

WS-I Compliance

BEA AquaLogic Service Bus provides Web Service Interoperability (WS-I) compliance in the run-time environment. The WS-I basic profile has the following goals:

- Disambiguate the WSDL and SOAP specifications wherever ambiguity exists.
- Define constraints that can be applied when receiving messages or importing WSDLs so that interoperability is enhanced. When messages are sent, construct the message so that the constraints are satisfied.

The WS-I basic profile is available at the following URL:

<http://www.ws-i.org/Profiles/BasicProfile-1.1.html>.

When you configure a proxy service or business service based on a WSDL, you can use the AquaLogic Service Bus Console to specify whether you want AquaLogic Service Bus to enforce WS-I compliance for the service. For more information on how to do this, see “Adding a Proxy Service” in [Proxy Services](#) in *Using the AquaLogic Service Bus Console*.

When you configure WS-I compliance for a proxy service, checks are performed on inbound request messages received by that proxy service. When you configure WS-I compliance for an invoked service, checks are performed when any proxy receives a response message from that invoked service. BEA recommends that you create an error handler for these errors, since by default, the proxy service SOAP client receives a system error handler-defined fault. For more information on creating fault handlers, see “Error Messages and Handling” in [Proxy Services](#) in *Using the AquaLogic Service Bus Console*.

For messages sent from a proxy service, whether as outbound request or inbound response, WS-I compliance checks are not explicitly performed. That is because the pipeline designer is responsible for generating most of the message content. However, the parts of the message generated by AquaLogic Service Bus should satisfy all of the supported WS-I compliance checks. This includes the following content:

- Service invocation request message.
- System-generated error messages returned by a proxy service.
- HTTP status codes generated by a proxy service.

The Enforce WS-I Compliance checkbox is displayed as shown in [Figure 2-15](#):

Figure 2-15 Enforce WS-I Compliance Checkbox

Edit a Proxy Service - Operation Selection Configuration (Path - default)

Enforce WS-I Compliance

Selection Algorithm

- Transport Header
- SOAPAction Header
- WS-Addressing
- SOAP Header
- SOAP Body Type

<< Back Next >> Finish Cancel

WS-I Compliance Checks

Note: WS-I compliance checks require that the system knows what operation is being invoked on a service. For request messages received by a proxy service, that means that the context variable `$operation` should not be null. That depends upon the operation selection algorithm being configured properly. For response messages received from invoked services, the operation should be specified in the action configurations for Route, Publish, and Service Callout.

When you configure WS-I compliance checking for a proxy service or a business service, AquaLogic Service Bus carries out the following checks.

Table 2-7 AquaLogic Service Bus WS-I Compliance Checks

| Check | WS-I Basic Profile Details | AquaLogic Service Bus Description |
|---|---|---|
| 3.1.1 SOAP Envelope Structure | R9980 An Envelope must conform to the structure specified in SOAP 1.1, Section 4, “SOAP Envelope” (subject to amendment). | This check applies to request and response messages. If a response message is checked and the message does not possess an outer <code>Envelope</code> tag, a <code>soap:client</code> error is generated. If the message is an <code>Envelope</code> tag but possesses a different namespace, it is handled by the 3.1.2 SOAP Envelope Namespace. |
| 3.1.2 SOAP Envelope Namespace | R1015 A Receiver must generate an error if they encounter an envelope whose document element is not <code>soap:Envelope</code> . | This check applies to request and response messages and is related to the 3.1.1 SOAP Envelope Structure. If a request message has a local name of <code>Envelope</code> , but the namespace is not SOAP 1.1, a <code>soap:VersionMismatch</code> error is generated. |
| 3.1.3 SOAP Body Namespace Qualification | R1014 The child elements of the <code>soap:body</code> element in an Envelope must be namespace qualified. | This check applies to request and response messages. All request error messages generate a <code>soap:Client</code> error. |
| 3.1.4 Disallowed Constructs | R1008 An Envelope must not contain a Document Type Declaration. | This check applies to request and response messages. All request error messages generate a <code>soap:Client</code> error. |
| 3.1.5 SOAP Trailers | R1011 An Envelope must not have any child elements of <code>soap:Envelope</code> following the <code>soap:body</code> element. | This check applies to request and response messages. All request error messages generate a <code>soap:Client</code> error. |

Table 2-7 AquaLogic Service Bus WS-I Compliance Checks

| Check | WS-I Basic Profile Details | AquaLogic Service Bus Description |
|--|--|--|
| 3.1.9 SOAP attributes on SOAP 1.1 elements | <p>R1032 The <code>soap:Envelope</code>, <code>soap:header</code>, and <code>soap:body</code> elements in an Envelope must not have attributes in the namespace <code>http://schemas.xmlsoap.org/soap/envelope/</code></p> | This check applies to request and response messages. Any request error messages generate a <code>soap:client</code> error. |
| 3.3.2 SOAP Fault Structure | <p>R1000 When an Envelope is a fault, the <code>soap:Fault</code> element must not have element children other than <code>faultcode</code>, <code>faultstring</code>, <code>faultactor</code>, and <code>detail</code>.</p> | This check only applies to response messages. |
| 3.3.3 SOAP Fault Namespace Qualification | <p>R1001 When an Envelope is a Fault, the element children of the <code>soap:Fault</code> element must be unqualified.</p> | This check only applies to response messages. |
| 3.4.6 HTTP Client Error Status Codes | <p>R1113 An instance should use a “400 Bad Request” HTTP status code if a HTTP request message is malformed.</p> <p>R1114 An instance should use a “405 Method not Allowed” HTTP status code if a HTTP request message is malformed.</p> <p>R1125 An instance must use a 4xx HTTP status code for a response that indicates a problem with the format of a request.</p> | Only applies to responses for a proxy service where you cannot influence the status code returned due to errors in the request. |
| 3.4.7 HTTP Server Error Status Codes | <p>R1126 An instance must return a “500 Internal Server Error” HTTP status code if the response envelope is a fault.</p> | This check applies differently to request and response messages. For request messages, any faults generated have a 500 Internal Server Error HTTP status code. For response messages, an error is generated if fault responses are received that do not have a 500 Internal Server Error HTTP status code. |

Table 2-7 AquaLogic Service Bus WS-I Compliance Checks

| Check | WS-I Basic Profile Details | AquaLogic Service Bus Description |
|------------------------------|---|--|
| 4.7.19 Response Wrappers | R2729 An envelope described with an <code>rpc-literal</code> binding that is a response must have a wrapper element whose name is the corresponding <code>wSDL:operation</code> name suffixed with the string <code>Response</code> . | This check only applies to response messages. AquaLogic Service Bus never generates a non-fault response from a proxy service. |
| 4.7.20 Part Accessors | R2735 An envelope described with an <code>rpc-literal</code> binding must place the part accessor elements for parameters and return value in no namespace. R2755 The part accessor elements in a message described with an <code>rpc-literal</code> binding must have a local name of the same value as the name attribute of the corresponding <code>wSDL:part</code> element. | This check applies to request and response messages. Any request error messages generate a <code>soap:client</code> error. |
| 4.7.22 Required Headers | R2738 An envelope must include all <code>soapbind:headers</code> specified on a <code>wSDL:input</code> or <code>wSDL:output</code> of a <code>wSDL:operation</code> of a <code>wSDL:binding</code> that describes it. | This check applies to request and response messages. Any request error messages generate a <code>soap:client</code> error. |
| 4.7.25 Describing SOAPAction | R2744 A HTTP request message must contain a SOAPAction a HTTP header field with a quoted value equal to the value of the <code>soapAction</code> attribute of <code>soap:operation</code> , if present in the corresponding WSDL description. R2745 A HTTP request message must contain a SOAP action a HTTP header field with a quoted empty string value, if in the corresponding WSDL description, the SOAPAction of <code>soapbind:operation</code> is either not present, or present with an empty string as its value. | This check applies to request messages and a <code>soap:client</code> error is returned. |

Modeling Message Flow in AquaLogic Service Bus

Monitoring

BEA AquaLogic Service Bus provides the capability to monitor and collect run-time information for systems operations purposes. AquaLogic Service Bus aggregates run-time statistics that you can view on a customizable Dashboard. The Dashboard allows you to monitor the health of the system and alerts you to problems in your messaging services. With this information, you can quickly and easily isolate and diagnose problems as they occur.

This chapter includes the following sections:

- [“Monitoring Scenarios” on page 3-1](#)
- [“About Monitoring” on page 3-8](#)
- [“Service Summary” on page 3-14](#)
- [“Server Summary” on page 3-21](#)
- [“Alert Summary” on page 3-29](#)
- [“Alert Rules” on page 3-38](#)
- [“Statistics Associated With Different Resources” on page 3-42](#)
- [“Auditing” on page 3-44](#)

Monitoring Scenarios

The following sections describe some of the tools and functionality available in AquaLogic Service Bus Console to monitor messages and system operations. It includes:

- “Operational Health” on page 3-2
- “Monitoring Alerts” on page 3-3
- “Verifying Service Level Agreements” on page 3-4
- “Pipeline Alert Action” on page 3-4
- “Alert Destination” on page 3-5

Operational Health

The Dashboard page in the AquaLogic Service Bus Console provides the ability to view the state of all servers and monitored services immediately. The Dashboard displays two pie charts, a table, and several links. The Service Summary pie chart shows the percentage of alerts according to their severity for all services that were issued in the past 30 minutes. The Server Summary pie chart shows the current status of every server in the AquaLogic Service Bus domain. Additionally, from the Server Summary panel, you can drill down and view the domain logs, which are grouped according to severity.

In addition to the pie charts, these Summaries include a list of the most active services and critical servers. The list displays up to ten services, with fully qualified service names, in descending order of the most number of alerts. The most critical server list displays the ten most critical servers. This display is based on the health of the running servers, as defined by the WebLogic Diagnostic Service. For more information about the WebLogic Diagnostic Service, see [Configuring and Using the WebLogic Diagnostics Framework](#).

From each of the summaries, you can drill down into more detail by clicking a specific area on a pie chart or by clicking one of the links on the page.

The Alert Summary table lists the alerts that were issued in the past 30 minutes. This table contains the following fields:

- Alert Severity: This field specifies if the severity of the alert. By default the table is sorted by alert severity, in the following order:
 - a. Fatal
 - b. Critical
 - c. Major
 - d. Minor

- e. Warning
 - f. Normal
 - Time Stamp: This field gives the details of when the error occurred in MM/DD/YY HH:MM format
 - Service: This field gives the path and the name of the service.
 - Alert Rule Name: This field gives the name of the alert rule configured for the service.
- Note:** This is available only in case of SLA violations

You can customize the layout of this table.

Monitoring Alerts

When you log into the AquaLogic Service Bus Console, you may see a list of alerts on the Dashboard. This display is dynamically refreshed. These alerts could be the result of SLA violations or pipeline alerts. Service Level Agreements (SLAs) are agreements that define the precise level of service expected from the AquaLogic Service Bus business and proxy services.

Each row of the table displays the information that you have configured, such as the severity, timestamp, and associated service. Clicking the severity link will display more details about the alert to help analyze the cause of the alert.

Console, e-mail, JMS, reporting or SNMP traps are the various alert destinations that can be configured for the alert. For example, you can choose e-mail or JMS as additional or replacement destination for alert notification.

Monitoring Statistics

Monitoring Statistics helps you know how many messages in a particular service have processed successfully and how many have failed. To access this information, from the Dashboard, you access the Service Monitoring Summary page and filter the display for the relevant service. Besides displaying the number of messages that have been processed successfully or failed, you can also see which project the service belongs to, the average execution time of message processing, and the number of alerts associated with the service. You can view monitoring statistics for the period of the current aggregation interval or for the period since you last reset statistics for this service or since you last reset statistics for all services.

You use the Global Settings page in the System Administration module of the AquaLogic Service Bus Console to reset statistics. When you do this, make sure you are not in a WebLogic session on the WebLogic Server Administration Console.

Clicking the name of the service brings you to that service's Service Monitoring Details page. This page provides additional information such as the minimum and maximum response times and the overall average time it takes for the service to execute a message, the success-failure ratio, the number of messages that have failed because of security or validation errors, and the number of messages associated with proxy service components (pipelines and route nodes). You can view this information for specific operations associated with the service. Again, you can view these statistics for the period of the current aggregation interval or you can display the statistics for the period since you last reset statistics for this service or since you last reset statistics for all services.

Verifying Service Level Agreements

Consider the following use case to verify the service level agreements:

Assume that a particular proxy service is generating a lot of SLA violation alerts due to slow response time. To investigate this problem further, you must log into the AquaLogic Service Bus Console and take a look at the detailed statistics for the proxy service. At this level, you will be able to identify that, a third-party web service invocation stage in the pipeline is taking a lot of time and is the actual bottleneck. After successfully renegotiating service-level characteristics with the third-party web service provider, you could configure alert metrics to track the web service provider's compliance with the new agreement terms. Thus you can use alerts as the basis for negotiating Service Level Agreements.

Pipeline Alert Action

You can also generate alerts inside a stage in the pipeline using the Alert action. For this you use the `Alert` action in the `Reporting` category of the Actions menu.

You define conditions under which a pipeline alert is triggered using the conditional constructs available in the Pipeline Editor such as Xquery Editor or an if-then-else construct. You can use the `Alert Destination` resource in an alert action to define the destination for alert. You will have complete control over the alert body including the pipeline, and context variables. Also you will be able to extract the portions of the message.

You can obtain an integrated view of all the alerts generated by a service on the Dashboard page in the AquaLogic Service Bus Console.

Note: For more information on adding Alert action in a stage, see [Proxy Services: Actions-Alert](#) in *Using the AquaLogic Service Bus Console*

Alert Destination

You can view the alerts by using one of the following alert destinations:

- [“AquaLogic Service Bus Console”](#) on page 3-5
- [“E-mail Alert Destination”](#) on page 3-5
- [“SNMP Traps”](#) on page 3-6
- [“JMS”](#) on page 3-8
- [“Reporting”](#) on page 3-8

AquaLogic Service Bus Console

The Dashboard shows the overall health related information of AquaLogic Service Bus. It provides an overview of the state of the system organized by server, services, and alerts.

After monitoring is enabled, the Service Monitoring Summary page in the AquaLogic Service Bus Console provides a view of the statistics collected for each service. It provides information about the alerts generated due to SLA violations or as a result of alert actions configured in the pipeline.

As previously mentioned, an SLA is an agreement that defines the precise level of service expected from business and proxy services in AquaLogic Service Bus. The SLA Manager, with the help of the AquaLogic Service Configuration module, allows users to configure SLA rule conditions and actions. The SLA Manager monitors SLA violations with the help of data provided by the Aggregator and sends notifications as configured in the alert rule actions. The SLA Manager is always deployed with the Aggregator and resides on only one managed server in cluster. The SLA Manager sends alerts to the Alert Log to store in the Alert Store.

E-mail Alert Destination

This is one of the destinations for the alerts. To configure this alert destination you have to first configure the SMTP global resource. This resource captures the address of the SMTP server corresponding to your e-mail destination, port number, and if required, the authentication credentials. The authentication credentials are stored inline and are not stored as a service account. The alert action makes use of the SMTP resource to send the outbound e-mail messages. You can also use the SMTP resource to send both pipeline alerts and SLA alerts. When an alert is delivered

over an e-mail the metadata consisting of the details about the alert is prefixed to the payload configured.

Note: For more information on SMTP Server resource, see [Overview of SMTP Servers](#) in *Using the AquaLogic Service Bus Console*

SNMP Traps

The Simple Network Management Protocol (SNMP) traps allows any third party software to interface monitoring Service Level Agreements (SLAs) within AquaLogic Service Bus. By enabling the notification of alerts using SNMP, Web Services Management (WSM) and the Enterprise Service Management (ESM) tools can monitor SLA violations by monitoring alert notifications.

Simple Network Management Protocol (SNMP) is an application-layer protocol which allows the exchange of information on the management of a resource across a network. It enables you to monitor a resource and if required, rectify it based on the data obtained from the resource. Both the SNMP version 1 and SNMP version 2 are supported in this version of the AquaLogic Service Bus. SNMP is made up of the following components:

- “Managed Resource” on page 3-6
- “Management Information Base(MIB)” on page 3-6
- “SNMP Agent” on page 3-7
- “SNMP Manager” on page 3-8
- “Network Management System (NMS)” on page 3-8

Managed Resource

This is the resource, which is being monitored. The resource and its attributes are added to the Management Information Base(MIB).

Management Information Base(MIB)

The Management Information Base (MIB) is a hierarchical data structure that stores all the resources to be monitored, in a hierarchical manner. It also stores the attributes of the resources, which are monitored. Each resource is given a unique identifier called the Object Identifier(OID). You can use the SNMP commands to retrieve the information on the management of a resource. The following section gives an illustration of the WebLogic Server MIB.

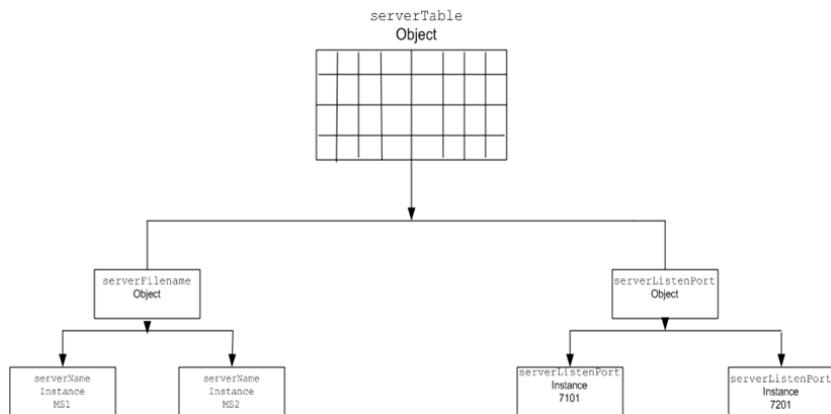
An Illustration of WebLogic Server MIB

The Weblogic Server installer creates a copy of the MIB in the following location:

```
<BEA_HOME>/weblogic92/server/lib/BEA-WEBLOGIC-MIB.asn1
```

where <BEA_HOME> is the directory in which you installed the WebLogic Server. WebLogic Server exposes thousands of data points in its management system. To organize this data it provides a hierarchical data model that reflects the collection of services and resources that are available in a domain. [Figure 3-1](#) illustrates the hierarchy of objects in the MIB.

Figure 3-1 Hierarchy of Objects in MIB



For example, if you created two managed servers, MS1 and MS2, in a domain, then MIB contains one object `serverTable`, which in turn contains one `serverName` object. The `serverName` object in turn contains two instances containing values MS1 and MS2. The MIB assigns a unique number called an object identifier (OID) to each managed object. Once assigned the you cannot change the OID. Each OID consists of a sequence of integers. This sequence defines the location of the object in the MIB tree. Each node in the path has both a number and a name associated with it.

For more information on WebLogic Server MIBs see WebLogic Server documentation at [WebLogic Server® 9.2 MIB Reference](#).

SNMP Agent

Each managed resource uses an SNMP agent to update the relevant information in the MIB. For this you should configure the SNMP agent to detect certain conditions within a managed resource and send trap notification (report) to the SNMP manager. You can configure the SNMP agent to generate traps in one of the following ways:

- Automatically: You can configure the SNMP agent to generate traps for events such as server startup or server shut down.
- Using log messages: Using filters, you can configure the SNMP agent to detect specific log messages and generate traps.
- Monitoring traps: You can create JMX API clients to monitor the changes in the attributes and notify SNMP agent to generate traps. You can also configure the SNMP agents to monitor the changes in the attribute.

SNMP Manager

The SNMP manager manages the SNMP agents. SNMP is also it is the primary interface to the Network Management System.

Network Management System (NMS)

The Network Management System forms the interface with the user. It gathers data using the SNMP manager and presents it to the user.

JMS

Java Messaging Service (JMS) is another destination for a pipeline alert and a SLA alert. You will have use a JNDI URL for the JMS destination for alerts. When you configure an alert rule to post a message to a JMS destination, you must create a JMS connection factory and a queue or topic, and target them to the appropriate JMS server in the WebLogic Server Administration Console. For information on how to do this, see “Configuring a JMS Connection Factory” and “JMS Resource Naming Rules for Domain Interoperability” in [Configuring JMS System Resources](#) in *Configuring and Managing WebLogic JMS*. When you define the JMS alert destination you can either use a destination queue or a destination topic. The message type can be bytes or text. For more information on how to configure JMS alert destination see [Alert Destinations](#) in *Using the AquaLogic Service Bus Console*.

Reporting

This is another process to monitor and analyze both pipeline alerts and SLA alerts. This process of monitoring is discussed in detail in [Chapter 5, “Reporting”](#)

About Monitoring

This section contains information on the following topics:

- [“Aggregation Interval” on page 3-9](#)
- [“Monitoring Architecture” on page 3-10](#)
- [“Monitoring Services” on page 3-11](#)
- [“Refresh Rate of Monitored Information” on page 3-11](#)
- [“Dashboard” on page 3-12](#)

Aggregation Interval

In AquaLogic Service Bus, the monitoring subsystem collects statistical information, such as message-count, execution time, over an aggregation interval. The aggregation interval is the time period over which statistical data is collected and displayed in the AquaLogic Service Bus Console.

Following is an illustration of how the aggregation interval works:

Consider a proxy service you have configured for processing a purchase order, for which you have enabled with an aggregation interval of 10 minutes. When you send the first message through the proxy service, monitoring is started. Until the first ten minutes elapse, the Service Summary page displays the partially computed data. At this time the system does not have 10 minutes of data. After the first 10 minutes of data aggregation, the system always displays the last 10 minutes of data. For example, at the 14th minute, the Dashboard displays minutes 4 through 14. If no messages are processed after the 15th minute, on the 25th minute, the Service does not display any data. For more information about how aggregation interval affects the display of monitored information, see [“Alert Rules” on page 3-38](#).

You must explicitly enable monitoring for any business or proxy service that you create; monitoring is disabled by default. After you have enabled monitoring and set the aggregation interval for your individual services, you can enable or disable monitoring for all those services from the Global Settings page in the System Administration module. For more information, see [“Monitoring Services” on page 3-11](#).

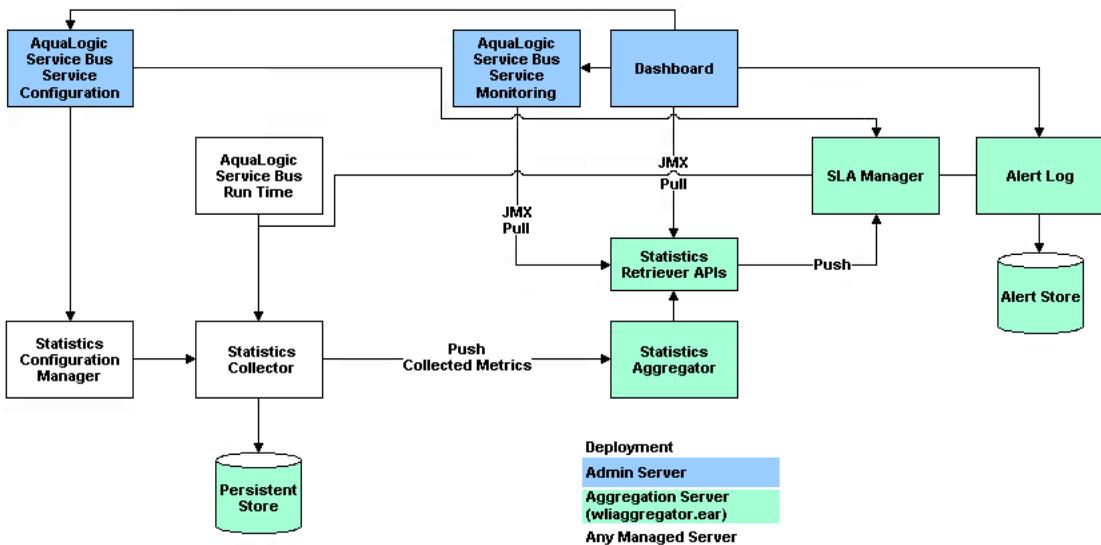
SLA alerts are automated responses to Service Level Agreements (SLAs) violations or occurrences, which are displayed on the Dashboard. You define alert rules to specify unacceptable service performance according to your business and performance requirements. Each alert rule allows you to specify the aggregation interval for that rule when configuring the alert rule. This aggregation interval is not affected by the aggregation interval set for the service. Alert rules also allow you to send notifications to the configured alert destinations on topic about

the violation. For information on defining alert rules, see [Creating Alert Rules](#) in *Using the AquaLogic Service Bus Console*

Monitoring Architecture

The following diagram shows the architecture of AquaLogic Service Bus monitoring.

Figure 3-2 Monitoring Architecture



The Statistics Configuration Manager stores and manages the statistics configuration for each operational resource. An operational resource is defined as the unit for which statistical information can be collected by the monitoring subsystem. An operational resource includes a proxy service, service operations, and pipelines. The Statistics Configuration Manager is notified about changes in the service definition, such as adding, updating, or deleting a pipeline.

Each managed server in a cluster hosts a Statistics Collector. The Statistics Collector collects statistics on operational resources as directed by the Statistics Configuration Manager. The Statistics Collector also keeps samples history within the aggregation interval for the collected statistics. At every system-defined checkpoint interval, the Statistics Collector stores a snapshot of current statistics into a persistent store for recovery purposes and sends the information to the Statistics Aggregator.

One of the managed servers in a cluster, called the *Aggregating Server* or *Aggregator*, is designated as the aggregator for cluster-wide statistics. At system-defined checkpoint intervals, each managed server in the cluster sends a checkpoint snapshot of its contributions to the Aggregator. The Aggregator then combines this information to offer cluster-wide statistics to its clients through Retriever APIs. The clients of Aggregator are the Dashboard, SLA Manager, and Service Monitoring modules.

To contribute a data point to the system, an operational resource in the system, such as a run-time proxy service pipeline, calls a method on the Statistics Collector, and identifies itself, the statistic, and the data point.

Monitoring Services

When you create a business or proxy service, monitoring is disabled by default for that service. Enable monitoring as follows:

- To enable monitoring for an individual service, select the Enable Monitoring checkbox on the Manage Monitoring page. Then set the aggregation interval for the service by selecting the interval times from the hour and minute drop-down lists. For information on how to do this, see “Viewing the Dashboard Statistics” in [Monitoring](#) in the *Using the AquaLogic Service Bus Console*.
- To enable monitoring for all services, select the Enable Monitoring checkbox on the Global Settings page. For information on how to do this, see “Enabling Monitoring” in [System Administration](#) in the *Using the AquaLogic Service Bus Console*.

Note: The Enable Monitoring option permits you to enable or disable monitoring of all services that have individually been enabled for monitoring. If monitoring for a particular service has *not* been enabled, you must first enable it and set the aggregation interval on the **Manage Monitoring** page before the system starts collecting statistics for that service.

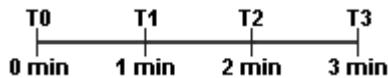
When creating alert rules, you must enable monitoring before you create the rule. For more information, see “[Alert Rules](#)” on page 3-38 and “Create an Alert Rule” in [Monitoring](#) in the *Using the AquaLogic Service Bus Console*.

Refresh Rate of Monitored Information

At run time, the default refresh rate for the Dashboard page is one minute. However, it may take up to three minutes for the information to be displayed on the Dashboard. This delay occurs because of the time gaps between when the messages are processed by the proxy service, when the metrics are collected, and the refresh rate of the Dashboard. The system works as follows:

1. Every minute the Statistics Collector sends the current snapshot to the aggregator.
2. Every minute, the aggregator merges all the documents it has received from the managed servers within the last minute.
3. The AquaLogic Service Bus Console refreshes every minute; that is, it runs a query on the aggregated document and then displays the results.

Figure 3-3 Aggregation Time Line



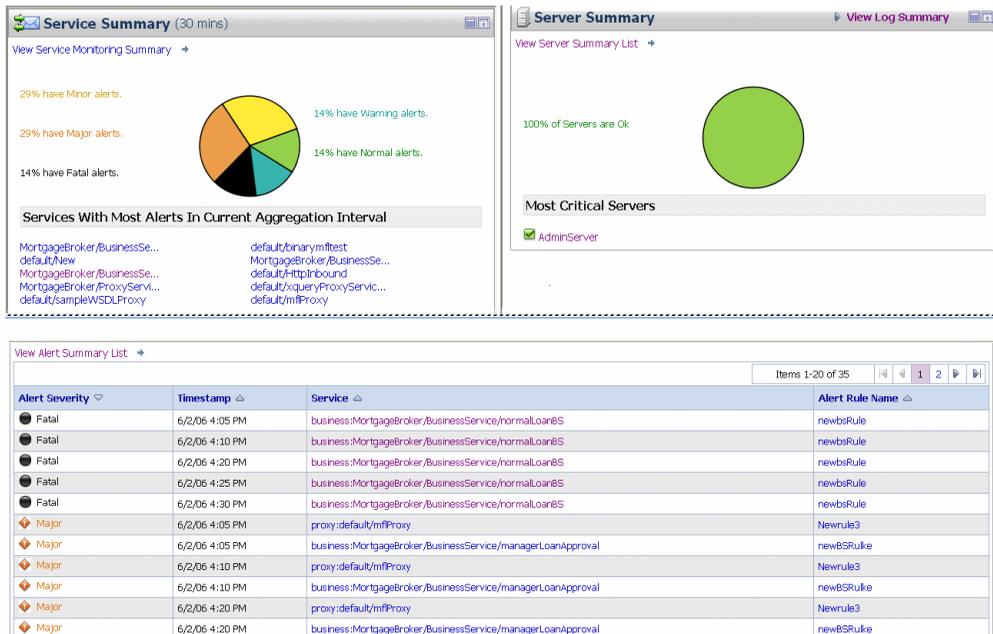
For example, a proxy service starts sending data in T1, as shown in [Figure 3-3](#). At T2—that is, the second minute—the Statistics Collector sends the data to the aggregator. However, if an aggregation cycle has just occurred, the aggregator does not merge this data until the next aggregation cycle, which occurs after one minute, or a maximum of two minutes from the previous aggregation cycle. When the data is merged, it is now available for the AquaLogic Service Bus Console. Since the console refreshes every minute, if the refresh cycle has just passed, but the console displays the alerts after a maximum time of three minutes.

You can change the Dashboard polling interval in the System Administration module in the AquaLogic Service Bus Console. For information on how to do this, see “Setting the Dashboard Polling Interval Refresh Rate” in [System Administration](#) in the *Using the AquaLogic Service Bus Console*.

Dashboard

When you log onto the AquaLogic Service Bus Console, the Dashboard is automatically displayed. The Dashboard shows the monitoring information for the last 30 minutes. It provides an overview of the state of the system—organized by server, services, and alerts, as shown in the following figure.

Figure 3-4 AquaLogic Service Bus Dashboard



As shown in the previous figure the Dashboard displays the following information:

- **Services Summary:** This summarizes the alert status for both proxy and business services if alerts have been configured. Alerts notify you of any violations in the service level agreements or if any alert action condition, which is defined in a pipeline is met.
- **Servers Summary:** displays the status of the servers.
- **Alerts Summary:** if alerts have been configured, displays which alert rules have been triggered.

From the Dashboard, you can drill-down into the system and easily find specific information, such as the average execution time of a service, the date and time an alert occurred, or the duration for which server has been running.

You configure the Dashboard and monitoring in the AquaLogic Service Bus Console, which is described in the [Monitoring](#) and [System Administration](#) sections of *Using the AquaLogic Service Bus Console*.

Service Summary

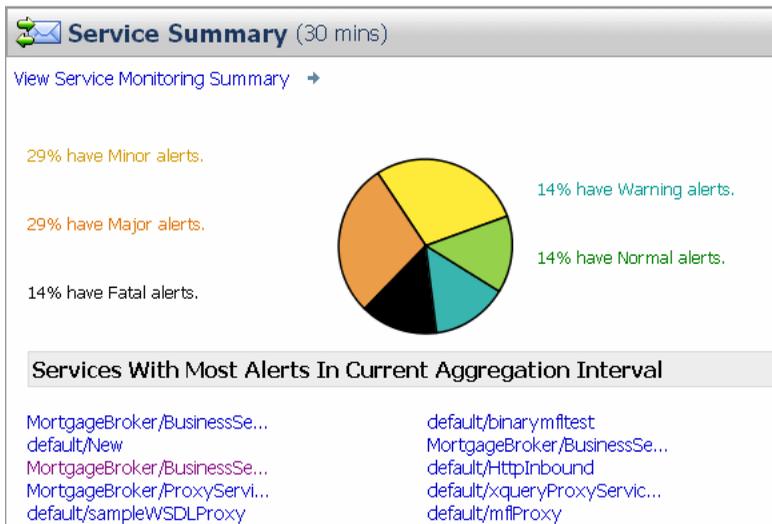
This section provides information on the following topics:

- “About the Service Summary” on page 3-14
- “Service Monitoring Summary” on page 3-15
- “Service Monitoring Details” on page 3-17

About the Service Summary

The Service Summary panel provides an overview of the state of the services. The Service Summary pie chart shows the percentage of alerts according to their severity for all services for which alerts are defined and monitoring is enabled for the last 30 minutes. The severity level of alerts is user configurable and has no absolute meaning. Severity types include Fatal, Critical, Major, Minor, Warning, and Normal. The services having the most number of alerts are listed beneath the pie chart, as shown in the following figure. Up to ten services are listed in descending order of services with the most alerts.

Figure 3-5 Services Summary Pane



From the Service Summary panel, you can access more information about alerts by clicking the following:

- A specific area on a pie chart: displays the Alert History page for alerts for the given level of severity.
- The name of a service under Services With Most Alerts In Current Aggregation Interval: displays the Service Monitoring Details page for that service.
- View Service Monitoring Summary: displays the Service Monitoring Summary page. To help you locate specific services, you can filter the services by different criteria.

Each of these pages is fully described in the sections that follow.

WARNING: When a service (or its component; for example, a pipeline node) is renamed or relocated, its statistical data is lost.

For information on how to access detailed alert information, see “Viewing the Dashboard Statistics” in [Monitoring](#) in the *Using the AquaLogic Service Bus Console*.

Service Monitoring Summary

The Service Monitoring Summary page provides two views of service monitoring statistics, as shown in the following figures.

The first is a dynamic view of statistical data collected by each service. This view is available when you select Current Aggregation Interval in the Show Metrics For field. The aggregation interval displayed in this view determines the statistics that are displayed. For example, if the aggregation interval of a particular service is 20 minutes, that service’s row displays the data collected in the last 20 minutes.

Figure 3-6 Service Monitoring Summary Page—Current Aggregation Interval

The screenshot shows the 'Service Monitoring Summary' page. At the top, there is a dropdown menu for 'Show Metrics For' set to 'Current Aggregation Interval'. Below this are search filters for 'Name' and 'Path', and radio buttons for 'Has Alerts', 'Has Errors', and 'Invoked by proxy:'. There are 'Search' and 'View All' buttons. The main content is a table with the following data:

| Name | Path | Aggregation Interval | Avg. Exec Time | Message Count | Error Count | Alert Count |
|-------------------|---------|-------------------------|----------------|---------------|-------------|-------------|
| Business Service1 | default | 0 Hour(s) and 1 Minutes | 0 | 0 | 0 | 0 |
| Proxy Service 1 | default | 0 Hour(s) and 1 Minutes | 0 | 0 | 0 | 4 |

Navigation controls at the bottom include 'Back' and 'Refresh' buttons, and pagination for 'Items 1-2 of 2'.

The second view is a running count of the metrics. This view is available when you select Since Last Reset in the Show Metrics For field. The statistics displayed in each row are for the period

since you last reset the statistics for an individual service or since you last reset the statistics for all services on the Global Settings page in the System Administration module.

Figure 3-7 Service Monitoring Summary Page—Since Last Reset

Service Monitoring Summary

Metrics For :

Search Name: Path:

Has Alerts

Has Errors

Invoked by proxy:

Items 1-15 of 15 1

| | Path ▲ | Avg. Exec Time(ms) ▲ | Message Count ▲ | Error Count ▲ | Alert Count ▲ |
|--------------|---|---|--|--|--|
| ky | ALSR Project | 0 | 0 | 0 | 0 |
| nftest | default | 0 | 0 | 0 | 0 |
| atingService | MortgageBroker/BusinessServices | 0 | 0 | 0 | 0 |
| bound | default | 0 | 0 | 0 | 38 |
| teway1 | MortgageBroker/ProxyService | 0 | 0 | 0 | 0 |

Items 1-5 of 5 1

As shown in the top section of the preceding figures, you can filter the display of information using the following criteria:

- Name—the name of the proxy service or business service.
- Path—the project folder in which the proxy service or business service resides.
- Has Alerts—by services that have alert messages.
- Has Errors—by services that have failed messages.
- Invoked by proxy—the name and path of the proxy service.

The Service Monitoring Summary table displays the following information:

- **Name**—the name of the proxy or business service. The name is a link to the **Service Monitoring Details** page. See “[Service Monitoring Details](#)” on page 3-17.
- **Path**—the project folder in which the service resides. The path is a link to the **Project View** or Folder View page, depending on whether the service resides in the top level of a project or in a folder.
- **Aggregation Interval**—the time period over which data points for specific statistics is collected and displayed for the service. This information is displayed only when you have selected **Current Aggregation Interval** in the **Show Metrics For** field.
- **Average Execution Time**—the average time it has taken the service to process a message for the period of the current aggregation interval or for the period since the last reset. This is measured in milli-seconds.
- **Message Count**—the total number of messages processed by the service for the period of the current aggregation interval or for the period since the last reset.
- **Error Count**—the number of messages that have failed for the period of the current aggregation interval or for the period since the last reset.
- **Alert Counts**—the number of alerts raised by alert rule occurrences and violations for the period of the current aggregation interval or for the period since the last reset.

Note: An Action column is displayed when you have selected Since Last Reset in the Show Metrics For field. In this column, you can click the Reset Statistics icon for a specific service to reset the statistics for that service. When you confirm that you want to do this, the system deletes all monitoring statistics that were collected for the service since the last time you clicked the Reset Statistics icon or the last time you clicked Reset Statistics on the Global Settings page. However, the system does not delete the statistics being collected during the Current Aggregation Interval for the service. Once you click the **Reset Statistics** icon, the system immediately starts collecting monitoring statistics for the service again.

Service Monitoring Details

The Service Monitoring Details page provides you with two views of detailed information about a specific service, as shown in the following figures.

The first is a dynamic view of the statistical data collected by the service. This view is available when you select Current Aggregation Interval in the Show Metrics For field. The aggregation interval displayed in this view determines the statistics that are displayed. For example, if the

aggregation interval of this service is 20 minutes, the view displays the data collected in the last 20 minutes.

Figure 3-8 Service Monitoring Details Page—Current Aggregation Interval

Service Monitoring Details - MortgageBroker/BusinessService/normalLoanBS

Show Metrics For : Current Aggregation Interval

Alert Status : Normal at 11:31:38 AM June 30, 2006

Aggregation Interval : 0 Hour(s) and 1 Minutes

Alerts For Current Aggregation Interval : 0 [Alert History](#)

Location Path : MortgageBroker/BusinessService

Display Metrics For : xbusServer

| Operations | | | | | | Performance | |
|----------------|---------------|-------------|--------------------|--------------------|--------------------|----------------------------------|-------|
| Items 1-1 of 1 | | | | | | | |
| Operations | Message Count | Error Count | Min Resp. Time(ms) | Max Resp. Time(ms) | Avg. Exec Time(ms) | | |
| processLoanApp | 0 | 0 | 0 | 0 | 0 | | |
| Items 1-1 of 1 | | | | | | | |
| | | | | | | Min Response Time(ms): | 0 |
| | | | | | | Max Response Time(ms): | 0 |
| | | | | | | Overall Avg. Execution Time(ms): | 0 |
| | | | | | | Total Number of Messages: | 0 |
| | | | | | | Messages With Errors: | 0 |
| | | | | | | Failover Count: | 0 |
| | | | | | | Success Ratio(%): | 100.0 |
| | | | | | | Failure Ratio(%): | 0.0 |
| | | | | | | Number of WS Security Errors: | 0 |
| | | | | | | Number of Validation Errors: | N/A |

Back Refresh

The second view is a running count of the metrics. This view is available when you select **Since Last Reset** in the **Show Metrics For** field. The statistics displayed are for the period since you last reset statistics for this particular service or since you last reset statistics for all services on the Global Settings page in the System Administration module.

Figure 3-9 Service Monitoring Details Page—Since Last Reset

Service Monitoring Details - default/New

Show Metrics For :

Alerts Since Last Reset : 13 [Alert History](#)

Location Path : default

Display Metrics For :

| Operations | | | | | | Performance | |
|----------------|---------------|-------------|--------------------|--------------------|--------------------|----------------------------------|-------|
| Items 1-1 of 1 | | | | | | | |
| Operations | Message Count | Error Count | Min Resp. Time(ms) | Max Resp. Time(ms) | Avg. Exec Time(ms) | Min Response Time(ms): | 0 |
| processLoanApp | 0 | 0 | 0 | 0 | 0 | Max Response Time(ms): | 0 |
| Items 1-1 of 1 | | | | | | Overall Avg. Execution Time(ms): | 0 |
| | | | | | | Total Number of Messages: | 0 |
| | | | | | | Messages With Errors: | 0 |
| | | | | | | Success Ratio(%): | 100.0 |
| | | | | | | Failure Ratio(%): | 0.0 |
| | | | | | | Number of WS Security Errors: | 0 |
| | | | | | | Number of Validation Errors: | 0 |

Flow Components

| Items 1-2 of 2 | | | |
|----------------------------|---------------|-------------|--------------------|
| Component Name | Message Count | Error Count | Avg. Exec Time(ms) |
| PipelinePairNode1_request | 0 | 0 | 0 |
| PipelinePairNode1_response | 0 | 0 | 0 |
| Items 1-2 of 2 | | | |

Back Refresh

The Service Monitoring Details Page displays the following set of information:

- Service Monitoring Details
 - Alert Status—the current alert status, which is displayed only when you have selected Current Aggregation Interval in the Show Metrics For field.
 - Aggregation Interval—the time period over which data points for specific statistics are collected and then displayed for the service. This information is displayed only when you have selected Current Aggregation Interval in the Show Metrics For field.
 - Alerts for Current Aggregation Interval—the total number of alerts associated with this service during the current aggregation interval. This information is displayed only when you have selected **Current Aggregation Interval** in the **Show Metrics For** field.
 - Alerts Since Last Reset—the total number of alerts associated with this service since you last reset statistics for the service or since you last reset statistics for all services on the Global Settings page. This information is displayed only when you have selected Since Last Reset in the Show Metrics For field.

Monitoring

- Alert History—a link to the Customized System Alerts History page. See [“System Alerts History” on page 3-32](#).
- Location Path—the project and folder where the service resides.
- Display Metrics For—displays the metrics for a server. For a single node, only one item is displayed.
- Operations
 - Operations—the operations associated with the service, if any exist.
 - Message Count—the number of messages associated with each operation during the current aggregation interval or within the period since the last reset.
 - Minimum Response Time—the minimum time this operation has taken to execute messages during the current aggregation interval or within the period since the last reset.
 - Maximum Response Time—the maximum time this operation has taken to execute messages during the current aggregation interval or within the period since the last reset.
 - Average Execution Time—the average time the operation has taken to execute messages during the current aggregation interval or within the period since the last reset.
- Performance
 - Minimum Response Time—the minimum time this service has taken to execute messages during the current aggregation interval or within the period since the last reset.
 - Maximum Response Time—the maximum time this service has taken to execute messages during the current aggregation interval or within the period since the last reset.
 - Overall Average Execution Time—the overall average time that the service has taken to execute messages during the current aggregation interval or within the period since the last reset.
 - Total Number of Messages—the total number of messages, including failed messages, during the current aggregation interval or within the period since the last reset.
 - Messages With Errors—the number of messages that failed during the current aggregation interval or within the period since the last reset. In two-way messaging, if

the response message fails, but the request message was processed, only the failed response message is counted.

- Failover Count—for business services only, the number of failover messages during the current aggregation interval or within the period since the last reset.
- Success Ratio—the percentage of successfully processed messages during the current aggregation interval or within the period since the last reset.
- Failure Ratio—the percentage of messages that failed to process during the current aggregation interval or within the period since the last reset.
- Number of WS Security Errors—the number of messages that failed due to security reasons, such as authentication errors, security policy violations, or authorization errors, during the current aggregation interval or within the period since the last reset.
- Number of Validation Errors—the number of messages that failed when a validate action compared one or more parts of a message against an XSD schema or WSDL resource, during the current aggregation interval or within the period since the last reset. This is displayed for proxy services only.
- Flow Components for proxy services
 - Component Name—the name of pipeline or node in the message flow.
 - Message Count—the number of messages associated with each component during the current aggregation interval or within the period since the last reset.
 - Error Count—the number of failed messages associated with each component during the current aggregation interval or within the period since the last reset.
 - Average Execution Time—the average time the component has taken to execute a message during the current aggregation interval or within the period since the last reset.

Server Summary

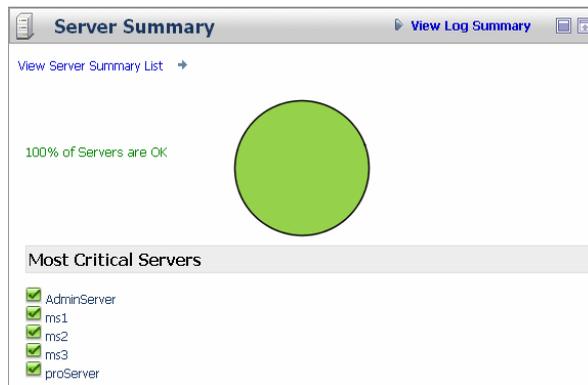
This section provides information on the following topics:

- [About the Server Summary](#)
- [Log Summary](#)
- [Server Summary](#)
- [Server Details](#)

About the Server Summary

The Server Summary panel provides an overview of the state of the servers. The pie chart shows the status of each server in the domain. The status for each server is derived from the WebLogic Diagnostic Service (see *Configuring and Using the WebLogic Diagnostics Framework*). The five most critical servers are displayed, as shown in Figure 3-10.

Figure 3-10 Server Summary Pane



The displayed statuses have the following meanings:

- Fatal—the server has failed and must be restarted.
- Critical—server failure likely; something must be done immediately to prevent failure. For more details, check the server logs and the corresponding `RuntimeMBean`.
- Warning—the server could have problems in the future. For more details, check the server logs and the corresponding `RuntimeMBean`.
- OK—the server is functioning without any problems.
- Overloaded—the server has more work assigned to it than the configured threshold; it might refuse more load.

Log Summary

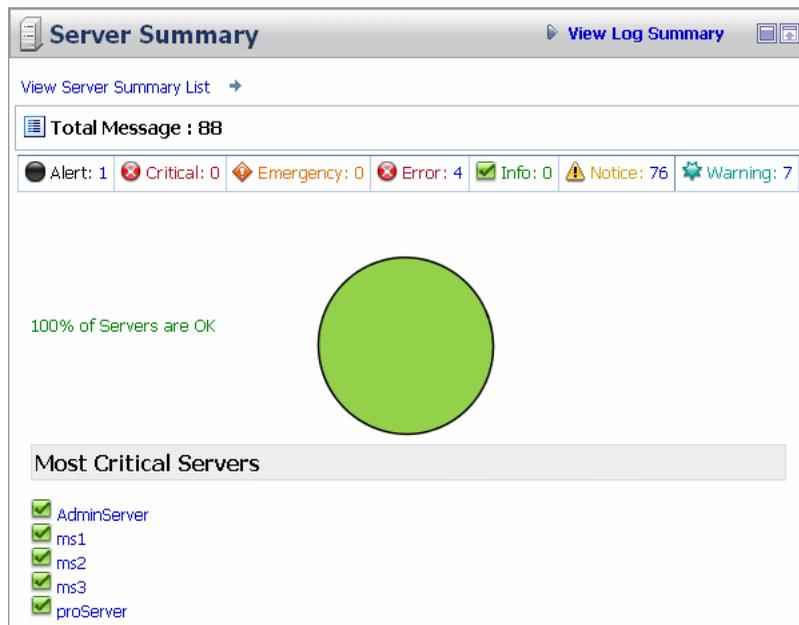
The AquaLogic Service Bus Console allows you to view the WebLogic Server domain log. The domain log file provides a central location from which to view the overall status of the domain. Each server instance forwards a subset of its messages to a domain-wide log file. By default, servers forward only messages of severity level NOTICE or higher. You can modify the set of

messages that are forwarded. For more information, see [Understanding WebLogic Logging Services](#) in *Configuring Log Files and Filtering Log Messages*.

If you configure the logging action in a pipeline, the log is forwarded to the server log. Unless you configure WebLogic Server to forward these messages to the domain log, you cannot view this log from AquaLogic Service Bus Console. For information in how to do this, see [Create Log Filters](#) in the *WebLogic Server Administration Console Online Help*.

To see the number of messages currently raised by the system, click the View Log Summary link in the Server Summary panel. A table is displayed that contains the number of messages grouped by severity, as shown in the following figure.

Figure 3-11 Log Summary



The displayed message statuses have the following meanings:

- **Alert**—a particular service is in an unusable state while other parts of the system continue to function. Automatic recovery is not possible; immediate attention of the administrator is required to resolve the problem.
- **Critical**—a system or service error has occurred. The system can recover but there might be a momentary loss or permanent degradation of service.

- **Emergency**—the server is in an unusable state. This severity indicates a severe system failure.
- **Error**—a user error has occurred. The system or application can handle the error with no interruption. Limited degradation of service may occur.
- **Info**—reports normal operations; a low-level informational message.
- **Notice**—an informational message with a higher level of importance than Info messages.
- **Warning**—a suspicious operation or configuration has occurred. However, normal operations may not be affected.

This display is based on the health state of the running servers, as defined by the WebLogic Diagnostic Service. For more information about the WebLogic Diagnostic Service, see [Configuring and Using the WebLogic Diagnostics Framework](#).

To view the domain log for a particular type of message, click the number corresponding with the type of message. The following figure shows an example of a domain log file displayed in the AquaLogic Service Bus Console.

Figure 3-12 Domain Log File Entries

Dashboard > Project Explorer > Proxy Services > Dashboard

This page shows you the latest contents of the domain log file.

Domain Log File Entries

Showing 1 - 8 of 8 Previous | Next

| Date | Subsystem | Severity | Message ID | Message |
|------------------------------|------------------|----------|------------|--|
| Jun 13, 2005 10:05:07 AM MDT | BEA-AlertManager | Error | BEA-394000 | Exception on executeAction, com.bea.wli.sb.transports.TransportException: com.bea.wli.sb.transports.TransportException: javax.mail.SendFailedException: Sending failed; nested exception is: javax.mail.MessagingException: Unknown SMTP host: host; nested exception is: java.net.UnknownHostException: host host com.bea.wli.sb.transports.TransportException: com.bea.wli.sb.transports.TransportException: javax.mail.SendFailedException: Sending failed; nested exception is: javax.mail.MessagingException: Unknown SMTP host: host; nested exception is: java.net.UnknownHostException: host com.bea.wli.sb.transports.email.EmailTransportProvider.sendMessageAsync(EmailTransportProvider.java:87) at jrockit.reflect.VirtualNativeMethodInvoker.invoke(Ljava.lang.Object;[Ljava.lang.Object;)Ljava.lang.Object;(Unknown Source) at jrockit.reflect.InitialMethodInvoker.invoke(Ljava.lang.Object;[Ljava.lang.Object;)Ljava.lang.Object;(Unknown Source) at java.lang.reflect.Method.invoke(Ljava.lang.Object;[Ljava.lang.Object;)Ljava.lang.Object;(Unknown Source) at com.bea.wli.sb.transports.TransportManagerImpl\$1.invoke(TransportManagerImpl.java:1195) at \$Proxy9.sendMessageAsync(Lcom.bea.wli.sb.transports.TransportMessageContext;Lcom.bea.wli.sb.transports.TransportSendListener;)V(Unknown Source) at com.bea.wli.sb.transports.TransportManagerImpl.sendMessageWithoutService(TransportManagerImpl.java:542) at com.bea.wli.sb.transports.TransportManagerImpl.sendMessageAsync(TransportManagerImpl.java:468) at com.bea.wli.monitoring.alert.action.emailAlert.EmailActionProvider.executeAction(EmailActionProvider.java:186) at com.bea.wli.monitoring.alert.AlertManager._invokeActions(AlertManager.java:593) at com.bea.wli.monitoring.alert.AlertManager.evaluateSingleRule(AlertManager.java:538) at |

The following information is displayed:

- **Date**—the date and time the entry was logged in a format that is specific to the local time zone and format.

- Subsystem—the WebLogic Server subsystem that was the source of the message, such as the EJB container or Java Messaging Service (JMS).
- Severity—indicates the degree of impact or seriousness of the event.
- Message ID—the unique six-digit identification for the message.
- Message—a description of the event or condition.

For more information, see “Message Attributes” in [Understanding WebLogic Logging Services](#) in *Configuring Log Files and Filtering Log Messages*.

To display details of a single log file on the page, select the radio button for the appropriate log, then click the View button.

Server Summary

The Server Summary page provides a customizable table of servers, as shown in the following figure.

Figure 3-13 Server Summary Page

The screenshot shows the 'Server Summary' page. At the top, there is a title bar with 'Server Summary' and a 'View as Graph' link. Below this is a summary bar with message counts: Fatal: 0, Critical: 0, Warning: 0, Ok: 1, Overloaded: 0, and Unknown: 0. A 'Customize table' link is also present. The main part of the page is a table with the following data:

| Status | Server | Cluster Name | Machine Name | State | Uptime |
|--------|-------------|--------------|--------------|---------|--------------------------------|
| Ok | AdminServer | | | RUNNING | 3 hours 1 min 29 secs 906 msec |

At the bottom of the table, there are 'Refresh' and 'Back' buttons.

As shown in the upper section of the [Figure 3-13](#), the Server Summary Page displays the number of messages currently raised by the system. For information about the meaning of each type of status message, see “[Log Summary](#)” on [page 3-22](#).

The server table displays the following information:

Monitoring

- Status—the status of the server:
 - Fatal—the server has failed and must be restarted.
 - Critical—server failure likely; something must be done immediately to prevent failure. For more details, check the server logs and the corresponding `RuntimeMBean`.
 - Warning—the server could have problems in the future. For more details, check the server logs and the corresponding `RuntimeMBean`.
 - OK—the server is functioning without any problems.
 - Overloaded—the server has more work assigned to it than its configured threshold; it cannot take on more load.
- Server—the name of the server. The name is a link to the View Server Details page. See [“Server Details” on page 3-27](#).
- Cluster Name—if the server is part of a cluster, the name of the cluster.
- Machine Name—the name of the computer associated with the server.
- State—the state of the server:
 - RUNNING
 - FAILED
 - SHUTDOWN
- Uptime—the duration for which this server has been running.

To view this information in the table as a pie or bar chart, click [View as a Graph](#).

To filter the display of servers, click [Customize Table](#) above the server table. The available filtering is shown in the following figure.

Figure 3-14 Server Summary Table Filter

 **Server Summary Table Filter**

| <input type="checkbox"/> Severity | All | | | | | | | | | | | | | | |
|---|--|-----------|--------|--|--------|--|--------|--|--------------|--|--------------|--|-------|--|--------|
| <input type="checkbox"/> Server | All | | | | | | | | | | | | | | |
| <input type="checkbox"/> Cluster Name | All | | | | | | | | | | | | | | |
| <input type="checkbox"/> Machine Name | All | | | | | | | | | | | | | | |
| <input type="checkbox"/> State | All | | | | | | | | | | | | | | |
| Columns Display | <table border="1"> <thead> <tr> <th>Available</th> <th>Chosen</th> </tr> </thead> <tbody> <tr> <td></td> <td>Status</td> </tr> <tr> <td></td> <td>Server</td> </tr> <tr> <td></td> <td>Cluster Name</td> </tr> <tr> <td></td> <td>Machine Name</td> </tr> <tr> <td></td> <td>State</td> </tr> <tr> <td></td> <td>Uptime</td> </tr> </tbody> </table> | Available | Chosen | | Status | | Server | | Cluster Name | | Machine Name | | State | | Uptime |
| Available | Chosen | | | | | | | | | | | | | | |
| | Status | | | | | | | | | | | | | | |
| | Server | | | | | | | | | | | | | | |
| | Cluster Name | | | | | | | | | | | | | | |
| | Machine Name | | | | | | | | | | | | | | |
| | State | | | | | | | | | | | | | | |
| | Uptime | | | | | | | | | | | | | | |
| Number of rows displayed per page | 20 | | | | | | | | | | | | | | |
| Maximum Results Returned | 10 | | | | | | | | | | | | | | |
| <input type="button" value="Apply"/> <input type="button" value="Reset"/> | | | | | | | | | | | | | | | |

For information about how to use the Server Summary Table Filter, see “Customize Your View of the Server Summary” in [Monitoring](#) in the *Using the AquaLogic Service Bus Console*.

Server Details

You can access the View Server Details page by clicking the name of a server under Most Critical Servers or by clicking the name of a server in the Servers Summary page.

The View Server Details page enables you to view more server monitoring details, as shown in the following figure.

Figure 3-15 Server Details Page—General Tab

Dashboard

General | Channels | Performance | Threads | Timers | Workload | Security | JMS | JTA

This page provides general runtime information about this server.

| | | |
|----------------------------|---|---|
| State: | RUNNING | core.server.servermonitoringgeneral.state.label.inlinehelp More Info... |
| ActivationTime: | Wed Jun 15 10:55:47 MDT 2005 | core.server.servermonitoringgeneral.activationtime.label.inlinehelp More Info... |
| ▼ Advanced | | |
| Weblogic Version: | WebLogic Server 9.0 Mon Jun 13 20:43:42 PDT 2005 585258 - internal build by sjbuild on client qs.build.auntie | core.server.servermonitoringgeneral.weblogicversion.label.inlinehelp More Info... |
| Java Vendor: | BEA Systems, Inc. | core.server.servermonitoringgeneral.javavendor.label.inlinehelp More Info... |
| Java Version: | 1.5.0_03 | core.server.servermonitoringgeneral.javaversion.label.inlinehelp More Info... |
| OSName: | Windows XP | core.server.servermonitoringgeneral.osname.label.inlinehelp More Info... |
| OSVersion: | 5.1 | core.server.servermonitoringgeneral.osversion.label.inlinehelp More Info... |

The information displayed on this page is a subset of the Monitoring tab in the AquaLogic Service Bus Console Server Settings page. The details available are:

- **General**—provides general run-time information about the server. Click **Advanced** to view more information, such as WebLogic Server version or operating system name.
- **Channels**—displays monitoring information about each channel.
- **Performance**—displays performance information about the server.
- **Threads**—displays current run-time characteristics and statistics for the server’s active executable queues.
- **Timers**—displays information about the timer used by the server.
- **Workload**—displays statistics for work managers, constraints, and policies configured on the server.
- **Security**—allows you to monitor user-lockout management statistics for the server.

- JMS—allows you to monitor JMS information about the server.
- JTA—displays the summary of all transaction information for all resource types on the server.

For more information, see [WebLogic Server Administration Console Online Help](#).

Alert Summary

This section provides information on the following topics:

- [“About the Alert Summary” on page 3-29](#)
- [“System Alerts History” on page 3-32](#)
- [“System Alert Details” on page 3-35](#)
- [“View Alert Rule Details” on page 3-36](#)

About the Alert Summary

In AquaLogic Service Bus there are two types of alerts that can occur. They are:

- [“Pipeline Alerts” on page 3-29](#)
- [“Service Level Agreement Alerts \(SLA\)” on page 3-30](#)

Pipeline Alerts

The alerts triggered when alert actions, configured within a pipeline are executed, are called as the pipeline alerts. You can use actions grouped under the reporting category. The actions available under the **Report** category are:

- Alert
- Log
- Report

For more information, see [Proxy Service: Actions](#) in *Using the AquaLogic Service Bus Console* The alerts are monitored using the alert destinations.

Service Level Agreement Alerts (SLA)

The Service Level Agreement (SLA) alerts are generated when the service violates the service level agreement or a predefined condition. The Alert Summary panel contains a customizable table displaying information about violations or occurrences of events in the system. These violations and occurrences are based on SLAs. AquaLogic Service Bus provides various SLA monitors that you can configure to monitor proxy and business services. Some examples of SLA monitors are maximum execution time and authorization failure. You configure these monitors by creating alert rules. When a rule evaluates to true, it raises an alert. This alert can be sent to console, SNMP trap, reporting stream, e-mail recipients or JMS queue/topic. These destinations for the alert are configured using the alert destination resource.

Note:

- When you configure an alert rule to post a message to a JMS destination, you must create a JMS connection factory and a queue or topic, and target them to the appropriate JMS server in the WebLogic Server Administration Console. For information on how to do this, see “Configuring a JMS Connection Factory” and “JMS Resource Naming Rules for Domain Interoperability” in [Configuring JMS System Resources](#) in *Configuring and Managing WebLogic JMS*.
- For more information on how to define conditions see [To Define Alert Rule Conditions](#).

The AquaLogic Service Bus Console provides several ways to view and find alerts, such as by severity and by service. You can also view alerts graphically. For information on how to do this, see “Listing and Locating Alerts” and “Viewing a Chart of Alerts” in [Monitoring](#) in *Using the AquaLogic Service Bus Console*.

The following figure shows the View Alert Summary List:

Figure 3-16 View Alert Summary List

View Alert Summary List →

| Alert Severity ▾ | Timestamp ▲ | Service ▲ | Alert Rule Name ▲ |
|------------------|----------------|---|-------------------|
| Fatal | 6/2/06 4:05 PM | business:MortgageBroker/BusinessService/normalLoanBS | newbsRule |
| Fatal | 6/2/06 4:10 PM | business:MortgageBroker/BusinessService/normalLoanBS | newbsRule |
| Fatal | 6/2/06 4:20 PM | business:MortgageBroker/BusinessService/normalLoanBS | newbsRule |
| Fatal | 6/2/06 4:25 PM | business:MortgageBroker/BusinessService/normalLoanBS | newbsRule |
| Fatal | 6/2/06 4:30 PM | business:MortgageBroker/BusinessService/normalLoanBS | newbsRule |
| Major | 6/2/06 4:05 PM | proxy:default/mflProxy | Newrule3 |
| Major | 6/2/06 4:05 PM | business:MortgageBroker/BusinessService/managerLoanApproval | newBSRule |
| Major | 6/2/06 4:10 PM | proxy:default/mflProxy | Newrule3 |
| Major | 6/2/06 4:10 PM | business:MortgageBroker/BusinessService/managerLoanApproval | newBSRule |
| Major | 6/2/06 4:20 PM | proxy:default/mflProxy | Newrule3 |
| Major | 6/2/06 4:20 PM | business:MortgageBroker/BusinessService/managerLoanApproval | newBSRule |
| Major | 6/2/06 4:25 PM | proxy:default/mflProxy | Newrule3 |
| Major | 6/2/06 4:25 PM | business:MortgageBroker/BusinessService/managerLoanApproval | newBSRule |
| Major | 6/2/06 4:30 PM | proxy:default/mflProxy | Newrule3 |
| Major | 6/2/06 4:30 PM | business:MortgageBroker/BusinessService/managerLoanApproval | newBSRule |
| Minor | 6/2/06 4:05 PM | business:MortgageBroker/BusinessService/CreditLoan | newAlertCS |
| Minor | 6/2/06 4:05 PM | proxy:default/New | Newalert4 |
| Minor | 6/2/06 4:10 PM | business:MortgageBroker/BusinessService/CreditLoan | newAlertCS |
| Minor | 6/2/06 4:10 PM | proxy:default/New | Newalert4 |
| Minor | 6/2/06 4:20 PM | business:MortgageBroker/BusinessService/CreditLoan | newAlertCS |

Items 1-20 of 35

The Alert Summary panel shows alerts for the last 30 minutes. It contains the following details:

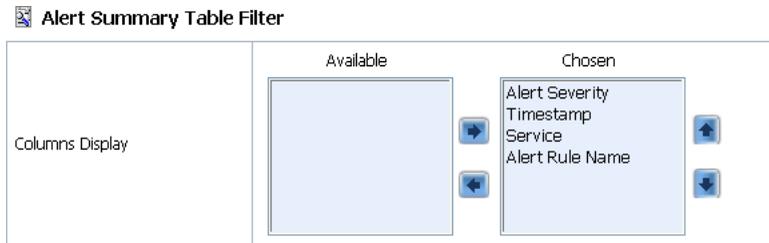
- **Alert Severity**—the user-defined severity of the alert. The Severity is a link to the Alert Details page. See [“System Alert Details” on page 3-35](#).
- **Timestamp**—the date and time when the alert occurred.
- **Alert Rule Name**—the name assigned to the alert. The name is a link to the View Alert Rule Details page. See [“View Alert Rule Details” on page 3-36](#).

Note: The Alert Rule Name content acts as a link only for SLA alerts (which possess rule configuration information) and not so for pipeline alerts.
- **Service/Project Name**—the name of the service and project associated with the alert. The name is a link to the Service Monitoring Details page. See [“Service Monitoring Details” on page 3-17](#).

To view a complete list of alerts, click View Alert Summary List. See [“System Alerts History” on page 3-32](#).

To customize the information displayed in the Alert Summary Panel, click Customize table above the summary table. The available filtering is shown in the following figure.

Figure 3-17 Alert Summary Table Filter



To customize the sort order of the displayed alerts, click the sort icons beside the column headers.

System Alerts History

To access the Customized System Alerts History page, in the Alert Summary panel, click View Alert Summary List. The Customized System Alerts History page enables you to view all the alerts by paging through the table (see [Figure 3-18](#)) or by filtering the display of the alerts (see [Figure 3-19](#)).

Figure 3-18 Customized System Alerts History

| Alert History | | | | | | | | | | | | | |
|--------------------------|------------------|----------------|-------------------|-------------|--|----------|--|------------|--|----------------------|--|-----------------|--|
| Fatal: 76 | | Critical: 38 | | Major: 38 | | Minor: 0 | | Warning: 0 | | Normal: 0 | | | |
| | | | | | | | | | | Items 141-152 of 152 | | 1 2 3 4 5 6 7 8 | |
| <input type="checkbox"/> | Alert Severity ▾ | Timestamp ▲ | Alert Rule Name ▲ | Service ▲ | | | | | | | | | |
| <input type="checkbox"/> | Major | 8/4/06 3:55 PM | newRule12 | default/New | | | | | | | | | |
| <input type="checkbox"/> | Major | 8/4/06 4:00 PM | newRule12 | default/New | | | | | | | | | |
| <input type="checkbox"/> | Major | 8/4/06 4:05 PM | newRule12 | default/New | | | | | | | | | |
| <input type="checkbox"/> | Major | 8/4/06 4:11 PM | newRule12 | default/New | | | | | | | | | |
| <input type="checkbox"/> | Major | 8/4/06 4:15 PM | newRule12 | default/New | | | | | | | | | |
| <input type="checkbox"/> | Major | 8/4/06 4:20 PM | newRule12 | default/New | | | | | | | | | |
| <input type="checkbox"/> | Major | 8/4/06 4:37 PM | newRule12 | default/New | | | | | | | | | |
| <input type="checkbox"/> | Major | 8/4/06 4:45 PM | newRule12 | default/New | | | | | | | | | |
| <input type="checkbox"/> | Major | 8/4/06 4:50 PM | newRule12 | default/New | | | | | | | | | |
| <input type="checkbox"/> | Major | 8/4/06 4:55 PM | newRule12 | default/New | | | | | | | | | |
| <input type="checkbox"/> | Major | 8/4/06 5:00 PM | newRule12 | default/New | | | | | | | | | |
| <input type="checkbox"/> | Major | 8/4/06 5:05 PM | newRule12 | default/New | | | | | | | | | |
| | | | | | | | | | | Items 141-152 of 152 | | 1 2 3 4 5 6 7 8 | |
| Delete | | Refresh | | Back | | | | | | | | | |

You can customize the table shown in the [Figure 3-18](#) and provides the following details:

- **Alert Severity**—the severity level of alerts is user configurable and has no absolute meaning. The field is a link to the System Alert Details page. See [“System Alert Details” on page 3-35](#).
- **Timestamp**—the date and time when the alert occurred.
- **Alert Rule Name**—the name assigned to the alert. The name is a link to the View Alert Rule Details page. See [“View Alert Rule Details” on page 3-36](#).

Note: The Alert Rule Name content acts as a link only for SLA alerts, which are configured with an alert rule and not for pipeline alerts.

- **Service**—the name of the service and project associated with the alert. The name is a link to the Service Monitoring Details page. See [“Service Monitoring Details” on page 3-17](#).

To view a pie or bar chart of the alerts, click View Graph in the table.

To search for a specific alert, you can filter the display of alerts by clicking Customize Table in the Customized System Alerts History table. The filtering is shown options are available in the following figure.

Figure 3-19 System Alerts Table Filter

System Alerts Table Filter

Time: 0 days 0 hours 30 mins
 June 11 2006 4 05 PM
 June 11 2006 4 35 PM

Severity: All

Service: All

Alert Rule Name: All

Columns Display:

| Available | Chosen |
|-----------|-----------------|
| | Alert Severity |
| | Timestamp |
| | Alert Rule Name |
| | Service |

Number of rows displayed per page: 20

Maximum Results Returned: Show All

Apply Reset

For information about how to use the Alerts Table Filter, see “Customizing Your View of Alerts” in [Monitoring](#) in the *Using the AquaLogic Service Bus Console*.

Note: When an alert is raised in your configuration, a message is sent to your domain log, which resides at the following location:

```
<BEA_HOME>\servers\server_name\logs\domain_name.log
```

where

- *BEA_HOME* represents the name of the BEA home directory.
- *domain_name* represents the name you assigned your AquaLogic Service Bus domain when you created it.

The message is logged as an alert and has this message ID: BEA-394015

The message body is a string that consists of the following details:

- Alert ID
- Alert Rule ID
- Alert Rule Name

- Severity
- Timestamp
- Service Name
- Service Path
- Body

System Alert Details

The System Alert Details page displays complete information about the alert and allows you to add an annotation to the alert, as shown in the following figure.

Figure 3-20 System Alert Details Page

"newbsRule" Details

Save Cancel Delete

| | |
|-----------------|---|
| Alert Rule Name | newbsRule |
| Description | |
| Timestamp | Fri Jun 30 12:40:52 IST 2006 |
| Severity | Fatal |
| Service | MortgageBroker/BusinessService/normalLoanBS |
| Annotation | <input type="text"/> |

Save Cancel Delete

The following details are displayed:

- Alert Rule Name— alert rule name for SLA alerts and alert summary for pipeline alerts. Acts as a link to the View Rule Details page for SLA alerts.
- Description—rule description of the alerts for the SLA and alert payload for pipelines.
- Timestamp—the date and time of the alert.
- Severity—the user-defined severity of the alert.

- Service—the name of the service associated with the alert. The name is a link to the Service Monitoring Details page. See “[Service Monitoring Details](#)” on page 3-17.
- Annotation—use this field to add notes to the alert.

You access this page from the dashboard by clicking Alert Severity in the Alert Summary table. This page also allows you to delete the alert.

View Alert Rule Details

The View Alert Rule Details page displays complete information about a specific alert rule, as shown in the following figure.

Figure 3-21 View Alert Rule Details Page

View Alert Rule Details - binaryAlertRule [default/binarymftest]

| | | |
|------------------|-----------------|--|
| Last Modified By | weblogic | Description - no description - |
| Last Modified On | 6/23/06 3:24 PM | |
| References | 1 | |
| Referenced By | 0 | |

General Configuration

| | |
|-----------------------------------|----------------------|
| Rule Name | binaryAlertRule |
| Alert Summary | |
| Alert Destination | default/NewalertDest |
| Start Time (HH:MM) | 12:00 AM |
| End time (HH:MM) | 12:00 AM |
| Rule Expiration Date (MM/DD/YYYY) | |
| Rule Enabled | true |
| Alert Severity | Critical |
| Alert Frequency | Every Time |
| Stop Processing More Rules | false |

Conditions

| | |
|----------------------|---|
| Condition Expression | Aggregation Interval :0 Hour(s) and 10 Minutes Success Ratio (%) < 100 |
|----------------------|---|

Back

Edit

The following information is displayed:

- Last Modified By—Specifies who modified the alert rule.
- Last Modified On—Specifies about when the alert rule was modified.
- References—Specifies the number of resources it refers to.
- Referenced By—Specifies the number of resources, which refer to the rule.
- General Configuration
 - Rule Name—the name of the alert rule. The value in this field will be used as the subject for an e-mail alert.
 - Alert Summary—The summary to describe the purpose of the alert rule. This is also used as the subject line for the e-mail message if this alert rule is configured with an e-mail destination.
 - Alert Destination—You associate the alert rule with the Alert Destination. By this you set the destinations for the alert notifications for the alert rule. You have to set the alert destination in order to determine the distribution of severity of the alerts.
Note: Although an alert is detected and counted even if alert destination is not set, you cannot determine the severity of the alert and hence it will not be reflected on the dashboard.
 - Start Time (HH:MM)—specifies the starting time during which the rule is active on each day prior to the expiration date.
 - End Time (HH:MM)—specifies the ending time during which the rule is active on each day prior to the expiration date.
 - Rule Expiration Date (MM/DD/YY)—the expiration date of the rule. The rule expires at 12.01am on the specified date. If you do not specify a date, the rule never expires.
 - Rule Enabled—indicates whether the rule is enabled or not.
 - Alert Severity—the user-defined severity of the alert.
 - Alert Frequency—indicates whether notifications should be issued to the configured alert destinations every time the alert rule evaluates to true or issued once when the rule evaluates to true within a given aggregation interval.
 - Stop Processing More Rules—when multiple rules associated with a service exist, this flag indicates whether subsequent rules associated with the service must be processed if the current rule evaluates to true.

- Conditions
 - Condition Expression—displays the condition that triggers the alert rule and aggregation interval details of the alert.

For information about how to define alert rules, see “Create an Alert Rule” in [Monitoring](#) in the *Using the AquaLogic Service Bus Console*.

Alert Rules

This section provides information on the following topics:

- [“About Alert Rules” on page 3-38](#)
- [“Some Uses for Alerts” on page 3-39](#)
- [“Understanding Alert Rules” on page 3-39](#)

About Alert Rules

As mentioned earlier, alerts are automated responses to SLAs violations, which are displayed on the Dashboard. You define alert rules to specify unacceptable service performance according to your business and performance requirements. Each alert rule allows you to specify the aggregation interval for that rule when configuring the alert rule. The alert aggregation interval is not affected by the aggregation interval set for the service.

On the Alert Rule page, if you set the Alert Frequency to `Every Time`, the notifications are issued every time the alert rule evaluates to true. If you set the Alert Frequency to `Once When Condition Is True` the notifications are issued the first time the rule evaluates to true, and no more notifications are generated until the condition resets itself and evaluates to true again.

In the case where the Alert Frequency is set to `Every Time`, the number of times an alert rule is fired depends on the aggregation interval and the sample interval associated with that rule. For example, if the aggregation interval is set to 5 minutes, the sample interval is 1 minute. Rules are evaluated each time 5 samples of data are available. Therefore, the rule is evaluated for the first time approximately 5 minutes after it is created and every minute thereafter.

In the case where the Alert Frequency is set to `Once When Condition is True`, after an alert is fired the first time in an aggregation interval, it is not fired again in the same aggregation interval.

Creating an alert rule involves three parts:

- **General Configuration**—defines the name, description, summary, duration, severity, frequency, state of the enabled alert rule and other general characteristic.
- **Define Condition**—defines one or more conditions that trigger the alert rule. Additionally, you can define the aggregation interval for the condition on this page.

Note: Rules can only be created for services that are enabled for monitoring.

For more information about creating an alert rule is located in “Create an Alert Rule” in [Monitoring](#) in the *Using the AquaLogic Service Bus Console*.

Some Uses for Alerts

The following are some uses for alerts:

- Monitoring and e-mail notification of WS-Security errors.
- Monitoring the number of messages passing through a particular pipeline.
- E-mail notification when the average execution time exceeds 5 seconds during stock exchange hours.

Understanding Alert Rules

The information in this section is presented in question-answer format.

Question 1: I created a service with an alert rule that has the following condition expression:

```
Aggregation Interval:0 Hours(s) and 1 Minutes
Message Count = 0
```

It has been 10 minutes and I have not received any alerts.

Answer: Monitoring statistic collection for each statistical attribute, such as message count and error count, associated with a service begins when a change in the value of that statistic occurs. Data collection for the Message Count attributes begins when the first message is processed by the service and the Message Count attribute is incremented. Similarly, collection of data for the Error Count statistic starts only when the service encounters its first error and the Error Count attribute is incremented. If the service is idle, no monitoring information is collected for that service and subsequently no alert rules are triggered. After the first message is processed, monitoring data for that service is continually collected even if the service does not receive any further requests. Check to see if the service has received any requests.

Question 2: I defined a new alert rule with an aggregation interval that did not exist before and that rule does not seem to raise any alerts. All other rules created prior to this one are working correctly.

Answer: The cause is the same as for Question 1; the service needs to process at least one request after a rule with a new aggregation interval is created to trigger the alert rule. The other rules defined with different aggregation interval values are not affected by the alert rule.

Question 3: I restarted the server and none of my services have processed any requests. Why do I see alerts being generated?

Answer: Once the Monitoring subsystem has started collecting data for services, stopping and restarting a server does not abort the collection process. The data collected is persisted and statistic collection picks up from where it left off.

Question 4: I have an alert rule with the following definition:

```
Aggregation Interval:0 Hours(s) and 5 Minutes  
Success Rate < 80%
```

The Service Monitoring Summary page shows the following values:

```
Message Count: 4  
Error Count: 1
```

Why am I being alerted in this case? Shouldn't the success rate be 80% in this case?

Answer: No, the message count value displayed is the total of all messages processed by the service, including the ones that generated an error. Subsequently, in this case, the success rate is 75%.

Question 5: I created a service with an aggregation interval of 10 minutes that sends a JMS message. I could see the message on the Service Monitoring Summary page, but some time later the message count for my service shows as zero.

Answer: The Service Monitoring Summary page displays dynamic statistics. In this case, it shows the message count in the last 10 minutes. Because no messages were processed by the system in the last 10 minutes, the message count is displayed as zero.

Question 6: I changed the aggregation interval of a service from 10 minutes to 5 minutes. The Service Monitoring Summary page shows all statistics as zero. One of the alerts in this server was configured to a statistical element with a 2 minute aggregation interval, which did not fire the next minute.

Answer: Changing the aggregation interval for a service removes the statistical information for all the services and alerts associated with that service. The alert initializes again and triggers an alert at the end of aggregation interval expiry.

Question 7: I have a business service with multiple endpoints with an alert rule defined as `Failover-count > 0`. When one of the endpoints goes down, the alert is triggered. However, when a service has only one endpoint, the `Failover-count` is not incremented for this service. Instead, an error is generated.

Answer: Set the Retry count to a number greater than zero. For information about setting the Retry count, see “Adding a Business Service” in [Business Services](#) in *Using the AquaLogic Service Bus Console*.

Question 8: I see that an alert is generated on the Dashboard but the value for the Alerts for `Current Aggregation Interval` field on the Service Monitoring Details page displays zero.

Answer: Alert rules are evaluated after the completion of the interval, which occurs after a checkpoint completion. If a rule evaluates to true, the rule’s actions are triggered, a log is generated, and the interval-count statistic attribute (Alerts for Current Aggregation Interval) is incremented. The updated value of this counter is processed in the next checkpoint, 60 seconds later. The Monitoring Details page displays the updated count approximately one minute after the alert is generated.

Question 9: How does the active time for rules that span midnight work?

Answer: Consider the case where the active time for a rule is specified as 22:00 to 09:00.

On a given date, say June 7, the rule will be active and inactive as follows:

```

June 6, 10:00 P.M. to June 7, 9:00 A.M. - Active
June 7, 9:01 A.M. to June 7, 9:59 P.M. - Inactive
June 7, 10:00 P.M. to June 8, 9:00 A.M. - Active

```

The `ServerStatistics` are sent to the dashboard. The `ServerStatistics` represents the monitoring runtime data for that minute. In other words, it contains the statistics information for the services that have been enabled.

The monitoring system aggregates the data received every minute makes it available for the retriever sub system. The aggregator thread is behind by 15 seconds with respect to the `Statistics Collector` checkpoint thread.

If you disable monitoring for the domain, you disable the collection of statistics for that domain. The monitoring data is no longer collected from the next minute, which means there is no data returned if you attempt to retrieve it. The same applies when you enable monitoring for the

domain. The system initially does not show any data. However, after a maximum of two minutes, the Service Summary page displays the results of monitoring.

Statistics Associated With Different Resources

The following section provides more information on different statistics associated with:

- “SERVICE” on page 3-42
- “FLOW_COMPONENT” on page 3-43
- “WEBSERVICE_OPERATION” on page 3-43

SERVICE

A service has an inbound endpoint or an outbound endpoint that is registered with the Service Directory of the AquaLogic Service Bus. Such services are associated with other resources such as WSDL, and security settings. The statistics reported for this resource type is listed in [Table 3-1](#). It also give you the type of the statistics.

Table 3-1 Statistics Reported for SERVICE

| Statistic | Type |
|-------------------|----------|
| message-count | count |
| error-count | count |
| failover-count | count |
| response-time | interval |
| validation-errors | count |
| severity-warning | count |
| severity-major | count |
| severity-minor | count |
| severity-normal | count |
| severity-fatal | count |
| severity-critical | count |

Table 3-1 Statistics Reported for SERVICE

| Statistic | Type |
|--------------|-------|
| severity-all | count |
| failure-rate | count |
| wss-error | count |
| success-rate | count |

FLOW_COMPONENT

Statistics are collected for two FLOW_COMPONENT types, namely, Pipeline-pair node and Route node. For more details on Pipeline-pair node and route node see [Table 2-1](#) of [Chapter 2](#), “[Modeling Message Flow in AquaLogic Service Bus](#)”. The statistics reported for FLOW_COMPONENT are listed in [Table 3-2](#)

Table 3-2 Statistics Reported For FLOW_COMPONENT

| Statistic | Type |
|---------------|----------|
| elapsed-time | interval |
| message-count | count |
| error-count | count |

WEBSERVICE_OPERATION

The statistics pertaining to the WEBSERVICE_OPERATION such as WSDLs are collected and stored in a runtime XML file. The statistics reported for this type of resource are listed in [Table 3-3](#)

Table 3-3 Statistics Reported for WEBSERVICE_OPERATION

| Statistics | Type |
|--------------|----------|
| elapsed-time | interval |

Table 3-3 Statistics Reported for WEBSERVICE_OPERATION

| Statistics | Type |
|---------------|-------|
| message-count | count |
| error-count | count |

Auditing

Auditing helps you to keep track of changes in the configuration of the AquaLogic Service Bus(ALSB). The three types of auditing you can perform are briefly described in:

- [“Configuration Change Auditing” on page 3-44](#)
- [“Runtime Auditing of Messages” on page 3-44](#)
- [“Security Auditing” on page 3-44](#)

Configuration Change Auditing

When you perform configurational changes in AquaLogic Service Bus console a track record of the changes is generated and history of all the configurational changes is maintained. Only the previous image of the object is maintained. You can view or access the history of configurational changes and the list of resources that have been changed during the session only through the console. However, in order to access all the information on configuration you have to activate the session.

Runtime Auditing of Messages

Auditing the entire message flow pipeline during is tedious. However, you can use the reporting action to perform selective auditing of the message flow pipeline during run time. You insert the reporting action at required points in the message flow pipeline and extract the required information. The extracted information may be then stored in a database or sent to the reporting stream in order to write the auditing report.

Security Auditing

When a message is sent to the proxy service and there is a breach in the transport level authentication or the security of the Web Services, WebLogic server generates an audit trail. You have to configure the WebLogic server to generate this audit trail. Using this you can audit all

security violations that occur in the message flow pipeline. It also generates an audit trail whenever it authenticates a user. For more information on security auditing, see [Configuring the WebLogic Security Framework: Main Steps](#) in *AquaLogic Service Bus Security Guide*.

Monitoring

Using the Test Console

The BEA AquaLogic Service Bus Test Console is a browser-based test environment used to validate and test the design of your system. It is an extension of the AquaLogic Service Bus Console. You can configure the object of your test (proxy service, business service, XQuery, XSLT, MFL resource), execute the test, and view the results in the console. In some instances you can trace through the code and examine the state of the message at specific trace points. Design time testing helps isolate design problems before you deploy a configuration to a production environment. The test console can test specific parts of your system in isolation and it can test your system as a unit.

The Test Console can be invoked to test any proxy service or business service and certain resources used by these services. You can also do in-line XQuery testing.

You can invoke the test console in a number of ways in the AquaLogic Service Bus Console, depending on what part of your process you want to test. You can invoke the test console from:

- The Project Explorer
- The Resource Browser
- The XQuery Editor

You can run and test a proxy service that makes a call to another proxy service or business service and vice versa. You can test the resources used by your services. When testing services you must be aware of the information that is passing from the test console to the service and vice versa.

Features

The test console supports the following features:

- Testing proxy services
- Testing business services
- Testing resources
- Testing in-line XQueries
- Tracing the message through the message flow (for proxy services only)

Prerequisites

To use the test console:

- You must have AquaLogic Service Bus running and you must have activated the session that contains the resource you want to test.
- You must disable the pop-up blockers in your browser for the inline XQuery testing to work. Note that if you have toolbars in the Internet Explorer browser, this may mean disabling pop-up blockers from under the options menu as well as for all toolbars that are configured to block them. Inline XQuery testing is done only in the design time environment (in an active session).
- If you want the test console to generate and send SAML tokens to a proxy service, you must configure the proxy service to require SAML tokens *and* to be a relying party. For more information on creating a SAML relying party, see [Create a SAML Relying Party](#) in *WebLogic Server Administration Console Online Help*.

Note: When creating a SAML relying party:

- Only WSS/Sender-Vouches and WSS/Holder-of-Key SAML profiles are applicable to a proxy service.
- When you are configuring the relying party, for the Target URL value provide the *URI* of the proxy service. You can view the URI of the proxy service by clicking on the proxy service name in the AquaLogic Service Bus Console's Project Explorer module. The URI displays in the Endpoint URI row of the Transport Configuration table.

Testing Proxy Services

You must have activated a session to test a proxy service. You can test a proxy service from the Resource Browser or Project Explorer. You can test the following types of proxy services:

- WSDL Web Service
- Messaging Service
- Any Soap Service
- Any XML Service

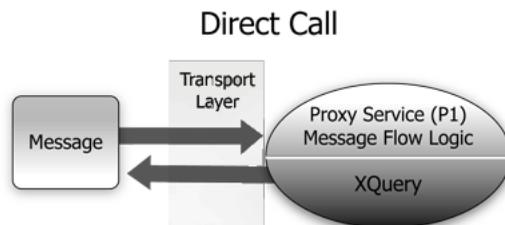
Direct Calls

A Direct Call is used to test a proxy service that is collocated in the AquaLogic Service Bus domain. Using the Direct Call option, messages are sent directly to the proxy service, bypassing the transport layer. When you employ the Direct Call option, tracing is turned on by default, allowing you to diagnose and troubleshoot a message flow in the test console. By default, testing of proxy services is done using the Direct Call option.

When you use the Direct Call option to test a proxy service, the configuration data you input to the test console must be that which is expected by the proxy service from the client that invokes it. In other words, the test console plays the role of the client invoking the proxy service. Also when you do a direct call testing you bypass the monitoring framework for the message and

The following figure illustrates a direct call. Note that the message bypasses the transport layer; it is delivered directly to the proxy service (P1).

Figure 4-1 Direct Call to Test a Proxy Service



A Direct Call strategy is best suited for testing proxy services' internal message flow logic. Your test data should simulate the expected message state at the time it is dispatched. Use this test

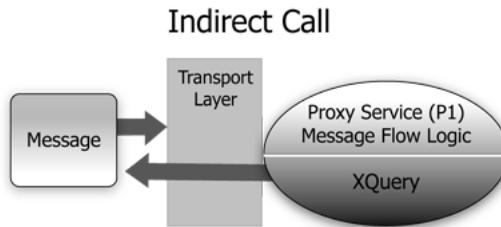
approach in conjunction with setting custom (inbound) transport headers in the test console’s Transport section to accurately simulate the service call.

Indirect Calls

When you test a proxy service with an *indirect* call (that is, when the Direct Call option is not checked), the message is sent to the proxy service through the transport layer. The transport layer performs manipulation of message headers or metadata as part of the test. The effect is to invoke a proxy service to proxy service invocation run-time path.

The following figure illustrates an indirect call. Note that the message is first processed through the transport layer and is subsequently delivered to the proxy service (P1).

Figure 4-2 Indirect Call to Test a Proxy Service



This testing strategy is recommended when testing a proxy service to proxy service interface when both services run in the same JVM. Use this test approach in conjunction with setting custom (outbound) transport headers in the test console’s Transport panel to accurately simulate the service call. For more information on Transport settings in the test console, see [“Test Console Transport Settings” on page 4-21](#).

Using the *indirect call*, the configuration data you input to the test is the data being sent from a proxy service (for example from a Route Node or a Service Callout action of another proxy service). In the *indirect call* scenario, the test console plays the role of the proxy service that routes to, or makes a callout to, another service.

HTTP Requests

When you test proxy services, the test console never sends a HTTP request over the network, therefore transport-level access control is not applied.

(This transport-level access control is achieved through the Web Application layer—in other words, even in the case that an indirect call is made through the AquaLogic Service Bus Console

transport layer, an HTTP request is not sent over the network and this transport-level access control is not applied.) For information about the AquaLogic Service Bus Console architecture, see [Overview](#) in *AquaLogic Service Bus Concepts and Architecture*.

For information about transport settings, see [“Understanding How the Run Time Uses the Transport Settings in the Test Console”](#) on page 4-23.

Testing Business Services

You must have activated a session to test services. You can test the following types of business services:

- WSDL Web Service
- Transport Typed Service
- Messaging Service
- Any Soap Service
- Any XML Service

When testing business services, the messages are always routed through the transport layer. The [“Direct Calls”](#) on page 4-3 option is not available. The configuration data that you provide to the test console to test the service is that which represents the state of the message that is expected to be sent to that business service—for example from a Route Node or a Service Callout action of a proxy service. The test console is in the role of the caller proxy service when you use it to test a business service.

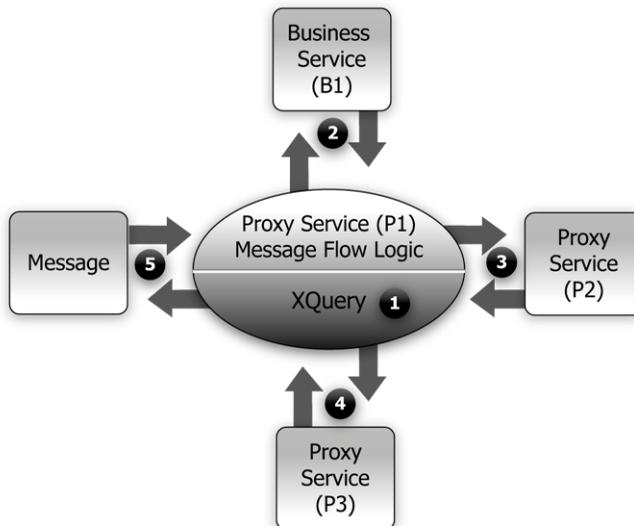
Transport Security

When using the test console to test HTTP(S) business services with BASIC authentication, the test console authenticates with the user name-password from the service account of the business service. Similarly, when testing JMS, e-mail, or FTP business services that require authentication, the test console authenticates with the service account associated with the business service.

Recommended Approaches to Testing Proxy and Business Services

In the scenario depicted in the following figure, a client invokes the proxy service (P1). The message flow invokes business service B1, then proxy service P2, then proxy service P3 before returning a message to the client. Interfaces are identified by number.

Figure 4-3 Test Scenario Example



There are many valid test strategies for this scenario. The following are recommended test strategies:

- It is recommended that you complete the testing of interfaces other than the client interface to a given proxy service before you test the client call. In the sample scenario illustrated in the preceding figure, this means that you complete the testing of interfaces 1 through 4 first, then test interface 5. In this way, the message flow logic for the proxy service (P1) can be iteratively changed and tested (via interface 5) knowing that the other interfaces to the proxy service function correctly.
- It is recommended that all the XQuery expressions in a message flow be validated and tested prior to a system test. In the preceding figure, interface 1 refers to XQuery expression tests.

- Proxy service to business service (interface 2 in the preceding figure) is tested using a *indirect call*. In other words, the messages are routed through the transport layer.
- Proxy service to proxy service tests (Interfaces 3 and 4 in the preceding figure) are tested using an *indirect call*. In other words, disable the Direct Call option, which means that during the testing, the messages are routed through the transport layer.
- Your final *system* test simulates the client invoking the proxy service P1. This test is represented by interface 5 in the preceding figure.

Test interface 5 with a Direct Call. In this way, during the testing, the messages bypass the transport layer. Tracing is automatically enabled with a Direct Call.

- It is recommended that the message state be saved after executing successful interface tests to facilitate future troubleshooting efforts on the system. Testing interface 5 is in fact a test of the complete system and knowing that all other interfaces in the system work correctly helps narrow the troubleshooting effort when system errors arise.

Tracing Proxy Services Using the Test Console

Tracing the message through a proxy service involves examining the message context and outbound communications at various points in the message flow. The points at which the messages are examined are predefined by AquaLogic Service Bus. AquaLogic Service Bus defines tracing for stages, error handlers and route nodes.

For each stage, the trace includes the changes that occur to the message context and all the services invoked during the stage execution. The following information is provided by the trace:

- New variables—  `added`—the names of all new variables and their value (values can be seen by clicking +)
- Deleted variables—  `deleted`—the names of all deleted variables
- Changed variables—  `changed`—the names of all variables for which the value changed. The new value is visible by clicking on the + sign).
- Publish—every publish call is listed. For each publish call, the trace includes the name of the service invoked, and the value of the `outbound`, `header`, `body` and `attachment` variables.
- Service Callout—every Service Callout is listed. For each Service Callout, the trace includes the name of the service that is invoked, the value of the `outbound` variable, the

value of the `header`, `body`, and `attachment` variables for both the request and response messages.

The trace contains similar information for Route Nodes as for stages. In the case of Route Nodes, the trace contains the following categories of information:

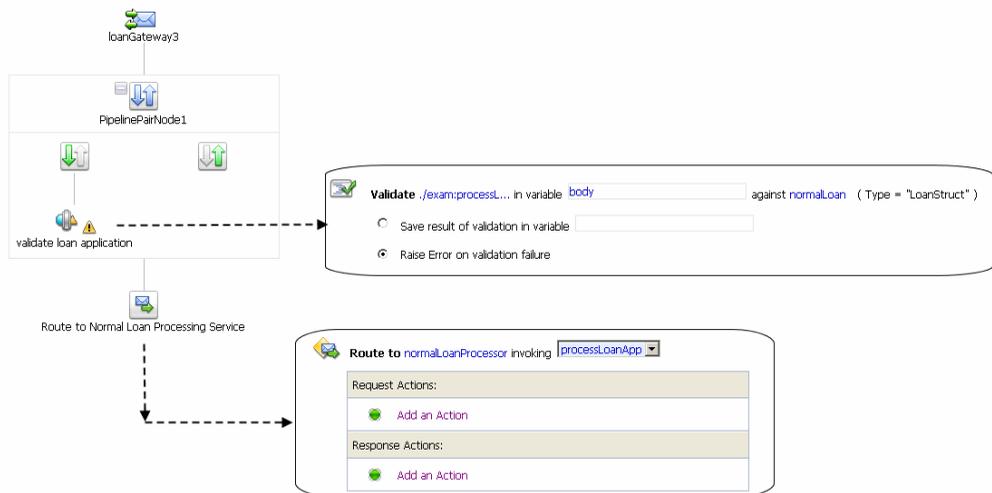
- The trace for service invocations on the request path
- The trace for the Routed Service
- The trace for the service invocations on the response path
- Changes made to the message context between the entry point of the route node (on the request path) and the exit point (on the response path)

Example: Testing and Tracing a Proxy Service

This example uses one of the proxy services in the example AquaLogic Service Bus domain as a basis of instruction.

For more information on how to start the example domain and run the examples provided there, see [BEA AquaLogic Service Bus Samples](#). This example scenario uses the proxy service named `loanGateway3`, associated with the *Validating a Loan Application* example.

The message flow for `loanGateway3` is represented in the following figure. The figure is annotated with the configuration for the *validate loan application* stage and the configuration for the route node.

Figure 4-4 Message Flow for Proxy Service (LoanGateway3)

To test this proxy service in the AquaLogic Service Bus examples domain using the test console, complete the following procedure:

1. Start the AquaLogic Service Bus examples domain and load the samples data, as described in [BEA AquaLogic Service Bus Samples](#).
2. Log in to the AquaLogic Service Bus Console, then select **Project Explorer** and locate the LoanGateway3 proxy service.
3. Select the *Launch Test Console* icon  for the LoanGateway3 proxy service. The *Proxy Service Testing - LoanGateway3* page is displayed. Note that the **Direct Call** and the **Include Tracing** options are selected.
4. Edit the test XML provided to send the following message for the test.

Listing 4-1 Test Message for LoanGateway3

```
<loanRequest xmlns:java="java:normal.client">
  <java:Name>Name_4</java:Name>
  <java:SSN>SSN_11</java:SSN>
  <java:Rate>4.9</java:Rate>
</loanRequest>
```

Using the Test Console

```
<java:Amount>2500</java:Amount>  
<java:NumOfYear>20.5</java:NumOfYear>  
<java:Notes>Name_4</java:Notes>  
</loanRequest>
```

5. Click **Execute**.

The results page is displayed. Scroll to the bottom of the page to see the tracing results in the **Invocation Trace** panel.

Figure 4-5 Invocation Trace for a Proxy Service (LoanGateway3) Test

Invocation Trace

(receiving request) ▾

Initial Message Context

- + added \$body ▾
- + added \$header ▾
- + added \$inbound ▾
- + added \$messageID ▾

↕ PipelinePairNode1

validate loan application ▾

Message Context Changes

- ▲ changed \$body ▾
- ▲ changed \$inbound ▾

⚠ Stage Error Handler ▾

\$fault:

```
<con:fault xmlns:con="http://www.bea.com/wli/sb/context">
  <con:errorCode>BEA-382505</con:errorCode>
  <con:reason>ALSB Validate action failed validation</con:reason>
  <con:details>
    <con1:ValidationFailureDetail xmlns:con1="http://www.bea.com/wli/sb/stages/transform/config">
      <con1:message>
        Decimal fractional digits (1) of value '20.5' does not match fractionDigits facet (0) for xs:int
      </con1:message>
      <con1:xmlLocation>
        <java:NumOfYear xmlns:java="java:normal.client">20.5</java:NumOfYear>
      </con1:xmlLocation>
    </con1:ValidationFailureDetail>
  </con:details>
  <con:location>
    <con:node>PipelinePairNode1</con:node>
    <con:pipeline>PipelinePairNode1_request</con:pipeline>
    <con:stage>validate loan application</con:stage>
  </con:location>
</con:fault>
```

Compare the output in the trace with the nodes in the message flow shown in [Figure 4-4](#).

The trace indicates the following:

- **Initial Message Context**—Shows the variables initialized by the proxy service when it is invoked. To see the value of any variable, click the + associated with the variable name.
- **Changed Variables**—\$body and \$inbound changed as a result of the processing of the message through the validate loan application stage. These changes are seen at the end of the message flow.

- The contents of the `fault` context variable (`$fault`) is shown as a result of the Stage Error Handler handling the validation error. (The non-integer value (**20.5**) you entered for the `<java:NumOfYear>` element in [Listing 4-1](#) caused the validation error in this case.)

For more information about this loan application scenario, see [Tutorial 3: Validating a Loan Application](#) in *AquaLogic Service Bus Tutorials*.

It is left as an exercise to the reader to test the service using different input parameters, or to change the behavior of the message flow in the AquaLogic Service Bus Console Project Explorer, and run the test again to view the results.

Testing Resources

You can test resources inside an active session or from outside a session. You can test the following resources:

- “MFL” on page 4-12
- “XSLT” on page 4-14
- “XQuery” on page 4-14

MFL

A Message Format Language (MFL) document is a specialized XML document used to describe the layout of binary data.

MFL resources support the following transformations:

- XML to Binary - there is one required input (XML) and one output (Binary).
- Binary to XML - there is one required input, Binary, and one output, XML.

Each transformation only accepts one input and provides a single output.

The following example describes an XML input file to be tested in the test console. When you invoke the test console to test the MFL file, sample XML data is generated. Execute the test using the sample XML—in this case, a successful test results in the transformation of the message content of the input XML document in to binary format. The following [“Example” on page 4-13](#) section describes the MFL, the test XML, and the data resulting from the test.

Example

The following listing is an example MFL file.

Listing 4-2 Contents of an MFL File

```
<?xml version='1.0' encoding='windows-1252'?>
<!DOCTYPE MessageFormat SYSTEM 'mfl.dtd'>
  <MessageFormat name='StockPrices' version='2.01'>
    <StructFormat name='PriceQuote' repeat='*'>
      <FieldFormat name='StockSymbol' type='String' delim=':'
codepage='windows-1252' />
      <FieldFormat name='StockPrice' type='String'
delim='|' codepage='windows-1252' />
    </StructFormat>
  </MessageFormat>
```

The XML input generated by the test console to test the MFL file in the [Listing 4-2](#) is described in the following listing.

Listing 4-3 Test Console XML Input

```
<StockPrices>
  <PriceQuote>
    <StockSymbol>StockSymbol_31</StockSymbol>
    <StockPrice>StockPrice_17</StockPrice>
  </PriceQuote>
</StockPrices>
```

In the test console, click **Execute** to run the test—the result is the Stock symbol and the stockPrice in binary format as shown in the following listing.

Listing 4-4 MFL Test Console Results

```
00000000:53 74 6F 63 6B 53 79 6D 62 6F 6C 5F 33 31 3A 53 StockSymbol_31:S  
00000010:74 6F 63 6B 50 72 69 63 65 5F 31 37 7C .. .. .. StockPrice_17|...
```

XSLT

Extensible Stylesheet Language Transformation (XSLT) describes XML-to-XML mappings in AquaLogic Service Bus. You can use XSL Transformations when you edit XQuery expressions in the message flow of proxy services

To test an XSLT resource, you must supply an input XML document. The test console displays the output XML document as a result of the test. You can create parameters in your document to assist with a transformation. XSLT parameters accept either primitive values or XML document values. You cannot identify the types of parameters from the XSL transformation. In the Input and parameters section of the XSLT Resource Testing page in the test console, you must provide the values to bind to the XSLT parameters defined in your document.

XQuery

XQuery uses the structure of XML intelligently to express queries across different kinds of data, whether physically stored in XML or viewed as XML.

An XQuery transformation can take multiple inputs and returns one output. The inputs expected by an XQuery transformation are variable values to bind to each of the XQuery external variables defined. The value of an XQuery input variable can be a primitive value (string, integer, date), an XML document, or a sequence of the previous types. The output value can be primitive value (string, integer, date), an XML document, a sequence of the previous types.

XQuery is a typed language—every external variable is given a type. The types can be categorized into the following groups:

- Simple/primitive type—string, int, float, and so on.
- XML nodes
- Untyped

In the test console, a single-line edit box is displayed if the expected type is a simple type. A multiple-line edit box is displayed if the expected data is XML. A combination input is used when

the variable is not typed. The test console provides the following field in which you can declare the variable type: [] as XML. Input in the test console is rendered based on the type. This makes it easy to understand the type of data you must enter.

For example, the following figure shows an XQuery with three variables: int, XML, and undefined type.

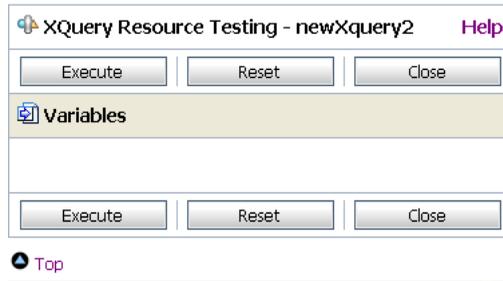
Figure 4-6 Input to the XQuery Test

| default/newXquery2 | | Description |
|--------------------|-------------------|--------------------|
| Last Modified By | weblogic | - no description - |
| Last Modified On | 06/05/06 11:57 AM | |
| References | 0 | |
| Referenced By | 1 | |

| XQuery | |
|--------|--|
| | <pre> <xs:element name="routing"> <xs:annotation> <xs:documentation>This is a simple routing table</xs:documentation> </xs:annotation> <xs:complexType> <xs:sequence> <xs:element name="local" maxOccurs="unbounded"> <xs:complexType> <xs:sequence> <xs:element name="logical"/> <xs:element name="physical"/> </xs:sequence> </xs:complexType> </xs:element> <xs:element name="remote" minOccurs="0" maxOccurs="unbounded"> <xs:complexType> <xs:sequence> <xs:element name="logical"/> <xs:element name="physical"/> </xs:sequence> </xs:complexType> </xs:element> </xs:sequence> </xs:complexType> </xs:element> </pre> |

In the Test Console, all three variables are listed in the Variables section. By default, for the untyped variable XML is checked as it is the most usual case. You must configure these variables.

Figure 4-7 Configuring the XQuery Variables in the Test Console



You can also test an XQuery expression from the XQuery Editor.

Performing In-line XQuery Testing

You must disable the pop-up blockers in your browser for the inline XQuery testing to work. Note that if you have toolbars in the Internet Explorer browser, you may need to disable pop-up blockers from under the browser's Options menu as well as for all toolbars that are configured to block them.

When performing in-line XQuery testing with the test console, you can use the Back button to return to the page from where you can execute a new test. But if you want to execute a new test after making changes to the in-line XQuery, you must close and re-open the test console for the changes to take effect.

Testing Services With Web Service Security

The test console supports testing proxy services and business services protected with Web Service Security (WSS). A SOAP service is protected with WSS if it has WS-Policies with WS-Security assertions assigned to it. Specifically, a service operation is protected with WS-Security if the operation's effective request and/or response WS-Policy includes WS-Security assertions. WS-Policies are assigned to a service by a mechanism called WS-PolicyAttachment. See "[Web Services Policy Attachment](#)" on page 3-33. Note that an operation may have both a request policy and a response policy.

When an operation has a request WS-Policy or response WS-Policy, the message exchange between the test console and the service is protected by the mechanisms of WS-Security. According to the operation's policy, the test service digitally signs and/or encrypts the message (more precisely, parts of the message) and includes any applicable security tokens. The input to the digital signature and encryption operations is the clear-text SOAP envelope specified by the

user as described in “Configuring Proxy Service Test Data” and “Configuring Business Service Test Data” in [Test Console](#) in the *Using the AquaLogic Service Bus Console*.

Similarly, the service processes the response according to the operation’s response policy. The response may be encrypted or digitally signed. The test service then processes this response and decrypts the message and/or verifies the digital signature.

The test console (**Security** panel) displays fields used for testing services with WS-Security: **Service Provider**, **Username** and **Password**.

Figure 4-8 Security Panel in Test Console



If you specify a proxy service provider in the test console, all client-side PKI key-pair credentials required by WS-Security are retrieved from the proxy service provider. You use the user name and password fields when an operation’s request policy specifies an Identity assertion and user name Token is one of the supported token types. For more information, see [Web Service Policy](#).

The Service Provider, user name, and Password fields are displayed whenever the operation has a request or response policy. Whether the values are required depends on the actual request and response policies.

The following table describes the different scenarios.

Table 4-1 Digital Signature and Encryption Scenarios

| Scenario | Is Proxy Service Provider Required? |
|---|--|
| The request policy has a Confidentiality assertion. | <p>No. The test service encrypts the request with the service’s public key. When testing a proxy service, the test service automatically retrieves the public key from the encryption certificate assigned to the proxy service provider of the proxy service.</p> <p>When testing a business service, the encryption certificate is embedded in the WSDL of the business service. The test service automatically retrieves this WSDL from the WSDL repository and extracts the encryption certificate from the WSDL.</p> |

Table 4-1 Digital Signature and Encryption Scenarios

| | |
|---|---|
| <p>The response policy has a Confidentiality assertion.</p> | <p>Yes. In this scenario, the operation policy requires the client to send its certificate to the service. The service will use the public key from this certificate to encrypt the response to the client. A proxy service provider <i>must</i> be specified and <i>must</i> have an associated encryption credential.</p> <p>If both request and response encryption are supported, different credentials must be used.</p> |
| <p>The request policy has an Integrity assertion.</p> | <p>Yes. The client must sign the request. A proxy service provider <i>must</i> be specified and <i>must</i> have an associated digital signature credential.</p> <p>Furthermore, if this is a SAML holder-of-key integrity assertion, a user name and password is needed in addition to the proxy service provider.</p> |
| <p>The response policy has an Integrity assertion.</p> | <p>No. In this case, the policy specifies that the service must sign the response. The service signs the response with its private key. The test console simply verifies this signature.</p> <p>When testing a proxy service, this is the private key associated to the proxy service provider’s digital signature credential for the proxy service.</p> <p>When testing a business service, the service signing key-pair is configured in a product-specific way on the system hosting the service.</p> <p>In the case that the current security realm is configured to do Certificate Lookup and Validation, then the certificate that maps to the proxy service provider must be registered valid in the certificate lookup and validation framework.</p> <p>For more information on Certificate Lookup and Validation, see "Configuring the Credential Lookup and Validation Framework" in Configuring WebLogic Security Providers in <i>Securing WebLogic Server</i>.</p> |

Table 4-2 Identity Policy Scenarios. It is Assumed the Policy has an Identity Assertion

| Supported Token Types¹ | Description | Comments |
|--|---|--|
| UNT | The service only accepts WSS user name tokens | The user must specify a user name and password in the security section. |
| X.509 | The service only accepts WSS X.509 tokens | The user must specify a proxy service provider in the security section and the proxy service provider must have an associated WSS X.509 credential. |
| SAML | The service only accepts WSS SAML tokens | The user must specify a user name and password in the security section <i>or</i> a user name and password in the transport section. If both are specified, the one from the security section is used as the identity in the SAML token. |
| UNT, X.509 | The service accepts UNT or X.509 tokens | The user must specify a user name and password in the security section <i>or</i> a proxy service provider in the security section with an associated WSS X.509 credential. If both are specified, only a UNT token is generated. |
| UNT, SAML | The service accepts UNT or SAML tokens | The user must specify a user name and password in the security section <i>or</i> a user name and password in the transport section. If both are specified, only a UNT token is sent. |
| X.509, SAML | The service accepts X.509 or SAML tokens | The user must specify one of the following: <ul style="list-style-type: none"> • a user name and password in the security section • a user name and password in the transport section • a proxy service provider with an associated WSS X.509 credential |
| UNT, X.509, SAML | The service accepts UNT, X.509 or SAML tokens | The user must specify one of the following: <ul style="list-style-type: none"> • a user name and password in the security section • a user name and password in the transport section • a proxy service provider with an associated WSS X.509 credential. |

1. From the Identity Assertion inside the request policy.

Limitations for Services and Policies

The following limitations exist for testing proxy services with SAML policies and business services with SAML holder-of-key policies:

- Testing of proxy services with inbound SAML policies is not supported
- Testing business services with a SAML holder-of-key policy is a special case.

The SAML holder-of-key scenario can be configured in two ways:

- as an integrity policy (this is the recommended approach)
- as an identity policy

In both cases the user must specify a user name and password—the SAML assertion will be on behalf of this user. If SAML holder-of-key is configured as an integrity policy, the user must also specify a proxy service provider. The proxy service provider must have a digital signature credential assigned to it. This case is special because this is the only case where a user name and password must be specified even if there is not an identity policy.

Note: After executing a test in the test console, the envelope generated with WSS is not always a valid envelope—the results page in the test console includes white spaces for improved readability. That is, the secured SOAP message is displayed printed with extra white spaces. Because white spaces can affect the semantic of the document, this SOAP message cannot always be used as the literal data. For example, digital signatures are white-space sensitive and can become invalid.

Test Console Transport Settings

The transport panel in the test console provides the functionality to specify the metadata and transport headers for messages in your test system. The following figure shows an example of a Transport panel on the test console.

Figure 4-9 Transport Panel in the Test Console

The Transport panel is a form with a title bar labeled "Transport" and a close button. It is organized into several sections:

- Authentication:** Username: Password:
- Metadata:** encoding: relative-URI: query-string: client-host: client-address:
- Headers:** Accept: Accept-Encoding: Accept-Language: Connection: Content-Encoding: Content-Length: Content-Type: Host: SOAPAction: User-Agent:
- User Headers:** name: value:

At the bottom of the panel are three buttons:

[Top](#)

The preceding figure displays an example of the transport panel for a given service—in this case, a WSDL-based proxy service.

You can set the metadata and the transport headers in the message flow of a proxy service. In doing this, you influence the actions of the outbound transport. You can test the metadata, the message, and the headers so that you can see the output you get in the pipeline. The fields that are displayed in the Transport panel when testing a proxy service represent those headers and metadata that are available in the pipeline. The test console cannot filter the fields it presents depending on the proxy service. The same set of transport parameters are displayed on the page for every HTTP-based request.

The **Username** and **Password** fields are used to implement basic authentication for the user that is running the proxy service. The **Username** and **Password** fields are not specifically transport related.

Metadata fields are grouped in the **Transport** panel, below the **Username** and **Password** fields and above the group of transport header fields. The fields displayed are based on the transport type of the service. Certain fields are pre populated in the test console depending on the operation selection algorithm you selected for the service when you defined it.

For example, in the case of the transport panel displayed in [Figure 4-9](#), the `SOAPAction` header field is populated with “`http://example.orgprocessLoanApp`”. This value was taken from the service definition (the selection algorithm selected for this proxy service was `SOAPActionHeader`). For more information about the selection algorithms, see “Adding a Proxy Service” in [Proxy Services](#) in *Using the AquaLogic Service Bus Console*.

When you specify values for fields in the transport panel, be aware whether you opted to test the service using a direct or indirect call—see “[Direct Calls](#)” on [page 4-3](#) and “[Indirect Calls](#)” on [page 4-4](#)—and specify the values according to whether the message will be processed through the transport layer or not.

When testing a proxy service with a direct call, the test data must represent the message as if it had been processed through the transport layer. That is, the test data should represent the message in the state expected at the point it leaves the transport layer and enters the service. When testing a proxy or business service, using an indirect call, the test data represents the data that is sent from a route node or a service callout. The test message is processed through the transport layer.

For information about specific headers and metadata and how they are handled by the test framework, see [Understanding How the Run Time Uses the Transport Settings in the Test Console](#).

About Security and Transports

- When using the test console to test HTTP(S) business services with BASIC authentication, the test console authenticates with the user name and password from the service account of the business service. Similarly, when testing JMS, e-mail, or FTP business services that require authentication, the test console authenticates with the service account associated with the business service.
- When you test proxy services, the test console never sends a HTTP request over the network. Therefore transport-level access control is not applied.

Understanding How the Run Time Uses the Transport Settings in the Test Console

The test console allows you to specify header values and metadata. However, when the message is sent out, some headers and metadata may be modified or removed, and the underlying transport may in turn, ignore some of the headers and use its own values when the test is executed.

The following table describes the headers and metadata for which there are limitations when using the test console.

Table 4-3 Limitations to Transport Header and Metadata Values You Specify in the Test Console When Testing a Service

| Transport | Testing this Service Type | Description of Limitation | Transport Headers Affected |
|----------------------|---------------------------|--|---|
| HTTP(S) ¹ | Proxy Service | All transport headers and other fields you set are preserved at run time. This is true whether or not the Direct Call option is set. | All |
| | Business Service | The AquaLogic Service Bus run time overrides any values you set for these parameters | <ul style="list-style-type: none"> • Content-Length • Content-Type • relative-URI • client-host • client-address |

Table 4-3 Limitations to Transport Header and Metadata Values You Specify in the Test Console When Testing a Service

| Transport | Testing this Service Type | Description of Limitation | Transport Headers Affected |
|-------------|---------------------------|--|---|
| JMS | Proxy Service | Direct Call When the Direct Call option is used, all transport headers and other fields you set are preserved at run time | All |
| | | X Direct Call When the Direct Call option is not used, the same limitations apply as for a transport header action configuration | See the limitations for JMS transport headers described in “Transport Headers” in Proxy Services: Actions in <i>Using the AquaLogic Service Bus Console</i> . |
| | Business Service | The same limitations apply as for a transport header action configuration | See the limitations for JMS transport headers described in “Transport Headers” in Proxy Services: Actions in <i>Using the AquaLogic Service Bus Console</i> . |
| | E-Mail | Proxy Service | No limitations. In other words, any transport headers and other fields you set are honored by the run time. This is true whether or not Direct Call is specified. |
| | Business Service | The AquaLogic Service Bus run time overrides any values you set for these parameters | <ul style="list-style-type: none"> Content-Type |
| File | Proxy Service | No limitations. In other words, any transport headers and other fields you set are honored by the run time. ² | |
| | Business Service | | |

Table 4-3 Limitations to Transport Header and Metadata Values You Specify in the Test Console When Testing a Service

| Transport | Testing this Service Type | Description of Limitation | Transport Headers Affected |
|------------|---|---|----------------------------|
| FTP | Proxy Service <hr/> Business Service | No limitations. In other words, any transport headers and other fields you set are honored by the run time. | |

1. When you test proxy services, the test console never sends a HTTP request over the network, therefore transport-level access control is not applied.
2. For example, in the case of FileName (Transport metadata)—the value you assign is used to append to the output file name. For example, 1698922710078805308-b3fc544.1073968e0ab.-7e8e-{\$FileName}

Using the Test Console

Reporting

BEA AquaLogic Service Bus delivers message data and alerts to one or more reporting providers. Message data can be captured from the body of the message and from other variables associated with the message, such as header or inbound variables. Alert data contains information about Service Level Agreement (SLA) violations that you can configure to monitor proxy services. You can use the message or alert data delivered to the reporting provider for functions such as tracking messages or regulatory auditing.

AquaLogic Service Bus includes a JMS Reporting Provider for message reporting. The Reporting module in the AquaLogic Service Bus Console displays the information captured from this reporting provider. If you do not wish to use the JMS Reporting Provider that is provided with your AquaLogic Service Bus installation, you can untarget it and create your own reporting provider using the Reporting Service Provider Interface (SPI). If you configure your own reporting provider for messages, no information is displayed in the AquaLogic Service Bus Console and you will need to create your own user interface. If you wish to capture SLA data, you will need to create a reporting provider for alerts.

This chapter contains information on the following topics

- [“Reporting Scenarios” on page 5-2](#)
- [“Reporting Framework” on page 5-3](#)
- [“JMS Reporting Provider” on page 5-5](#)
- [“How to Enable Message Reporting” on page 5-6](#)
- [“Removing, Stopping, or Untargeting a Reporting Provider” on page 5-15](#)

Reporting Scenarios

The following scenarios describe some of the ways in which you can use AquaLogic Service Bus to track messages:

- [“Message Tracking” on page 5-2](#)
- [“Search for a Particular Message” on page 5-2](#)
- [“Logging for Regulatory Auditing” on page 5-2](#)
- [“Alert Reporting Provider” on page 5-3](#)

Message Tracking

In the AquaLogic Service Bus Console, when you filter messages to display the Post-Trade Processing proxy service. When you drill down into some of the messages, you discover that the pipeline errors are due to message transformation errors and that the mangled messages are coming out of the portal associated with the proxy service. To solve the problem, you can add a new transformation to the pipeline for all messages originating from that portal site.

Search for a Particular Message

Customer Service calls Operations with a complaint that the customer who submitted a trade did not receive a trade confirmation. You can log into the AquaLogic Service Bus Console and search for the trade number. The search shows that two messages were processed in the request pipelines, but no response messages. You can ask Customer Service to contact the customer and assure the customer that the trade was successfully processed.

Logging for Regulatory Auditing

At the end of the month, the Mortgage Processing team must provide their compliance department with information about the mortgages they have processed. To fulfill this requirement, the Application Development team updates the applicable proxy service pipelines to capture the relevant message information. Specifically, data is extracted from the messages for the customer ID and name, mortgage ID, mortgage amount, property address, and the date the loan application was submitted.

Alert Reporting Provider

By configuring a reporting provider for alerts, you can receive an alert notification outside of the AquaLogic Service Bus Console and process the alert according to your business needs. For example, you could develop an alert reporting provider that utilizes the reporting stream for alerts and then display the alerts on a custom console, such as HP OpenView, or Tivoli.

Reporting Framework

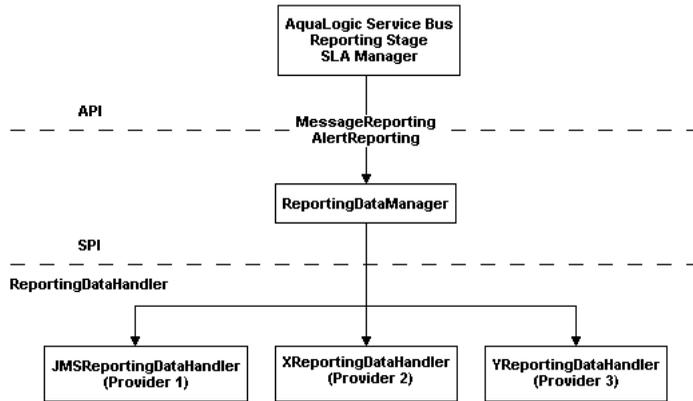
AquaLogic Service Bus contains an extensible framework for creating one or more reporting providers for messages or alerts.

To enable message reporting you must first create a Report action in the message flow for the proxy service. The Report action allows you to extract information from each message and write it to the AquaLogic Service Bus Reporting Data Stream. You do not need to configure a report action for alert reporting. Alert data is always available in the Reporting Data Stream. For more information, see [“How to Enable Message Reporting” on page 5-6](#).

All the information you need in order to create your own reporting provider is located in `com.bea.wli.reporting` in the [Javadoc for AquaLogic Service Bus](#). The Javadoc provides information about what you need to do to implement a reporting provider, including how to package it, where it goes, how to deploy it, and the order of deployment. The location of the reporting schema (`MessageReporting.xsd`) is `<BEA_HOME>/weblogic92/servicebus/lib/sb-schemas.jar`, where `BEA_HOME` is the directory in which you installed BEA products.

The following figure shows the reporting framework.

Figure 5-1 Reporting Framework



As shown in the [Figure 5-1](#), both report messages and alerts are exported to reporting data streams. In the Report stage, information is extracted by the Report action from each message and written to the Reporting Data Stream with metadata that adheres to `MessageReporting.xsd`. Similarly, the SLA Manager uses Reporting Data Manager APIs to write to the Alert Reporting Stream with metadata that adheres to the `AlertReporting.xsd`. If you want to develop a reporting provider for alerts or your own message reporting provider, you need to implement an interface called `ReportingDataHandler` and use `ReportingDataManager` class.

The `ReportingDataHandler` interface takes the reporting or alert data stream and processes it. It can either process or store, or both this stream in a relational database, file, JMS queue, and so on. Depending on which stream you want to use, you need to implement the appropriate handle methods to process the data stream:

- **Message Reporting Stream**—the report action of AquaLogic Service Bus run time uses the following two `handle` methods to write to the Message Reporting Stream:

```

handle(com.bea.xml.XmlObject metadata, String s)
handle(com.bea.xml.XmlObject metadata, com.bea.xml.XmlObject data)
  
```

- **Alert Reporting Stream**—the Alert Manager uses the following `handle` method to write to the Alert Reporting Stream:

```

handle(com.bea.xml.XmlObject metadata, com.bea.xml.XmlObject data)
  
```

The `ReportingDataManager` is a server-local object that keeps a registry of reporting providers. Reporting providers implement the `ReportingDataHandler` interface. The `ReportingDataManager` provides operations to do the following:

- Add and remove reporting data handlers
- Export reporting data stream using various handle operations.

JMS Reporting Provider

The JMS Reporting Provider provides a pluggable architecture to capture the reporting information from each message via a Report action. All messages across the cluster are aggregated and stored in the JMS Reporting Provider Data Store in a database specific format. When you use the JMS Reporting Provider, which is provided with AquaLogic Service Bus installation the Reporting module in the AquaLogic Service Bus Console displays information from the JMS Reporting Provider Data Store.

Note: The JMS Reporting Provider is automatically configured when you create an AquaLogic Service Bus domain. If you do not wish to use this reporting provider, you must untarget it. For more information, see [“Removing, Stopping, or Untargeting a Reporting Provider”](#) on page 5-15.

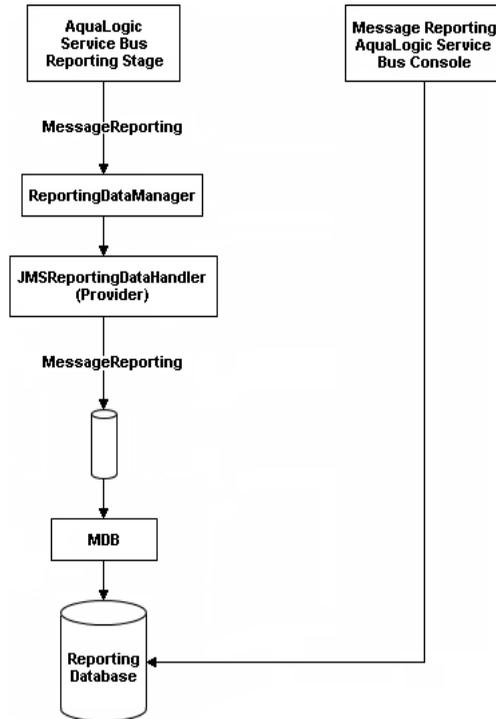
This section contains information on the following topics:

- [“About the JMS Reporting Provider”](#) on page 5-5
- [“How to Enable Message Reporting”](#) on page 5-6
- [“Using the Reporting Module”](#) on page 5-8

About the JMS Reporting Provider

The JMS Reporting Provider consists of a producer and a consumer, which are decoupled to improve scalability. The producer is a JMS producer and the Message Driven Bean (MDB) acts as the JMS consumer, as shown in the following diagram.

Figure 5-2 JMS Reporting Provider



The Reporting stage contains the Report actions that collect the reporting information and dispatch the reporting stream to JMS Reporting Provider through various `handle` operations in the `ReportingDataManager`. The `JMSReportingDataHandler` is the JMS producer of the reporting provider. The `JMSReportingDataHandler` takes the reporting stream and logs the information to a JMS queue. The MDB listens to the JMS reporting queue, which processes the message asynchronously and stores the data in the JMS Reporting Provider Data Store.

How to Enable Message Reporting

To receive report messages from either the JMS Reporting Provider, which is provided with AquaLogic Service Bus installation, or your reporting provider, you must first create a Report action in the message flow for the proxy service. The Report action allows you to extract information from each message and write it to the AquaLogic Service Bus Reporting Data Stream. In the Report action, you must specify the information you want to extract from the message and add to the AquaLogic Service Bus Reporting Data Stream.

You do not need to configure a report action for alert reporting. Alert data is always available in the Reporting Data Stream.

When configuring a Report action, you use key values to extract key identifiers from the message. You can configure multiple keys. Information can be captured not only from the body of the message but any other variable associated with the message, such as header or inbound variables. For more information about message variables, see [Message Context](#) in *Using the AquaLogic Service Bus Console*.

You can use any XML elements as a key:

```
<?xml version="1.0" encoding="utf-8"?>
  <poIncoming>
    <areacode>408</areacode>
    <item-quantity>100</item-quantity>
    <item-code>ABC</item-code>
    <item-description>Medicine</item-description>
  </poIncoming>
```

For example, you can specify the key as the `itemcode`, the value as `./item-code` (an XPath expression), and the variable as message body (`body`), as shown in the following figure.

Figure 5-3 Key Name and Value



Report `$body` with search keys:

| Key Name | Key Value | Options |
|----------|---|---------|
| itemCode | <code>./itemCode</code> in variable <code>body</code> | |

If you are using the JMS Reporting Provider, which is provided with AquaLogic Service Bus installation, the keys and associated values are displayed in the Report Index column of the Summary of Messages table. If you configure multiple keys, the key-value pairs are displayed in Report Index Column with each key-value separated by a semicolon, as shown in [Figure 5-4](#).

Figure 5-4 Keys and Associated Values Display

| Summary of Messages | | |
|--|--|---|
| Report Index  | DB TimeStamp  | Inbound Service  |
| Customer ID=EA-3822883 | 6/10/06 8:11 PM | ProxyService\$MortgageBroker\$ |
| Customer Name=John Smith | 6/10/06 8:11 PM | ProxyService\$MortgageBroker\$ |
| Property Address=2315 North St, San Jose, CA 95131 | 6/10/06 8:11 PM | ProxyService\$MortgageBroker\$ |
| Date=6/20/05 | 6/10/06 8:11 PM | ProxyService\$MortgageBroker\$ |
| Mortgage Amount=\$778,900,Interest Rate=05.00% | 6/10/06 8:11 PM | ProxyService\$MortgageBroker\$ |

For information on how to create a Report action or on how to view the Summary of Messages page, see the following in *Using the AquaLogic Service Bus Console*:

- Report in [Proxy Services: Actions](#).
- Listing and Locating Messages in [Reporting](#).

Using the Reporting Module

The reporting module in the AquaLogic Service Bus Console displays the information collected by the JMS Reporting Provider Data Store. The first page of the Reporting module, called the Summary of Messages, displays a table containing the extracted information and other information, such as the time the message was written to the database and the service with which the message is associated. You can customize the display of information on this page by filtering and sorting the data. You can also drill down to view detailed information about specific messages, including error information.

The Reporting module provides a purge function to help you manage your message data. You can purge all of the messages from the reporting datastore or base the purge on a time-range.

The JMS Reporting Provider Data Store requires a database. An evaluation version of the PointBase database is installed with WebLogic Server. You can use PointBase for a development environment but not for production. AquaLogic Service Bus also supports databases from other vendors. Be sure to apply standard database administration practices to the database hosting the JMS Reporting Provider Data Store. For more information, see [“Configuring a Database for the JMS Reporting Provider Store”](#) on page 5-14.

For more information on how to use the reporting module is located in the [Using the AquaLogic Service Bus Console](#).

This section includes information on the following topics:

- [“Summary of Messages” on page 5-9](#)
- [“View Message Details” on page 5-10](#)
- [“Purging Messages” on page 5-13](#)

Summary of Messages

When you click Reporting in the navigation panel, the Summary of Messages page is displayed. This page contains a table that provides a list of report messages sorted by the database timestamp.

Figure 5-5 Summary of Messages

| Summary of Messages Search | | | |
|---|-----------------------------|---|---------------------------|
| Report Index △ | DB TimeStamp △ | Inbound Service ▽ | Error Code △ |
| errorCode=BEA-382000 | 6/5/06 2:35 PM | ProxyService\$MortgageBroker\$ProxyServices\$loanGateway3 | BEA-382000 |
| errorCode=BEA-382000 | 6/5/06 2:35 PM | ProxyService\$MortgageBroker\$ProxyServices\$loanGateway3 | BEA-382000 |
| errorCode=BEA-382000 | 6/5/06 2:35 PM | ProxyService\$MortgageBroker\$ProxyServices\$loanGateway3 | BEA-382000 |
| errorCode=BEA-382000 | 6/5/06 2:40 PM | ProxyService\$MortgageBroker\$ProxyServices\$loanGateway3 | BEA-382000 |

If the messages are not filtered, the Summary of Messages table displays up to 100 of the latest messages based on the database timestamp. If you filter the messages, up to 1000 messages are displayed.

Note: After you filter the message, the filter remains in effect until you update it.

The table shown in the preceding figure provides the following information:

- **Report Index:** displays the key-value pairs extracted from the message context variables or the message payload. For more information, see [“About the JMS Reporting Provider” on page 5-5](#).
- **DB TimeStamp:** the time when the message was written to the database.
- **Inbound Service:** the inbound service associated with the message. The service is a link to the View Proxy Service Details page.

- **Error Code:** the error code associated with this message, if it exists. For more information about error codes, see “Error Messages and Handling” in [Proxy Services: Error Handlers](#) in *Using the AquaLogic Service Bus Console*.

To search for specific messages, you can filter the display of messages by clicking Filter in the Summary of Messages Table. The available filtering is shown in the following figure.

Figure 5-6 Summary of Messages Search

The screenshot shows a web form titled "Summary of Messages". It contains several search filters: "Start Date" and "End Date" (both set to June 16, 2005, 4:53 PM), "For the Last" (0 days, 0 hours, 00 mins), "Inbound Service Name", "Error Code", and "Report Index" (with a "+" button). There are "Search" and "Reset" buttons at the bottom.

As shown in the [Figure 5-6](#), you can filter report messages for a specified period of time, by the name of a service, by error code, and by report index. After you filter the messages, the title of the page changes to Summary of Filtered Messages. For information on how to use the Summary of Messages filter, see “Listing and Locating Messages” in [Reporting](#) in *Using the AquaLogic Service Bus Console*.

To view more information about a report message, click the name of the message in the Report Index column. The View Message Details page is displayed.

View Message Details

The View Message Details page displays complete information about the report messages, as shown in the following figure.

Figure 5-7 Report Message Detail Page

| View Message uuid:c91ab9e973f25cc1:17db5bc9:10480ef1244:-7fd9 Details | |
|---|---|
| OK | |
| General Configuration | |
| Message ID | uuid:c91ab9e973f25cc1:17db5bc9:10480ef1244:-7fd9 |
| Database Timestamp | Wednesday, June 07 2006 5:05:00 PM MDT |
| Time at point of Logging | Wednesday, June 07 2006 5:05:00 PM MDT |
| Server Name | xbusServer |
| State | ERROR |
| Node Name | PipelinePairNode1 |
| Pipeline Name | PipelinePairNode1_request |
| Stage Name | validate loan application |
| Inbound Service | |
| Name | ProxyService\$MortgageBroker\$ProxyServices\$loanGateway3 |
| URI | /loan/gateway3 |
| Operation | processLoanApp |
| Outbound Service | |
| Name | |
| URI | |
| Operation | |
| Report Index | |
| Report Index Text | errorCode=BEA-382000 |
| Fault | |
| Error Code | BEA-382000 |
| Message | Decimal fractional digits (1) of value '10.1' does not match fractionDigits (2) of format '###,###.##' in environment 'default' in class 'java.lang.NumberFormatException'. |

The page shows the following information:

- General Configuration
 - Message ID: the unique identification for this message.
 - Database Timestamp: the time when the message was written to the database.
 - Time at point of Logging: the date and time, on the server machine, that the message was reported.
 - Server name: the name of the server from which this message was generated.

Reporting

- State: state of the pipeline from which this message was generated, as follows:
 - REQUEST: indicates that the reporting action was executed in a request pipeline.
 - RESPONSE: indicates that the reporting action was executed in a response pipeline.
 - ERROR: the action was running in the service-level error handler.
- Node Name: the pipeline node from which this message was generated.
- Pipeline Name: the pipeline from which this message was generated.
- Stage Name: the stage from which this message was generated.
- Inbound Service
 - Name: the inbound proxy service associated with this message. An inbound proxy service exchanges messages with client applications. The name is a link to the View Proxy Service Details page. For more information about this page, see “Viewing and Changing Proxy Services” in [Proxy Services](#) in *Using the AquaLogic Service Bus Console*.
 - URI: the URI associated with the proxy service.
 - Operation: the inbound operation associated with this message. Operations are the tasks performed by a pipeline or route node in the message flow associated with the service.
- Outbound Service
 - Name: the outbound business service associated with this message. An outbound business service exchanges messages with an AquaLogic Service Bus proxy service. The name is a link to the View Business Service Details page. For more information about this page, see “Viewing and Changing Business Services” in [Business Services](#) in *Using the AquaLogic Service Bus Console*.
 - URI: the URI to the outbound business service end point.
 - Operation: name of the operation invoked on the outbound service. Operations are the tasks performed by a pipeline or route node in the message flow associated with the service.
- Report Index
 - Report Text Index: displays the key-value pairs extracted by a Report Action from the message context variables or the message payload. For more information, see “[About the JMS Reporting Provider](#)” on page 5-5.

- Fault
 - Error Code: the code associated with the error, if any. For more information, see “Error Messages and Handling” in [Proxy Services: Error Handlers](#) in *Using the AquaLogic Service Bus Console*.
 - Reason: the reason for the error.
 - Detail: The fault details associated with the error. These details, if present, are typically a stack trace of where a particular fault occurred. The stack trace may be truncated due to a size limitation in the database. The limit is 2048 characters.
- Report Body
 - Detail: opens a browser window that displays the report body in a browser. You can use an XQuery expression in a Report Action to capture the report body text. For more information, see “Report” in [Proxy Services: Actions](#) and “Using the Inline XQuery Expression Editor” in [Proxy Services: XQuery Editors](#) in *Using the AquaLogic Service Bus Console*.

Purging Messages

You can purge all of the messages from the reporting datastore or base the purge on a range of time. Message purging is an asynchronous process that occurs in the AquaLogic Service Bus Console. This feature enables you to work with the Summary of Messages page in the AquaLogic Service Bus Console while the purge occurs in the background.

Figure 5-8 Purging Messages Page

The screenshot shows a web form titled "Purge Messages". At the top left is a list icon. Below the title are two radio buttons: "Purge All Messages" (unselected) and "Purge From" (selected). The "Purge From" section contains a row of dropdown menus: "January", "1", "1970", "12", "00", and "AM". The "Purge To" section contains a row of dropdown menus: "June", "11", "2006", "5", "31", and "PM". At the bottom of the form is a "Submit" button.

The duration of time it takes a purge to complete depends on how many messages are in the purge queue. The deletion of messages is slowed if you search for reporting messages during the purge process. Moreover, the Summary of Messages page may display incorrect data as some data may not yet be purged.

Because the purge process is asynchronous and occurs in the background, the AquaLogic Service Bus Console does not display any messages to indicate that a purge is in process. However, if another user attempts to start a purge when a purge is in progress, the following message is displayed:

A Purge job is already running. Please try later.

Configuring a Database for the JMS Reporting Provider Store

AquaLogic Service Bus requires a database for the JMS Reporting Provider Data Store. The PointBase database that is installed with WebLogic Server is for evaluation purposes only and not intended for a production environment. Non-evaluation development or other use of the PointBase Server requires that you obtain a separate PointBase license.

In a production environment you must use one of the supported databases. For the latest information about supported databases, see “Supported Databases and Drivers” in [Supported Configurations for WebLogic Platform](#) in *Supported Configurations for AquaLogic Service Bus*.

This section provides information on the following topics:

- “Configuring a Database in a Development Environment” on page 5-14
- “Configuring a Database for Production” on page 5-15

Configuring a Database in a Development Environment

When you create an AquaLogic Service Bus domain, the Configuration Wizard does not create database tables automatically. In a development environment, the JMS Reporting Provider, which is provided with AquaLogic Service Bus installation checks whether tables exist for the specified database at run time. If tables do *not* exist, the Reporting Provider creates them; if they do exist, the Reporting Provider uses them.

Note: If you are using Pointbase, you do not need to specify a database in the Configuration Wizard.

You can specify which database is used by the JMS Reporting Provider in one of the following ways:

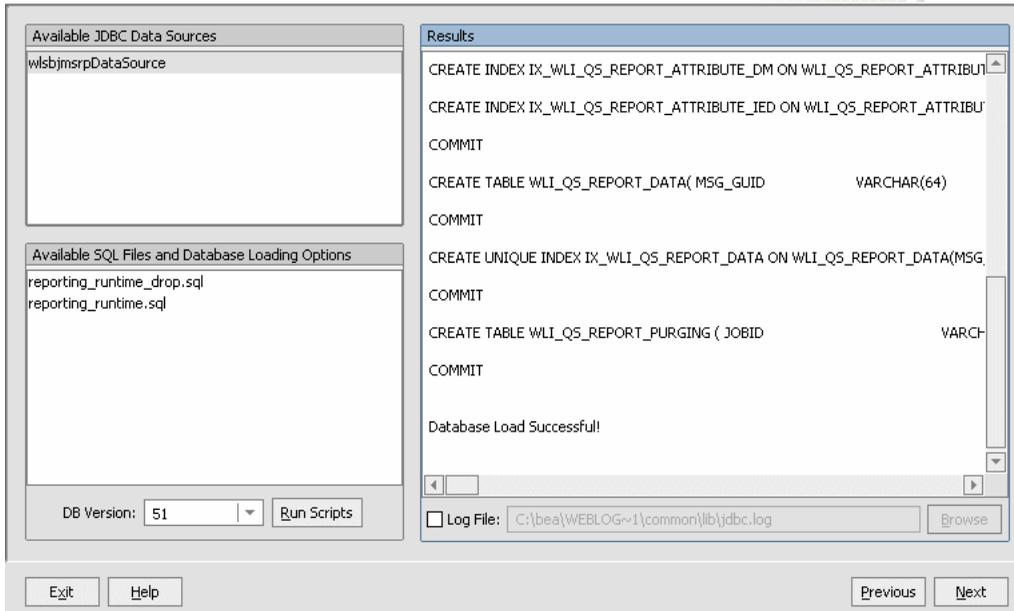
- Run the reporting SQL scripts in your AquaLogic Service Bus domain. The scripts are located in `<BEA_HOME>/weblogic92/integration/common/dbscripts`, where `BEA_HOME` represents the location in which you installed WebLogic products.

- When you create a domain in the Configuration Wizard, customize the JDBC Settings on the Run Database Scripts page (see [Figure 5-9](#)). For more information, see [Creating WebLogic Domains Using the Configuration Wizard](#).

Figure 5-9 Run Database Scripts in the Configuration Wizard

Run Database Scripts

If your connections test OK, select the database version, then click Run Scripts to create database objects and load necessary data.



Configuring a Database for Production

Complete information about configuring a database for production is located in the *AquaLogic Service Bus Deployment Guide* in the following chapters:

- [Configuring a Single-Server Deployment](#)
- [Configuring a Cluster Deployment](#)

Removing, Stopping, or Untargeting a Reporting Provider

As previously mentioned, the JMS Reporting Provider, which is provided with AquaLogic Service Bus installation is automatically configured when you create an AquaLogic Service Bus

domain. If you do not wish to use this reporting provider or any reporting provider, you must untarget it.

Note: If no reporting provider exists, you can still define a Report action. However, no data will be written.

The following sections provide information on how to stop or untarget any reporting provider:

- [“Stopping a Reporting Provider when the Server is Running” on page 5-16](#)
- [“Untargeting a Reporting Provider when the Server is Running” on page 5-18](#)
- [“Untargeting the JMS Reporting Provider—Server Not Running” on page 5-19](#)

Stopping a Reporting Provider when the Server is Running

If you wish to stop a reporting provider when the server is running in the AquaLogic Service Bus domain, do the following steps:

1. Start the WebLogic Server Administration Console. For more information, see “Starting the Administration Console” in [Overview of the Administration Console](#) in *Introduction to WebLogic Server and WebLogic Express*.
2. After logging into the WebLogic Server Administration Console, in the Domain Structure, click **Deployments**. The Summary of Deployments page is displayed.
3. In the Deployments table, select the checkbox beside the reporting provider you wish to stop.

Figure 5-10 Stopping a Reporting Provider

Summary of Deployments

Control **Monitoring**

This page displays a list of J2EE Applications and stand-alone application modules that have been installed to this domain. Installed applications and modules can be started, stopped, updated (redeployed), or deleted from the domain by first selecting the application name and using the controls on this page.

To install a new application or module for deployment to targets in this domain, click the Install button.

Deployments

Install Update Delete Start Stop Showing 21 - 30 of 52 Previous N

| <input type="checkbox"/> | Name | State | Type | Deployment Order |
|-------------------------------------|-----------------------------|--------|------------------------|------------------|
| <input checked="" type="checkbox"/> | ftpttransport-l10n(2.5,2.5) | Active | Library | 150 |
| <input type="checkbox"/> | IMS Reporting Provider | Active | Enterprise Application | 125 |
| <input type="checkbox"/> | jsf-myfaces(1.1,1.1.1.1) | Active | Library | 1 |
| <input type="checkbox"/> | jsf-ri(1.1,1.1.1.1) | Active | Library | 1 |
| <input type="checkbox"/> | jstl(1.1,1.1.1.2) | Active | Library | 1 |
| <input type="checkbox"/> | Message Reporting Purger | Active | Enterprise Application | 126 |
| <input type="checkbox"/> | sp13n-app-lib(9.2.0,9.2.0) | Active | Library | 1 |
| <input type="checkbox"/> | publish | Active | Web Application | 100 |
| <input type="checkbox"/> | sbconsole-l10n(2.5,2.5) | Active | Library | 400 |
| <input type="checkbox"/> | ServiceBus_Console | Active | Web Application | 810 |

Install Update Delete Start Stop Showing 21 - 30 of 52 Previous N

4. Click **Stop** and after the list is displayed, choose the appropriate command.
5. After the Stop Application Assistant page is displayed, click **Yes**. The Deployments table shows that the state of the reporting provider is now Prepared.

Untargeting a Reporting Provider when the Server is Running

If you wish to untarget a reporting provider when the server is running in the AquaLogic Service Bus domain, do the following:

1. Start the WebLogic Server Administration Console. For more information, see “Starting the Administration Console” in [Overview of the Administration Console](#) in *Introduction to WebLogic Server and WebLogic Express*.
2. After logging into the WebLogic Server Administration Console, in the Change Center, click **Lock & Edit**.
3. From the left panel, under Domain Structure, click **Deployments**. The Summary of Deployments page is displayed.
4. In the Deployments table, click the reporting provider you wish to untarget. The Settings page for the Reporting Provider is displayed.
5. Click the **Targets** tab.
6. Clear the appropriate checkbox.

Figure 5-11 Untargeting a Reporting Provider



7. Click **Save**. A message is displayed indicating that the settings have been successfully updated.

8. After you untarget the reporting provider, untarget the data source used by the reporting provider, as follows:

Note: This step is only required for reporting providers that use their own data sources. If you are untargeting the JMS Reporting Provider, which is provided with AquaLogic Service Bus installation you must perform the following steps.

- a. In the left panel, under Domain Structure, select **Services**→**JDBC**→**Data Sources**.
- b. In the Summary of JDBC Data Source page, click the name of the data source you wish to untarget. The Settings page for the data source is displayed.
- c. Click the **Targets** tab.
- d. Clear the appropriate checkbox.
- e. Click **Save**. A message is displayed indicating that the settings have been successfully updated.
- f. To activate the changes, in the Change Center, click **Activate Changes**.

Untargeting the JMS Reporting Provider—Server Not Running

If the server is not running in the AquaLogic Service Bus domain, you can use the WebLogic Scripting Tool (WLST) to remove the JMS Reporting Provider from the AquaLogic Service Bus domain. For more information about WLST, see [WebLogic Scripting Tool](#) in the WebLogic Server documentation.

To untarget a reporting provider, complete the following steps:

1. If you have not already set up your environment to use WLST, see “Main Steps for Using WLST” in [Using the WebLogic Scripting Tool](#) in *WebLogic Scripting Tool*.

2. Open a UNIX shell or terminal window.

3. Invoke WLST Offline.

```
C:>java com.bea.plateng.domain.script.jython.WLST_offline
```

4. Read the domain that was created using the Configuration Wizard. For example:

```
wls:/offline>readDomain("C:/bea/user_projects/domains/base_domain")
```

5. Untarget the reporting provider data source. For example:

```
wls:/offline/base_domain>unassign("JdbcSystemResource",
"wlslbjmsrpDataSource", "Target", "AdminServer")
```

6. Untarget the reporting provider application. For example:

```
wls:/offline/base_domain>unassign("AppDeployment", "JMS Reporting
Provider", "Target", "AdminServer")
```

7. Update the domain:

```
wls:/offline/base_domain>updateDomain()
```

8. Close the domain:

```
wls:/offline/base_domain>closeDomain()
```

9. Exit from the WLST command prompt:

```
wls:/offline>exit()
```

After the AquaLogic Service Bus JMS reporting provider is untargeted, the Reporting module in the AquaLogic Service Bus Console will indicate that the reporting provider is not deployed, as shown [Figure 5-12](#).

Figure 5-12 Reporting Provider Not Deployed



Note: In a cluster, the JMS Reporting Provider is targeted to Cluster. Therefore in a cluster, to view and purge messages, you must configure at least one managed server to run with the Administration server. If no managed servers are running, AquaLogic Service Bus Console displays the message shown in the previous figure.

Tracing

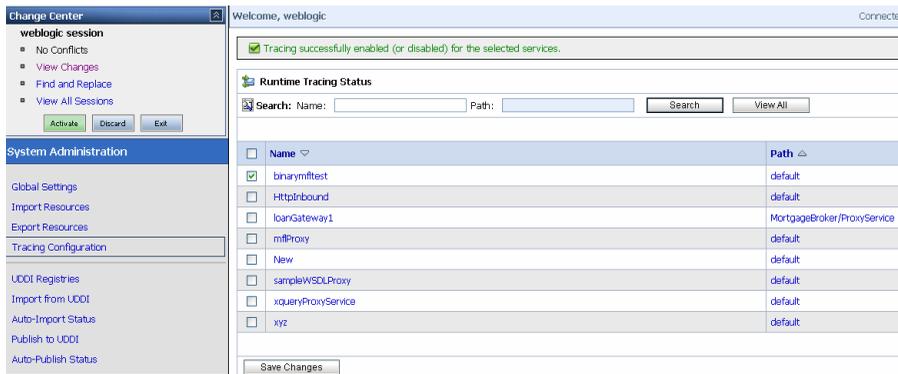
BEA AquaLogic Service Bus supports to trace messages without having to shut down the server. This feature is useful in both a development and production environment. Tracing allows administrators, support engineers, and systems engineers to troubleshoot and diagnose a message flow in one or more proxy services.

For example, if one of your proxy services is failing and you want to find out at which stage the problem exists, you can enable tracing for that proxy service. After tracing is enabled, the system logs various details extracted from the message flow such as stage name, name of the pipeline, and route node name. The entire message context is also printed, including headers and message body. When a fault occurs in the message flow, additional details such as error code and reason are logged. Tracing occurs at the beginning and end of each component in the message flow, which includes stages, pipelines, and nodes. Actions are not traced individually.

To Enable Tracing

You can enable tracing in the System Administration module of the AquaLogic Service Bus Console, as shown in the following figure.

Figure 6-1 Tracing Configuration



As shown in the preceding figure, the Tracing Configuration page displays the tracing status of the proxy services. If the checkbox adjacent to the name of the proxy service is selected, tracing is enabled for that service. The Runtime Tracing Status table displays the following information:

- Name: the name of the proxy service. The name is a link to the View Proxy Service Details page.
- Path: the project name and the name of the folder in which the proxy service resides. It is a link to the Project Details or Folder Details page.

Information about the pages referenced from the Runtime Tracing Status table is available in the *Using the AquaLogic Service Bus Console*, as follows:

- View Proxy Service Details: “Viewing and Changing Proxy Services” in [Proxy Services](#)
- Project Details: “Viewing Project Details” in [Project Explorer](#).
- Folder Details: “Viewing Folder Details” in [Project Explorer](#).

For information on how to use the AquaLogic Service Bus Console to enable tracing, see “Enabling Runtime Tracing Status of Proxy Services” in [System Administration](#) in *Using the AquaLogic Service Bus Console*.

Note: Remember to activate the session to start logging. Once the session is activated, the trace setting is persisted along with the other details of the proxy service configuration.

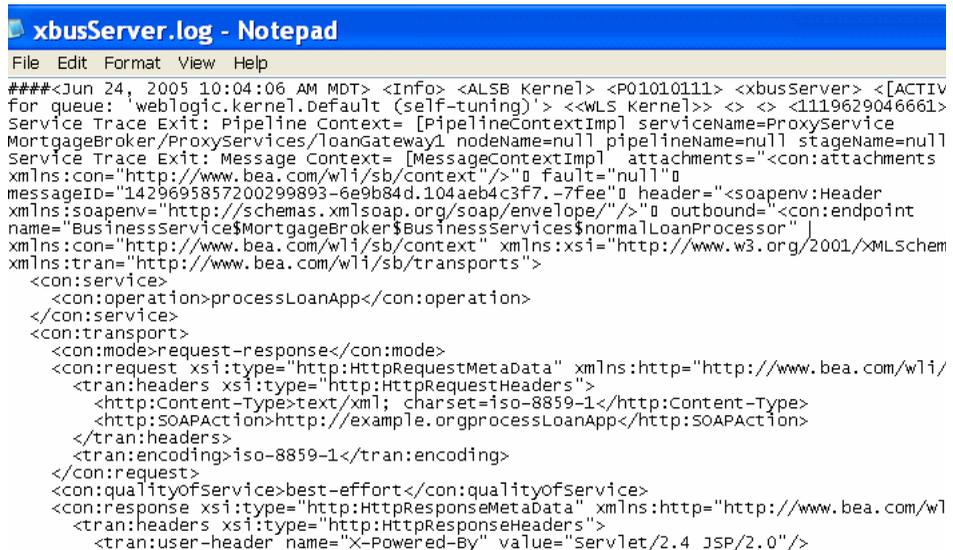
The tracing information is stored in the server directory logs. For example, in the AquaLogic Service Bus Examples, if tracing is enabled for the services before they are tested, the tracing information is logged in the following log file.

```
<BEA_HOME>\weblogic92\samples\domains\servicebus\servers\xbusServer\logs\xbusServer.log
```

where, BEA_HOME is the directory in which you installed your BEA product.

The following figure shows a sample of the tracing log.

Figure 6-2 Tracing Log Example



```
xbusServer.log - Notepad
File Edit Format View Help
####<Jun 24, 2005 10:04:06 AM MDT> <Info> <ALSB Kernel> <P01010111> <xbusServer> <[ACTIV
for queue: 'weblogic.kernel.Default (self-tuning)']> <<WLS Kernel>> <> <> <1119629046661>
Service Trace Exit: Pipeline Context= [PipelineContextImpl serviceName=ProxyService
MortgageBroker/ProxyServices/loanGateway1 nodeName=null pipelineName=null stageName=null
Service Trace Exit: Message Context= [MessageContextImpl attachments="<con:attachments
xmlns:con="http://www.bea.com/wli/sb/context"/>"0 fault="null"0
messageID="1429695857200299893-6e9b84d.104aeb4c3f7.-7fee"0 header="<soapenv:Header
xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/">"0 outbound="<con:endpoint
name="BusinessServiceMortgageBroker$BusinessServices$normalLoanProcessor"
xmlns:con="http://www.bea.com/wli/sb/context" xmlns:xsi="http://www.w3.org/2001/XMLSchema
xmlns:tran="http://www.bea.com/wli/sb/transport">
  <con:service>
    <con:operation>processLoanApp</con:operation>
  </con:service>
  <con:transport>
    <con:mode>request-response</con:mode>
    <con:request xsi:type="http:HttpRequestMetadata" xmlns:http="http://www.bea.com/wli/
      <tran:headers xsi:type="http:HttpRequestHeaders">
        <http:Content-Type>text/xml; charset=iso-8859-1</http:Content-Type>
        <http:SOAPAction>http://example.org/processLoanApp</http:SOAPAction>
      </tran:headers>
      <tran:encoding>iso-8859-1</tran:encoding>
    </con:request>
    <con:qualityofservice>best-effort</con:qualityofservice>
    <con:response xsi:type="http:HttpResponseMetadata" xmlns:http="http://www.bea.com/wl
      <tran:headers xsi:type="http:HttpResponseHeaders">
        <tran:user-header name="X-Powered-By" value="Servlet/2.4_JSP/2.0"/>

```

Note: The tracing pattern in the server log is identical to the tracing in the test console. For more on tracing in the test console, see [Tracing Proxy services—Test Console](#) in *Using the AquaLogic Service Bus Console*

Tracing

UDDI

This section contains the information on the following topics:

- [“Overview of BEA AquaLogic Service Bus and UDDI” on page 7-1](#)
- [“Using AquaLogic Service Bus and UDDI” on page 7-8](#)
- [“Configuring a Registry” on page 7-9](#)
- [“Publishing a Proxy Service to a UDDI Registry” on page 7-10](#)
- [“Using Auto-Publish” on page 7-11](#)
- [“Importing a Service from a Registry” on page 7-11](#)
- [“Using Auto-Import” on page 7-13](#)
- [“Mapping AquaLogic Service Bus Proxy Services to UDDI Entities” on page 7-15](#)
- [“Canonical tModels Supporting AquaLogic Service Bus Services”](#)
- [“Example” on page 7-25](#)

Overview of BEA AquaLogic Service Bus and UDDI

Universal Description, Discovery and Integration (UDDI) registries are used in an enterprise to share Web services. Using UDDI services helps companies organize and catalog these Web services for sharing and reuse in the enterprise or with trusted external partners.

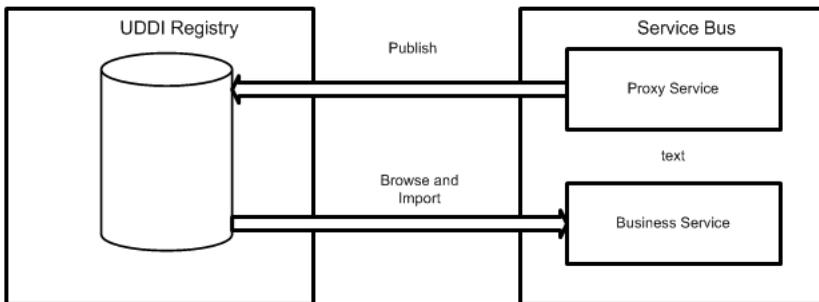
A UDDI registry service for Web services is defined by the UDDI specification available at:

<http://www.oasis-open.org/committees/uddi-spec/doc/tcspecs.htm#uddiv3>

UDDI registries are based on this specification, which provides details on how to publish and locate information about Web services using UDDI. The specification does not define run-time aspects of the services (it is only a directory of the services). UDDI provides a framework in which to classify your business, its services, and the technical details about the services you want to expose.

Publishing a service to a registry requires knowledge of the service type and the data structure representing that service in the registry. A registry entry has certain properties associated with it and these property types are defined when the registry is created. You can publish your service to a registry and make it available for other organizations to *discover* and use. Proxy services developed in BEA AquaLogic Service Bus can be published to a UDDI registry. AquaLogic Service Bus can interact with any UDDI 3.0 compliant registry. BEA provides the AquaLogic Service Registry.

Figure 7-1 AquaLogic Service Bus integration with UDDI



AquaLogic Service Bus' Web-based interface to AquaLogic Service Registry makes the registry accessible and easy to use. In working with UDDI, AquaLogic Service Bus promotes the reuse of standards based Web services. In this way, AquaLogic Service Bus registry entries can be searched for and discovered and used by a multiple domains. Web services and UDDI are built on a set of standards, so reuse promotes the use of acceptable, tested Web services and application development standards across the enterprise. The Web services and interfaces can be catalogued by type, function, or classification so that they can be discovered and managed more easily.

Basic Concepts of the UDDI Specification

UDDI is based upon several established industry standards, including HTTP, XML, XML Schema Definition (XSD), SOAP, and WSDL. The latest version of the UDDI specification is available at:

<http://www.oasis-open.org/committees/uddi-spec/doc/tcspecs.htm#uddiv3>

An UDDI specification describes a registry of Web services and its programmatic interfaces. UDDI itself is a set of Web services. The UDDI specification defines services that support the description and discovery of:

- Businesses, organizations, and other Web services providers
- The Web services they make available
- The technical interfaces that can be used to access and manage those services

Benefits of Using a UDDI Registry with AquaLogic Service Bus

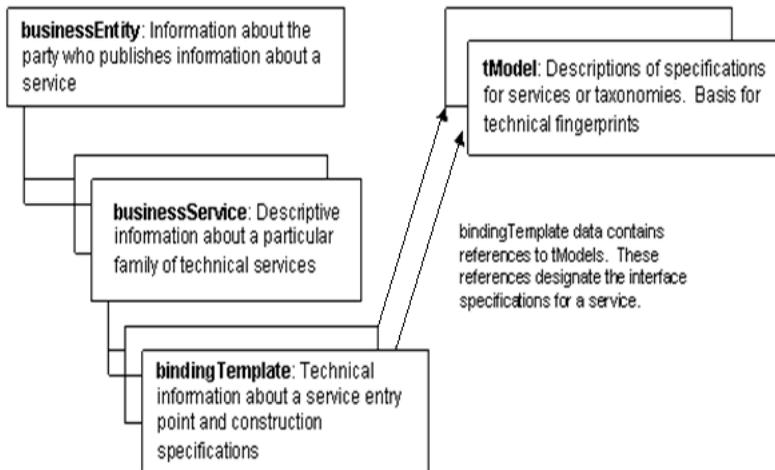
A UDDI registry stores data and metadata about business services. It is a standards-based library of catalogued and managed information about Web services for discovery and reuse by other applications. UDDI offers several benefits to IT managers at both design time and run time, including increasing code reuse. UDDI also provides benefits to developers, including the following:

- UDDI improves infrastructure management by publishing information about proxy services to the registry and categorizes the services for discovery. Thus growing a portfolio of services making it easier to understand and manage relationships among services, component versioning, and dependencies.
- UDDI services can be imported from a registry to configure the parameters required to invoke the Web service and the necessary transport and security protocols.
- UDDI promotes the use of standards-based Web services and business services development in business applications and provides a link to a library of resources for Web services developers. This decreasing the development lifecycle and improves productivity. It also increases the prospect of interoperability between business applications by sharing standards-based resources.
- UDDI provides a user friendly interface for searching and discovering Web services. You can search on criteria specified by you.

Introduction to UDDI Entities

UDDI uses a specific data model to represent entities that define organizations and services. [Figure 7-2](#) shows the relationship between different UDDI entities.

Figure 7-2 UDDI Entities Representing Organizations and Services



A high-level overview of the UDDI entities is provided.

Table 7-1 High-Level Description of UDDI Entities

| | |
|------------------|--|
| Business Entity | An organization or group of people who own and provide the services. It can be described by a set of names, descriptions, contact details for the service provider, a set of categories that represent the business entity’s features, unique identifiers, discovery URLs. |
| Business Service | A business service represents functionality or resources provided by a business entity. It is described by a name, a description, and a set of categories that represent the function of the service. It is not necessarily a Web service. |

Table 7-1 High-Level Description of UDDI Entities

| | |
|------------------|--|
| Binding Template | A binding template represents the technical details of how to invoke a business service. A business service can contain one or more binding templates. It is described by an Access Point representing the service endpoint (the endpoint URI and protocol specification), tModel instance information, and categories to reference specific features of the binding template. |
| tModel | This is the technical model describing how services must be represented in the UDDI registry. The description of a service includes a name, a description, an overview document (a reference to a document specifying the purpose of the tModel), a category, and an identifier (to uniquely identify the tModel). |

For more information on the UDDI data model and entities used in UDDI, see [Introduction to BEA AquaLogic Service Registry](#) in *BEA AquaLogic Service Registry 2.1 User's Guide*.

Prerequisites

Before using AquaLogic Service Bus with a UDDI registry you must perform the following tasks:

- AquaLogic Service Registry must be installed and running. For information on how to install BEA AquaLogic Service registry 2.1, see in [BEA AquaLogic Service Registry Installation Guide](#).
- AquaLogic Service Bus must be installed and running. For information on how to install BEA AquaLogic Service Bus 2.5, see, in [BEA AquaLogic Service Bus Installation Guide](#).

Certification

AquaLogic Service Bus works with any UDDI registry that is fully compliant with the version 3 implementation of UDDI (Universal Description, Discovery and Integration).

AquaLogic Service Registry 2.1 is a version 3 UDDI-compliant registry and is certified to work with AquaLogic Service Bus.

Features

The AquaLogic Service Bus Console provides you with access to any version 3 implementation of UDDI registry once it has been set up to work with AquaLogic Service Bus. The following features are available:

- Configure AquaLogic Service Bus to work with one or more version 3 UDDI compliant registries.
- The import feature allows you to search for specific services in a registry or list all services available. You can search on business entity, service name pattern, or both, and then make your selection from the results list.
- Import selected business services from a registry.
- Publish selected AquaLogic Service Bus proxy services to the registry.

For more information on how to configure and search the registry, import business services to AquaLogic Service Bus, and how to publish proxy services to a UDDI registry, see the following topics in [System Administration](#) in *Using the AquaLogic Service Bus Console*:

- [Configuring a UDDI Registry](#)
- [Importing a Business Service from UDDI Registry](#)
- [Publishing a Proxy Service to a UDDI Registry](#)

What is the BEA AquaLogic Service Registry?

BEA AquaLogic Service Registry is a compliant fully with version 3 implementation of UDDI and is a key component of a Service Oriented Architecture (SOA).

Note: AquaLogic Service Registry is not provided with AquaLogic Service Bus. In order to use AquaLogic Service Registry you have to buy a separate licence from BEA.

A UDDI registry provides a standards-based foundation infrastructure for locating services, invoking services, and managing metadata about services (security, transport or quality of service). Using the Registry Console you can browse and publish registry content. The Registry Console is the primary console for administrators to perform registry management. You can launch the AquaLogic Service Registry console in a Web browser by opening the following URL: `http://hostname:port/uddi/web`, where hostname and port are defined when AquaLogic Service Registry is installed. The default port is 8080. For more information on the management of AquaLogic Service Registry, particularly configuring the registry and managing permissions, approval, and replication, see [BEA AquaLogic Service Registry Administrator's Guide](#).

Sample Business Scenario for AquaLogic Service Bus and UDDI

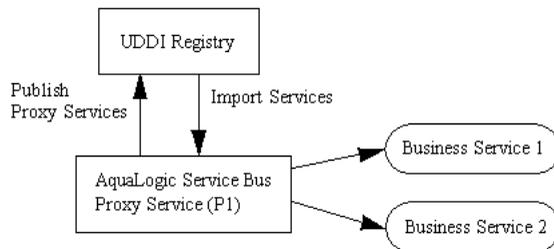
The following are two sample business scenario that highlight the benefit of using UDDI.

Basic Proxy Service Communication with a UDDI Registry

This scenario describes how you can use AquaLogic Service Bus to import services from a registry and then publish the services back to a registry as part of an AquaLogic Service Bus proxy service.

AquaLogic Service Bus imports business services from a UDDI registry. Proxy services are configured to communicate with the business services in the Message Flow. The proxy services themselves can be published back to the registry and made available for use by other domains.

Figure 7-3 Proxy Service Communication with a UDDI Registry



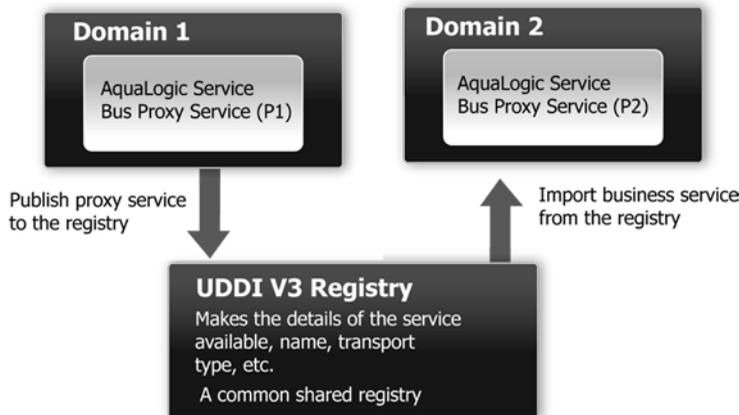
Cross-Domain Deployment in AquaLogic Service Bus

This scenario describes cross-domain deployment using AquaLogic Service Bus. An AquaLogic Service Bus application in one domain requires to access another AquaLogic Service Bus service in another domain at run time.

An instance of AquaLogic Service Bus is deployed in each of two domains. The AquaLogic Service Bus Proxy service (P1) is configured in domain (D1). The AquaLogic Service Bus Proxy service (P2) in domain (D2) requires to access proxy service (P1). As the domains can not communicate directly with each other, P2 in D2 can not discover P1 in D1. The AquaLogic Service Bus import and export feature does not support run-time discovery of services in different domains, but publishing the service to a publicly available UDDI registry allows for the discovery of the service in any domain. Once P1 is made available in the UDDI registry it can be invoked at run time (for example, get a stock quote) and imported as a business services in another AquaLogic Service Bus proxy service.

When importing and exporting from different domains you should have network connectivity. A proxy service might reference schemas located in the repository of a different domain, in which case it needs HTTP access to the domain to import using the URL. In the absence of connectivity an error message will be returned.

Figure 7-4 Sample Business Case of Cross-Domain Deployment



Using AquaLogic Service Bus and UDDI

You can use the AquaLogic Service Bus Console to

- Publish information about any proxy service to a registry, including the following service types: WSDL, messaging, any SOAP, and any XML.
- Search a registry for information about a service and discover the service.
- Configure a registry to allow users to publish services and import services.
- Import Web services and integrate them with your application.

UDDI Workflow

The typical workflow for using UDDI with AquaLogic Service Bus is as follows:

- Install AquaLogic Service Bus. For more information on the installation, see [AquaLogic Service Bus Installation Guide](#).
- Install AquaLogic Service Registry. For information on installation, see [AquaLogic Service Registry Installation Guide](#).
- Configure the registry in the ALSB console. For more information, see “Detaching a Service” in [System Administration](#) in *Using the AquaLogic Service Bus Console*.

- Set the default registry. For more information see [Setting Up Default Registry](#) in *Using the AquaLogic Service Bus Console*.

Configuring a Registry

You can configure a UDDI registry, make it available in AquaLogic Service Bus, and then publish AquaLogic Service Bus proxy services to it or import business services from the registry to be used in a proxy service. In an active AquaLogic Service Bus session in the AquaLogic Service Bus Console to configure the registry.

The following table describes the fields required for configuring a UDDI registry.

Table 7-2 UDDI Registry Configuration Settings

| | |
|----------------------------|---|
| Name* | This is the name of the registry. The name is assigned to it when it is first published. Select the registry name to edit the details for the registry. You can edit the Inquiry URL, Publish URL, Security URL, and the Service Account, but not the name of the registry. This field is required. |
| Inquiry URL* | This is the URL used to locate and import a service. The Inquiry URL is required. To read from a registry, you only need to specify a name and Inquire URL. This field is required. |
| Publish URL* | This is the URL used to publish a service. When publishing a service you must also specify a security URL and specify the service account associated with the registry. This field is required. |
| Security URL | This URL is used to get an authentication token so that you can publish to the registry. You must specify a publish URL and a security URL if you have a service account defined. This field is required. |
| Subscription URL | This URL is used to subscribe to changes from the corresponding service in the registry. You will use this URL to synchronize the service in the AquaLogic Service Bus console with the changes in the corresponding service in the registry using Auto-Import. This field is mandatory |
| User name | You will have to supply the <code>user name/password</code> corresponding to the user name for the registry console. This is required for authentication into the registry console. |
| Password/(Change Password) | You will have to supply the <code>user name/password</code> . This is required for authentication into the registry console. |
| Options | Delete is the only option defined for the registry. |

When publishing services to AquaLogic Service Registry, to gain access to the registry, you must be authenticated for which you should have a valid user name and password. The user name and password combination is implemented as a service account resource in AquaLogic Service Bus. Service accounts must be defined before configuring proxy services so that the authentication criterion is set up to work with a service during the configuration of the proxy service.

You can set up registries with multiple user name and passwords allowing different users to have different permissions based on the service account. Permissions in AquaLogic Service Registry are such that administrators can manage users' privileges in BEA AquaLogic Service Registry and create views into the registry, specific to the needs of the different user types. User permissions set in AquaLogic Service Bus govern access to the registries, their content, and the functionality available to you.

Publishing a Proxy Service to a UDDI Registry

You can use the AquaLogic Service Bus Console to publish proxy services to AquaLogic Service Registry. You should have an account set up in AquaLogic Service Registry to do this. You can publish any proxy service to a UDDI registry except proxy service using the local transport. The service types and transports are listed in [Table 7-3](#).

Table 7-3 Service Types and Transports for a Proxy Service

| Service Type | Transports |
|--------------|---|
| WSDL | HTTP(S), JMS |
| Any SOAP | HTTP(S), JMS |
| Any XML | HTTP(S), JMS, E-mail, File, FTP, Tuxedo |
| Messaging | HTTP(S), JMS, E-mail, File, FTP, Tuxedo |

Note: Messaging services can have different content for requests and responses, or can have no response at all (one-way messages). E-mail, File, and FTP should be one-way.

You can select the Business Entity under which a service is to be published. Business Entity Administration (including creation, removal, update, and deletion of entities) is done using the management console provided by the registry vendor (the Business Service Console in the case of AquaLogic Service Registry). The first time you publish to a registry you must load the `tModels` to that registry. This is done at the time you configure the publishing details in the AquaLogic Service Bus Console.

For more information on how to publish to a UDDI registry, see “Publishing a Proxy Service to a UDDI Registry” in [System Administration](#) in *Using the AquaLogic Service Bus Console*.

AquaLogic Service Bus works with any UDDI version 3 compliant registry but it has been certified to work with AquaLogic Service Registry only.

Using Auto-Publish

When you create a proxy service you can publish it to the default registry automatically. In order to do this you have to first set a default registry. Default is the registry to which the proxy services are published to, when you create or modify them. You can use the checkbox beside **Publish To Registry** in the **Create a Proxy Service-General Configuration** page to enable or disable the Auto-Publish feature for individual proxy services. For more information on setting up a default registry see, [Setting up a Default Registry](#) in *Using the AquaLogic Service Bus Console*.

When you enable the **Publish To Registry** in the **Create a Proxy Service-General Configuration** page the proxy service is published to the default registry. The services are automatically published to the registry when you activate the session, only when the **Publish to Registry** checkbox is selected for the proxy service. If the Registry is unavailable, the publish is retried in the background. Any further changes to the proxy service resets the retry attempts. When a proxy service is republished to UDDI, all taxonomies and categorizations, which are defined in UDDI for the proxy service are preserved.

When you change the default registry all the proxy services that are enabled to auto-publish will be published to the new default registry. Synchronization will now take place with the current default registry. When a proxy service is not synchronized, the AquaLogic Service Bus Console

console displays the  icon beside the proxy service.

Note: When you have a default registry and you import a `sbconfig.jar`, which has a default registry set with the same logical name during the import, it is possible that the default registry will have incorrect value for the business entity. You may now see errors in the **Auto Publish Status** page, if there are any auto-published proxy services. You can correct this by selecting the default registry again.

Importing a Service from a Registry

You can import services from a registry as AquaLogic Service Bus business services. When importing a WSDL-based service, if multiple UDDI binding templates are encountered, AquaLogic Service Bus creates a different business service for each binding template.

Establish access to registries in AquaLogic Service Bus by someone who has AquaLogic Service Bus system administration privileges for the UDDI registries. The registry entries automatically appear in the Import page of the AquaLogic Service Bus Console. When importing, you make a selection from the list of available registries. To discover a service in a registry you must query a specific registry. Entries in registries are unique. This query is performed when you specify what registry you want to use for importing a service.

You can import the following business services types from a UDDI registry into AquaLogic Service Bus:

- WSDL over HTTP binding. When multiple UDDI binding templates are present, a business service is create for each binding template.
- SOAP or XML binding over HTTP, or HTTP(S).
- Services that are categorized as AquaLogic Service Bus services. These are AquaLogic Service Bus proxy services that are published to a UDDI registry. This feature is primarily used in multi-domain AquaLogic Service Bus deployments where proxy services from one domain need to discover and route to proxy services in another domain.

For information on how to use the AquaLogic Service Bus Console to import services from a UDDI registry, see “Importing a Business Service from a UDDI Registry” in [System Administration](#) in Using the AquaLogic Service Bus Console.

When a service is updated, you must re-import the service from the registry to get the most recent version.

Services have documents associated with them and these documents can include a number of other documents (schemas, policies, and so on). On import, the UDDI registry points to the document location based on the inquiry URL of the service. When a document that includes or references other resources is located, all of the referenced information and each included item is added as a separate resource in AquaLogic Service Bus.

Business Entity and pattern are the criteria used to search for a service in a registry. For example, you can enter `foo%`, when searching for a service. Services published by AquaLogic Service Bus have specific `tmodel` keys identifying the services that are used when searching for the service in the registry.

Import automatically tries to connect to a registry when you attempt to get the list of business entities from the registry. The Business Entity is the highest level of organization in the registry, though you can use other search criteria, such as business, application type, and so on. If you require authentication, then you need a user name and password which you must get from your systems Administrator.

Related References

- Technical Notes can be found at <http://www.oasis-open.org/committees/uddi-spec/doc/tns.htm>. The note on *s Using WSDL in a UDDI Registry* is important.
- UDDI product and development tool information is available at the OASIS UDDI Solutions page at <http://uddi.org/solutions.html>.
- The UDDI specifications The specification defines the following:
 - SOAP APIs that applications use to query and to publish information to a UDDI registry
 - XML Schema schemata of the registry data model and the SOAP message formats
 - WSDL definitions of the SOAP APIs
 - UDDI registry definitions (`tModels`) of various identifier and category systems that may be used to identify and categorize UDDI registrations

Using Auto-Import

You can use the Auto-Import feature to synchronize the business services, which are imported from the AquaLogic Service Registry, with the corresponding services in the registry. For more information on Using Auto-Import see, [Auto-Import](#) in Using the AquaLogic Service Bus Console. You can use Auto-Import to do the following:

- “Synchronize” on page 7-13
- “Detach” on page 7-14

Synchronize

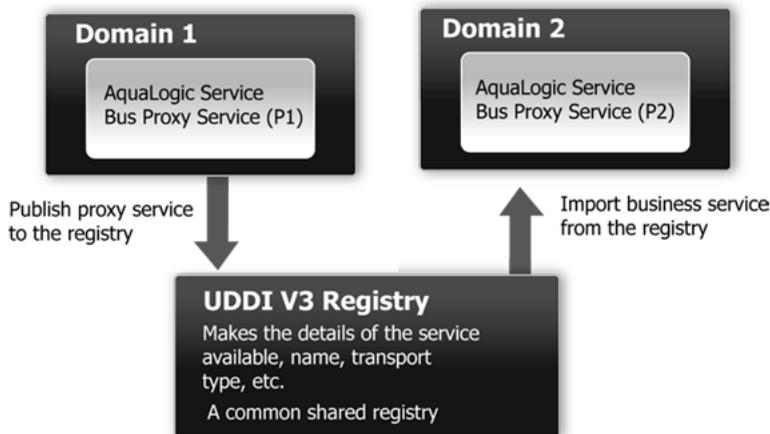
You can synchronize the services you have imported from the registry. If the services in the registry change, you can synchronize services in the AquaLogic Service Bus Console with those in the registry. The following use case illustrates the process of synchronization. If the business service is not detached from the registry, AquaLogic Service Bus automatically subscribes to any

changes to the service in the registry. If the service changes, the  icon in the resource browser and project explorer indicates the service needs to be synchronized. In addition, the **Auto Import**

Status page shows this service and provides the options to synchronize the service or detach it from the registry. Under certain circumstances, synchronizing the service might result in semantic validation errors that shows up in the view conflicts page. These will have to be fixed manually fixed before the session is activated.

When a service is synchronized, the service is updated only with fields that are obtained from UDDI. Other fields in the service definition will preserve their values if modified since last import.

Figure 7-5 Sample Business Case of Cross-Domain Deployment



Consider a scenario where you publish services from Domain1(see [Figure 7-5](#)) to a registry. You then import these services to another domain, Domain2. When you make changes to the corresponding service in Domain1 and update it in the registry. You can update the services in Domain2 by synchronizing it with the registry using Auto-Import.

Detach

When you do not want the service in the AquaLogic Service Bus Console synchronized with the corresponding service in the registry then, you can avoid synchronization by detaching it from the registry. For more information on using Detach, see “Detaching a Service” in [System Administration](#) in *Using the AquaLogic Service Bus Console*.

Mapping AquaLogic Service Bus Proxy Services to UDDI Entities

AquaLogic Service Bus proxy service attributes must be mapped to the data model supported by the UDDI registry to allow a proxy service to be published as a UDDI business entity. The following table shows the service types, message types, and transports relevant to the UDDI registry mapping for an AquaLogic Service Bus proxy service.

Table 7-4 Proxy Service Attributes and Service Types

| Service Type | Message Content Type | Transports |
|--------------|---------------------------------|---|
| WSDL | SOAP or XML (with attachment) | HTTP(S), JMS |
| Any SOAP | Untyped SOAP (with attachment) | HTTP(S), JMS |
| Any XML | Untyped XML (with attachment) | HTTP(S), JMS, E-mail, File, FTP, and Tuxedo |
| Messaging | Binary, Text, MFL, XML (schema) | HTTP(S), JMS, E-mail, File, FTP, and Tuxedo |

Note: Optional parts are listed in parentheses. Messaging services can have different content for requests and responses, or can have no response at all (one-way messages). E-mail, File, and FTP should be one-way.

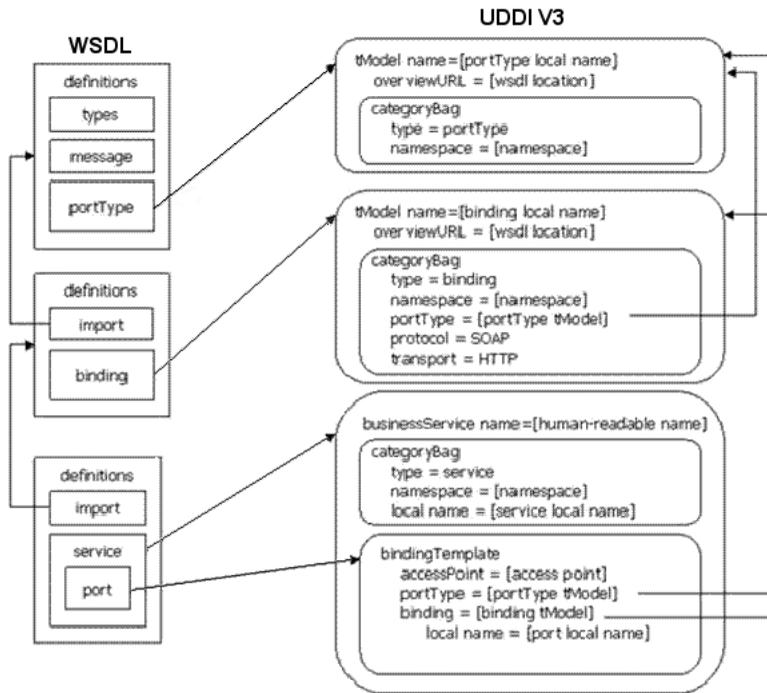
Proxy services have attributes in common and also attributes that are specifically defined by the transport protocols used by the service and the type of service. Each proxy service can deliver messages of a certain type.

The primary relevant entities in UDDI are:

- `businessService`: this represents the service as a whole and contains high-level general information about the service.
- `bindingTemplate`: this contains information for accessing the service.
- `tModels`: tModels are used to supply the individual attributes for categorizing and defining the service.

Figure 7-6 shows how WSDL-based services are mapped to UDDI business entities.

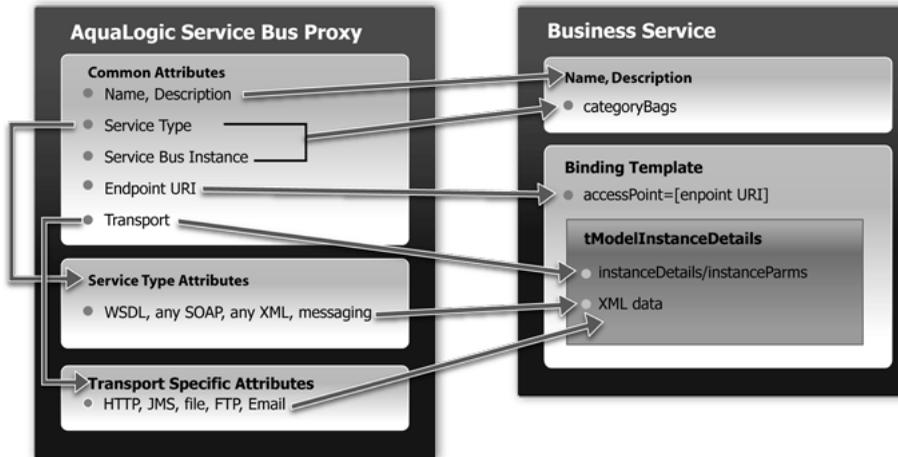
Figure 7-6 WSDL Service to UDDI Mapping



The technical note on *Using WSDL in a UDDI registry, version 2.0.2*, at <http://www.oasis-open.org/committees/uddi-spec/doc/tns.htm>, is used as the basis for publishing WSDL-based proxy services to the UDDI registry. This document is also used as a reference point for publishing non-WSDL based services. The document and the base UDDI specification describe the canonical technical models (tModels) used to describe UDDI entities. To publish AquaLogic Service Bus proxy services as entities in the UDDI registry, you must add additional canonical tModels to support some of the constructs specific AquaLogic Service Bus. Not all attributes of an AquaLogic Service Bus proxy service are useful when searching for a service, for example service type and transport details. These attributes do not categorize the service. tmodels are configuration details of the service once it has been discovered. These configuration details are mapped to the business service binding template tmodelinstanceDetails section. Other attributes specifically identify a service and can be used as the search criteria for the service. These attributes are mapped using keyed references to tModels with values in the categoryBag of the binding template.

An example of how AquaLogic Service Bus maps to UDDI is shown in [Figure 7-7](#).

Figure 7-7 AquaLogic Service Bus to UDDI Mapping



UDDI Mapping Details for an AquaLogic Service Bus Proxy Service

AquaLogic Service Bus high-level proxy service information maps into the Business Service as follows:

- Name and Description map to `businessService` elements.
- There is a Special keyed Reference Group for AquaLogic Service Bus properties. An example of a key is `uddi:bea.com:attributes:aqualogicservicebus`.
- AquaLogic Service Bus type (WSDL, SOAP, XML, and Mixed) and Instance are mapped to `keyedReferences` in the service category. An example of a key is `uddi:bea.com:serVICetype`.
- An AquaLogic Service Bus Instance maps to a `keyedReference` in the AquaLogic Service Bus `keyedReferenceGroup` (Name = "AquaLogicServiceBus", Values = URL of the AquaLogic Service Bus instance).

This instance serves two purposes:

- To indicate that this service is in fact hosted by an AquaLogic Service Bus server.
- To contain the URL of the AquaLogic Service Bus instance.

[Listing 7-1](#) shows a mapping of high-level proxy service information to a business service.

Listing 7-1 Sample Proxy Service to Business Service Mapping

```
<keyedReferenceGroup tModelKey="uddi:bea.com:servicebus:properties">
  <keyedReference tModelKey="uddi:bea.com:servicebus:servicetype"
    keyName="Service Type"
    keyValue="SOAP"/>
  <keyedReference tModelKey="uddi:bea.com:servicebus:instance"
    keyName="Service Bus Instance"
    keyValue="http://FO002.amer.bea.com:7001"/>
</keyedReferenceGroup>
```

Note: The key for the businessService created when a proxy service is published is a publisher assigned key name. It is derived from the AquaLogic Service Bus domain name, the path of the proxy service, and the proxy service name. It takes the following form:

```
uddi:bea.com:servicebus:<domainname>:<path>:<servicename>.
```

For example, AnonESBan, which is a domain in AquaLogic Service Bus, contains a project named Proxy, which contains a folder named Accounting, which in turn contains a proxy service called PayoutProxy. When PayoutProxy is published to UDDI, its businessService is created with the following key:

```
uddi:bea.com:servicebus:AnonESB:Proxies:Accounting:PayoutProxy.
```

AquaLogic Service Bus detailed proxy service information maps into the binding template as follows:

- The Endpoint URI maps to the access point.
- The Marker tModel for each transport maps to tModelInstanceDetails.
 - Transport tModels for HTTP, JMS, File, FTP, E-mail. New tModels are packaged with AquaLogic Service Bus to support JMS and file transports.
 - Detailed AquaLogic Service Bus configuration information maps to instanceParms.
- The Market tModel for each service type maps to the tModelInstanceDetails. This includes the following:

- Protocol `tModels` for WSDL, any SOAP, any XML, Messaging. New `tModels` are packaged with AquaLogic Service Bus to support anySOAP, anyXML, and Messaging.
- WSDL maps via WSDL to UDDI technology note.
- Messaging has detailed configuration information that maps to `InstanceParms`.

[Listing 7-2](#) shows a detailed information mapping to the binding template.

Listing 7-2 Sample Detailed Mapping to the Binding Template

```
<bindingTemplate bindingKey="uddi:" serviceKey="uddi:">
  <accessPoint useType="endPoint">file:///c:/temp/in3</accessPoint>
  <tModelInstanceDetails>
    <tModelInstanceInfo tModelKey="uddi:uddi.org:transport:file">
      <InstanceDetails>
        <InstanceParms><ALSBInstanceParms xmlns="http://www.bea.com/wli/sb/uddi">
          <property name="fileMask" value="*.*"/>
          <property name="sortByArrival" value="false"/> </ALSBInstanceParms>
        </InstanceParms>
      </InstanceDetails>
    </tModelInstanceInfo>
    <tModelInstanceInfo tModelKey="uddi:bea.com:servicebus:protocol:
      messagingservice">
      <InstanceDetails>
        <InstanceParms><ALSBInstanceParms xmlns="http://www.bea.com/wli/sb/uddi">
          <property name="requestType" value="XML"/>
          <property name="RequestSchema" value="http://domain.com:7001
            /sbresource?SCHEMA%2FDJS%2FOAGProcessPO"/>
          <property name="RequestSchemaElement"
            value="PROCESS_PO"/>
        </InstanceParms>
      </InstanceDetails>
    </tModelInstanceInfo>
  </tModelInstanceDetails>
</bindingTemplate>
```

```

        <property name="responseType" value="None"/></ALSBIInstanceParms>
    </InstanceParms>
</InstanceDetails>
</tModelInstanceInfo>
</tModelInstanceDetails>
</bindingTemplate>

```

Transport Attributes

Each of the transport types in the `uddi:uddi.org:transport: *` group has a different set of detailed metadata. (See [Table 7-4, “Proxy Service Attributes and Service Types,”](#) on page 7-15.) This metadata provides configuration details of the transport for the proxy service. It is neither useful for characterizing the service nor useful in querying the service. However, after the service has been discovered, this data is needed to access the service. The metadata is represented by an XML string and is located in the `instanceParms` field in `tModelInstanceInfo`.

If you are mapping a proxy service that uses the HTTP transport, and as part of the HTTP configuration you need to describe some detailed configuration details, including the required client authorization and the request and response character encoding, the following [Listing 7-3](#) provides an example of what must appear in the bindingTemplate `tModelInstanceDetails`.

Listing 7-3 Example of tModelInstanceDetails

```

<tModelInstanceDetails>
  <tModelInstanceInfo tModelKey="uddi:uddi.org:transport:http">
    <instanceDetails>
      <instanceParms>
        <ALSBIInstanceParms xmlns="http://www.bea.com/wli/sb/uddi">
          <property name="basic-auth-required" value="true"/>
          <property name="request-encoding" value="iso-8859-1"/>
          <property name="response-encoding" value="utf-8"/>
          <property name="Scheme" value="http"/>
        </ALSBIInstanceParms>
      </instanceParms>
    </instanceDetails>
  </tModelInstanceInfo>
</tModelInstanceDetails>

```

Note: For each transport, the service endpoint is always stored in the bindingTemplate's `accessPoint` field.

Table 7-5 is organized by transport type and lists the `tModelKey` and `instanceParms` used by each of the transports.

Table 7-5 Transport Attributes

| Transport | tModelKey | InstanceParms |
|---------------------|--|--|
| HTTP | <code>uddi:uddi.org:transport:http</code> | <ul style="list-style-type: none"> • Client Authentication [None, Basic, Client Cert (HTTP(S) only)] • Request encoding • Response encoding |
| JMS | <code>uddi:uddi.org:transport:jms</code> | <ul style="list-style-type: none"> • Destination Type [Queue, Topic] • Response required, Response URI • Response Message Type [Bytes, Text] • Request encoding • Response encoding |
| File | <code>uddi:uddi.org:transport:file</code> | <ul style="list-style-type: none"> • File Mask • Sort by Arrival [Boolean] • Request Encoding |
| FTP | <code>uddi:uddi.org:transport:ftp</code> | <ul style="list-style-type: none"> • File Mask • Sort by Arrival [Boolean] • Transfer Mode [Text, Binary] • Request Encoding |
| E-mail ¹ | <code>uddi:uddi.org:transport:smtp</code> | <ul style="list-style-type: none"> • Attachment supported [Boolean] • Request Encoding |
| Tuxedo | <code>uddi:bea.org:transport:tuxed</code> ○ | <ul style="list-style-type: none"> • Response required • Access point ID • Buffer type • Buffer subtype • Classes jar • Field table classes • View classes |

1. The `accessPoint` in the Binding Template for an E-mail Transport uses the standard `mailto` URL format:

```
mailto:name@some_server.com
```

This is different from the one configured for the proxy service in AquaLogic Service Bus, which is a URL oriented toward reading e-mail. It is not possible to derive this `mailto` URL from the proxy service definition as the server name is not known. For example, if the proxy service is defined to read from a POP3 server, it might be defined with a URL such as `mailfrom:pop3.bea.com`. When publishing such a proxy service, a dummy server is added. In the above example, the published URL will take the form `mailto:some_name@some_server.com`.

Service Type Attributes

Table 7-6 provides a high-level description of each of the service types.

Table 7-6 Service Type Attributes

| Service | Description |
|----------|---|
| WSDL | WSDL based proxies map to UDDI based on the <i>Using WSDL in a UDDI Registry, version 2.0.2</i> technical note at URL: http://www.oasis-open.org/committees/uddi-spec/doc/tn/uddi-spec-tc-tn-wsdl-v202-20040631.htm . |
| Any SOAP | A simple marker protocol in the <code>tModel</code> in the <code>bindingTemplate</code> 's <code>tModelInstanceDetails</code> , as well as in the <code>categoryBag</code> , defines the Any Soap attributes. |

Table 7-6 Service Type Attributes

| Service | Description |
|--------------------|--|
| Any XML | A simple marker protocol <code>tModel</code> within the <code>bindingTemplate</code> 's <code>tModelInstanceDetails</code> , as well as in the <code>categoryBag</code> defines the Any XML attributes. This is a new detailed <code>tModel</code> . |
| Messaging Services | A simple marker protocol <code>tModel</code> in the <code>bindingTemplate</code> 's <code>tModelInstanceDetails</code> , defines the messaging services attributes. This is a new detailed <code>tModel</code> . Unlike the other service types, messaging services have additional configuration information associated with them, which provide details about the request and response messages. The configuration details are represented as XML data in the <code>InstanceParms</code> data for the following <code>tModel</code> reference in the <code>tModelInstanceInfo</code> : <ul style="list-style-type: none"> • Input message format (XML, Text, Binary, MFL) • URL of input message Schema in AquaLogic Service Bus (optional, if input message is XML) • URL of input message MFL in AquaLogic Service Bus (if input message is MFL) • Output message format (none, XML, Text, Binary, MFL) • URL of output message Schema in AquaLogic Service Bus (optional, if output message is XML) • URL of output message MFL in AquaLogic Service Bus (if output message is MFL) |

Canonical tModels Supporting AquaLogic Service Bus Services

The AquaLogic Service Bus-UDDI mapping introduces a number of new canonical `tModels` that are used to represent AquaLogic Service Bus metadata and relationships. These `tModels` must be registered in the UDDI registry to support this mapping. You can create these `tModels` in AquaLogic Service Registry under the administrator ID.

The following table provides a summary of the new `tModels`.

Table 7-7 AquaLogic Service Bus tModels

| Name | Value | Description |
|---|--------------------------------------|--|
| CategorizationGroup tModel Types | | |
| bea-com:servicebus:properties | | Describes very specific attributes of an AquaLogic Service Bus service. In the data model it is used in the business service categoryBag. |
| Categorization tModel Types | | |
| bea-com:servicebus:serviceType | WSDL, SOAP, XML, Messaging Service | Describes the service type of the AquaLogic Service Bus service. |
| bea-com:servicebus:instance | URL of AquaLogic Service Bus Console | Describes the service instance in AquaLogic Service Bus responsible for publishing the service to UDDI. |
| Transport tModel Types | | |
| uddi-org:jms | | Describes the type of transport used by the service. A reference to it is found in the accessPoint attribute of the business service binding template. |
| uddi-org:file | | Describes the type of transport used to invoke the service. A reference to it is found in the accessPoint attribute of the business service binding template. |
| Protocol tModel Types | | |
| bea-com:servicebus:anySoap | | Describes the type of protocol used to access the service. It designates services that have a SOAP message but not defined by a WSDL or schema. The message body content is determined dynamically by the application. |

Table 7-7 AquaLogic Service Bus tModels

| Name | Value | Description |
|---|-------|--|
| bea-com:servicebus:anyXML | | Describes the type of protocol used to access the service. It designates services having an XML message but not defined by a WSDL or schema. The message body content is determined dynamically by the application. |
| bea-com:servicebus:messaging Service | | Describes the type of protocol used to access the service. It designates services where the request message can be any XML (with or without schema), text, binary, or MFL and whose response message can be any of the above or none. The message body content is determined dynamically by the application. |

Example

The following is an example of the mapping for a Messaging Service, configured with JMS transport, the request being XML with a Schema and the response being a text message.

Listing 7-4 Sample Messaging Service Mapping

```
<businessService
  serviceKey="uddi:bea.com:servicebus:Domain:Project:JMSMessaging"
  businessKey="uddi:9cb77770-57fe-11da-9fac-6cc880409fac"
  xmlns="urn:uddi-org:api_v3">
  <name>JMSMessagingProxy</name>
  <bindingTemplates>
  <bindingTemplate
    bindingKey="uddi:4c401620-5ac0-11da-9faf-6cc880409fac"
    serviceKey="uddi:bea.com:servicebus:
      Domain:Project:JMSMessaging">
  <accessPoint useType="endPoint">
    jms://server.com:7001/weblogic.jms.XAConnectionFactory/
      ReqQueue
  </accessPoint>
  <tModelInstanceDetails>
    <tModelInstanceInfo tModelKey="uddi:uddi.org:transport:jms">
      <instanceDetails>
        <instanceParms>
```

```

    <ALSBIInstanceParms
      xmlns="http://www.bea.com/wli/sb/uddi">
      <property name="is-queue" value="true"/>
      <property name="request-encoding"
        value="iso-8859-1"/>
      <property name="response-encoding"
        value="utf-8"/>
      <property name="response-required"
        value="true"/>
      <property name="response-URI"
        value="jms://server.com:7001/
          .jms.XAConnectionFactory/
          RespQueue"/>
      <property name="response-message-type"
        value="Text"/>
      <property name="Scheme" value="jms"/>
    </ALSBIInstanceParms>
  </instanceParms>
</instanceDetails>
</tModelInstanceInfo>
<tModelInstanceInfo
  tModelKey="uddi:bea.com:servicebus:
    protocol:messaging-service">
  <instanceDetails>
    <instanceParms>
      <ALSBIInstanceParms xmlns=
        "http://www.bea.com/wli/sb/uddi">
        <property name="requestType" value="XML"/>
        <property name="RequestSchema"
          value="http://server.com:7001/
            sbresource?SCHEMA%2FDJS%2FOAGProcessPO"/>
        <property name="RequestSchemaElement"
          value="PROCESS_PO_007"/>
        <property name="responseType" value="Text"/>
      </ALSBIInstanceParms>
    </instanceParms>
  </instanceDetails>
</tModelInstanceInfo>
</tModelInstanceDetails>
</bindingTemplate>
</bindingTemplates>
<categoryBag>
<keyedReferenceGroup tModelKey="uddi:bea.com:servicebus:properties">
  <keyedReference tModelKey="uddi:bea.com:servicebus:serVICetype"
    keyName="Service Type"
    keyValue="Mixed" />
  <keyedReference tModelKey="uddi:bea.com:servicebus:instance"
    keyName="Service Bus Instance"
    keyValue="http://cyberfish.bea.com:7001" />

```

```
</keyedReferenceGroup>  
</categoryBag>  
</businessService>
```

UDDI

Transports

You can configure BEA AquaLogic Service Bus proxy services and business services to use different transport protocols. The transport protocol used depends on the type of service, the type of authentication required, the type of the invoking service and so on. This section describes the of transport protocols supported by AquaLogic Service Bus. They include:

- [“E-mail” on page 8-1](#)
- [“EJB” on page 8-4](#)
- [“File” on page 8-4](#)
- [“FTP” on page 8-6](#)
- [“HTTP” on page 8-9](#)
- [“HTTP\(S\)” on page 8-11](#)
- [“JMS” on page 8-13](#)
- [“Local” on page 8-18](#)
- [“Tuxedo” on page 8-18](#)

E-mail

You can choose the e-mail transport protocol when you configure a Messaging Type or Any XML Service type of proxy service or business service. The following sections describe

- [“Configuring Proxy Services using E-mail Transport Protocol” on page 8-2](#)
- [“Configuring Business Services using E-mail Transport Protocol” on page 8-3](#)

Configuring Proxy Services using E-mail Transport Protocol

When you configure a proxy service using the e-mail transport protocol, you must specify an endpoint URI in the following format:

```
mailfrom:<mailserver-host:port>
```

where

- `mailserver-host`: is the name of the host mail server
- `port`: is the port used by the mailserver host

You can configure the following parameters for an e-mail transport proxy service:

- **Service Account:** This is a mandatory parameter. This is the service account resource. The service account consists of a `user name/password` combination required to access the e-mail account.
- **Polling Interval:** This is a mandatory parameter. This parameter specifies the interval in milli-seconds. The default value is `60 ms`.
- **E-mail protocol:** This is a mandatory parameter. There are two types of protocol you can choose from, namely `imap` and `pop3`. The default protocol is set to `pop3`.
- **Read Limit:** This is a mandatory parameter. This specifies the number of files to be read in each poll. The default value is `10`.
- **Pass By Reference:** If this parameter is enabled, the file is staged in the archive directory and passed as a reference in the message headers.
- **Post Read Action:** This is a mandatory parameter. This specifies whether the files should be deleted, moved, or archived after being read by the service. By default the files are to be deleted after reading.
- **Attachments:** This is a mandatory parameter. This parameter specifies if the attachments are to be archived or ignored. By default this parameter is set to `ignore`.
- **IMAP Move Folder:** This is the destination of the messages if the `Post Read Action` is set to `move`.

Note: You must configure this field only if `Post Read Action` is set to `move`.

- **Download Directory:** This is a mandatory parameter. It specifies the file system directory path to download the message.
- **Archive Directory:** This is a mandatory parameter. A file URI that points to the directory where the files are archived. This field is active only when `Post Read Action` parameter is set to `archive`.
- **Error Directory:** This is a mandatory parameter. This URI that points to a directory, in which the contents of the file will be stored in case of a error.
- **Request Encoding:** This is an optional parameter. This parameter specifies the type of encoding to read the request message. The default encoding is `iso-8859-1`.

For more information on how to configure e-mail services, see [Adding a Proxy Service: Transport Configuration](#) in *Using the AquaLogic Service Bus Console*.

Configuring Business Services using E-mail Transport Protocol

When you configure a business service using e-mail transport protocol, you must specify the endpoint URI in the following format:

```
mailto:<name@domain_name.com>
```

where `<name@domain_name.com>` is the e-mail destination.

You can configure the following parameters for an e-mail transport business service:

- **SMTP Server:** You must choose the SMTP Server from the drop-down list.

Note: You must first create an SMTP Server resource before you can choose it from the drop-down list.
- **Mail Session:** This parameter is optional. It is the JNDI name of the configured mail session. You can choose mail sessions from the drop-down list.

Notes:

- You must first configure mail sessions in the WebLogic Server Console. For more information on configuring a mail session, see [Create a Mail Session](#) in WebLogic Server Administration Console.
- Also you should set either the `Mail Session` parameter or the `SMTP Server` parameter.

- **From Name:** This is an optional parameter. This parameter specifies the name from which the reply should be sent.
- **From Address:** This is an optional parameter. This parameter specifies the e-mail address from which the e-mail message should be sent.
- **Reply To Name:** This is an optional parameter. This parameter specifies the name to which the reply should be sent.
- **Reply To Address:** This is an optional parameter. This parameter specifies the e-mail address, to which the reply should be sent.
- **Connection Timeout:** This is an optional parameter. You can use this parameter to specify time in milli-seconds after which the connection to the SMTP server times out.
- **Request Encoding:** This is an optional parameter. This parameter specifies the type of encoding to read the request message. The default encoding is `iso-8859-1`.

For more information on how to configure this transport, see [Adding a Business Service: Transport Configuration](#) in *Using the AquaLogic Service Bus Console*.

EJB

An EJB business service is a `Transport Typed Service`, that is, the type of the transport is determined by the configuration of the service. The EJB transport supports native Remote Method Invocation (RMI) of Stateless Session Beans deployed on WebLogic Server. The EJB transport can also be leveraged to directly expose an EJB as a Web service through AquaLogic Service Bus. For information about the EJB transport, see [Chapter 9, “EJB Transport.”](#)

File

You can choose file transport protocol when you configure a Messaging Type or Any XML Service type of proxy service and the endpoint URI is of the form:

```
file:///<root-dir/dir1>
```

where `root-dir/dir1` is the absolute path to the destination directory.

The following sections describe:

- [“Configuring Proxy Services using File Transport Protocol”](#) on page 8-5
- [“Configuring Business Services using File Transport Protocol”](#) on page 8-6

Configuring Proxy Services using File Transport Protocol

To configure the File transport for a proxy service you must specify the following fields:

- **File Mask:** This is an optional parameter. This specifies the files that should be polled by the proxy service. If the URI is a directory and `*.*` is specified, then the service will poll for all the files in the directory.
- **Polling Interval:** This is a mandatory parameter. This specifies the value for the polling interval in milli-seconds. The default value is `60 ms`.
- **Read Limit:** This is a mandatory parameter. This specifies the number of files to be read in each poll. The default value is `10`.

Note: If `'0'` is specified, all the files are read.

- **Sort By Arrival:** This is an optional parameter. This parameter indicates the sequence of events raised in the order of the arrival of files. The default value for this parameter is `False`.
- **Scan Subdirectories:** This is optional. If enabled, the sub-directories are also scanned.
- **Pass By Reference:** If this parameter is enabled, the file is staged in the archive directory and passed as a reference in the headers.
- **Post Read Action:** This parameter is mandatory. This specifies whether the files should be deleted or archived after being read by the service. By default the files are to be deleted after reading.
- **Stage Directory:** This is a mandatory parameter. This file URI that points to the staging directory.
- **Archive Directory:** This is a mandatory parameter. This file URI that points to the directory where the files are archived. This field is active only when `Post Read Action` parameter is set to `archive`.
- **Error Directory:** This is a mandatory parameter. This URI that points to a directory, in which the contents of the file will be stored in case of a error.
- **Request Encoding:** This is an optional parameter. This parameter specifies the type of encoding to read the request message. The default encoding is `utf-8`.

For more information on how to configure this transport, see [Adding a Proxy Service: Transport Configuration](#) in *Using the AquaLogic Service Bus Console*.

Configuring Business Services using File Transport Protocol

When you configure a business service using e-mail transport protocol you must specify the endpoint URI in the following format:

```
file:///<root-dir/dir1>
```

where `root-dir/dir1` is the absolute path to the destination directory.

When you use this type of transport to configure a business service you must configure the following fields:

- **Prefix:** This is an optional parameter. This parameter specifies the prefix to be attached to the filename.
- **Suffix:** This is an optional parameter. This parameter specifies the suffix to be attached to the filename.
- **Request Encoding:** This is an optional parameter. This specifies the type of encoding to read the message. The default encoding which will be used is `utf-8`.

For more information on how to configure this transport, see [Adding a Business Service: Transport Configuration](#) in *Using the AquaLogic Service Bus Console*.

FTP

You can choose FTP transport protocol when you configure a Messaging Type or Any XML Service type of proxy service and the endpoint URI is of the form:

```
ftp://<hostname:port/directory>
```

where

- `hostname:` is the name of the host on which the destination directory is stored.
- `port:` is the port number at which the `ftp` connection is made.
- `directory:` is the destination directory.

The following sections describe

- [“Configuring Proxy Services using FTP Transport Protocol”](#) on page 8-7
- [“Configuring Business Services using FTP Transport Protocol”](#) on page 8-8

Configuring Proxy Services using FTP Transport Protocol

To configure the FTP transport for a proxy service you must specify the following fields:

- **User Authentication:** You must choose one of the following types of `User Authentication`:
 - **anonymous:** If you choose this type of `User authentication`, you do not require any login credentials to login to the ftp server. But you optionally supply your e-mail ID for identification.
 - **external user:** If you choose this type of `User authentication`, you have reference a `Service Account` resource, which contains your `user name/password` for the ftp server.
- **Pass By Reference:** This is an optional parameter. If this parameter is enabled, the file is staged in the archive directory and passed as a reference in the headers.
- **Remote Streaming:** This is an optional parameter. Setting this parameter to `True` will poll ftp files directly from the remote server at processing time.
- **File Mask:** This is a mandatory parameter. This specifies the files that should be polled by the proxy service. If the URI is a directory and `*.*` is specified, then the service will poll for all the files in the directory.
- **Polling Interval:** This is a mandatory parameter. This specifies the value for the polling interval in milli-seconds. The default value is `60 ms`.
- **Read Limit:** This is a mandatory parameter. This specifies the value for the polling interval in milli-seconds. The default value is `60 ms`.
- **Post Read Action:** This is a mandatory parameter. This specifies whether the files should be deleted or archived after being read by the service. By default the files are to be deleted after reading.
- **Transfer Mode:** This parameter specifies whether the mode of file transfer is `binary` or `ascii`. By default it is an `binary` transfer
- **Stage Directory:** This is a mandatory parameter. This file URI that points to the staging directory.
- **Archive Directory:** This is a mandatory parameter. This file URI that points to the directory where the files are archived. This field is active only when `Post Read Action` parameter is set to `archive`.

- **Error Directory:** This is a mandatory parameter. This URI that points to a directory location, where the contents of the file will be stored in case of a error.
- **Request Encoding:** This is an optional parameter. This parameter specifies the type of encoding to read the request message. The default encoding is `utf-8`.
- **Advanced Settings:** To configure the advance settings click on  icon on the right hand side to expand the advanced settings section. Configuring parameters in this section is optional.
 - **Scan Subdirectories:** This is optional. If enabled the sub-directories are also scanned.
 - **Sort By Arrival:** This is an optional parameter. This parameter indicates the sequence of the events being raised in the order of the arrival of files. This default value for this parameter is `False`.
 - **Timeout:** This is an optional parameter. This parameter specifies the ftp timeout interval, in seconds, before the connection is dropped. The default value for this parameter is `0`.
 - **Retry:** This is an optional parameter. This parameter specifies the maximum number of retries for the ftp connection failures.

For more information on how to configure this transport, see [Adding a Proxy Service: Transport Configuration](#) in *Using the AquaLogic Service Bus Console*.

Configuring Business Services using FTP Transport Protocol

You can choose the FTP transport protocol when you configure a Messaging Type or Any XML Service type of business service and the endpoint URI is of the form:

```
ftp://<hostname:port/directory>
```

where

- `hostname`: is the name of the host on which the destination directory is stored.
- `port`: is the port number at which the `ftp` connection is made.
- `directory`: is the destination directory.

To configure the FTP transport for a business service you must specify the following fields:

- **User Authentication:** You must choose one of the following types of `User Authentication`:
 - **anonymous:** If you choose this type of `User authentication`, you do not require any login credentials to login to the ftp server. But you optionally supply your e-mail ID for identification.
 - **external user:** If you choose this type of `User authentication`, you have reference a `Service Account` resource, which contains your `user name/password` for the ftp server.
- **Prefix for destination filename:** This is a mandatory parameter. This parameter specifies the prefix to be attached to the filename.
- **Suffix for destination filename:** This is a mandatory parameter. This parameter specifies the suffix to be attached to the filename.
- **Request Encoding:** This is an optional parameter. This parameter specifies the encoding for the request message.

For more information on how to configure this transport, see [Adding a Business Service: Transport Configuration](#) in *Using the AquaLogic Service Bus Console*.

HTTP

The following sections describe:

- [“Configuring Proxy Services using HTTP Transport Protocol”](#) on page 8-9
- [“Configuring Business Services using HTTP Transport Protocol”](#) on page 8-10

Configuring Proxy Services using HTTP Transport Protocol

You can choose HTTP as the transport protocol when you configure any type of proxy service and the endpoint URI is of the form:

```
/<someService>
```

where `someService` is the name of proxy service or a business service

You can configure the HTTP transport for a proxy service you must specify the following fields:

- **Basic Authentication Required:** If you enable this, basic authentication is required to access this service.

- **Dispatch Policy:** You must configure work managers in the WebLogic Server console in order to have other dispatch policies in addition to the `default` dispatch policy. For more information on how to configure a work manager, see [Create a Global Work Manager](#) in WebLogic Server Administration Console.
- **Request Encoding:** This parameter specifies the character set encoding for the request messages.
- **Response Encoding:** This parameter specifies the character set encoding for the response messages.

For more information on how to configure this transport, see [Adding a Proxy Service: Transport Configuration](#) in *Using the AquaLogic Service Bus Console*.

Configuring Business Services using HTTP Transport Protocol

You can choose HTTP as the transport protocol when you configure any type of business service and the endpoint URI is of the form:

```
http://<host:port/someService>  
where
```

- `host`: is the name of the system that hosts the service.
- `port`: is the port number at which the connection is made.
- `someService`: is a target service.

To configure the HTTP transport for a business service you must specify the following fields:

- **Timeout:** This parameter specifies the HTTP timeout interval, in seconds, before the connection is dropped. The default value for this parameter is 0.
- **HTTP Request Method:** This parameter enables you to choose the `get` method or the `post` method for the HTTP requests.
- **Basic Authentication Required:** If you enable this, basic authentication is required to access this service.
- **Service Account:** This resource contains the login credentials required for the `Basic Authentication`.
- **Follow HTTP redirects:** Enabling this parameter allows the business to follow the `HTTP` redirects.

- **Dispatch Policy:** You must configure work managers in the WebLogic Server console in order to have other dispatch policies in addition to the `default` dispatch policy. For more information on how to configure a work manager, see [Create a Global Work Manager](#) in WebLogic Server Administration Console.
- **Request Encoding:** This parameter specifies the character set encoding for request messages.
- **Response Encoding:** This parameter specifies the character set encoding for response messages.

For more information on how to configure this transport, see [Adding a Business Service: Transport Configuration](#) in *Using the AquaLogic Service Bus Console*.

HTTP(S)

You can choose HTTP(S) transport protocol when you configure any type of proxy service and the endpoint URI is of the form:

/<someService>

where `someService` is the name of a proxy service or a business service. Following sections describe:

- [“Configuring Proxy Services using HTTP\(S\) Transport Protocol”](#) on page 8-11
- [“Configuring Business Services using HTTP\(S\) Transport Protocol”](#) on page 8-12

Configuring Proxy Services using HTTP(S) Transport Protocol

You can configure the HTTP(S) transport for a proxy service you must specify the following fields:

- **Client Authentication:** The Client Authentication method provides you with three options:
 - **None:** This option enables one-way SSL. No client authentication is required.
 - **Basic:** This option enables one-way SSL. But this requires `user/password` client authentication.
 - **Client Certificate:** This option enables one-way SSL. But this requires `user/password` client-side and server-side authentication. BEA recommends this method of client authentication.

- **Dispatch Policy:** You must configure work managers in the WebLogic Server console in order to have other dispatch policies in addition to the `default` dispatch policy. For more information on how to configure a work manager, see [Create a Global Work Manager](#) in WebLogic Server Administration Console.
- **Request Encoding:** This parameter specifies the character set encoding for the request messages.
- **Response Encoding:** This parameter specifies the character set encoding for the response messages.

For more information on how to configure this transport, see [Adding a Proxy Service: Transport Configuration](#) in *Using the AquaLogic Service Bus Console*.

Configuring Business Services using HTTP(S) Transport Protocol

You can choose the HTTP(S) transport protocol when you configure any type of business service and the endpoint URI is of the form:

```
http(s)://<host:port/someService>  
where
```

- `host`: is the name of the system that hosts the service.
- `port`: is the port number at which the connection is made.

To configure the HTTP(S) transport for a business service you must specify values for the following fields:

- **Timeout:** This parameter specifies the HTTP timeout interval, in seconds, before the connection is dropped. The default value for this parameter is 0.
- **HTTP Request Method:** This parameter enables you to choose the `get` method or the `post` method for the HTTP requests.
- **Basic Authentication Required:** Enable this basic authentication is required to access this service.
- **Service Account:** This resource contains the login credentials required for the Basic Authentication.
- **Follow HTTP redirects:** Enabling this parameter allows business services to follow HTTP redirects.

- **Dispatch Policy:** You must configure work managers in the WebLogic Server console in order to have other dispatch policies in addition to the `default` dispatch policy. For more information on how to configure a work manager, see [Create a Global Work Manager](#) in WebLogic Server Administration Console.
- **Request Encoding:** This parameter specifies the character set encoding for the request messages.
- **Response Encoding:** This parameter specifies the character set encoding for the response messages.

For more information on how to configure this transport, see [Adding a Business Service: Transport Configuration](#) in *Using the AquaLogic Service Bus Console*.

JMS

You can choose JMS as the transport protocol for all the types of proxy services. AquaLogic Service Bus is certified with the following JMS implementations:

- WebLogic Server 9.x JMS
- IBM WebSphere MQ/JMS release 5.3
- TIBCO Enterprise Message Service™ release 4.2

The proxy services and business services must be configured to use the JMS transport as described in [Adding a Proxy Service: Transport Configuration](#) and [Adding a Business Service: Transport Configuration](#) sections of *Using the AquaLogic Service Bus Console*.

For more information on the JMS transport, see [Interoperability Solutions for JMS and WebSphere MQ](#). The following sections describe:

- [“Configuring Proxy Services using JMS Transport Protocol”](#) on page 8-13
- [“Configuring Business Services using JMS Transport Protocol”](#) on page 8-16

Configuring Proxy Services using JMS Transport Protocol

You can choose the HTTP(S) transport protocol when you configure any type of business service and the endpoint URI is of the form:

```
jms://<host:port[,host:port]*/factoryJndiName/destJndiName>
```

where

- `host`: is the name of the system that hosts the service.
- `port`: is the port number at which the connection is made.
- `[,host:port]*`: indicates that you can configure multiple hosts with corresponding ports.

Note: To target a target a JMS destination to multiple servers, use the following format of the URI:

```
jms://host1:port,host2:port/QueueConnectionFactory/destJndiName
```

where `QueueConnectionFactory` is name of the Connection factory Queue.

For more information on how to define a Connection factory queue, see [Configure resources for JMS system modules](#) in *Administration Console Online Help*.

- `factoryJndiName`: The name of the JNDI Connection Factory. For more information on how to define a Connection factory queue, see [Configure resources for JMS system modules](#) in *Administration Console Online Help*.
- `destJndiName`: is the name of the JNDI destination.

To configure a proxy service, using JMS transport protocol you must specify values for the following fields:

- **Destination Type:** You must specify the destination to be of one of the following:
 - **Queue:** This defines a point-to-point destination type. You can use this destination type for peer asynchronous communications. A message, which is delivered to a queue is distributed to only one destination.
 - **Topic:** This defines a publish or a subscribe destination type. You can use this destination type for peer asynchronous communications. A message, which is delivered to a topic is distributed to all the subscribers of the topic.
- **Is Response Required:** This specifies if you expect a response to the outbound message. If you expect a response you must specify the following parameters for the response message:
 - **Response Correlation Pattern:** This configures the design pattern for JMS response message to be one of the following:
 - **JMS Correlation ID:** You must choose `JMS Correlation ID` design pattern for your message if you want to correlate by JMS Correlation ID and send the response to the URI configured in the `Response URI` field.

Note: The `Response URI` field is active only when you choose `JMS Correlation ID` design pattern for your response message.

- **JMS Message ID:** You can JMS Message ID design pattern for your message if you want to correlate by JMS Message ID and send the response to the `JMSReplyTo Destination` by configuring the `Response Connection Factory` field.

Note: The `Response Connection Factory` field is active only when you choose `JMS Message ID` design pattern for your response message.

- **Response Message Type:** You can set type of the response message to `Bytes` or `Text`.
- **Response Encoding:** You can set the encoding for the response message. The default encoding is `UTF-8`.
- **Client Response Timeout:** This parameter specifies the response timeout interval, in seconds.
- **Request Encoding:** You can set the encoding for the request message. The default encoding is `UTF-8`.
- **Dispatch Policy:** This specifies the dispatch policy for the endpoint. You must configure work managers in the WebLogic Server console in order to have other dispatch policies in addition to the default dispatch policy. For more information on how to configure a work manager, see [Create a Global Work Manager](#) in WebLogic Server Administration Console.

- **Advanced Settings:** To configure the advance settings click on  icon on the right hand side to expand the advanced settings section. Configuring parameters in this section is optional.

You can set the following parameters in the section:

- **Use SSL:** This specifies if the connections can be made over SSL or not.
- **Message Selector:** You can use this field to specify the criteria for selecting messages.
- **Durable Subscription:** You can check if the subscription is durable only if the destination type is `Topic`.
- **Retry Count:** In this field you can configure the maximum number of retries for the connection.
- **Retry Interval:** In this field you can configure the time interval in milli-seconds between consecutive retries.
- **Error Destination:** In this field you can configure the name of the target destination for the messages, which have reached the maximum number of retry count.

- Expiration Policy: In this field you can specify the message expiration policy to be used when you encounter an expired message at a WebLogic Server or a JMS destination.
- JMS Service Account: In this field you can specify the name of the service account resource to be used to make the connection over the secured socket layer.

Configuring Business Services using JMS Transport Protocol

You can choose the HTTP(S) transport protocol when you configure any type of business service and the endpoint URI is of the form:

```
jms://<host:port[,host:port]*/factoryJndiName/destJndiName >
```

where

- `host`: is the name of the system that hosts the service.
- `port`: is the port number at which the connection is made.
- `[,host:port]*`: indicates that you can configure multiple hosts with corresponding ports.

To target a target a JMS destination to multiple servers, use the following format of the URI:

```
jms://host1:port,host2:port/QueueConnectionFactory  
/destJndiName
```

where `QueueConnectionFactory` is name of the Connection factory Queue. For more information on how to define a Connection factory queue, see [Configure resources for JMS system modules](#) in *Administration Console Online Help*.

- `factoryJndiName`: The name of the JNDI Connection Factory. For more information on how to define a Connection factory queue, see [Configure resources for JMS system modules](#) in *Administration Console Online Help*.
- `destJndiName`: is the name of the JNDI destination.

To configure a business service, using JMS transport protocol you must specify values for the following fields:

- Destination Type: You can specify the destination to be of one of the following:
 - Queue: This defines a point-to-point destination type. You can use this destination type for peer asynchronous communications. A message, which is delivered to a queue is sent to only one destination.

- Topic: This defines a publish or a subscribe destination type. You can use this destination type for peer asynchronous communications. A message, which is delivered to a topic is distributed to all the subscribers of the topic.
- Is Response Required: You can specify if you expect a response to the outbound message. If you expect a response you must specify the following parameters for the response message:
 - Response Correlation Pattern: You can configure the design pattern for JMS response message to be one of the following:
 - JMS Correlation ID: You must choose `JMS Correlation ID` design pattern for your message if you want to correlate by JMS Correlation ID and send the response to the URI configured in the `Response URI` field.

Note: The `Response URI` field is active only when you choose `JMS Correlation ID` design pattern for your response message.
 - JMS Message ID: You must choose `JMS Message ID` design pattern for your message if you want to correlate by JMS Message ID and send the response to the `JMSReplyTo Destination` by configuring the `Response Connection Factory` field.

Note: The `Response Connection Factory` field is active only when you choose `JMS Message ID` design pattern for your response message.
 - Response Message Type: You can set type of the response message to `Bytes` or `Text`.
 - Response Encoding: You can set the encoding for the response message. The default encoding is `UTF-8`.
 - Client Response Timeout: This parameter specifies the response timeout interval, in seconds.
- Request Encoding: You can set the encoding for the request message. The default encoding is `UTF-8`.
- Dispatch Policy: You can specify the dispatch policy for the endpoint. You must configure work managers in the WebLogic Server console in order to have other dispatch policies in addition to the `default` dispatch policy. For more information on how to configure a work manager, see [Create a Global Work Manager](#) in WebLogic Server Administration Console.
- Advanced Settings: To configure the advance settings click on  icon on the right hand side to expand the advanced settings section. Configuring parameters in this section is optional.

You can set the following parameters in the section:

- Use SSL: This specifies if the connections can be made over SSL or not.
- Expiration: This specifies the time interval in milli-seconds after which the message will expire. Default value is 0, which means that the message never expires.
- Unit Of Order: This is a value added feature of WebLogic, which enables a message producer to group messages into a single unit with respect to the order of processing. This single unit is called the `Unit of Order`. All the messages in a unit must be processed sequentially in the same order they were created.

Note: This is supported by WebLogic Server 9.0.

- JNDI Service Account: You can choose the service account resource to be used for JNDI lookups.
- JMS Service Account: You can choose the service account resource to be used for the JMS server connection.

Local

Every proxy service is associated with a protocol that determines the level of communication used by the clients to send requests to the proxy service. In AquaLogic Service Bus there two categories of proxy services—the proxy services of first category are invoked directly by the clients; those of the second category are invoked by other proxy services in the message flow. The proxy services of the second category use a new transport called the *local transport*. For more information on Local Transport, see [Chapter 10, “Local Transport.”](#)

Tuxedo

BEA AquaLogic Service Bus and BEA Tuxedo can inter-operate to use the services each of them offer. The Tuxedo transport is secure, reliable, high performing, and provides bi-directional access to the Tuxedo domain from AquaLogic Service Bus. You can access domains in AquaLogic Service Bus from Tuxedo. Also you can access tuxedo domains from AquaLogic Service Bus. For more information about Tuxedo transport, see [Interoperability Solution for Tuxedo](#).

You can configure both proxy services and business services in AquaLogic Service Bus. For more information on configuring a proxy service, see [Adding a Proxy Service: Transport Configuration](#) and for more information on configuring a business service, see [Adding a Business Service: Transport Configuration](#) sections of *Using the AquaLogic Service Bus Console*.

EJB Transport

Using the EJB Transport, AquaLogic Service Bus supports native RMI invocation of Stateless Session Beans deployed on WebLogic Server 8.1, 9.0, 9.1 or 9.2. It allows transactional and secure communications. The EJB transport can also be leveraged to expose an EJB as a Web service through AquaLogic Service Bus.

This section includes the following topics:

- [“Introduction” on page 9-1](#)
- [“Invoking EJBs from AquaLogic Service Bus” on page 9-3](#)
- [“Exposing EJBs as Web Services” on page 9-11](#)
- [“Advanced Topics” on page 9-11](#)
- [“Troubleshooting” on page 9-15](#)

Introduction

You can design business services in AquaLogic Service Bus to use the EJB transport. The EJB transport is fully integrated into the AquaLogic Service Bus configuration, management, monitoring, and test consoles. Business services built with the EJB transport can be used for Publish, Service Callout, and service invocations. You cannot create proxy services that use the EJB transport.

An EJB can be exposed as a Web service, without the need for tools or the modification of the legacy code on the application server that hosts the EJB.

The EJB transport provides the following capabilities:

Transactional Integrity

You can call EJB Business service in the context of a global transaction. The EJB Transport can also suspend or start a global transaction before invoking an EJB.

Security Propagation

The security context established at the beginning of a message flow, from an AquaLogic Service Bus client is propagated to the other system. In other words, an incoming SOAP over HTTP request to AquaLogic Service Bus that requires authentication is authenticated by AquaLogic Service Bus and the authenticated subject can then be propagated to the EJB server.

HTTP Tunneling and Encrypted Communication

You can access EJBs that are behind a fire wall with HTTP tunneling. For additional security, you can use SSL to encrypt all of the communications with the EJB Server.

JNDI Provider

EJB transport leverages the JNDI provider—an AquaLogic Service Bus resource. The JNDI provider defines communication protocols and security credentials for accessing remote servers. A JNDI provider can be reused by multiple EJB business services. This provides a centralized way for administrators to manage remote EJB server configurations.

For information about JNDI provider resources, see [System Administration](#) in *Using the AquaLogic Service Bus Console*.

High Performance Caching

The EJB transport is built on high performance cache. This allows the reuse of established connections and minimizes EJB stubs lookups.

Failover and Load Balancing

The EJB transport can take advantage of scenarios in which the same EJB is deployed in multiple domains or on a cluster for load balancing or failover or both.

Advanced XML to Java Binding Capabilities

The EJB transport leverages the WebLogic Server JAX-RPC stack to perform Java to XML bindings. The JAX-RPC stack is a high performance engine that supports advanced Java objects such as XML Beans. If the Java type is not recognized by the stack, an extension mechanism is provided to facilitate support of these Java types. For information about this extension mechanism (using the converter classes), see [Supported Types and Converter Class](#).

Intelligent Retries

The EJB transport makes retry decisions based on the nature of the failure that can occur during the invocation of an EJB.

Invoking EJBs from AquaLogic Service Bus

Before you can configure a business service in AquaLogic Service Bus, you must register a JNDI provider resource and a client JAR resource. This section describes how to design and configure an EJB transport business service in AquaLogic Service Bus. It includes the following topics:

- [“Register a JNDI Provider Resource” on page 9-3](#)
- [“Register an EJB Client JAR Resource” on page 9-4](#)
- [“Create an EJB Business Service” on page 9-5](#)
- [“Invoking EJB Business Services” on page 9-11](#)

Register a JNDI Provider Resource

A JNDI Provider resource allows you to specify the communication protocols and security credentials used to retrieve EJB stubs bound in the JNDI tree of remote WebLogic 8.1 or 9.x domains. (For more information how to setup a JNDI tree, see [Programming WebLogic JNDI](#) in the [BEA WebLogic Server documentation](#).)

Typically, the target EJB is not located in the same domain as AquaLogic Service Bus. In this case, you must register a JNDI Provider resource. When the EJB is located in the same domain, you can define a provider to specify credentials and take advantage of stubs caching, although it is optional in this case.

The JNDI provider has a high performance caching mechanism for remote connections and EJB stubs. The preferred communication protocol from AquaLogic Service Bus to a WebLogic Server domain is `t3` or `t3s`. If messages need to go through a fire wall, you can use HTTP tunneling. For more information about HTTP tunneling, see [HTTP Tunneling and Encrypted Communication](#).

Notes:

- Although it is possible to use a WebLogic Server foreign JNDI Provider, BEA recommends that you do not.
- 2-way SSL is not supported.

- EJB transport provider does not support CLIENT CERT to look-up JNDI tree or access a method on an EJB.

Adding a JNDI Provider

For information about registering and configuring a JNDI provider resource in AquaLogic Service Bus, see “Adding a JNDI Provider” in [System Administration](#) in *Using the AquaLogic Service Bus Console*.

Register an EJB Client JAR Resource

A client JAR must be registered as a resource in AquaLogic Service Bus. It is therefore part of the AquaLogic Service Bus configuration and can be exported from and imported into a project.

An EJB client JAR file must contain the interfaces and classes needed by AquaLogic Service Bus to access an EJB. This includes the remote and home interfaces and any dependent types to which the client is exposed, such as method parameter types or application exceptions. If your business service requires converter classes, they also need to be included in the JAR file. For information about converter classes, see [Converter Classes](#).

Consider the following guidelines when using EJB client JARs:

- Adding Home and remote interfaces in the system classpath is bad practice and is not supported by AquaLogic Service Bus.
- BEA recommends that you keep the client JAR size small, include a single home interface per JAR and not register the entire ejb-jar file.
- You can use WebLogic Workshop to obtain a client JAR for EJBs deployed on WebLogic Server 8.1 or 9.x.
- Client-jars compiled with JDK 1.4 or later are supported.

Adding a Client JAR

For information about registering and configuring a JAR resource in AquaLogic Service Bus, see “Adding a JAR” in [JARs](#) in *Using the AquaLogic Service Bus Console*.

Create a Service Account (Optional)

If the EJB methods are protected, you can specify the credentials you want to use for the invocations. Those credentials are often different than the credentials used by the JNDI provider.

For information about adding and using service accounts, see [Service Accounts](#) in *Using the AquaLogic Service Bus Console*.

Locate an EJB in the JNDI Tree

If you do not know the JNDI name for an EJB, you can browse the EJB Server JNDI tree. For information about browsing the JNDI tree using the WebLogic Server Administration Console, see:

- [JNDI](#) in the *WebLogic Server 8.1 Administration Console Online Help* (for WebLogic Server 8.1)
- [View objects in the JNDI tree](#) in the *WebLogic Server 9.2 Administration Console Online Help* (for WebLogic Server 9.x)

Create an EJB Business Service

This section provides information about creating a business service that uses the EJB transport. It includes the following topics:

- [“General Configuration”](#) on page 9-5
- [“EJB Transport-Specific Configuration”](#) on page 9-7
- [“EJB Business Service Interface Configuration”](#) on page 9-9

General Configuration

1. Open the AquaLogic Service Bus Console and in an active session, select Project Explorer from the left navigation panel. The **Project View** page is displayed.
2. Select the project in which you want to create the business service. A page in which you can create a business service is displayed—create a new business service.
3. On the General Configuration page, as shown in the following figure, enter a name for the business service and select the **Transport Typed Service** as the **Service Type**.

An EJB business service is a *Transport Typed Service*, meaning that the type of the transport is determined by the configuration of the service. The EJB transport is currently the only such transport type. You can add other transports by using the AquaLogic Service Bus Transport SDK.

An entry in the **Description** field is optional.

Figure 9-1 Create a Business Service - General Configuration

Create a Business Service - General Configuration (default/)

Service Name*

Description

Service Type*

Create a New Service

WSDL Web Service
 (port or binding)

Transport Typed Service

Messaging Service

Any SOAP Service

Any XML Service

Create From Existing Service

Business Service

Proxy Service

| |

4. Click **Next** to open the transport-specific configuration page.
5. In the Transport Configuration page, select **ejb** as the **Protocol**.

Figure 9-2 Create a Business Service—Transport Configuration

Create a Business Service - Transport Configuration (default/EJB Business Service)

Protocol*

Load Balancing Algorithm

Endpoint URI* Format: `ejb:<provider>:jndi_name`

| EXISTING URIS | OPTIONS |
|----------------------------------|--|
| ejb:WLS81_Provider:alsb.ejb.Ejb1 | <input type="button" value="↑"/> <input type="button" value="↓"/> <input type="button" value="🗑"/> |

Retry Count

Retry Interval

| | |

6. Enter the Endpoint URI and add it to the list of EXISTING URIs. To build the URI you need the name of the provider you used in [“Adding a JNDI Provider” on page 9-4](#) and the location of the EJB home interface in the JNDI tree you determined in [“Locate an EJB in the JNDI Tree” on page 9-5](#):

```
ejb:provider:jndi_name
```

If the EJB is deployed locally, you need not provide a JNDI provider name. In this case, the URI format is:

```
ejb::jndi_name
```

7. As for any AquaLogic Service Bus transport, you can also specify the Load Balancing Algorithm, Retry count, Retry Interval and specify multiple URIs for failover. See [“Retries and Failover” on page 9-13](#).
8. Click **Next** to open the EJB transport-specific configuration page.

EJB Transport-Specific Configuration

After completing the general configuration of the business service you specify EJB transport-specific information such as the Home and Remote interfaces. The EJB Transport Configuration page in the AquaLogic Service Bus Console is shown in the following figure.

Figure 9-3 Create a Business Service— EJB Transport Configuration

| Create a Business Service - EJB Transport Configuration (default/EJB Business Service) | |
|--|---|
| Service Account | <input type="text"/> <input type="button" value="Browse..."/> |
| Supports Transaction | <input checked="" type="checkbox"/> |
| Client Jar* | <input type="text"/> <input type="button" value="Browse..."/> |
| Home Interface* | <input type="text" value="v"/> |
| Remote Interface* | <input type="text"/> |
| <input type="button" value=" << Prev."/> <input type="button" value=" Next >>"/> <input type="button" value=" Finish"/> <input type="button" value=" Cancel"/> | |

To Configure the EJB Transport

1. Optionally select a Service Account.

If the EJB methods are protected and you defined a service account as described in [“Create a Service Account \(Optional\)” on page 9-4](#), click **Browse** to locate the appropriate Service Account.

- By default, the **Supports Transaction** option is selected. This specifies that the EJB supports transaction. If you do not want to propagate transactions, or if the EJB does not support transactions, deselect **Supports Transaction**.

For information about transaction processing with the EJB Transport, see [“Transaction Processing, Retries, and Errors Handling” on page 9-12](#).

- Select the Client JAR—browse and select the Client JAR you registered previously, as described in [“Adding a Client JAR” on page 9-4](#).

When you select a Client JAR, a list of its Home Interface is displayed on this page.

- Select the Home Interface from the list of interfaces provided in the **Home Interface** field.

Notice that the Remote Interface is automatically deduced from the Home Interface and the configuration page is refreshed. The **Remote Interface** field is populated and other options are provided that allow you to control the interface of the service and the WSDL generated when you finish configuration of this business service.

Figure 9-4 Create a Business Service— EJB Transport Configuration after Selecting the Home Interface

| Create a Business Service - EJB Transport Configuration (default/EJB Business Service) | | | | | | | |
|--|---|----------------------------------|-------------------------------------|----------|----------------------------------|-------------------------------------|------|
| Service Account | <input type="text"/> <input type="button" value="Browse..."/> | | | | | | |
| Supports Transaction | <input checked="" type="checkbox"/> | | | | | | |
| Client Jar* | <input type="text" value="ejb/ejb client.jar"/> <input type="button" value="Browse..."/> | | | | | | |
| Home Interface* | <input type="text" value="com.bea.wli.sb.ejb.SampleEjbHome"/> <input type="button" value="v"/> | | | | | | |
| Remote Interface* | <input type="text" value="com.bea.wli.sb.ejb.SampleEjb"/> | | | | | | |
| TargetNamespace* | <input type="text" value="http://www.openuri.org/"/> | | | | | | |
| Style | <input checked="" type="radio"/> Document Wrapped <input type="radio"/> RPC | | | | | | |
| Encoding | <input checked="" type="radio"/> Literal <input type="radio"/> Encoded | | | | | | |
| Methods | <table border="1"> <tr> <td><input type="button" value="+"/></td> <td><input checked="" type="checkbox"/></td> <td>sayHello</td> </tr> <tr> <td><input type="button" value="+"/></td> <td><input checked="" type="checkbox"/></td> <td>sort</td> </tr> </table> | <input type="button" value="+"/> | <input checked="" type="checkbox"/> | sayHello | <input type="button" value="+"/> | <input checked="" type="checkbox"/> | sort |
| <input type="button" value="+"/> | <input checked="" type="checkbox"/> | sayHello | | | | | |
| <input type="button" value="+"/> | <input checked="" type="checkbox"/> | sort | | | | | |
| <input type="button" value=" << Prev."/> <input type="button" value=" Next >>"/> <input type="button" value=" Finish"/> <input type="button" value=" Cancel"/> | | | | | | | |

EJB Business Service Interface Configuration

An EJB business service is a Transport Typed Service, which means the type of the transport is determined by the configuration of the service.

The type of an EJB business service is equivalent to a SOAP XML service—in other words, you can use an EJB business service like any other SOAP XML business service. A WSDL is generated when you save the EJB Transport Configuration.

The WSDL is generated based on the interface of the EJB. The EJB transport configuration page provides configuration options for you to control the interface of the service and the WSDL that is generated. To do so, complete the configuration on the EJB Transport Configuration page as shown in the preceding figure:

1. **TargetNamespace**—Specify the target namespace of the WSDL.
2. **Style**—You can select Document Wrapped or RPC.
3. **Encoding**—Select Literal or Encoded.
4. The methods displayed are those of the EJB Remote Interface you selected. For example, the following figure displays two methods: `sayHello` and `sort`.

Figure 9-5 Create a Business Service— EJB Transport Configuration, Expanded Methods Configuration

| Methods | | | | | | | | | | | | | |
|---|--|----------------------|------|-----------|--------|-----------------------------------|----------------------|------|------|-----------|--------|-------------------------------------|----------------------|
| [-] <input checked="" type="checkbox"/> | <p>sayHello</p> <p>Operation* <input type="text" value="sayHello"/></p> <p>Parameters* <table border="1"> <thead> <tr> <th>Type</th> <th>Name</th> <th>Converter</th> </tr> </thead> <tbody> <tr> <td>String</td> <td><input type="text" value="arg0"/></td> <td><input type="text"/></td> </tr> </tbody> </table> </p> <p>Return* <table border="1"> <thead> <tr> <th>Type</th> <th>Name</th> <th>Converter</th> </tr> </thead> <tbody> <tr> <td>String</td> <td><input type="text" value="return"/></td> <td><input type="text"/></td> </tr> </tbody> </table> </p> | Type | Name | Converter | String | <input type="text" value="arg0"/> | <input type="text"/> | Type | Name | Converter | String | <input type="text" value="return"/> | <input type="text"/> |
| Type | Name | Converter | | | | | | | | | | | |
| String | <input type="text" value="arg0"/> | <input type="text"/> | | | | | | | | | | | |
| Type | Name | Converter | | | | | | | | | | | |
| String | <input type="text" value="return"/> | <input type="text"/> | | | | | | | | | | | |
| [-] <input checked="" type="checkbox"/> | <p>sort</p> <p>Operation* <input type="text" value="sort"/></p> <p>Parameters* <table border="1"> <thead> <tr> <th>Type</th> <th>Name</th> <th>Converter</th> </tr> </thead> <tbody> <tr> <td>List</td> <td><input type="text" value="arg0"/></td> <td><input type="text"/></td> </tr> </tbody> </table> </p> <p>Return* <table border="1"> <thead> <tr> <th>Type</th> <th>Name</th> <th>Converter</th> </tr> </thead> <tbody> <tr> <td>List</td> <td><input type="text" value="return"/></td> <td><input type="text"/></td> </tr> </tbody> </table> </p> | Type | Name | Converter | List | <input type="text" value="arg0"/> | <input type="text"/> | Type | Name | Converter | List | <input type="text" value="return"/> | <input type="text"/> |
| Type | Name | Converter | | | | | | | | | | | |
| List | <input type="text" value="arg0"/> | <input type="text"/> | | | | | | | | | | | |
| Type | Name | Converter | | | | | | | | | | | |
| List | <input type="text" value="return"/> | <input type="text"/> | | | | | | | | | | | |

5. You can exclude the methods you do not want to expose by unchecking the checkbox associated with the method names.
6. You can change the default operation name for a given method. (By default, the operation name is the method name.) If an EJB contains methods with same name, you must change the operation names so that they are unique—WSDLs require unique operation names.
7. You must exclude the methods with parameters or return types that are not supported by the JAX-RPC stack or you must associate such arguments with “[Converter Classes](#)” on [page 9-15](#).

The following figure shows an example of a custom methods configurations.

Figure 9-6 Create a Business Service— EJB Transport Configuration, Customized Methods Configuration

The screenshot shows a configuration window titled "Methods" with two entries:

- sayHello** (checkbox unchecked):
 - Operation*: sayHello
 - Parameters*:

| Type | Name | Converter |
|--------|---------|------------|
| String | message | [Dropdown] |
 - Return*:

| Type | Name | Converter |
|--------|----------|------------|
| String | response | [Dropdown] |
- sort** (checkbox checked):
 - Operation*: sort
 - Parameters*:

| Type | Name | Converter |
|------|-------|---------------------------------|
| List | names | ArrayToListConverter [Dropdown] |
 - Return*:

| Type | Name | Converter |
|------|-------------|---------------------------------|
| List | sortedNames | ListToArrayConverter [Dropdown] |

8. Click **Next**. Save the service and activate the session.

Note: If the credentials or transaction settings are different between the methods for a given EJB, you can leverage the ability to customize the methods for a given business service, and create a business service per method. This gives you fine-grained control over transactions and credentials.

Invoking EJB Business Services

An EJB business service can be used as a SOAP XML business service. You can publish to, route to, or callout to an EJB business service. If you need transaction support, set the QoS to *Exactly-Once*. See [“Transaction Processing, Retries, and Errors Handling” on page 9-12](#).

You can also use the test console to validate your configuration and to help you to determine the shape of the XML request.

Exposing EJBs as Web Services

You can leverage the EJB transport to easily expose EJBs as Web Services.

Note: You cannot create a proxy service from an existing EJB business service—you must first get the WSDL generated from the EJB business service, and then create the proxy service based on that WSDL. To do so, complete the following steps:

1. Create an EJB business service pointing to the EJB you want to expose, as described in [“Create an EJB Business Service” on page 9-5](#).
2. From the service details page on the AquaLogic Service Bus Console, get the WSDL for the EJB business service.

The WSDL is contained in a JAR file. You can obtain the WSDL only if there is no pending session.

3. Extract the WSDL from the JAR and register it as a WSDL resource. For information about creating WSDL resources, see [WSDLs](#) in *Using the AquaLogic Service Bus Console*.

If the configuration of the business service changes, a new WSDL is generated. If that happens, you must get the new WSDL and re-register it as a WSDL resource.

4. Create a SOAP XML proxy service based on the WSDL.
5. Edit the proxy service pipeline and route to the EJB business service.

You can now invoke the EJB as a Web Service with no need for purchasing an expensive Web Service toolkit or carrying out intrusive actions on the EJB server.

Advanced Topics

This section includes information about EJB transport that will help you understand how EJB business services behave at run time depending on how they are configured at design time. It includes the following topics:

- [“Transaction Processing, Retries, and Errors Handling” on page 9-12](#)
- [“Supported Types and Converter Class” on page 9-14](#)

Transaction Processing, Retries, and Errors Handling

Transactions

The EJB transport can create, suspend, and propagate transactions. The transaction between AquaLogic Service Bus and the EJB server are XA transactions. If you use transactions with HTTP tunneling or have a dedicated communication channel and the EJBs are deployed on 8.1 servers, you must set the *security interoperability* mode for the transaction manager to *performance*. For information about setting the security interoperability mode and other transaction configurations, see [Configuring Transactions](#) in *Programming WebLogic JTA*.

To determine the behavior of the EJB business service, considerations include whether the proxy service pipeline has a transactional context, and what qualities of service (QoS) settings are specified in the pipeline when invoking the service:

QoS Best-Effort

If *Best Effort* QoS is specified in the pipeline, no transaction is propagated to the EJB—any ongoing transaction is suspended before invocation, and resumed after invocation.

QoS Exactly-Once

If *Exactly Once* QoS is specified in the pipeline, and

If the EJB does not support transactions (that is, if the **Supports Transaction** option on the EJB transport configuration page is unchecked), no transaction is propagated to the EJB. As in the case of *Best Effort*, any ongoing transaction is suspended before invocation and resumed afterwards.

or

If the EJB supports transactions (that is, if the **Supports Transaction** option on the EJB transport configuration page is checked), the EJB is invoked in the context of a transaction—any ongoing transaction is propagated to the EJB. If no transaction is present, a transaction is created before invocation and committed afterwards.

For more information about QoS in AquaLogic Service Bus services, see [“Quality of Service” on page 2-55](#).

Retries and Failover

Assuming that the EJB business service is configured for retries or failovers, the EJB transport distinguishes the following types of exceptions:

- Runtime Exceptions or Remote Exceptions—typically unexpected fatal errors or communication exceptions
- Exception raised by the JAX-RPC engine—exceptions that occur during the XML to Java conversion
- EJB Checked Exceptions—exceptions declared in the EJB method signature specific to the EJB implementation; also called Business Exceptions

Retries and failover are based on the type of errors and also in the QoS:

QoS Best-Effort

If a run-time or remote exception is thrown, the EJB transport attempts retries or failovers.

If an exception occurs in the JAX-RPC engine, an error is raised to the pipeline and no retries or failover attempts are made.

If an EJB Checked Exception is thrown, an error is raised to the pipeline and no retries or failover attempts are made.

QoS Exactly-Once

If a run-time or remote exception is thrown and the ongoing transaction has been set as *rollback only* (likely by the EJB container), it means the EJB container has been reached and a fatal error either occurred within the EJB container or the EJB. In this case, no retries or failover attempts are made and an error is raised to the pipeline.

If a runtime or remote exception is thrown but the ongoing transaction has not been set as *rollback only*, it means an error occurred before the invocation of the EJB container and the EJB transport will attempt retries or failovers. Note that in this case, the EJB transport still respects the *exactly-once* semantic.

If an exception occurs in the JAX-RPC engine, the EJB transport sets the ongoing transaction to *rollback only* and an error is raised to the pipeline; no retries or failover attempts are made.

If an EJB Checked Exception is thrown, an error is raised to the pipeline and no retries or failover attempts are made.

See [“Transactions” on page 9-12](#) for other repercussions of QoS specifications for an EJB business service.

Error Handling

When throwing a checked exception, according to the EJB specifications, the ongoing transaction can be specified as *rollback only*.

If the ongoing transaction is set as *rollback only* by the EJB developer, the transaction is eventually rolled back by its creator (most likely the proxy service).

If the ongoing transaction is not set to *rollback only*, and a checked exception is raised, it is important to catch EJB checked exceptions in the pipeline with an error handler. If those exceptions are not caught, the pipeline errors are propagated back to the proxy service. The proxy service, in turn, is likely to rollback the ongoing transaction (depending of the transport implementation)—this may not be the intended result.

For example, assume you have an EJB with the following method:

```
public void withdrawFunds(float amount) throws RemoteException,
    InsufficientFundsException {...}
```

Also assume that when an `InsufficientFundsException` exception is thrown, the EJB does not set the current transaction as *rollback only*. In most scenarios, it is wrong to allow the proxy service to roll back the transaction—you may need to configure an error handler in the pipeline to catch the error and avoid this scenario.

Supported Types and Converter Class

The EJB transport is responsible for the XML \leftrightarrow Java conversion. The conversion is performed by the WebLogic Server JAX-RPC engine.

The EJB transport natively supports the following types:

- Primitive types
- XmlObject (both Apache and BEA versions)
- Schema generated XMLBeans (both Apache and BEA versions)
- JavaBean classes

For the full list of natively supported types, see [Data Types and Data Binding](#) in *Programming Web Services for WebLogic Server*.

An EJB method can use parameters/return types that are either not supported by the JAX-RPC engine (an error is reported at design time), or that do not map directly to XML (errors occur at run time). The most commonly used unsupported types are:

- “Object”, “Object[]”
- Java Collections as they are not strongly-typed (for example, List, Set)
- Java classes that do not follow the JavaBean pattern (for example, Map)

You can write a custom converter class than converts those types into types more suitable for XML \leftrightarrow Java conversions. The EJB transport supports custom converter classes.

Converter Classes

A Converter class is a Java class that implements and conforms to the contract defined by the `com.bea.wli.sb.transports.ejb.ITypeConverter` Java interface of the AquaLogic Service Bus public API. For information about the `ITypeConverter` Java interface and other AquaLogic Service Bus APIs, see the AquaLogic Service Bus [Javadoc](#).

To use a converter class for an EJB business service, you must:

1. Create a converter class by implementing and compiling the interface.
2. Add the converter class to the client JAR. (See [“Adding a Client JAR” on page 9-4](#)).
3. When customizing the method configuration during the creation of an EJB business service, navigate to one of the parameter/return types and select the desired converter. See [step 7 in “EJB Business Service Interface Configuration” on page 9-9](#)—the AquaLogic Service Bus Console displays a list of the converters available in the Client-JAR that can be applied to a particular parameter/return type.

Troubleshooting

The information in this section is provided to help you troubleshoot problems when designing or running an EJB business service.

Enabling Debug Mode

The EJB transport uses the same logger as other AquaLogic Service Bus transports. To enable the debug mode, before starting the server, edit the `wlidebug.xml` file in the domain directory and set the category `wli-sb-transports-debug` to `true`. For more information about the `wlidebug.xml` file and the debug flags, see [Appendix B, “Debugging AquaLogic Service Bus.”](#)

Temp Directories

During design time, the EJB transport generates files in the subfolder `alsbejbtransport` and subfolders prefixed with `appgen_` in the `temp` directory. It is safe to delete those folders and

files, and sometimes may be useful to check them to determine what went wrong during activation.

Deployed Application

When an EJB business service is created an application is deployed on the AquaLogic Server. You can use the WebLogic Server Administration Console to monitor and tune this application. The name of EJB business service applications is prepended with `ALSB_EJB`, which is followed by the WSDL type and an auto generated suffix.

Errors

The following items may help in the event that you need to troubleshoot a problem with an EJB business service:

- The following error when creating a business service is due to a Windows operating system limitation—paths containing more than 255 characters are not supported:

```
The system cannot find the path specified):Probably the string length of the path of the file being extracted was too long
```

You can try to reduce the path length by creating a shorter path to the AquaLogic Service Bus domain, or you can use the following option to override the WebLogic Server `temp` directory when starting the server:

```
-Dweblogic.j2ee.application.tmpDir=$desired_short_dir
```

- If you get an XML marshalling error when invoking an EJB business service and you believe the request to be valid against the service WSDL, you probably need to write a converter class. For information, see [“Converter Classes” on page 9-15](#).
- If the EJB interfaces and stubs are changed on the remote server, the first time you try to invoke the new EJB, an error is thrown. Those changes on the remote server are not visible to AquaLogic Service Bus—it tries to invoke the cached EJB stubs, which are no longer valid. However, when the invocation error occurs, the transport assumes that those stubs are now invalid, and remove them from the cache—in this way, the error is prevented on subsequent attempts to invoke the EJB. To avoid this first-time error, you can reset the JNDI Provider in the AquaLogic Service Bus Console.

- For HTTP tunneling between WebLogic Server 9.2 and WebLogic Server 8.1 to work, you must set the `t3-server-abbrev-table-size` element to 255 in the `config.xml` file in the AquaLogic Service Bus domain, as shown in the following code snippet:

```
<server>
  <name>AdminServer</name>
  <ssl>
    <name>AdminServer</name>
    <enabled>true</enabled>
  </ssl>
  <t3-server-abbrev-table-size>255</t3-server-abbrev-table-size>
  <listen-address></listen-address>
</server>
```

EJB Transport

Local Transport

This chapter provides information about the AquaLogic Service Bus local transport. It includes the following topics:

- [“Introduction” on page 10-1](#)
- [“Features and Characteristics of Local Transport Proxy Services” on page 10-2](#)
- [“Usage of Local Transport Proxy Services” on page 10-3](#)
- [“Limitations” on page 10-4](#)

Introduction

Commonly, service bus architectures include complex message flows, in which messages are routed through multiple proxy services that are organized into larger multiple proxy service flows. Individual proxy services in these multiple proxy service scenarios route, publish, or callout to the next proxy service in the flow. The reason for a multiple proxy services design is to support modularity and compartmentalization of the various components of the end-to-end message flow. The individual proxy services in a multiple proxy service flows need to:

- Communicate efficiently and securely.
- Allow transactions and transactional behavior to be propagated
- Allow security context to be propagated so that the identity can be propagated end-to-end. The security context propagation also allows the client of the first proxy service in a multiple service flow to be authorized by the proxy services that are subsequently invoked

in the flow—thus supporting fine-grained access control generic headers in the local transport.

Using the local transport for proxy services ensures support for these capabilities.

Features and Characteristics of Local Transport Proxy Services

Local transport-based proxy services can only be invoked by other proxy service, not by other clients. The invocation is optimized by AquaLogic Service Bus. Local proxy services do not have an URI. However, there are no constraints on the service and interface types supported by local transport proxy services. The one exception is that SAML is only supported in a pass through scenario.

If the quality of service (QoS) for an invoking proxy service is defined as Exactly Once, the transaction of that service is propagated to the local transport proxy service.

In other words, the invoked local transport proxy service inherits the transactional behavior of the invoking proxy service. A proxy service can authenticate at the transport level or the message level. If it is enabled, the effective client is the message-level authenticated client. If the message-level authenticated client is not enabled, then the transport-level authenticated client is the effective client (if that is enabled). If neither the message-level nor the transport-level authenticated client is enabled, the anonymous client becomes the effective client.

When a proxy service invokes a local transport proxy service, the effective client of the invoking proxy service becomes the transport-level client of the invoked local proxy service. A local transport proxy service can authorize this client for access with an access control policy. In this way, it is possible to propagate the client of the first proxy service to all the subsequent proxy services in the overall end-to-end message flow.

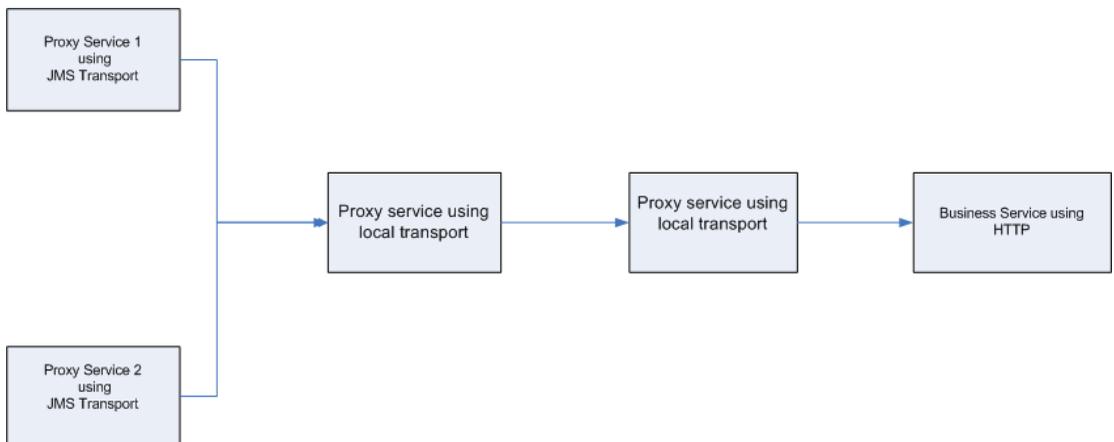
Local transport proxy services support user-defined transport headers. Consider a scenario in which a proxy service uses the HTTP transport; it routes to a local proxy service and the HTTP proxy service passes headers to the local proxy service using the Transport Header action. In this scenario, if the HTTP proxy service received the Content-Type header, this header is available as a user header in the local transport and is therefore accessible through the standard user header, instead of as a typed transport header.

You can invoke a local transport proxy service from the AquaLogic Service Bus test browser. Metrics are collected for a local transport proxy service in the same way as they are any other service. However, local transport proxy services cannot be published to UDDI.

Usage of Local Transport Proxy Services

A common scenario that can be supported using local transport proxy services is one in which a proxy service needs to be invoked using different transports. This can be achieved by putting a set of front-end proxy services (one service per transport) in front of a local transport proxy service in the path of the message flow. These front-end proxy services simply route messages to the local transport proxy service. The following figure illustrates this scenario.

Figure 10-1 Using Local Transport to Implement Convergence

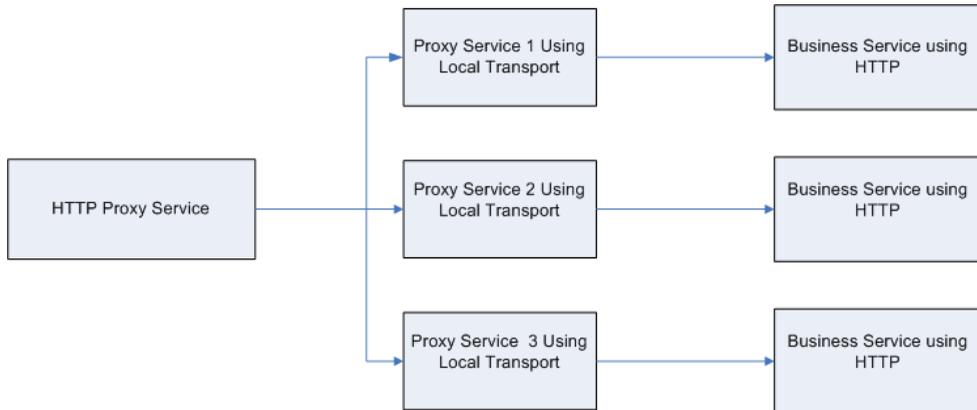


Another common scenario is one in which an Any SOAP or XML type proxy service acts as a front-end to different enterprise systems. This front-end proxy service can receive messages in a variety of formats and uses a technique common to all these messages (for example, a WS-Addressing SOAP header) to route the messages to an appropriate local transport proxy service. In this scenario, the front-end proxy service is acting as a generic router with little knowledge of the enterprise systems or the message formats and semantics. To further abstract the knowledge of the routing rules at design time, the front-end proxy service can use dynamic routing to route messages to the local transport proxy services. For an example of how dynamic routing is used in the proxy services, see [“Using Dynamic Routing” on page 2-32](#) in the users guide.

Each of the local transport proxy services to which messages are routed from the front-end service in turn acts as a front-end proxy service for a specific business service. The local transport proxy services are aware of the message format required by the business services to which they route. In this scenario, these local transport proxy services act as functional proxy services. The

roles of a functional proxy services are to enforce access control for invoking a particular business service, and to perform any transformation of the messages required to invoke the target business service correctly. The following figure illustrates this scenario.

Figure 10-2 Using Local Transport to Access Multiple Business Services



Limitations

The limitations of the local transport are:

- You can invoke the proxy service using the local transport only from another proxy service.
- The proxy services using the local transport cannot be published to the UDDI.
- A local transport proxy service cannot process inbound WS-Security SAML tokens.

Extensibility Using Java Callouts and POJOs

To allow you to extend the capabilities of AquaLogic Service Bus in your organization, you can invoke custom Java code from within proxy services. AquaLogic Service Bus supports a *Java exit mechanism* via a *Java Callout* action that allows you to call out to a Plain Old Java Object (POJO). Static methods can be accessed from any POJO. The POJO and its parameters are visible in the AquaLogic Service Bus Console at design time; the parameters can be mapped to message context variables.

For information about configuring a Java Callout to a POJO, see “Java Callout” in [Proxy Services: Actions](#) in *Using the AquaLogic Service Bus Console*.

Usage Guidelines

The scenarios in which you can use Java Callouts to POJOs in AquaLogic Service Bus include the following:

- Custom validation—Examples of custom validation include validation against a DTD, or doing cross-field semantic validation in Java.
- Custom transformation—Examples of custom transformations can include converting a binary document to base64Binary, or vice versa, or using a custom Java transformation class.
- Custom authentication and authorization—Examples of custom authentication and authorizations include scenarios in which a custom token in a message needs to be authenticated and authorized. However, the authenticated user’s identity cannot be

propagated by AquaLogic Service Bus to the services or POJOs subsequently invoked by the proxy service.

- Lookups for message enrichment—For example, a file or Java table can be used to look up any piece of data that can enrich a message.
- Binary data access—You can use a Java Callout to a POJO to sniff the first few bytes of a binary document to deduce the MFL type. The MFL type returned is used for a subsequent NonXML-to-XML transformation using the [MFL Transform](#) action.
- Implementing custom routing rules or rules engines.

Note: The input and return types for Java Callouts are limited to primitives and XmlObject. Enterprise JavaBeans (EJBs) also provide a Java exit mechanism. The use of EJBs is recommended over the use of POJOs in the following cases:

- When you already have an EJB implementation.
- When you require read access to a JDBC database—Although POJOs can be used for this purpose, EJBs were specifically designed for this and provide better support for management of, and connectivity to, JDBC resources.
- When you require write access to a JDBC database or other J2EE transactional resource—EJBs were specifically designed for transactional business logic and they provide better support for proper handling of failures. However, transaction and security context propagation is supported with POJOs and they can be used for this purpose.

For outbound messaging, BEA recommends that you write a custom transport instead of using POJOs or EJBs.

Best Practices

POJOs are registered as JAR resources in AquaLogic Service Bus. For information about JAR resources, see [JARs](#) in *Using the AquaLogic Service Bus Console*.

In general, BEA recommends that the JARs are small and simple—any large bodies of code that a JAR invokes or large frameworks that are made use of are best included in the system classpath. Note that if you make a change to the system classpath, you must reboot the server.

BEA recommends that you put dependent and overlapping classes in the same JAR resource; put them in different JARS if they are naturally distinct. Any change to a JAR causes all the services that reference it to be redeployed—this can be time consuming for your AquaLogic Service Bus

system. The same class can be located in multiple JAR resources without causing conflicts. The JARs are dynamically class loaded when they are first referenced.

A single POJO can be invoked by one or more proxy services. All the threads in the proxy services invoke the same POJO. Therefore, the POJO must be thread safe. A class or method on a POJO can be synchronized, in which case it serializes access across all threads in all of the invoking proxy services. Any finer-grained concurrency (for example, to control access to a DB read results cache and implement stale cache entry handling) must be implemented by the POJO code.

It is generally a bad practice for POJOs to create threads.

Extensibility Using Java Callouts and POJOs

Tuning AquaLogic Service Bus

This section provides AquaLogic Service Bus tuning tips.

- Whenever possible, set the logging level to *warning*. You set the logging level in the WebLogic Server Administration Console. For more information, see [Servers: Logging: General](#) in the WebLogic Server Administration Console *Online Help*. The following code displays the output server `config.xml` file when the logging level is set to warning. For more information on logging, see “Log” in [Proxy Services: Actions](#) in the *Using the AquaLogic Service Bus Console*.

```
<server>
  <name>AdminServer</name>
  <log>
    <file-min-size>5000</file-min-size>
    <log-file-severity>Warning</log-file-severity>
    <log-file-filter xsi:nil="true"></log-file-filter>
    <stdout-severity>Off</stdout-severity>
    <stdout-filter xsi:nil="true"></stdout-filter>
    <domain-log-broadcast-severity>Error</domain-log-broadcast-severity>
    <domain-log-broadcast-filter
xsi:nil="true"></domain-log-broadcast-filter>
    <memory-buffer-severity>Error</memory-buffer-severity>
    <memory-buffer-filter xsi:nil="true"></memory-buffer-filter>
  </log>
```

</server>

- Group JMS queues on different JMS servers based on message loads. Different JMS servers use different file stores, which you can distribute to separate disk volumes. For more information, “Adding a Business Service” in [Business Services](#) in the *Using the AquaLogic Service Bus Console* and pay particular attention to the *JMS* configuration information.
- If you are using an Oracle database as a JMS persistent store, it is recommended that you use a 10g database and ensure that it has sufficient JDBC connections. Create a JDBC store on a separate schema to use a separate tablespace.
- If you do not require monitoring for a proxy or business service, disable the monitoring capability. For more information, see “Overview of Monitoring” in [Monitoring](#) in the *Using the AquaLogic Service Bus Console*.
- If possible, set the routing data in the JMS message properties. AquaLogic Service Bus does not deserialize message content until the content is explicitly accessed in the pipeline. For example, if the content is an XML document, XML parsing does not happen until an XQuery or XSLT operation happens in the pipeline. For more information about working with the message context in the message flow, see [Message Context](#) in the *Using the AquaLogic Service Bus Console*.
- If you need to extract some of the inbound header elements for processing, you should specify that AquaLogic Service Bus retrieves specific header elements instead of all the elements.
- If you are using an Oracle database as a JMS persistent store BEA recommends that you should ensure that enough JDBC connections are available. This allows an administrator to tune the block size based upon the message size. It is also possible to create each tablespace on a separate datafile and put these datafiles on separate disks. Such an arrangement allows for greater concurrency at the hardware level when the datafiles are not stored on a RAID.
- Where possible, use the insert action instead of the assign action. The insert action uses “in-place” modification semantics enhancing the performance compared to the assign action. For information on configuring actions, see “Adding an Action” in [Proxy Services: Actions](#) in the *Using the AquaLogic Service Bus Console*.
- Use AquaLogic Service Bus clustering and WebSphere MQ clustering to achieve scalability.
- When you are configuring the **Accept Backlog** parameter in the WebLogic Server BEA recommends that you should first increase the default value by twenty five percent.If the

‘connection refused, socket exception’ is thrown, then continue increasing the value till the exception is not thrown. For more information on tuning the WebLogic Server see [Tuning WebLogic Server](#).

- If a front-end application invokes AquaLogic Service Bus synchronously, AquaLogic Service Bus can use the synchronous-asynchronous feature to communicate with WebSphere MQ synchronously. On the WebSphere MQ side, a request and a response queue is set up. AquaLogic Service Bus sends a request to the request queue and waits for a response from the response queue. To achieve improved performance, you can use a dedicate work manager for the response Message Driven Bean. You configure the dedicate work manager in the WebLogic Server Administration Console. For more information, see [Work Manager](#) in the *WebLogic Server Administration Console Online Help*. The following code displays the output server `config.xml` file after the dedicate work manager is configured.

```
<self-tuning>
  <min-threads-constraint>
    <name>minThreadsConstraint</name>
    <target>AdminServer</target>
    <count>20</count>
  </min-threads-constraint>
  <work-manager>
    <name>MQWorkManager</name>
    <target>AdminServer</target>
    <min-threads-constraint> minThreadsConstraint
  </min-threads-constraint>
    <ignore-stuck-threads>false</ignore-stuck-threads>
  </work-manager>
</self-tuning>
```

Tuning AquaLogic Service Bus

Debugging AquaLogic Service Bus

This section provides information about enabling debugging for different modules in AquaLogic Service Bus. You can enable and disable debugging by modifying the corresponding entries in the `wlidebug.xml` file, which is located in the root directory of the AquaLogic Service Bus domain. If the `wlidebug.xml` file is not in the root directory or if it has been deleted, it is created again without any contents when the server starts. The following listing provides an example of the contents of the `wlidebug.xml` file with debugging disabled for all modules (all entries set to `false`).

Listing B-1 `wlidebug.xml` File

```
<?xml version='1.0' encoding='UTF-8'?>
<java:wli-debug-logger xmlns:java="java:com.bea.wli.debug">
  <nl:name xmlns:nl="java:weblogic.diagnostics.debug">wlidebug</nl:name>
  <java:wli-management-debug>false</java:wli-management-debug>
  <java:wli-monitoring-debug>false</java:wli-monitoring-debug>
  <java:wli-management-dashboard-debug>false</java:wli-management-dashboard-de
bug>
  <java:wli-config-debug>false</java:wli-config-debug>
  <java:wli-config-transaction-debug>false</java:wli-config-transaction-debug>
  <java:wli-config-deployment-debug>false</java:wli-config-deployment-debug>
```

Debugging AquaLogic Service Bus

```
<java:wli-config-component-debug>>false</java:wli-config-component-debug>
<java:wli-sb-transport-debug>>false</java:wli-sb-transport-debug>
<java:wli-sb-pipeline-debug>>false</java:wli-sb-pipeline-debug>
<java:wli-alert-manager-debug>>false</java:wli-alert-manager-debug>
<java:wli-jms-reporting-provider-debug>>false</java:wli-jms-reporting-provider-debug>
<java:wli-monitoring-aggregator-debug>>false</java:wli-monitoring-aggregator-debug>
<java:wli-credential-debug >>false</java:wli-credential-debug >
<java:wli-management-common-debug >>false</java:wli-management-common-debug >
</java:wli-debug-logger>
```

Although debugging should be disabled during normal AquaLogic Service Bus operation, you may find it helpful to turn on certain debug flags while you are developing your solution and experimenting with it for the first time. For example, you may want to turn on the alert debugging flag when you are developing alerts and would like to investigate how the alert engine works.

Some of the available debug flags are:

- `wli-config-debug`—Provides information on general aspects of AquaLogic Service Bus configuration.
- `wli-config-deployment-debug`— Provides debug information on session creation, activation, and distribution of configuration in a cluster.
- `wli-config-transaction-debug`—Provides low level debug information about changes made to in-memory data structures and files. This alert flag also generates server startup recovery logs.
- `wli-config-component-debug`—Provides low level debug information about create, update, delete, and import operations.
- `wli-sb-transport-debug`—Provides transport related debug information, including transport headers, which is printed per-message.
- `wli-sb-pipeline-debug`—Prints errors that are generated within the pipeline.
- `wli-alert-manager-debug`—Prints an evaluation of alerts.

All other debug flags are self explanatory.

For all flags, debug information is logged to the server log at `{domaindir}/servers/{servername}/logs/{servername}.log`, except for the `wli-monitoring-aggregator-debug` flag. The `wli-monitoring-aggregator-debug` flag enables debugging for aggregator. This flag logs the aggregated document every minute and stores the log files in the `{domain}\monitoring` folder.

Note: Turning the `wli-monitoring-aggregator-debug` flag on generates large amounts of debug data. Therefore, you should only use this flag for debugging purposes for short periods of time.

Debugging AquaLogic Service Bus

XQuery Implementation

AquaLogic Service Bus uses the [BEA AquaLogic Data Services Platform](#) implementation of the XQuery engine which fully supports all of the language features that are described in the World Wide Web (W3C) specification for XQuery with one exception: modules.

It also supports the following XQuery functions:

- A robust subset of the XQuery functions that are described in W3C specification. For a list of the supported functions and a description of each function, see [BEA XQuery Implementation](#) in the *XQuery Developer's Guide*.
- A subset of the function extensions and language keywords that BEA AquaLogic Data Services Platform provides. See “[Supported Function Extensions from AquaLogic Data Services Platform](#)” on page C-1.
- Its own function extensions. See [Function Extensions from AquaLogic Service Bus](#).

All of the BEA function extensions use the following function prefix `fn-bea:` . For example, the full XQuery notation for an extended function is: `fn-bea: function_name`.

Supported Function Extensions from AquaLogic Data Services Platform

AquaLogic Service Bus supports all function extensions that BEA AquaLogic Data Services Platform provides except for the following:

- `fn-bea:is-access-allowed`

- `fn-bea:is-user-in-group`
- `fn-bea:is-user-in-role`
- `fn-bea:userid`
- `fn-bea:async`
- `fn-bea:timeout`
- `fn-bea:get-property`

BEA recommends that you do not use the following functions in AquaLogic Service Bus—they are better covered by other language features:

- `fn-bea:if-then-else`
- `fn-bea:QName-from-string`
- `fn-bea:sql-like`

For a list of all AquaLogic Data Services Platform function extensions and a description of each function, see [BEA XQuery Implementation](#) in the *XQuery Developer's Guide*.

Function Extensions from AquaLogic Service Bus

AquaLogic Service Bus provides the following XQuery function:

- [fn-bea:lookupBasicCredentials\(\)](#)
- [fn-bea:uuid\(\)](#)
- [fn-bea:serialize\(\)](#)

fn-bea:lookupBasicCredentials()

The `fn-bea:lookupBasicCredentials` function returns the user name and unencrypted password from a specified service account. You can specify any type of service account (static, pass-through, or user-mapping). See [Service Account](#) in *Using the AquaLogic Service Bus Console*.

Use the `fn-bea:lookupBasicCredentials` function as part of a larger set of XQuery functions that you use to encode a user name and password in a custom transport header or in an application-specific location within the SOAP envelope. You do not need to use this function if you only need user names and passwords to be located in HTTP Authentication headers or as WS-Security user name tokens. AquaLogic Service Bus already retrieves user names and

passwords from service accounts and encodes them in HTTP Authentication headers or as WS-Security user name tokens when required.

The function has the following signature:

```
fn-bea:lookupBasicCredentials( $service-account as xs:string ) as
UsernamePasswordCredential
```

where `$service-account` is the path and name of a service account in the following form:

```
project-name[/folder[...]]/service-account-name
```

The return value is an XML element of this form:

```
<UsernamePasswordCredential
  xmlns="http://www.bea.com/wli/sb/services/security/config">
  <username>name</username>
  <password>unencrypted-password</password>
</UsernamePasswordCredential>
```

You can store the returned element in a user-defined variable and retrieve the user name and password values from this variable when you need them.

For example, your AquaLogic Service Bus project is named `myProject`. You create a static service account named `myServiceAccount` in a folder named `myFolder1/myFolder2`. In the service account, you save the user name of `pat` with a password of `patpassword`.

To get the user name and password from your service account, invoke the following function:

```
fn-bea:lookupBasicCredentials (
myProject/myFolder1/myFolder2/myServiceAccount )
```

The function returns the following element:

```
<UsernamePasswordCredential
  xmlns="http://www.bea.com/wli/sb/services/security/config">
  <username>pat</username>
  <password>patpassword</password>
</UsernamePasswordCredential>
```

fn-bea: uuid()

The function `fn-bea:uuid()` returns a universally unique identifier. The function has the following signature:

```
fn-bea:uuid() as xs:string
```

You can use this function in the proxy pipeline to generate a unique identifier. You can insert the generated unique identifier into an XML document as an element. You cannot generate a unique

identifier to the system variable. You can use this mainly to modify the payload. Following use case illustrates the usage of the function to generate a unique identifier:

Suppose you want to generate a unique identifier to add to a message for tracking purpose. You could use this function to generate a unique identifier. The function returns a string, and add it to the SOAP header.

fn-bea:serialize()

You can use the `fn-bea:serialize()` function if you need to represent an XML Document as a string instead of as an XML element. For example, you may want to exchange an XML document through an EJB interface and the EJB method takes String as argument. The function has the following signature:

```
fn-bea:serialize($input as item()) as xs:string
```