



BEA WebLogic® Integration

Best Practices for WLI Application Lifecycle

Copyright

Copyright © 1995-2007 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 © 1995-2007 BEA Systems, Inc. All Rights Reserved. BEA, BEA JRockit, 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

1. Introduction

Features of WLI	1-2
---------------------------	-----

2. Understanding Requirements

3. Architecture and Design of the WLI Application

Modeling Business Processes and Services	3-3
Defining Business Processes	3-3
Identifying Process Objectives and Goals	3-3
Identifying Key Performance Indicators for the Process	3-4
Identifying Process Actors or Participants	3-4
Identifying Public and Private Processes	3-4
Identifying Initiator and Participant Processes	3-5
Keeping Processes Modular	3-5
Enabling end-to-end, Cross-functional Processes	3-6
Separating Production and Monitoring Processes	3-6
Designing and Developing Services	3-6
Service Classification	3-7
Identify Services	3-7
Define Service Contract	3-8
Define Input and Output Service Messages	3-8
Define Pre and Post Conditions for Services	3-8
Decide the Service Calling Paradigm	3-8

Decide the Service Granularity	3-9
Define Quality of Service Requirements	3-9
Design Service With a Service Proxy	3-9
Design Reusable Services	3-9
Design Loosely Coupled Services	3-10

4. Composing and Developing WLI Applications

Naming Standard for WLI Application Artifacts and Variables	4-2
Modular JPD Design	4-3
Modular XML Document	4-4
Parallel Node	4-4
JPD Exceptions	4-4
Event Handlers for Process Events	4-6
JPD Transactions and Compensation Management	4-7
Transaction Boundaries	4-7
Transactions for Synchronous and Asynchronous Processes	4-7
Transactions and Controls	4-7
JPD State Management	4-10
JPD Versioning.	4-11
Singleton JPD	4-12
Race Condition With Dynamic Subscription	4-13
Dead Letter Channel Subscription	4-13
High Quality of Service JPD.	4-13
SLA Threshold for JPDs	4-14
Monitor JPDs	4-15
Security Policy for JPDs	4-16
Interoperable JPDs	4-17
Communication Between JPDs.	4-18

Using Controls	4-19
Using Dynamic Properties and Annotation for Controls	4-20
Buffering Service Control Methods During Asynchronous Calls	4-21
Using Control Factory to Manage Multiple Instances of a Control	4-21
Data Transformation	4-21
Canonical Data Model	4-22
Runtime Selection of a Transformation	4-22
Developing a Task Plan	4-23
Task Plan for Exception Management	4-23
Integrating Custom Logic With a Task Plan	4-23

5. Deploying and Maintaining WLI Applications

Deploying WLI Application During Runtime	5-1
Deploying WLI Application in a Cluster	5-1
Configuring Trading Partner Integration Resources	5-2
Changing Cluster Configurations and Deployment Requests	5-2
Load Balancing in a WLI Cluster	5-3
HTTP Functions in a Cluster	5-3
JMS Functions in a Cluster	5-3
Synchronous Clients and Asynchronous Business Processes	5-3
RDBMS Event Generators	5-3
Application Integration Functions in a Cluster	5-4

6. Core Implementation Patterns for WLI Applications

Core Implementation Patterns for a JPD	6-1
Pattern 1: Basic Synchronous Stateless two-way Service	6-3
Pattern 2: Basic Asynchronous Stateless two-way Service	6-4
Pattern 3: Basic Asynchronous Stateless one-way Service	6-5

Pattern 4: Basic Asynchronous Stateful two-way Service.	6-6
Pattern 5: Basic Asynchronous Stateful one-way Service.	6-7
Pattern 6: Composite Synchronous Stateless two-way Service.	6-8
Pattern 7: Composite Synchronous Stateful two-way Service.	6-9
Pattern 8: Composite Asynchronous Stateless two-way Service.	6-10
Pattern 9: Composite Asynchronous Stateless one-way Service.	6-11
Pattern 10: Composite Asynchronous Stateful two-way Service	6-12
Pattern 11: Composite Asynchronous Stateful one-way Service.	6-13
Other Patterns.	6-13
Course-Grained Process Front-end for a Fine-Grained Process	6-15
Loosely Coupled Process With a Common Message Interface	6-16
Dynamic Property Driven Processes.	6-17

A. Understanding Requirements

Introduction

This document specifies the best practices, lessons learnt, and key considerations which help WebLogic Integration users such as application architects, developers, and operation staff to develop and run high quality WebLogic Integration (WLI) applications.

Note: It is recommended that you read the following documents before you use this document:

- [Introduction to WebLogic Integration](#)
- [Guide to Building Business Processes](#)
- [Tips and Tricks for WebLogic Integration](#)
- [Using the Worklist](#)
- [Using Integration Controls](#)

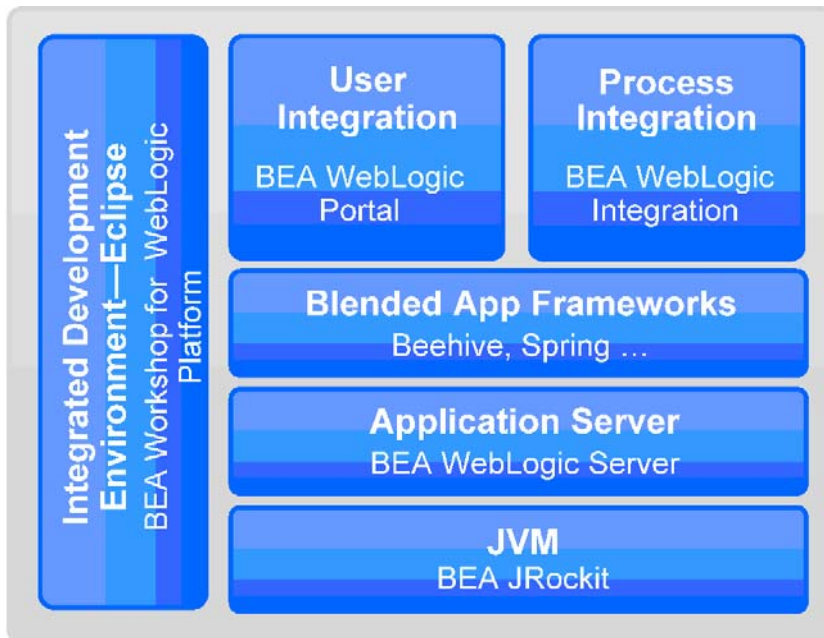
While reading this document, you may have to refer to relevant documents and information about WLI 9.2, available on the [WebLogic Integration documentation](#) page.

This document:

- Explains the basic [Features of WLI](#)
- Provides general best practices to help you with [Understanding Requirements, Architecture and Design of the WLI Application](#), and design of business processes and services.
- Describes best practices with WLI for:
 - [Composing and Developing WLI Applications](#)
 - [Deploying and Maintaining WLI Applications](#)

WebLogic Integration is specially designed for enterprise integration. [Figure 1-1](#) shows the architecture of WebLogic platform.

Figure 1-1 Architecture of WebLogic Platform



Features of WLI

WebLogic Integration provides you with the capability to design, develop, deploy and run integration-centric applications. The main features or capabilities of WebLogic Integration are:

- Creates system (Java Process Definition) and human-centric (Task Plan) business process applications using an Eclipse-based unified IDE.
- Defines system-centric transactional processes that have fine-grained control over the process using JPD graphical editor.
- Defines human-centric, multi-step, task-oriented, and long running business processes using Task Plan graphical editor.
- Uses standard-based Beehive controls to enable easier access to enterprise resources such as EJBs, JMS and web services. In most WLI applications, JPDs, controls, and Task Plans

work with each other in realizing overall business objectives. They can be easily integrated with each other and remain loosely coupled.

- Uses JPD controls and task plans that are property and annotation driven. Properties and annotations can be either static or dynamic. Dynamic properties can be changed during runtime which makes business processes that run with WLI applications agile and flexible.
- Provides extensive support of Business to Business (B2B) capabilities for inter-enterprise business processes using Trading Partner Management, Rosettanet, ebXML, and EDI.
- Creates loosely coupled, publish and subscribe architecture, style based process integration using event generators, message brokers, and event handlers.
- Provides extensive Enterprise Application Integration features such as adapters, and XQuery (2004) based data transformation. The large number of integration controls such as service and process controls make the task of integrating enterprise resources easier.
- Supports enterprise computing services support such as transaction management, clustering, security, and J2EE container services.

Figure 1-2 shows the functional overview of WLI.

Figure 1-2 Functional Overview of WLI



For more information on the features and functions of WLI, see [Introduction to WebLogic Integration](#).

Introduction

Understanding Requirements

The key to successful software development is that all stake holders develop a clear and uniform understanding of application requirements.

Software requirements can be broadly classified into two groups:

- Functional or problem domain requirements
- Non-functional or solution domain requirements

In a problem domain, the focus is on the functional or business requirements. It is recommended that you create a domain model of your functional requirements before you start thinking of the solution domain.

In a solution domain, we focus on how to deliver the solution for functional or business requirements.

Some of the important non-functional requirements of a WLI application are:

- Quality attributes such as security, high availability, scalability, performance, and reliability.
- User interface
- Integration
- Message format, transport, and protocol
- Data format and transformation
- Internationalization or i18n

Understanding Requirements

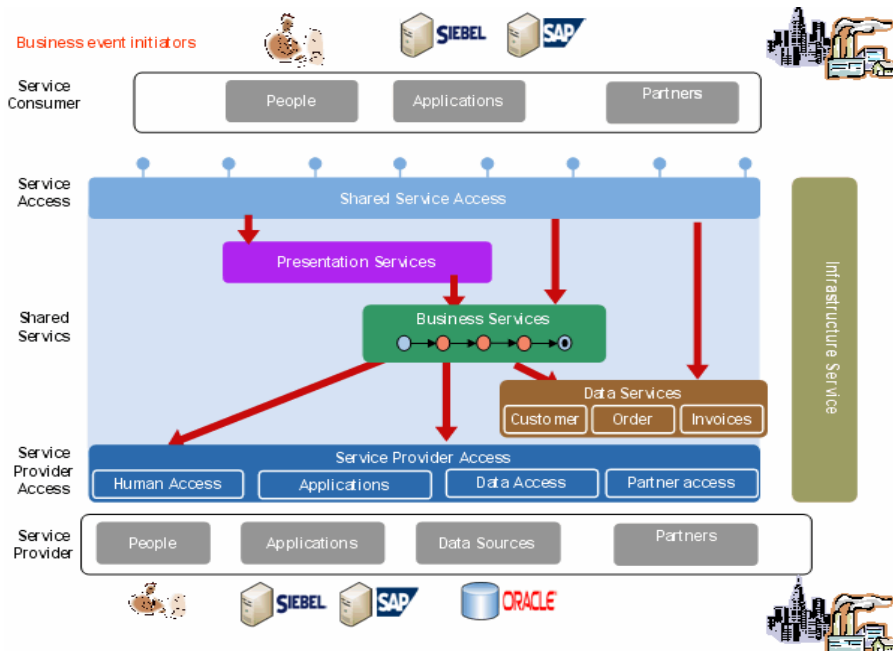
- Legal and compliance
- Runtime infrastructure
- Networking and communication
- Constraints such as the use of specific RDBMS System, protocols, and standards

For general guidelines on understanding requirements, see [Understanding Requirements](#).

Architecture and Design of the WLI Application

WLI is an enterprise class product, which can be used to implement process-driven Service Oriented Architecture (SOA) applications. [Figure 3-1](#) shows the BEA reference architecture for service oriented applications.

Figure 3-1 BEA Reference Architecture for Service Oriented Applications



It is recommended that you use BEA reference architecture for service oriented applications for a WLI application.

The reference architecture is generic and can be implemented with WLI standalone or with a combination of several products. The main components of the reference architecture are as follows:

- **Business event initiators:** The entities that initiate business actions or events. Business initiators are either human, or based on systems and applications.
- **Service consumers or composite applications:** The applications that are developed while using WLI. They handle business actions or events initiated by business entities.
- **Shared service access layer:** Provides access to shared business services. It is based on Validate, Enrich, Transform, Route, and Operate or invoke (VETRO) patterns. This layer can be implemented using WLI or an Enterprise Service Bus (ESB) such as AquaLogic Service Bus.
- **Shared services layer:** The shared and reusable services that are used in service orchestration while creating business processes.

The types of shared services are as follows:

- Presentation services that present the data to the user.
- Business services that represent core business capabilities. Business services can range from relatively simple to very complex. An example of a simple business service is credit card validation. An example of a complex business service is a cross-functional, inter-enterprise business process such as order fulfillment. Business services are task or activity-oriented. You can implement task-oriented services such as purchase order approval using task plans, and system-centric transactional services using JPDs. Most complex business services require both JPDs and task plans. JPDs and task plans can work independently and also be loosely coupled as required.
- Data services that are entity services which provide access to enterprise data. Data services have a Validate user, Create, Retrieve, Update, and Delete (CRUD) interface and can be implemented using WLI components or special data service enabling products such as AquaLogic Data Services Platform.
- Infrastructure services are non-functional services such as security, and audit.
- **Common service provider access layer:** Provides a common access to service providers. This layer can be implemented using WLI or an ESB. A service provider provides services and can be either human or based on systems and applications.

Modeling Business Processes and Services

The steps involved in modeling business processes and services are:

- [Defining Business Processes](#)
- [Designing and Developing Services](#)

Defining Business Processes

You can define a good business process by:

- [Identifying Process Objectives and Goals](#)
- [Identifying Key Performance Indicators for the Process](#)
- [Identifying Process Actors or Participants](#)
- [Identifying Public and Private Processes](#)
- [Identifying Initiator and Participant Processes](#)
- [Keeping Processes Modular](#)
- [Enabling end-to-end, Cross-functional Processes](#)
- [Separating Production and Monitoring Processes](#)

Identifying Process Objectives and Goals

It is important to understand the main objectives behind the decision to automate a business process. A business process is defined to:

- Improve operational efficiency and productivity
- Gain better information visibility and insight into business operations
- Improve customer service
- Improve human collaboration.

A few examples of process objectives are:

- An account receivable and payable process to improve visibility.
- A purchase order processing and fulfillment process to improve operational efficiency and productivity.

- A document approval process to improve human collaboration.
- A complaint resolution process in a call center for improving customer service and operational efficiency.

Identifying Key Performance Indicators for the Process

KPIs are the key indicators that provide the visibility for process performance. A few sample process KPIs are:

- The visibility of Days of Sales Outstanding (DSO) for an Accounts Receivable process.
- The time taken from order to delivery (efficiency), or the number of purchase orders completed in a day (productivity) for Purchase Order processing and fulfillment.

Identifying Process Actors or Participants

Identify process actors or participants before you define the process. In addition to human actors and applications, data sources and partners can also serve as actors or participants in a business process.

Identifying Public and Private Processes

Public and private processes have a special significance in the B2B domain. A public process is one that is visible to the external user. A private process is one that performs a task that is invisible to the external user within the application in the background. The private process is linked to the public process. When a public process is used as a front end for a private process, you can modify the private process without disturbing your public process clients.

It is a good practice to keep your public process modular. Public processes have higher requirements in terms of security, scalability and availability. Public processes always have special requirements for messaging, protocols, and standards, which need to be identified in advance. [Figure 3-2](#) illustrates a public and private process pattern.

One way to integrate public and private processes is to use loosely coupled publish and subscribe architecture between these processes using the message broker component of WLI.

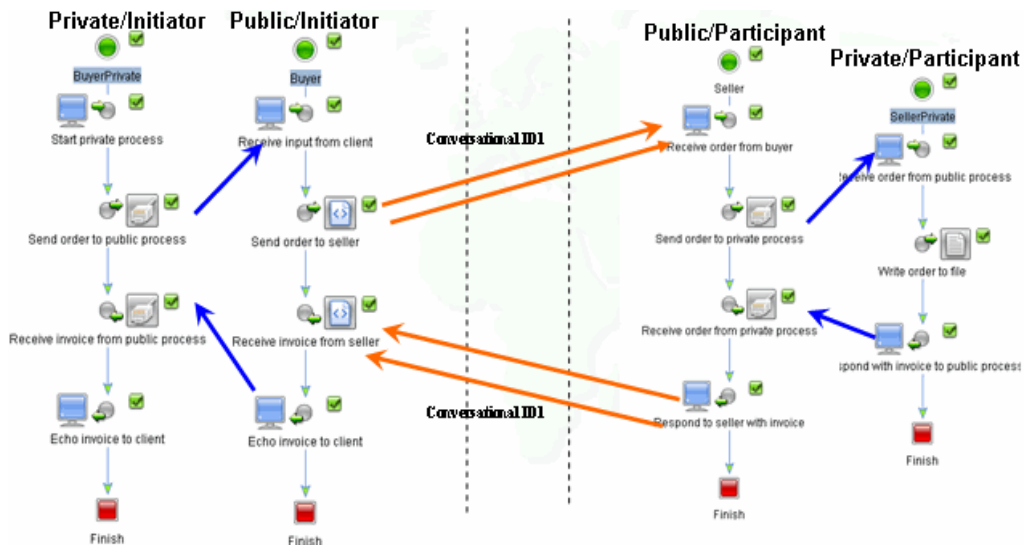
Another alternative to integrate public and private processes, is to have a process control between public and private processes, and use Java Web Services (JWS) as the front end for your public processes. Process Control (PC) is a better option if your public and private processes are in the same domain. See the [Communication Between JPDs](#) section for more details.

Identifying Initiator and Participant Processes

Many business processes, especially in the B2B domain, are conversational in nature. For example, a request quote service, where an agreement on the quoted price is achieved after several rounds of message exchange. An initiator uses an initiator process to initiate a business conversation. Participant processes respond to the request from an initiator process. An initiator process:

- Can have multiple conversations with a single participant process at the same time.
- Can also interact with multiple participant processes at the same time.
- Generates a unique conversation ID for each process. All the related participant processes, include this conversation ID in their response messages. Conversation IDs act as correlation identifiers across a conversation. Figure 3-2 shows a sample implementation of Public, Private, and Conversational patterns.

Figure 3-2 Public, Private, and Conversational Patterns



Keeping Processes Modular

It is a good practice to define small and modular processes. A modular process meets a specific business need. Modular process definition results in the reuse of business processes. It is easier to maintain many modular processes than a few long processes.

Enabling end-to-end, Cross-functional Processes

You can get the best results from Business Process Management (BPM) when you enable end-to-end cross-functional processes, to realize a specific business objective.

A few examples of end-to-end, cross-functional processes are as follows:

- Order to cash
- Request to quote
- Complaint to resolution

The ideal way to enable end-to-end processes is using loosely coupled integration between a number of modular business processes.

Separating Production and Monitoring Processes

It is always a good practice to separate your core production processes from your monitoring processes. Production processes have higher scalability, performance, and availability requirements than monitoring processes. An example of a core production process is the order fulfillment process. An example of monitoring processes are audit, compliance, and KPI calculation processes.

You can have loosely coupled integration between production and monitoring processes using publish and subscribe architecture with Message Broker. For more information, see [Design Loosely Coupled Services](#).

Note: It is a good practice to capture monitoring data from a production process and publish it to specific channels, which the monitoring process can then subscribe to. Subsequently, a monitoring process can process the monitoring data as per business needs.

Designing and Developing Services

The steps that you need to follow when you design and develop services are:

- [Service Classification](#)
- [Identify Services](#)
- [Define Service Contract](#)
- [Define Input and Output Service Messages](#)
- [Define Pre and Post Conditions for Services](#)

- Decide the Service Calling Paradigm
- Decide the Service Granularity
- Define Quality of Service Requirements
- Design Service With a Service Proxy
- Design Reusable Services
- Design Loosely Coupled Services

Service Classification

Table 3-1 lists the services that can be broadly classified into categories.

Table 3-1 Service Classification

Category	Service Category as per Reference Architecture	Characteristic	Example
Task-oriented	Business services	Long running and human-centric	Document-oriented such as Purchase Order approval
Activity	Business services	System-centric	<ul style="list-style-type: none"> • Purchase Order fulfillment • Credit card verification service
Entity service	Data services	CRUD service for business entities	Retrieve customer address

Identify Services

The first step in service design is to identify the candidate services. You can identify business services based upon business events. Table 3-2 contains a few examples of business events and related services.

Table 3-2 Sample Business Events and Related Services

Business Events	Services
Purchase Order is received	Purchase Order processing service
Purchase Order is validated	Purchase Order fulfillment service

Table 3-2 Sample Business Events and Related Services

Invoice is received	Invoice processing service
Purchase Order amendment request is received	Purchase Order amendment service

If you identify a service with its related business event, you have event-driven services, which can be easily orchestrated into long running business processes.

Define Service Contract

Once you have identified a service, the next step is to define a service contract. The service contract specifies the functions that a service provides in a format that the service consumer easily understands. A service contract defines how a service can be used and specifies the Quality of Service (QoS) parameters for a service. However, a service contract does not contain any implementation level details.

Define Input and Output Service Messages

After defining the service contract, the next most important step is to define service messages. A service receives an input message from service consumer and may or may not return a message to the consumer. Define an appropriate service message schema for each service message.

Define Pre and Post Conditions for Services

The next step is to define pre and post conditions for each service. Pre-condition represents the condition which should be satisfied, before a service is invoked. The service consumer has to ensure that the service is invoked, only when the pre-condition is satisfied. The service provider has to ensure that post conditions are satisfied after the service invocation is completed. If you develop your services using this method, you can perform better exception management for services.

Decide the Service Calling Paradigm

You need to decide whether the service is invoked in a synchronous or asynchronous mode. Small and data oriented services such as obtaining a customer address, can be designed in a synchronous mode. Long running services that require state management, should be designed as asynchronous services. Asynchronous services are more loosely coupled and scalable compared to synchronous services.

Decide the Service Granularity

Services should be more course-grained in comparison to API calls. You should decide on the actual granularity based upon your specific situation. A service call involves a round trip on the network. For performance reasons, you should try to minimize round trips on the network. In case you have too many fine-grained services, it is good practice to create a course-grained service by performing a light weight orchestration of the fine-grained services. Subsequently, you can expose the course-grained services to the user.

Define Quality of Service Requirements

Services are required to meet minimum quality requirements in terms of availability and performance. Service providers are required to have a Service Level Agreement (SLA) with service consumers. You should understand and specify QoS requirements well in advance.

One way to improve your QoS is to design your service as idempotent. Idempotent services do not change the state of the system, even if the service is invoked by the same input message several times. With idempotent services, you can easily use standards such as web service reliable messaging, which can re-transmit a message in case of failure. All read only services are idempotent.

However, even write services can be designed as idempotent. This design ensures that your services are more reliable and improves QoS. For example, when a service receives an invoice from a supplier, the invoice amount should be added to the total accounts receivable amount. In such a scenario, you can first assign a unique number to the invoice in the database and then increase the account payable amount. If the same message is sent again, the uniqueness constraint of the invoice number throws a database exception and discards the duplicate message.

Design Service With a Service Proxy

If you design your service with a proxy service as the front end, you can obtain several advantages using a proxy service as your main service. A proxy service can be used for authentication, enrichment, transformation, and versioning. The use of proxy services establishes a loose coupling between service providers and consumers.

Design Reusable Services

You should identify business functions that are reusable across different domains or departments. Such functions should be prime candidates for being exposed as standard based services.

Design Loosely Coupled Services

Services should be designed as loosely coupled. There should be minimum coupling between service consumers and service providers. The best practices to design loosely coupled services are as follows:

- **Message format and protocol coupling:** Use a proxy as a front end for your service to reduce the message format and protocol coupling.
- **Time coupling:** Design asynchronous services and reduce time coupling. You can also reduce time coupling by making your services highly available (7/24/365).
- **Type coupling:** Design document-oriented services to reduce type coupling. Document-oriented services are more loosely coupled than Remote Procedure Call (RPC) styled services. In an RPC styled service, the client makes a method call to a service. The client has to know the method name and data type of the method input as well as the output parameters. This leads to tight coupling between the service consumer and provider.

In the case of document-oriented services, the service consumer interacts with the service by sending documents that are meant to be processed as complete entities. These documents are written in XML, and defined in a commonly agreed upon schema between the service provider and consumer. For example, a document-oriented service is a purchase order processing service which receives a purchase order (in XML form), processes it, and returns a purchase order confirmation (XML document) to the client. JPDs are designed to be document-oriented. A JPD uses a document and processes it as an entity.

- **Location coupling:** You can reduce location coupling by making your service available. You can publish your service WSDL and meta data to a registry or repository. The service consumer can then discover the service from the registry and then invoke the service as required.
- **Version Coupling:** You can reduce version coupling by making your service backward compatible. Existing service consumers should not be affected if there is a new version of the service. If you use a proxy service as a front end, it receives the initial request and then directs it to the appropriate version of the service.

Composing and Developing WLI Applications

There are several best practices for composing and developing WLI applications. They are described in the following sections:

- [Naming Standard for WLI Application Artifacts and Variables](#)
- [Modular JPD Design](#)
- [Modular XML Document](#)
- [Parallel Node](#)
- [JPD Exceptions](#)
- [Event Handlers for Process Events](#)
- [JPD Transactions and Compensation Management](#)
- [JPD State Management](#)
- [Singleton JPD](#)
- [Race Condition With Dynamic Subscription](#)
- [Dead Letter Channel Subscription](#)
- [High Quality of Service JPD](#)
- [SLA Threshold for JPDs](#)
- [Monitor JPDs](#)

- [Security Policy for JPDs](#)
- [Interoperable JPDs](#)
- [Communication Between JPDs](#)
- [Using Controls](#)
- [Data Transformation](#)
- [Canonical Data Model](#)
- [Runtime Selection of a Transformation](#)
- [Developing a Task Plan](#)

For more information on this section, see [Guide to Building Business Processes](#) and [Using WebLogic Integration Controls](#).

Naming Standard for WLI Application Artifacts and Variables

When you develop WLI applications, you should follow a uniform and consistent naming standard for various WLI artifacts such as JPDs, task plans, controls, projects, and files. All controls, processes, methods, variables, and other WLI object names should follow standard object oriented and Java naming standards.

In WLI version 9.2, JPD, control, and data-transformation files have common .java extensions. You should add a suffix to each artifact so that they can be easily identified. [Table 4-1](#) contains an example of naming controls and the related suffix for each artifact.

Table 4-1 Example of Naming Controls

Type of Artifact	Suffix	Example
JPD projects	EAR, JPD	<Name>earjpd
	Web, JPD	<Name>webjpd
	Util, JPD	<Name>utiljpd
Task plan projects	EAR, JPD	<Name>eartp
	Web, JPD	<Name>webtp
	Util, JPD	<Name>utiltp

Table 4-1 Example of Naming Controls (Continued)

Type of Artifact	Suffix	Example
JPD	JPD	<name>jpd.java
Data transformation	DTF	<name>dtf.java
Custom control	None	None
Database control	DBC	<name>DBC.java
Web service	WSC	<name>WSC.java
EJB control	EJBC	<name>EJBC.java
JMS	JMSC	<name>JMSC.java
E-mail	EmailC	<name>EmailC.java
File	FileC	<name>FileC.java
HTTP	HTTPC	<name>HTTPC.java
MQ Series control	MQC	<name>MQC.java
Process control	PControl	<name>PControl.java
Task control	TASKC	<name>TASKC.java
Channels	Channel	<name>Channel.java
Event generators	EG	<name>EG

Modular JPD Design

If you have to define a large process, it is always a good practice to make it modular. A modular JPD meets a specific business objective. A few examples of modular JPDs are:

- Process purchase order
- Fulfill purchase order

You can divide a large JPD into number of smaller JPDs or sub-processes. Sub-processes can be invoked using a central or main JPD. You can use process control to communicate between a JPD and a sub-process (another JPD). You can also loosely couple two JPDs using Message Broker

or publish and subscribe architecture. See [Communication Between JPDs](#) and [Core Implementation Patterns for a JPD](#) for more details.

Modular XML Document

Every JPD has a start node and uses a XML document to start. Ensure that the XML document is modular and small for better performance. It is a good practice to have separate XSDs for each type of document. For example, you should have a separate XSDs for a Purchase Order and Invoice, instead of one single large schema for both the Purchase Order and the Invoice.

The modularity of a JPD and its associated document is a relative concept. If a JPD is too modular, it may be counter productive. You should take a balanced decision on modularity depending upon the specific scenario.

Parallel Node

The JPD provides a parallel node. Before you use this node, it is recommended that you understand how this node works. Parallel nodes can be used for managing multiple tasks which are related but not dependent. Parallel nodes are meant for business level parallelism. The actual execution does not take place in parallel.

At any given time, execution can start in one branch without any execution in other branches. For example, the branch executing a task may request for a quote from a partner and wait for a response. While this branch is waiting for a response, execution may start in another branch thus achieving business level parallelism.

There are two forms of parallel nodes – AND and OR mode. In the AND mode, all parallel paths should complete their execution, before the business process can proceed further. In the OR mode, the path that completes execution first is the winner. The processing of other path is stopped and the business process proceeds further.

The branches of a parallel node are isolated by transaction (default behavior). You can override this behavior to improve the performance of parallel nodes. Set the `continueTransaction` property to `true` in the source for each parallel element as follows:

```
<parallel continueTransaction="true">
```

JPD Exceptions

You can define exceptions and exception handlers in JPDs at nodes, groups, or process-levels. Exception handlers are executed in the following order:

1. Nodes
2. Groups
3. Processes

It is a good practice at the process-level to have a **catch all** process-level exception handler.

Exceptions that are not handled can cause process failure. You can set a **freezeOnfailure** property to avoid this scenario. If this property is set, any exceptions that are not handled cause the process to freeze. The process then rolls back to the last committed point and the administrator can restart the process from the last committed state.

[Table 4-2](#) contains a list of high-level design guidelines which you can follow while working on a JPD exception.

Table 4-2 High-Level Design Guidelines

High-Level Design Guideline	Description
Asynchronous two-way processes should establish a way to pass exceptions or errors back to the caller process, with a separate client response or publication node in the exception handler.	When an error or an exception occurs, it is possible to catch an exception within the process. However, asynchronous processes, by nature cannot immediately respond to the exception. However, the caller should be informed that the process did not complete successfully. Your callback process should have a success or failure path and the result in both cases should be communicated to the caller.
All synchronous stateless processes must throw an exception back to the caller.	If the caller is blocking, it must be notified of the failure, as it is the only component that understands what should be done next.
Process designers should define a process-level exception handler at the start of every process definition. This handler acts as a global exception handler and catches any undefined exceptions.	This is a catch-all process-level exception handler. Exceptions that are not handled cause the process to fail and no recovery is possible from this state. When you define a global exception handler, you can ensure that the process never goes to an aborted state.

Table 4-2 High-Level Design Guidelines (Continued)

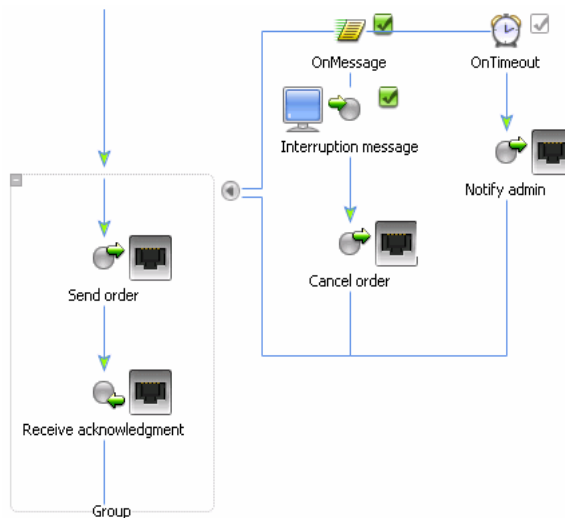
High-Level Design Guideline	Description
Set the freezeOnFailure property to true	If you set the freezeOnFailure property on the process, the process is rolled back to last commit state and persisted. The administrator can then fix the problem and re-activate the process.
Process designers should be especially careful about handling exceptions when processes call sub processes synchronously through the process control.	Exceptions that are not handled in the called sub process can set transactions to roll back. In this case, both the sub-process and caller process can be rolled back. You can define an appropriate exception handler in the sub process to avoid the roll back of the caller process.

Event Handlers for Process Events

Events handlers allow outside events to interrupt the process using **OnMessage** and **OnTimeout** event handlers. **OnMessage** event handlers can accept client requests. When a timer event is fired, control, receive, and message broker subscription message events are invoked.

Figure 4-1 shows an example of an **OnMessage** and **OnTimeout** event handler.

Figure 4-1 Example of OnMessage and OnTimeout Event Handler



JPD Transactions and Compensation Management

The following sections describe the best practices for various JPD transactions and compensation management:

- [Transaction Boundaries](#)
- [Transactions for Synchronous and Asynchronous Processes](#)
- [Transactions and Controls](#)
- [JPD Versioning](#)
- [Using Dynamic Properties and Annotation for Controls](#)
- [Buffering Service Control Methods During Asynchronous Calls](#)
- [Using Control Factory to Manage Multiple Instances of a Control](#)

Transaction Boundaries

Processes in WLI are transactional in nature. Every step of a process is executed within the context of a Java Transaction API (JTA) transaction. When you are building a process, implicit transaction boundaries are formed based upon the location of blocking elements such as a Control Receive or a Client Send. As you add process nodes, the transaction boundaries within a process keep changing.

You can also create explicit transaction boundaries. To do this, select contiguous nodes and declare them in a separate transaction to distinguish between them and the implicit nodes that the application creates. The transaction may also contain resources accessed by a process, depending on the nature of the resource and the control that provides the access.

Transactions for Synchronous and Asynchronous Processes

A stateless process is executed either in a client transaction or when a new transaction is started. Using JPD proxy, a Java client can invoke a JPD over RMI. In this scenario, the client transactions are propagated to the JPD. The caller transaction is not propagated when a JPD is invoked as a web service. The caller transaction is not propagated to a JPD for asynchronous processes as well.

Transactions and Controls

Transaction controls are of three types:

- [Transactional and XA Compliant](#)
- [Transactional and not XA Compliant Controls](#)
- [Non-transactional Controls](#)

Transactional and XA Compliant

If all controls used within a process are transactional and XA compliant, then the transaction of the process can be used to commit or terminate the underlying transaction branches. The process needs exception handlers to catch any issues and make the necessary transaction decisions. The developer needs to be aware where the transaction was started as any abort or rollback takes the control back to the starting point and this may not be within the process where the exception occurred.

Transactional and not XA Compliant Controls

XA is a protocol used to manage distributed transactions. WLI extends XA to allow non-XA resources to participate in distributed transactions, with the limitation that in a given transaction, only one transactional resource can be non-XA compliant. Therefore, if more than one transactional non-XA resource needs to be accessed in a process, then the access to these resources should be encapsulated in separate JPD sub-processes, which should be called asynchronously by the original process. Asynchronous invocation is necessary because synchronously called subprocesses run in the same transaction as the calling process. See [Using Integration Controls](#) for more details.

Non-transactional Controls

Non-transactional controls such as email controls do not support transactions. In the case of non-transactional controls, you need to have a strategy for exception handling. Use automatic rollback where controls are transactional. Use a caught exception for non-transactional controls to handle the error based on the business problem.

In addition to providing support for Atomicity Consistency Isolation and Durability (ACID) and XA transaction, JPDs also provide support for compensation. You should ensure that the transaction block for which you want to provide compensation, is not marked for **rollback only**. Define an exception handler path for the transaction block and enable the **execute on rollback** exception handler property. In such a situation, when a transaction fails, the exception handler path is executed first. You can define your undo or compensation logic in such a path. [Figure 4-2](#) shows an example of non-transactional controls.

Figure 4-2 Non-transactional Controls

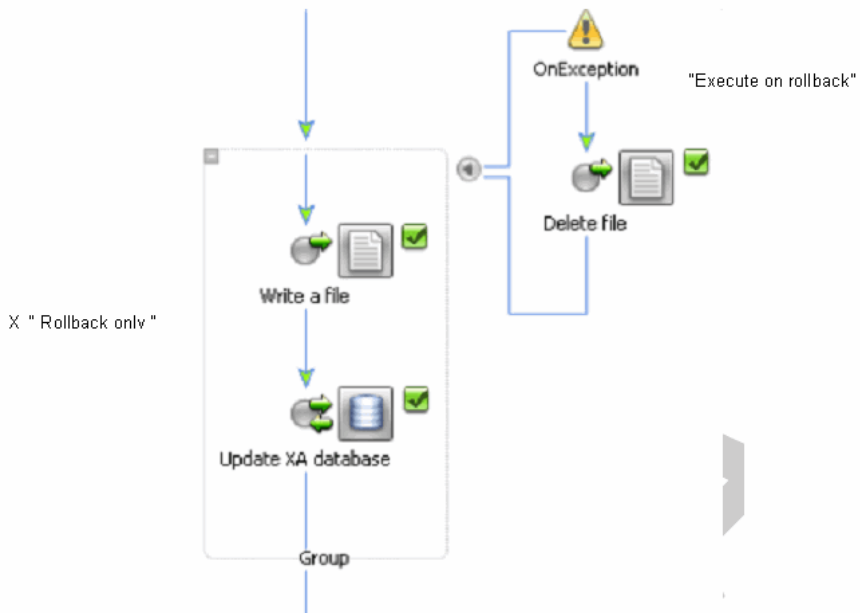


Table 4-3 contains a list of high-level guidelines to decide the transaction characteristic for your JPD.

Table 4-3 High-Level Design Guidelines for JPD Transaction Characteristics

High-Level Design Guidelines	Description
Implement the automatic execution of compensating transactions only with extreme caution.	At a design level it seems very simple to implement the reverse. The problems arise when the compensation fails, the question being the recovery from the failure. The solution to this situation is often unclear. The default design rule should be to raise an alert and manually decide on a solution.

Table 4-3 High-Level Design Guidelines for JPD Transaction Characteristics (Continued)

High-Level Design Guidelines	Description
A process needs to have exception handlers in place to find any issues and make the necessary transaction decisions.	You must not leave anything to default, always identify the exceptions and decide on a solution. This is applicable to transactional and non-transactional resources.
When more than one transactional non-XA resource has to be accessed in a process, the access to these resources should be encapsulated in separate JPD sub-processes.	The original process should call the sub-processes asynchronously. In this way, the sub-process runs in a separate transaction and is able to access the non-transactional resource.

JPD State Management

A stateless JPD is a process executed in memory only. Its state does not persist. All stateless processes are compiled into a stateless session bean. Stateless processes are intended to support business scenarios that involve short-running logic and have high performance requirements.

As the JPD does not persist its state to a database, it is optimized for lower-latency and higher-performance execution.

[Table 4-4](#) contains a list of high-level guidelines that you should follow, when working on stateless processes.

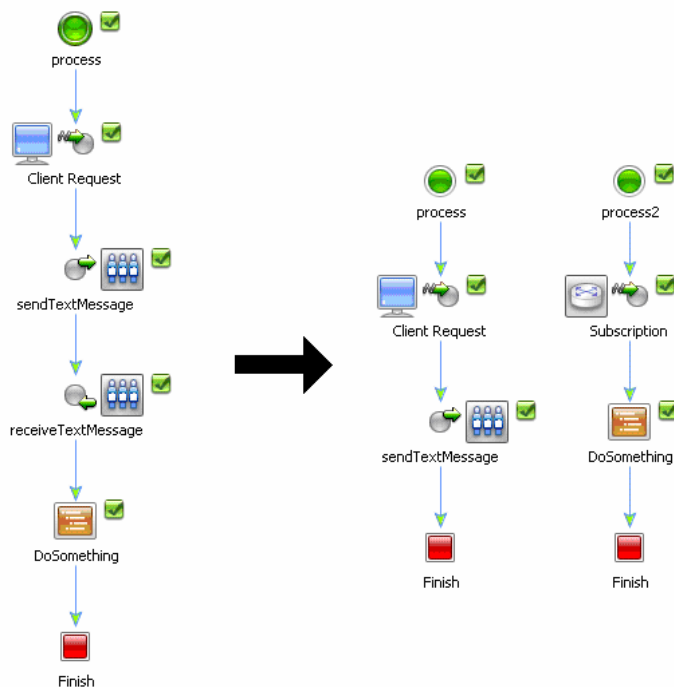
Table 4-4 High-Level Guidelines for Stateless Processes

High-Level Design Guidelines	Description
Do not use default values in global variables, unless the variable is either declared static or final. Initialize all global variables before use.	Stateless processes are implemented as stateless session beans. After a process is complete, subsequent process instances reuse the same stateless session bean instances, and therefore inherit the last known value of the global variables.
Set the on sync failure property on the process to re-throw for synchronously called processes where the requestor needs to handle transaction demarcation.	This property only applies to your process if it is configured to be a synchronous sub-process; it is ignored for any other business processes. If a synchronous sub-process fails, the default behavior is to mark it as rollback , which causes both the sub-process and the parent process to rollback. However, if the on sync failure property is set to re-throw , only the sub-process is rolled back.

A stateful process is a process that runs within the scope of more than one transaction. The process persists its state in the database. The state of the JPD survives even if the server crashes. The stateful JPD process is compiled into an entity bean. Stateful processes are intended to support business scenarios that involve complex, and long-running logic.

Stateful processes, in general, are slower than stateless processes. Use stateless processes, especially in scenarios where a state does not need to persist. In certain situations, you can split a stateful process into several stateless processes. [Figure 4-3](#) shows how you can split a stateful process into a stateless process.

Figure 4-3 Splitting a Stateful Process Into a Stateless Processes



JPD Versioning

WLI has a version feature that helps you change your business process without interrupting any instances of the process that are currently running. When you create a version of a business process, you are actually creating a child version of a business process that shares the same public interface as the parent business process. At runtime, the version of the process that is marked

active, is the process that external clients access using the public URI. Through the regular development cycle, new process versions are deployed with new versions of the application.

When the new version of a JPD is deployed, the existing instances run to completion on the same version of the JPD that they started with. You can version business processes, but not the individual controls associated with that process, or other business process related components such as schemas and transformations. When you version a business process, you must also version the sub-processes of that process, as they are not assigned a version automatically along with their parent process.

Table 4-5 contains a list of high-level guidelines that you can follow to set versions for JPDs.

Table 4-5 High-Level Guidelines to Version JPDs

High-Level Design Guidelines	Description
If you have to deploy a new version of a business process, you must set a version for your existing process before deploying the new version.	If you do not follow that process, you must let non-versioned instances run to completion before deploying the new versioned process.
Set the process version strategy according to the parent-child relationship of the business process.	<p>This describes how you can invoke sub-processes when different versions of the parent process exist. From the strategy drop-down menu:</p> <ul style="list-style-type: none">• Select loosely-coupled if you want the sub-process version to be set when the sub-process is invoked.• Select tightly-coupled if you want the sub-process version to be set at the time the parent process is invoked.

Singleton JPD

A JPD can subscribe to a message broker channel in two ways: Static or Dynamic.

If a process subscribes to a channel at the start node, this is called a static subscription. The subscription is known at the time the application is compiled and remains through the life time of the application.

When a process subscribes to a message broker channel during its flow using the message broker control, this is known as a dynamic subscription. The subscription starts and ends at runtime. A static subscription to a message broker channel can be specified as suppressible if you set the suppressible attribute for the subscription in the business process. The accepted values for the suppressible attribute are true and false (false is the default value).

A singleton JPD has only one instance of the JPD running at any time. To create a singleton JPD, first define a JPD with a static subscription and set the suppressible attribute to True. In this situation, the first message on the channel invokes the JPD and creates an instance of the JPD. Subsequent messages on the static channel do not lead to the creation of a new JPD instance. Singleton JPDs created in this way can continue to receive messages using dynamic subscription.

Race Condition With Dynamic Subscription

A race condition is possible when you use the message broker with a dynamic subscription. A message is lost if it is sent before the subscription is complete. This problem is evident when processes start waiting indefinitely for a response, coupled with messages appearing in the dead-letter channel.

If the request message is non-transactional, for example, in the case of a web service call, the message is sent immediately, whereas the subscription to the response occurs only when the transaction is committed. If the subscription appears before the message is sent in the process flow, it might occur later. In this case, ensure that the transaction is committed (for example, add an empty explicit transaction) after the subscription node and before the request message is sent.

Dead Letter Channel Subscription

If a message is sent to a channel without a subscriber after the filtering process is complete, the message goes to the dead-letter channel. Subscribe to the dead-letter channel to check if messages are published there, because in most cases, this is not a desired behavior.

High Quality of Service JPD

You need to take special steps to implement **at-least-once** or **once-only** Quality of Service. If these steps are not taken, the default is the **at-least-once** Quality of Service. The steps for the different process types are as follows:

- Invoking a Synchronous JPD: The quality of service is the responsibility of the caller.
- Invoking an Asynchronous JPD with JMS: When you use JMS as a calling mechanism to a process, the steps to achieve high Quality of Service are:
 - Configure re-delivery limit and re-delivery delay override for the associated JMS queue
 - JMS queues persist
 - JMS connection factories are transactional
 - JMS queue has an error queue configured

- Any exception that was not handled must be resent to the caller using the call back process. This step ensures that the message is delivered to the process. If an exception is not handled, the message is stored in the error queue and can be recovered.

Note: You cannot ensure the successful completion of every computer program, but a high Quality of Service means that the error is always recoverable.

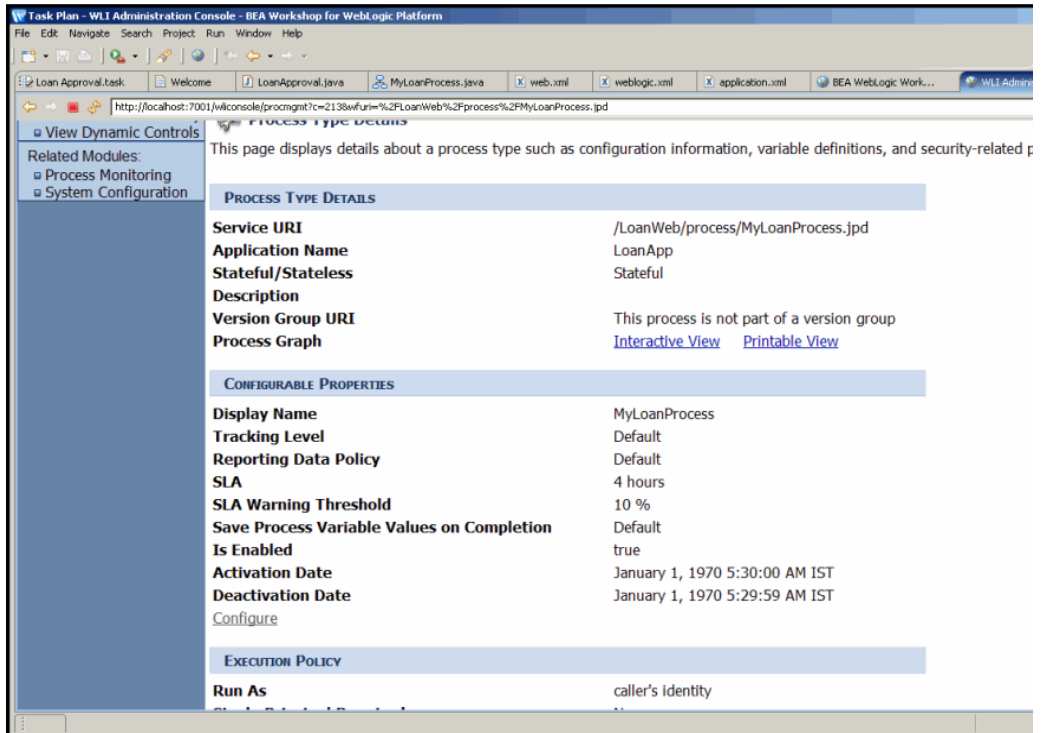
- The **once-only** Quality of Service cannot be implemented for event generator synchronous services such as files and e-mail that have source events that are not transactional.

SLA Threshold for JPDs

A service level agreement (SLA) specifies the performance target for a JPD. An internal or external commitment shows that a JPD is executed within a specified period of time. To help you achieve the SLA for a process, the WLI Administration Console allows you to set the following thresholds ([Figure 4-4](#)):

- SLA threshold, that represents the commitment applicable to the process type. For example, number of seconds, minutes, hours, or days.
- SLA warning threshold, which is a percentage of the total SLA.

Figure 4-4 SLA Threshold Details



The process status that is relative to these thresholds is tracked for each process instance as follows:

When the elapsed time for a process instance reaches the warning threshold, a warning is displayed on the Process Instance Summary and Detail pages. The amount of time remaining until the SLA threshold is reached is also displayed. When the elapsed time exceeds the SLA, a red flag is displayed. The time limit by which the SLA threshold has been exceeded is also displayed. This ability to set SLA thresholds allows you to easily identify processes that do not execute within the target time frame. You can then make the required changes to meet agreements between suppliers and customers, or to achieve your own performance goals.

Monitor JPDs

You can track a process at various levels. The system contains a default tracking level and then each process can override this. The WLI Administration Console and underlying MBeans

provide a monitoring interface to running instances and their variable values. Although variable values cannot be changed, it is possible to query the process using the instance ID or process label. You can set the process label in the instance by calling the JPD context `setProcessLabel`.

[Table 4-6](#) contains a list of high-level design guidelines to monitor JPDs.

Table 4-6 High-Level Design Guidelines for Monitoring JPDs

High-Level Design Guidelines	Description
Use tracking data only for operational support, and not for business or audit logging.	You need to purge tracking data regularly. Tracking can be set on and off dynamically for each process at runtime using the WLI Administration Console.
Disable process tracking for processes that require high performance.	Process tracking involves information that is written into the database several times. For performance reasons, tracking must be eliminated.
WLI system data should only be used for support and not business level audit.	The WLI system data should not be used for business level audit. A designed audit or logging framework must be implemented outside the WLI system data capture area.
Log every InstanceID.	You must record the InstanceID for any error, log, and audit message.
Schedule regular archiving of WLI data.	Archiver is a process that runs a Select query on the database, so it must be scheduled to run regularly on small amounts of data. Avoid scheduling the archiving process at peak hours.
Set the process label to relevant query values.	The process label is a preferred query string to design the values, rather than depending on the person who implements the values. For example, a support person finds an order number easily if it is represented on the process label.

Security Policy for JPDs

You can define the security policy for a JPD. The security policy controls the identity that the JPD uses to access external or backend systems. It allows the administrator to configure whether a JPD accesses an external system as the invoking application, or as an application that calls into

the process later. For example, if a process subscribes to a channel and then waits for a client request, the administrator can set the execution policy and use the identity from the client request while accessing backend resources.

The JPD security policy has four main components:

- **Process Execution Policy:** The execution policy specifies whether the operations in the process are run with start user or the caller's identity.
 - If the administrator specifies start user, each operation assumes the identity of the user that started the process.
 - If the administrator specifies caller's ID, the operation after the call assumes the identity of the calling component.

The policy also ascertains if a single principal is required or not. If a single principal is required, all incoming client requests must come from the same user.

- **Process authorization policy:** The administrator can configure the roles authorized to invoke the process method or client request. All methods in the process inherit the roles specified in the process authorization policy. If the process authorization policy is not defined, everyone is authorized.
- **Method authorization policy:** The administrator can configure the roles authorized to invoke the process methods or client requests. All methods inherit the roles specified in the process authorization policy. You can also add roles to the authorization policy for a method.
- **Callback authorization policy:** The administrator can configure the roles authorized to invoke the process callback. If the callback authorization policy is not defined, everyone is authorized to make callbacks.

Interoperable JPDs

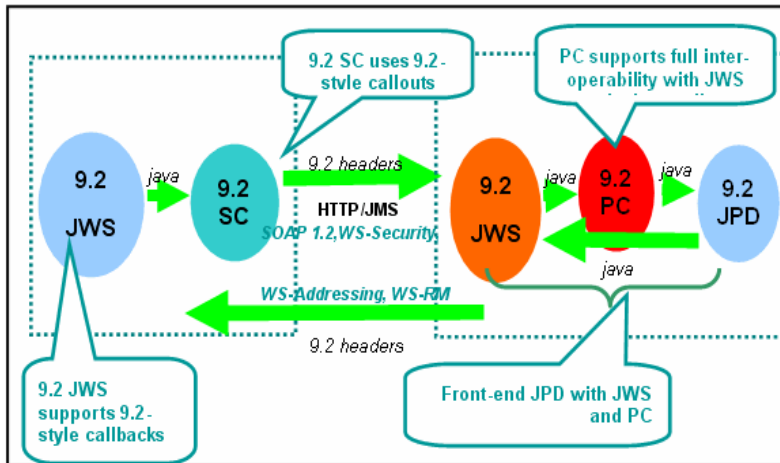
You can expose a JPD as a Java Web Service (JWS). There are limitations to developing interoperable JPDs, which include:

- JWS supports web service standards such as WS-Security and WS-Addressing.
- WebLogic Server 9.2 JWS does not support WLI 8.1 JPD style callbacks.
- WLI 9.2 JPDs support only WLI 8.1 conversations.
- WLS 9.2 Service Control (SC) can invoke WLI 8.1 JPDs.

- Process Control (PC) supports full interoperability with JWS via Java calls over RMI.

Figure 4-5 shows the recommended architecture, keeping the interoperability limitations in mind.

Figure 4-5 Front-end JWS with SC and JPD with a JWS and Process Control



Communication Between JPDs

A WLI application contains several JPDs that communicate with each other. A JPD that is invoked by other JPDs is called a sub-process. Table 4-7 lists the high-level guidelines that you can follow for JPD to JPD communication.

Table 4-7 High-Level Design Guidelines for JPD Communication

High-Level Design Guideline	Description
For asynchronous communications between processes in different domains, the use of JMS and the messaging bridge is recommended, together with the JMS event generator or web services over JMS.	The messaging bridge supports all of the QoS options for asynchronous messaging between domains. The store-and-forward capabilities of the messaging bridge insulate local processes from problems and provide access to remote providers.

Table 4-7 High-Level Design Guidelines for JPD Communication (Continued)

For asynchronous two-way communications between processes in the same domain, it is recommended that you use process control rather than the message broker.	You get approximately the same performance when you start a process using process control or message broker, but it is significantly faster to receive a process control callback than a message broker subscription. The message broker subscription filter mechanism uses a database to map the filter values to process instances. Process control callbacks are routed directly to process instances.
Use raw JMS where the message size is large.	JMS is more efficient than marshalling large messages in web services.
Use the process control, not service or service broker control for synchronous communications between processes in the same domain, if a transaction must be propagated between processes, or performance is an issue.	Process control is transactional and avoids SOAP marshalling. Service and service broker controls are not transactional.
For synchronous communications between processes in different domains, use service and service broker controls or the JPD proxy.	Process control can only be used within a domain. The JPD proxy provides transactional propagation at the cost of tight coupling. Service and service broker control do not provide transaction propagation but enable loose coupling.
Use the service broker control rather than the service or process control, if the data-dependent routing facilities are required.	The service broker control allows for configurable routing of service calls to different service implementations depending on data, but this facility is not transactional.

Using Controls

Controls are an integral part of a JPD. WLI 9.2 controls are based upon open Apache Beehive standards. They support annotations based on JSR-175 standards. Controls are reusable and provide easy access to enterprise resources. Controls can be used within a process definition to make calls to a backend system. For example, the database control allows a process to send SQL

to an RDBMS using a JDBC connection pool. [Table 4-8](#) lists the high-level guidelines that you can follow for using controls.

Table 4-8 High-Level Guidelines for Using Controls

High-Level Design Guideline	Description
Control reuse	All controls should be written so that they can be reused across process definitions.
You should develop custom controls keeping in mind native WLI 9.2 controls.	Use the controls available in WLI 9.2 as much as possible. Write custom controls only when absolutely necessary.
Use custom controls instead of custom Java code if reuse of the code is a consideration.	A control that is created for custom Java code means that the code can be reused across different process definitions.
Controls should be versioned in source code control.	Controls should be treated as Java code and versioned appropriately.
Each application project in BEA Workshop for WebLogic should have a separate control project.	Import controls from the component library into this project so that they become available to all the applications.

Using Dynamic Properties and Annotation for Controls

In many cases, control attributes are statically defined using annotations. However, some controls provide a Java API to dynamically change attributes. Dynamic controls, such as service broker and process controls provide the means to dynamically set control attributes. A dynamic or late binding process is used where attributes are determined at runtime using a combination of lookup rules and values. Controls that support dynamic binding are called dynamic controls.

Look-up rules are defined during design time. Look-up values can be defined in the `DynamicProperties.XML` file by changes during runtime. This file is a domain-wide file shared by all WLI applications in the domain. This feature allows the complete de-coupling of control attributes from the application. This file can be re-configured while an application is running, and you do not need to redeploy the application for the changes to take effect.

To achieve the dynamic binding of properties, use:

- Selectors
- The `setProperties()` API

- Setter methods for individual properties, such as `setEndPoint()`.

Use the `getProperties()` method to retrieve the current property settings.

Buffering Service Control Methods During Asynchronous Calls

When you call web service controls asynchronously from business processes, it is recommended that you buffer the asynchronous call to ensure that the message sent from the business process to the web service is enqueued. An asynchronous call to a resource marks the boundary of a transaction in your business process. A call to a resource is not enqueued until the transaction is committed.

By buffering the call to the resource, you ensure that the transaction is committed before there is any response from the resource. If you do not buffer the call, your business process must wait for the HTTP acknowledgement before the transaction is committed. In this situation, the resource may attempt to respond to the business process before the HTTP acknowledgement.

Using Control Factory to Manage Multiple Instances of a Control

The control factory feature of a control enables a JPD to interact with a multiple instances of the same JPD. You can implement File, e-mail, WLI JMS, Trading Partner Management, Service, and Worklist controls as control factories. For example, if a JPD is required to send a document, such as loan application, to multiple service providers, it can use a control factory to create multiple instances of the service control, and dispatch requests to the service control instances in parallel. If the control uses callbacks, a single parameterized callback handler in the calling JPD, can manage the callbacks received from all the control instances.

Data Transformation

Data transformation and manipulation is an integrated part of a business process. Service consumers and providers require varied data format and types. WLI provides the following tools for data transformation:

- XQuery(2004)
- XSLT
- Format Builder, Message Format Language (MFL)
- XML beans
- Data model

Note: WLI 9.2 supports the runtime execution of XQuery 2002 for the purpose of backward compatibility with WLI 8.1.

Canonical Data Model

It is recommended that you create a canonical data model for your application. A canonical data model maps the data to an agreed standard form. If correctly implemented, such a data model provides the services with an Enterprise Information System (EIS) neutral interface. It specifies how you can map the reference, static, and identifier data from an EIS data format to a standard format, de-coupling the data model of the host and the data model of the recipient. You should first create standards for the service interface, that in turn, enforces a common representation of data types and entities within the enterprise. You must maintain a consistent method of representing dates, numbers, post codes, and addresses.

Runtime Selection of a Transformation

A dynamic transformation control enables a business process to dynamically select and execute a transformation during runtime. It allows you to choose the XQuery, XSLT, or MFL file that is invoked at runtime. For example, if you have an integration hub that receives documents from various regional offices, you can use the dynamic transformation control to perform different transformations based on the area code of each regional office. [Table 4-9](#) lists the high-level design guidelines for data transformation.

Table 4-9 High-Level Design Guidelines for Data Transformations

High-Level Design Guideline	Description
Create new data transformations using XQuery rather than XSLT.	XQuery has several advantages over XSLT, including additional functionality, and better performance.
Use a data model (Enterprise Data Model (EDM)) only where there are clear benefits that offset the additional cost and complexity.	An enterprise requires a large amount of resources to use a data model. A data model should be used only when an organization is willing to commit the required resources.
Implementing a data model involves high cost and social challenges. You can implement standards at the interface to ensure a quick return on investment.	<p>This is to ensure that data is represented consistently across all service interfaces. For example, dates, post codes, and addresses are always represented in a specific style.</p> <p>Ensure that you represent information such as message headers in a standard format.</p>

Table 4-9 High-Level Design Guidelines for Data Transformations

High-Level Design Guideline	Description
The origin of the message should not be disclosed to the recipient.	If a data model or standard service interface is appropriately implemented, the recipient of a message is not aware of the data format of the message despatching system.
Wrap any re-usable transformation as a stateless process.	If a re-usable transformation is wrapped as a stateless process, other components can call it a service. This includes direct calls from a front-end system.

Developing a Task Plan

WLI 9.2 worklist has several enhanced features as follows:

- You can model multi-step tasks such as loan approval, and purchase order approval, using a simple drag-and-drop task plan editor. You do not need any prior knowledge of Java or J2EE to be able to use task plan editors.
- A task can be assigned to multiple human actors (one at a time).
- An enhanced and system generated web user interface allows you to work with and test a task plan. You can write your custom user interface using the APIs that are exposed by the task plan.
- The enhanced task assignment and load balancing facilities allow the efficient utilization of a human actor's time.
- The task plan can work without using a JPD. However, you can use a task plan control to interact with the JPD. JPDs can also subscribe to task plan events and process them as per business needs.

Task Plan for Exception Management

You can use the task plan to manage exceptions in a JPD. Users can work on a task plan that is invoked when an exception occurs.

Integrating Custom Logic With a Task Plan

Use the task plan event service to integrate your custom logic with the task plan.

To subscribe to task plan events, write your own custom event listeners and register them with the task plan. You can write custom code in your listener class. This custom code is executed when task plan events invoke the respective event listeners associated with the task.

You can use a custom assignment handler in the custom logic to change the default task assignment at runtime.

See [Using the Worklist](#) for detailed information.

Deploying and Maintaining WLI Applications

There are several best practices for deploying, running, and maintaining WLI applications, as explained in the following sections:

- [Deploying WLI Application During Runtime](#)
- [Deploying WLI Application in a Cluster](#)

Deploying WLI Application During Runtime

When you work in the development mode, you can use the WLI IDE to build and deploy your application. The IDE provides a feature that helps you generate an Ant script to create a build for production purposes. You can also execute these build scripts outside the IDE via the command prompt and generate a single EAR file. You can deploy this EAR file via the command prompt or using the WebLogic Server Console. For more information, see [Deploying WebLogic Integration Solutions](#).

Deploying WLI Application in a Cluster

A WebLogic Server cluster domain contains only one administration server, and one or more managed servers. The managed servers in a WLI domain can be grouped in a cluster. When you configure WLI resources that can be clustered, you target the resources to a named cluster. If you specify a cluster as the target for resource deployment, you can dynamically increase the capacity by adding managed servers to your cluster. The best practices that you can apply to a cluster are as follows:

- [Configuring Trading Partner Integration Resources](#)

- [Changing Cluster Configurations and Deployment Requests](#)
- [Load Balancing in a WLI Cluster](#)

Configuring Trading Partner Integration Resources

You must deploy Trading Partner Integration components homogeneously to a cluster. To avoid a single point of failure, ensure that Trading Partner Integration resources are deployed identically on every managed server.

The guidelines you can follow to configure Trading Partner Integration in a cluster are as follows:

- Specify the host and port numbers of the hardware or software routers as the HTTP end points in the binding of trading partners. This step protects the identity of your managed servers, which are behind a firewall, and allows managed servers to change operational status without impacting the external customer.
- Update Trading Partner Management configuration using the WLI Administration Console. A JMS broadcast mechanism propagates these changes to the managed servers. These changes are quick, but not instantaneous. There is a brief period during which the managed servers contain a mix of the old and new configuration information. To minimize the impact, it is recommended that you update configurations when the resources requiring updates are inactive.

Changing Cluster Configurations and Deployment Requests

You can change configurations for a cluster. For example, you can add new nodes to the cluster or modify Trading Partner Integration configuration only while the administration server of the cluster is active.

Requests to deploy or disable a cluster are interrupted if the administration server is inactive, but the managed servers continue to serve requests. If you can ensure that the required configuration files such as `msi-config.xml`, `SerializedSystemIni.dat`, and optionally `boot.properties` exist in each managed server's root directory, you can boot or reboot managed servers using an existing configuration.

Managed servers that work without an administrative server, operate in a Managed Server Independence (MSI) mode. For more information about MSI mode, see "*Understanding Managed Server Independence Mode*" sub-section in [Avoiding and Recovering from Server Failure](#) of *Managing WebLogic Server Start up and Shutdown*.

Load Balancing in a WLI Cluster

One of the goals of clustering in your WLI application is to achieve scalability. For a cluster to be scalable, each server must be fully utilized. Load balancing distributes the work load proportionally among all the servers in a cluster, so that each server can run at full capacity. Load balancing is required for various functional areas in a WLI cluster. The functions are:

- [HTTP Functions in a Cluster](#)
- [JMS Functions in a Cluster](#)
- [Synchronous Clients and Asynchronous Business Processes](#)
- [RDBMS Event Generators](#)
- [Application Integration Functions in a Cluster](#)

HTTP Functions in a Cluster

Web services (SOAP or XML over HTTP) and WebLogic Trading Partner Integration protocols can use HTTP load balancing. You can use the WebLogic `HttpClusterServlet`, a web server plug-in, or a hardware router for external load balancing.

JMS Functions in a Cluster

WLI or WLI applications most often utilize JMS queues that are configured as distributed destinations. The exception to this rule is that a JMS queue is targeted to a single managed server.

Synchronous Clients and Asynchronous Business Processes

If your WLI solution includes communication between a synchronous client and an asynchronous business process, you can enable server affinity for the `weblogic.jws.jms.QueueConnectionFactory`. This is the default setting.

WARNING: If you disable server affinity for a solution that includes communication between a synchronous client and an asynchronous business process in an attempt to tune JMS load balancing, the resulting load balancing behavior is unpredictable.

RDBMS Event Generators

The RDBMS event generator has a dedicated JMS connection factory called `wli.internal.egrdbms.XAQueueConnectionFactory`. Load balancing is enabled for this connection factory by default. You must disable load balancing and enable server affinity for

`wli.internal.egrdbs.XAQueueConnectionFactory` to disable load balancing for RDBMS events.

Application Integration Functions in a Cluster

Application Integration allows load balancing of synchronous and asynchronous services and events within a cluster. The usage of synchronous and asynchronous services are explained in detail as follows:

- [Synchronous Services](#)
- [Asynchronous Services](#)

Synchronous Services

Synchronous Services are implemented as method calls on a session EJB. They are load balanced within the cluster according to EJB load balancing rules. These EJBs are published at design time and each application view is represented as two session EJBs: one stateless, and one stateful.

In a standard operation, stateless session EJBs invoke the services, and load balancing occurs on a per-service basis. Every time you invoke a service on an application view, you may be routed to a different EJB on a different WebLogic managed server instance.

When you use the local transaction facilities of the application view during a local transaction, the stateful session EJB invokes the services. The stateful session EJB keeps the connection to the EIS open, so that the local transaction state can persist between service invocations. In this mode, service invocations are pinned to a single EJB instance on a single managed server within the cluster. Once the transaction is complete, either through a commit or rollback, the standard per-service load balancing is applicable.

Asynchronous Services

Asynchronous services are always invoked as method calls on a stateless session EJB. You cannot use the local transaction facility of the application view for asynchronous service invocations.

A single asynchronous service invocation translates to two method invocations on two different stateless session EJB instances. The load balancing for asynchronous service in this case occurs on two occasions, the first upon receipt of the request, and the second in the execution of the request and delivery of the response.

In addition, both the asynchronous service request and response are posted to a distributed JMS queue. JMS load balancing as a result, applies to both the request and the response. In this case,

the `invokeServiceAsync` method of the application view may be serviced on one managed server, the request delivered to a second managed server where the request is processed and the response generated, and the response delivered to a third server for retrieval by the client.

Deploying and Maintaining WLI Applications

Core Implementation Patterns for WLI Applications

There are several core implementation patterns for WLI applications, as explained in the following sections:

- [Core Implementation Patterns for a JPD](#)
- [Course-Grained Process Front-end for a Fine-Grained Process](#)
- [Loosely Coupled Process With a Common Message Interface](#)
- [Dynamic Property Driven Processes](#)

Core Implementation Patterns for a JPD

JPDs are of several types and are designed with the following intrinsic characteristics:

- **Basic complexity:** A basic process makes only one call out to another service or an external system. A basic process does not contain complex logic but may contain any number of steps for transformation and exception handling.
- **Composite complexity:** A composite process contains two or more calls out to external systems. A process that makes two calls out to a single back end service is defined as composite. Composite processes may contain complex logic such as parallel calls out, and client requests after the start node.

Note: Composite processes have more complex exception management because of the multiple calls.

- Synchronous or asynchronous calling paradigm

- Stateful or stateless JPD state
- One or two-way exchange paradigm

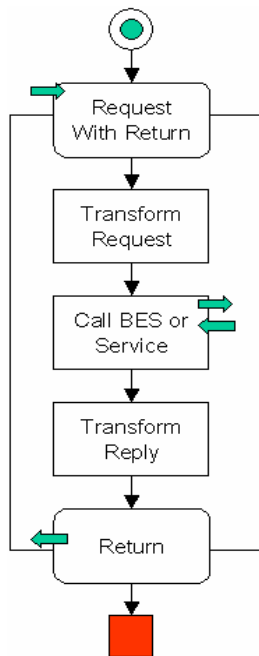
Several patterns of JPDs have resulted from combinations of these characteristics. The following sections describe these patterns in brief:

- [Pattern 1: Basic Synchronous Stateless two-way Service](#)
- [Pattern 2: Basic Asynchronous Stateless two-way Service](#)
- [Pattern 3: Basic Asynchronous Stateless one-way Service](#)
- [Pattern 4: Basic Asynchronous Stateful two-way Service](#)
- [Pattern 5: Basic Asynchronous Stateful one-way Service](#)
- [Pattern 6: Composite Synchronous Stateless two-way Service](#)
- [Pattern 7: Composite Synchronous Stateful two-way Service](#)
- [Pattern 8: Composite Asynchronous Stateless two-way Service](#)
- [Pattern 9: Composite Asynchronous Stateless one-way Service](#)
- [Pattern 10: Composite Asynchronous Stateful two-way Service](#)
- [Pattern 11: Composite Asynchronous Stateful one-way Service](#)
- [Other Patterns](#)

Pattern 1: Basic Synchronous Stateless two-way Service

Use this pattern illustrated in [Figure 6-1](#) to implement some of the fastest processes. The standard usage is for simple access to a backend system or to implement helper processes.

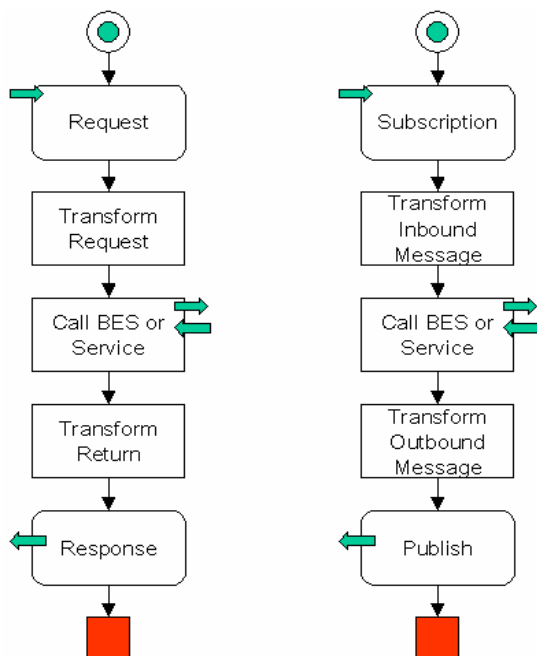
Figure 6-1 Pattern 1



Pattern 2: Basic Asynchronous Stateless two-way Service

Use this pattern illustrated in [Figure 6-2](#) to implement some of the fastest processes. The standard usage is for simple access to a backend system, when de-coupling with the process client is required.

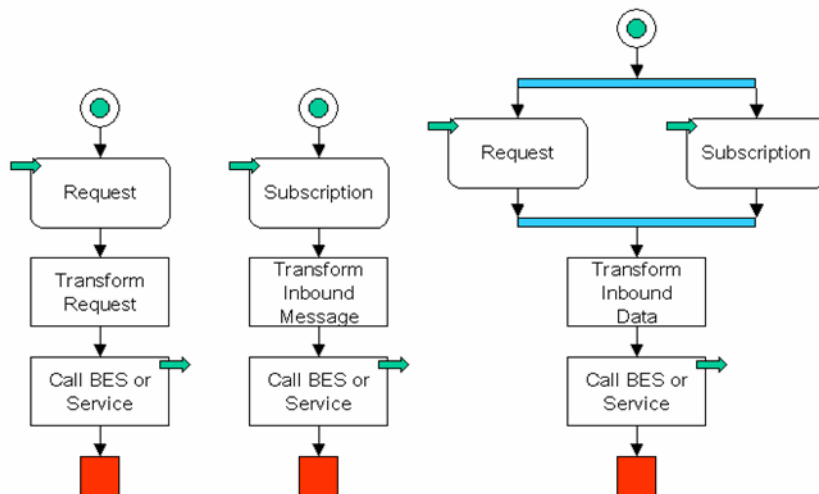
Figure 6-2 Pattern 2



Pattern 3: Basic Asynchronous Stateless one-way Service

Use this pattern illustrated in [Figure 6-3](#) to implement some of the fastest processes. The standard usage is to access backend systems in event-driven situations.


Figure 6-3 Pattern 3



Pattern 4: Basic Asynchronous Stateful two-way Service

This pattern is not as fast as its stateless equivalent. Use this pattern illustrated in [Figure 6-4](#) to access backend systems that provide an asynchronous interface.


Figure 6-4 Pattern 4



Pattern 5: Basic Asynchronous Stateful one-way Service

This pattern is very rare because most of the one-way services are stateless and so there is no need to wait for an answer. Use this pattern illustrated in [Figure 6-5](#) to quickly access a backend system in event-driven situation where a call that is waiting is required.


Figure 6-5 Pattern 5



Pattern 6: Composite Synchronous Stateless two-way Service

Use this pattern illustrated in [Figure 6-6](#) to implement integration logic that requires good performance and coupling with the client.


Figure 6-6 Pattern 6



Pattern 7: Composite Synchronous Stateful two-way Service

This pattern represents an unusual case, where the logic is implemented after the client has received a response. Use this pattern illustrated in [Figure 6-7](#) to enable a stateful process that runs for a period of time, and is started by a synchronous request or response. Once the synchronous reply is sent, the process moves to a traditional asynchronous model.


Figure 6-7 Pattern 7



Pattern 8: Composite Asynchronous Stateless two-way Service

This pattern illustrated in [Figure 6-8](#) is a standard method to implement integration logic that requires good performance while maintaining de-coupling from its client. It can also be used to access a backend system when de-coupling from the caller is required. It is not always possible to implement this pattern as it requires all the resources used to be stateless.

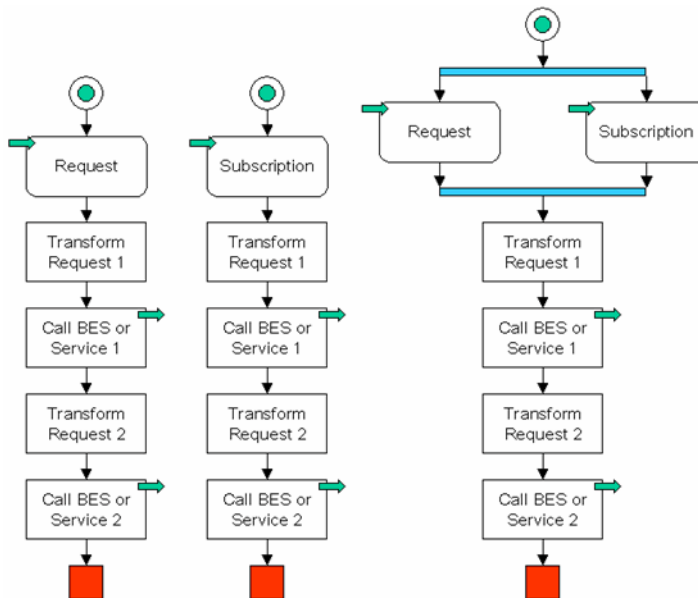
Figure 6-8 Pattern 8



Pattern 9: Composite Asynchronous Stateless one-way Service

Use this pattern illustrated in [Figure 6-9](#) to implement integration logic that requires good performance in event-driven situations.

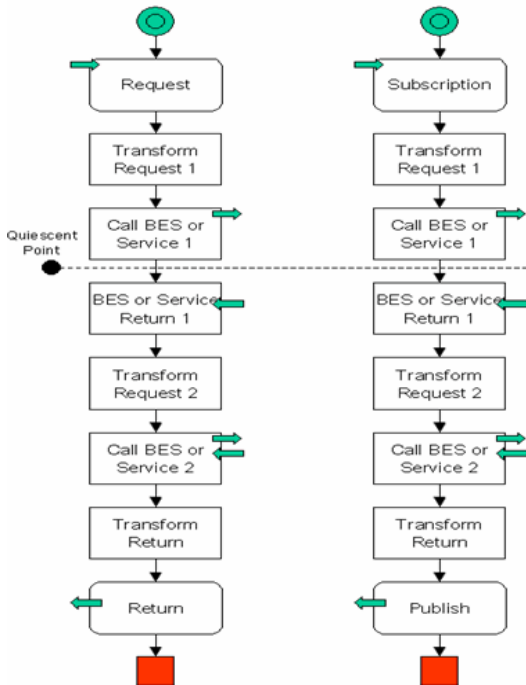
Figure 6-9 Pattern 9



Pattern 10: Composite Asynchronous Stateful two-way Service

Use this pattern illustrated in [Figure 6-10](#) when the stateless version of the same pattern is not possible, as at least one of the resources contains an asynchronous interface. It is also used as a standard pattern for long-running processes.

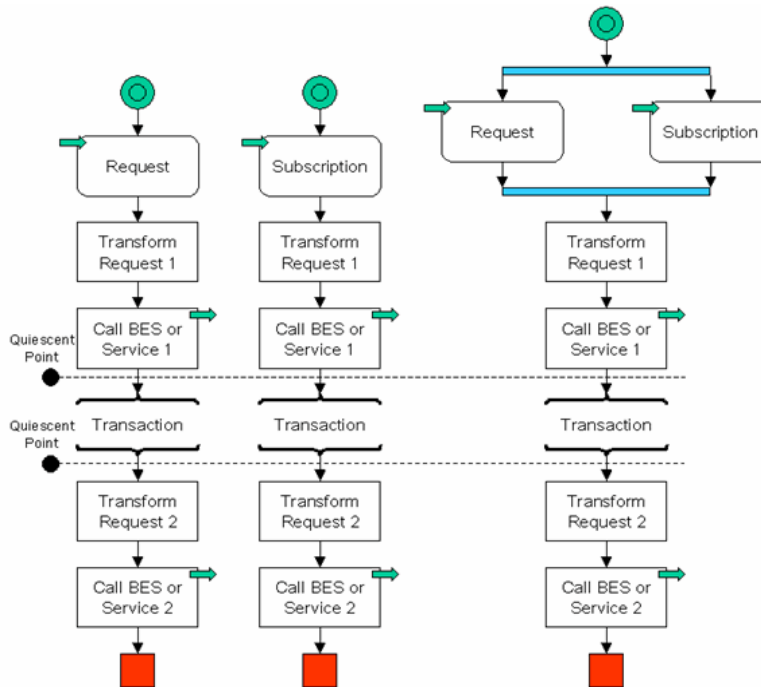
Figure 6-10 Pattern 10



Pattern 11: Composite Asynchronous Stateful one-way Service

This pattern is not common as most of the one-way services are usually stateless and there is no need to wait for an answer. Use this pattern illustrated in [Figure 6-11](#) to implement integration logic in event-driven scenarios, and where a call that is waiting is also required.

Figure 6-11 Pattern 11



Other Patterns

This section re-groups useful patterns, which can be used in combination with one of the core patterns.

SyncAsync Pattern

You can create a business process containing a client request node with a sync or async callback name attribute property, to enable synchronous clients to interact with business processes that have asynchronous interactions with resources. The client request node property holds the name of the callback method that the associated client response node uses. The client request and client response nodes delineate the activities (including asynchronous activities) that occur while the

client is blocking the process. After setting this property, generate the sync-to-async WSDL. The synchronous WSDL generation process replaces the SOAP address of the service with a modified SOAP address. The modified address causes the synchronous servlet to process the client request and subsequent return action. A sample of the generated service entry is as follows:

Normal WSDL

```
<service name="syncAsync">
  <port name="syncAsyncSoap" binding="s0: syncAsyncSoap">
    <soap: address
      location="http://localhost:7001/SyncAsyncWeb/processes/syncAsync.jpdl"/>
  </port>
```

Synchronous WSDL

```
<service name="syncAsync">
  <port name="syncAsyncSoap" binding="s0: syncAsyncSoap">
    <soap:address
      location="http://localhost:7001/sync2AsyncIM/SyncAsyncWeb/processes/syncAs
      ync.sync2JPD"/>
  </port>
```


De-Synchronizer Pattern

The de-synchronizer pattern allows you to asynchronously call a synchronous process. This is recommended if:

- The process needs to support both synchronous and asynchronous clients.
- De-coupling is required, in particular if a sub-process needs to run in its own transaction.
- The client is unable to call services synchronously.

Implement the process in [Figure 6-12](#) to provide an asynchronous interface to a synchronous process.

Figure 6-12 De-Synchronizer Pattern



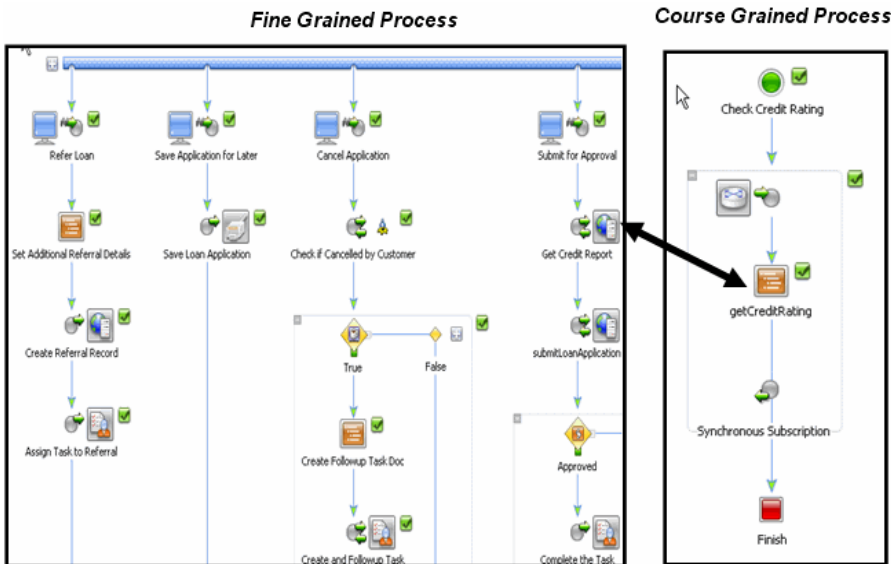
De-Synchronizer Service

This is a simple proxy process. The signature of the request and the response should remain the same as that of the original service. The de-synchronizer passes its input values to the sub-process and sends the return value to the caller of the sub-process.

Course-Grained Process Front-end for a Fine-Grained Process

[Figure 6-13](#) illustrates an example of how you can use a combination of course-grained and fine-grained processes. A course-grained process assigns specific work to a number of fine-grained processes.

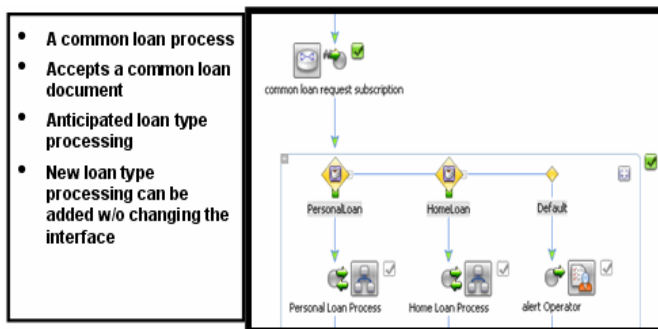
Figure 6-13 Front-end a Fine-Grained Process With a Course-Grained Process



Loosely Coupled Process With a Common Message Interface

This pattern illustrated in [Figure 6-14](#) shows how you can create a loosely coupled process using a common messaging interface through message brokers, based upon a Publish and Subscribe architecture.

Figure 6-14 Loosely Coupled Process Using Message Interface



Dynamic Property Driven Processes

You can use the following controls to define a dynamic, property driven process:

- Dynamic controls:
 - Dynamic Transformation control
 - XML Meta data control
- Dynamic Binding controls:
 - Service Broker control
 - Process control
- Dynamic subscription:
 - Message Broker Subscription control
- Runtime controls - Most controls have **setProperties(<type>ControlPropertiesDocument xml)**

Figure 6-15 illustrates an example of an agile and dynamic process with a Register and Table look-up.

Figure 6-15 Agile, Dynamic Process with a Register and Table Look-up



The features of this process are as follows:

- The target URL is determined at runtime and contains a:
 - Registry Look-up
 - Table Look-up
- The Service Broker control accepts the end point as the runtime property
- The Service Broker invokes the specified service at a location determined at runtime.

Understanding Requirements

You can use the following SMART criteria to evaluate functional and non-functional requirements:

- **Specific** - Is the requirement unambiguous, with consistent terminology, simple, and at the appropriate level of detail?
- **Measurable** - Is it possible to verify that this requirement has been met? What tests must be performed, or what criteria must be satisfied to verify that the requirement is met?
- **Attainable** - What is your professional judgment of the technical feasibility of the requirement?
- **Realistic** - Is the requirement realistic, given the resources? Do you have adequate staff? Do you have the skills? Do you have access to the requisite development infrastructure? Do you have access to the required runtime infrastructure? Do you have enough time?
- **Traceable** - Is the requirement linked from its conception through its specification to its subsequent design, implementation, and test?

When in a solution domain, use robustness analysis to categorize your use cases.

You can divide a use case into four kinds of objects using robustness analysis:

- **Actors** - Objects external to use cases. Actors could be human or other external objects such as systems, applications, or devices. Actors interact with the use case by sending or receiving messages.

- Boundary objects - The public face of the use case. Actors interact with use cases using boundary objects. The user interface element or the public API of the use case are examples of boundary objects.
- Entity objects - The objects with long lives in the use case. Examples of entity objects are purchase order, invoice, and customer.
- Controller objects - The glue between boundary elements and entity objects. They contain methods for specific functions in a use case such as Validate user, Create, Retrieve, Update, and Delete (CRUD) functions.

Figure A-1 shows how the use case is divided into objects.

Figure A-1 Use Case Realization





Figure A-2 illustrates the rules of robustness analysis.

Figure A-2 Rules of Robustness Analysis



The rules are as follows:

- You can have one or more actors, boundary objects, controllers, and entity objects in a use case.
- Actors can only talk to boundary objects.
- Boundary objects can only talk to controllers and actors.
- Entity objects can only talk to controllers.
- Controllers can talk to boundary and entity objects, other controllers, but not to actors.

Note: Dr. Ivar Jacobson developed the use case and robustness analysis technique. He is well known as the father of use cases and was also one of the authors of the original UML specifications.