

TIBCO iProcess[®] Engine

Architecture Guide

Software Release 11.4.1
April 2014

Two-Second Advantage[®]



Important Information

SOME TIBCO SOFTWARE EMBEDS OR BUNDLES OTHER TIBCO SOFTWARE. USE OF SUCH EMBEDDED OR BUNDLED TIBCO SOFTWARE IS SOLELY TO ENABLE THE FUNCTIONALITY (OR PROVIDE LIMITED ADD-ON FUNCTIONALITY) OF THE LICENSED TIBCO SOFTWARE. THE EMBEDDED OR BUNDLED SOFTWARE IS NOT LICENSED TO BE USED OR ACCESSED BY ANY OTHER TIBCO SOFTWARE OR FOR ANY OTHER PURPOSE.

USE OF TIBCO SOFTWARE AND THIS DOCUMENT IS SUBJECT TO THE TERMS AND CONDITIONS OF A LICENSE AGREEMENT FOUND IN EITHER A SEPARATELY EXECUTED SOFTWARE LICENSE AGREEMENT, OR, IF THERE IS NO SUCH SEPARATE AGREEMENT, THE CLICKWRAP END USER LICENSE AGREEMENT WHICH IS DISPLAYED DURING DOWNLOAD OR INSTALLATION OF THE SOFTWARE (AND WHICH IS DUPLICATED IN THE LICENSE FILE) OR IF THERE IS NO SUCH SOFTWARE LICENSE AGREEMENT OR CLICKWRAP END USER LICENSE AGREEMENT, THE LICENSE(S) LOCATED IN THE "LICENSE" FILE(S) OF THE SOFTWARE. USE OF THIS DOCUMENT IS SUBJECT TO THOSE TERMS AND CONDITIONS, AND YOUR USE HEREOF SHALL CONSTITUTE ACCEPTANCE OF AND AN AGREEMENT TO BE BOUND BY THE SAME.

This document contains confidential information that is subject to U.S. and international copyright laws and treaties. No part of this document may be reproduced in any form without the written authorization of TIBCO Software Inc.

TIBCO, Two-Second Advantage, TIBCO ActiveMatrix BusinessWorks, TIBCO Business Studio, TIBCO Enterprise Message Service, TIBCO Hawk, TIBCO iProcess, TIBCO iProcess Suite, and TIBCO Rendezvous are either registered trademarks or trademarks of TIBCO Software Inc. in the United States and/or other countries.

Enterprise Java Beans (EJB), Java Platform Enterprise Edition (Java EE), Java 2 Platform Enterprise Edition (J2EE), and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle Corporation in the U.S. and other countries.

All other product and company names and marks mentioned in this document are the property of their respective owners and are mentioned for identification purposes only.

THIS SOFTWARE MAY BE AVAILABLE ON MULTIPLE OPERATING SYSTEMS. HOWEVER, NOT ALL OPERATING SYSTEM PLATFORMS FOR A SPECIFIC SOFTWARE VERSION ARE RELEASED AT THE SAME TIME. SEE THE README FILE FOR THE AVAILABILITY OF THIS SOFTWARE VERSION ON A SPECIFIC OPERATING SYSTEM PLATFORM.

THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS DOCUMENT COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THIS DOCUMENT. TIBCO SOFTWARE INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS DOCUMENT AT ANY TIME.

THE CONTENTS OF THIS DOCUMENT MAY BE MODIFIED AND/OR QUALIFIED, DIRECTLY OR INDIRECTLY, BY OTHER DOCUMENTATION WHICH ACCOMPANIES THIS SOFTWARE, INCLUDING BUT NOT LIMITED TO ANY RELEASE NOTES AND "READ ME" FILES.

Copyright © 1994-2014 TIBCO Software Inc. ALL RIGHTS RESERVED.

TIBCO Software Inc. Confidential Information

Contents

Figures	ix
Tables	xi
Preface	xiii
Related Documentation	xiv
TIBCO iProcess Engine Documentation	xiv
Other TIBCO Product Documentation	xiv
Typographical Conventions	xvi
Connecting with TIBCO Resources	xix
How to Join TIBCOCommunity	xix
How to Access TIBCO Documentation	xix
How to Contact TIBCO Support	xix
Chapter 1 Introduction to TIBCO iProcess Products	1
TIBCO Product Overview 1	3
TIBCO Product Overview 2	4
TIBCO Business Studio	5
TIBCO BusinessWorks	6
TIBCO Enterprise Message Service	7
TIBCO Hawk	8
TIBCO iProcess Analytics	9
TIBCO iProcess Analytics Export	10
TIBCO iProcess Decisions Server	11
TIBCO iProcess Decisions Studio	12
TIBCO iProcess Engine	13
TIBCO iProcess Objects Server	15
TIBCO iProcess Objects Director	16
TIBCO iProcess Server Objects	17
TIBCO iProcess Technology Plug-ins	18
TIBCO iProcess Web Services Plug-in	20
TIBCO iProcess Workspace (Browser)	21

TIBCO iProcess Workspace Lite	22
TIBCO iProcess Workspace Plug-ins	24
TIBCO iProcess Workspace (Windows)	25
TIBCO Rendezvous	27
Chapter 2 Introduction to the TIBCO iProcess Engine	29
TIBCO iProcess Engine Architecture	30
The Role of the TIBCO iProcess Engine	31
iProcess Physical Architecture	32
Installing the iProcess Engine on a Single Server	32
Installing the iProcess Engine on a Node Cluster	32
TIBCO iProcess Workspace and TIBCO iProcess Engine Communication	33
TIBCO iProcess Engine Process Structure	35
Process Sentinels	35
Foreground Processes	36
Mbox Sets	36
Background Processes	37
Event Handling	37
Where is TIBCO iProcess Engine Case Data Stored?	39
How Do Work Items Appear in Work Queues?	40
Sending Instructions From the TIBCO iProcess Workspace to the TIBCO iProcess Engine	40
Mbox Sets	41
User Access to iProcess Engine Work Queues	42
24*7 TIBCO iProcess Engine Operation	44
The iProcess Engine and Hardware Clustering	44
Chapter 3 Using the TIBCO iProcess Suite in a Multilingual Environment	47
Overview	48
Globalization Options in the iProcess Suite	50
Advantages of UTF-8 Encoding	50
Issues with UTF-8 Encoding	51
iProcess Names	51
Recommendations	53
Configuring the iProcess Suite Using UTF-8	54
TIBCO iProcess Engine	54
iProcess Clients	55
The LDAPCONF Utility	56
TIBCO Business Studio	56
iProcess Plug-ins	56
Globalization Support Using Native Encoding	58

Using iProcess Suite in a Single-Byte Native Encoding Environment	58
Using iProcess Suite With a Multi-Byte Character Encoding Environment	59
Implementing iProcess Suite in an International Environment with Native Encoding	60
Chapter 4 TIBCO iProcess Engine Processes	61
Foreground Processes	62
Work Queue Server	63
Allocation of Work Queues to WIS Processes	65
RPC Pool Server	66
RPC Listeners	66
Work Item Server	67
WIS Mbox Daemon	67
Mbox Sets and Message Queues	70
Transaction Control of Messages	71
Background Processes	72
Background	73
Case Prediction Processor(s)	73
Database Queue Daemon	74
Deadline Manager	74
IAPJMS Process	76
RPC Background Process	76
Chapter 5 Introduction to Transactional Business Process Automation	77
Overview	78
What is a Local Transaction?	79
Example of a Local Transaction	80
What is a Distributed Transaction?	81
Transaction Scope	82
Oracle Server Transaction Scope	82
DB2 Transaction Scope	82
SQL Server Transaction Scope	82
Using Distributed Transactions with MSDTC	83
Using Enterprise Application Integration Steps in Procedures	84
What is MSDTC?	85
Examples of Transaction Control	86
Case Data Updates to the SQL Server using MSDTC	86
External Updates Using EAI Steps	87
Transaction Failures and Rollbacks	88
Poison Transactions	88

Chapter 6 iProcess Mbox Sets	89
Overview	90
What are iProcess Messages?	90
Definition of Mbox Sets	91
Configuring Mbox Sets	91
Transaction Management of Messages	92
UNIX Oracle Transaction Implementation	92
Windows SQL Server Transaction Implementation	92
UNIX DB2 Transaction Implementation	92
Chapter 7 Monitoring Activities	93
Overview	94
Activity Publishing	94
Work Queue Delta Publication	95
How Messages are Processed From the BG Process to the IAPJMS Process	96
How Activity Messages are Processed From the IAPJMS Process to the External Application	97
How Messages are Processed From the WIS Process to the IAPJMS Process	98
How Work Queue Delta Messages are Processed From the IAPJMS Process to the External Application	99
Understanding the Message Types	100
IAP Message Types	100
WQD Message Types	101
Chapter 8 Database Failure and Failover	103
Overview	104
TIBCO iProcess Engine Behavior	105
TIBCO iProcess Workspace Behavior	107
iProcess Objects and iProcess Server Objects Behavior	108
TIBCO iProcess Engine Configuration Requirements	109
Oracle (UNIX or Linux)	109
Oracle (Windows), SQL Server and DB2 (UNIX or Linux)	110
Chapter 9 Process Management	111
Responsibilities of the Process Sentinels	112
Distribution and Hierarchy of Process Sentinels	112
Master and Slave Responsibilities	113
How Processes are Controlled by the Process Sentinels	115
Starting the TIBCO iProcess Engine Processes	116
Determining Where Processes Run	118
Restarting Failed Processes	119

Restarting Failed Process Sentinels	119
Shutting Down Processes	121
Configuring the Process Sentinels	122
Chapter 10 Network Communication	125
TIBCO iProcess Workspace and TIBCO iProcess Engine Network Communication	126
Function of a Portmapper	126
The TIBCO iProcess Engine RPC Service	127
TCP/IP	127
TIBCO iProcess Engine to TIBCO iProcess Engine Network Communication in a Node Cluster	129
Using the TIBCO iProcess Engine in a Firewalled Environment	130
What is a Firewall?	130
iProcess RPC and Firewall Access	130
Port/RPC Number Resource Logging	131
Using Oracle Events Through a Firewall	131
Using JMX Through a Firewall	132
Index	133

Figures

Figure 1	A Local Transaction	79
----------	---------------------------	----

Tables

Table 1	General Typographical Conventions	xvi
Table 2	Syntax Typographical Conventions	xvii

Preface

This guide describes the architecture of the TIBCO iProcess Engine and how each server process works and interacts with other processes. It also provides information about how the TIBCO iProcess Engine interacts with other applications such as databases and transaction control applications. You can use this information to help plan and implement your system, debug system problems and trace performance problems to specific processes.

Topics

- [Related Documentation, page xiv](#)
- [Typographical Conventions, page xvi](#)
- [Connecting with TIBCO Resources, page xix](#)

Related Documentation

This section lists documentation resources you may find useful.

TIBCO iProcess Engine Documentation

The following documents form the TIBCO iProcess Engine documentation set:

- *TIBCO iProcess Engine Installation* Read this manual for instructions on site preparation and installation.
- *TIBCO iProcess Engine Release Notes* Read the release notes for a list of new and changed features. This document also contains lists of known issues and closed issues for this release.
- TIBCO iProcess Suite Documentation This documentation set contains all the manuals for TIBCO iProcess Engine and other TIBCO products in TIBCO iProcess[®] Suite. The manuals for TIBCO iProcess Engine are as follows:
 - *TIBCO iProcess Engine Architecture Guide*
 - TIBCO iProcess Engine Administrator's Guides:
 - TIBCO iProcess Engine Administrator's Guide*
 - TIBCO iProcess Objects Director Administrator's Guide*
 - TIBCO iProcess Objects Server Administrator's Guide*
 - TIBCO iProcess Engine Database Administrator's Guides:
 - TIBCO iProcess Engine (DB2) Administrator's Guide*
 - TIBCO iProcess Engine (Oracle) Administrator's Guide*
 - TIBCO iProcess Engine (SQL) Administrator's Guide*
 - *TIBCO iProcess swutil and swbatch Reference Guide*
 - *TIBCO iProcess Engine System Messages Guide*
 - *TIBCO iProcess User Validation API User's Guide*
 - *LDAPCONF Utility User's Guide*

Other TIBCO Product Documentation

You may find it useful to read the documentation for the following TIBCO products:

- TIBCO ActiveMatrix BusinessWorks[™]

- TIBCO Business Studio™
- TIBCO Enterprise Message Service™
- TIBCO Hawk®
- TIBCO Rendezvous®

Typographical Conventions

TIBCO iProcess Engine can be run on both Microsoft Windows and UNIX/Linux platforms. In this manual, the Windows convention of a backslash (\) is used. The equivalent pathname on a UNIX or Linux system is the same, but using the forward slash (/) as a separator character.



UNIX or Linux pathnames are occasionally shown explicitly, using forward slashes as separators, where a UNIX or Linux-specific example or syntax is required.

Any references to UNIX in this manual also apply to Linux unless explicitly stated otherwise.

The following typographical conventions are used in this manual..

Table 1 *General Typographical Conventions*

Convention	Use
<i>SWDIR</i>	<p>TIBCO iProcess Engine installs into a directory. This directory is referenced in documentation as <i>SWDIR</i>. The value of <i>SWDIR</i> depends on the operating system. For example,</p> <ul style="list-style-type: none"> on a Windows server (on the C: drive) <p>if <i>SWDIR</i> is set to the C:\swserver\staffw_nod1 directory, then the full path to the <code>swutil</code> command is in the C:\swserver\staffw_nod1\bin\swutil directory.</p> on a UNIX or Linux server <p>if <i>SWDIR</i> is set to the /swserver/staffw_nod1 directory, then the full path to the <code>swutil</code> command is in the /swserver/staffw_nod1/bin/swutil directory or the <code>\$SWDIR/bin/swutil</code> directory.</p> <p>Note: On a UNIX or Linux system, the environment variable <code>\$SWDIR</code> should be set to point to the iProcess system directory for the <i>root</i> and <i>swadmin</i> users.</p>
code font	<p>Code font identifies commands, code examples, filenames, pathnames, and output displayed in a command window. For example:</p> <p>Use <code>MyCommand</code> to start the foo process.</p>

Table 1 General Typographical Conventions (Cont'd)

Convention	Use
bold code font	<p>Bold code font is used in the following ways:</p> <ul style="list-style-type: none"> • In procedures, to indicate what a user types. For example: Type admin. • In large code samples, to indicate the parts of the sample that are of particular interest. • In command syntax, to indicate the default parameter for a command. For example, if no parameter is specified, MyCommand is enabled: MyCommand [enable disable]
<i>italic font</i>	<p>Italic font is used in the following ways:</p> <ul style="list-style-type: none"> • To indicate a document title. For example: See <i>TIBCO ActiveMatrix BusinessWorks Concepts</i>. • To introduce new terms. For example: A portal page may contain several portlets. <i>Portlets</i> are mini-applications that run in a portal. • To indicate a variable in a command or code syntax that you must replace. For example: MyCommand <i>PathName</i>
Key combinations	<p>Key name separated by a plus sign indicate keys pressed simultaneously. For example: Ctrl+C.</p> <p>Key names separated by a comma and space indicate keys pressed one after the other. For example: Esc, Ctrl+Q.</p>
	The note icon indicates information that is of special interest or importance, for example, an additional action required only in certain circumstances.
	The tip icon indicates an idea that could be useful, for example, a way to apply the information provided in the current section to achieve a specific result.
	The warning icon indicates the potential for a damaging situation, for example, data loss or corruption if certain steps are taken or not taken.

Table 2 Syntax Typographical Conventions

Convention	Use
[]	<p>An optional item in a command or code syntax.</p> <p>For example:</p> <p>MyCommand [optional_parameter] required_parameter</p>

Table 2 *Syntax Typographical Conventions (Cont'd)*

Convention	Use
	<p>A logical OR that separates multiple items of which only one may be chosen.</p> <p>For example, you can select only one of the following parameters:</p> <pre>MyCommand param1 param2 param3</pre>
{ }	<p>A logical group of items in a command. Other syntax notations may appear within each logical group.</p> <p>For example, the following command requires two parameters, which can be either the pair param1 and param2, or the pair param3 and param4.</p> <pre>MyCommand {param1 param2} {param3 param4}</pre> <p>In the next example, the command requires two parameters. The first parameter can be either param1 or param2 and the second can be either param3 or param4:</p> <pre>MyCommand {param1 param2} {param3 param4}</pre> <p>In the next example, the command can accept either two or three parameters. The first parameter must be param1. You can optionally include param2 as the second parameter. And the last parameter is either param3 or param4.</p> <pre>MyCommand param1 [param2] {param3 param4}</pre>

Connecting with TIBCO Resources

How to Join TIBCOmmunity

TIBCOmmunity is an online destination for TIBCO customers, partners, and resident experts. It is a place to share and access the collective experience of the TIBCO community. TIBCOmmunity offers forums, blogs, and access to a variety of resources. To register, go to <http://www.tibcommunity.com>.

How to Access TIBCO Documentation

You can access TIBCO documentation here:

<http://docs.tibco.com>

How to Contact TIBCO Support

For comments or problems with this manual or the software it addresses, contact TIBCO Support as follows:

- For an overview of TIBCO Support, and information about getting started with TIBCO Support, visit this site:

<http://www.tibco.com/services/support>

- If you already have a valid maintenance or support contract, visit this site:

<https://support.tibco.com>

Entry to this site requires a user name and password. If you do not have a user name, you can request one.

Introduction to TIBCO iProcess Products

This chapter provides an overview of the products in the TIBCO iProcess Suite and other TIBCO products that interact with the TIBCO iProcess Suite. It gives a brief description of each product and explains how they interact with each other.



All Java-based TIBCO iProcess Suite components (for example, TIBCO iProcess Java Server Plug-in and TIBCO iProcess BusinessWorks Server Plug-in) must be run using a 32-bit JVM.

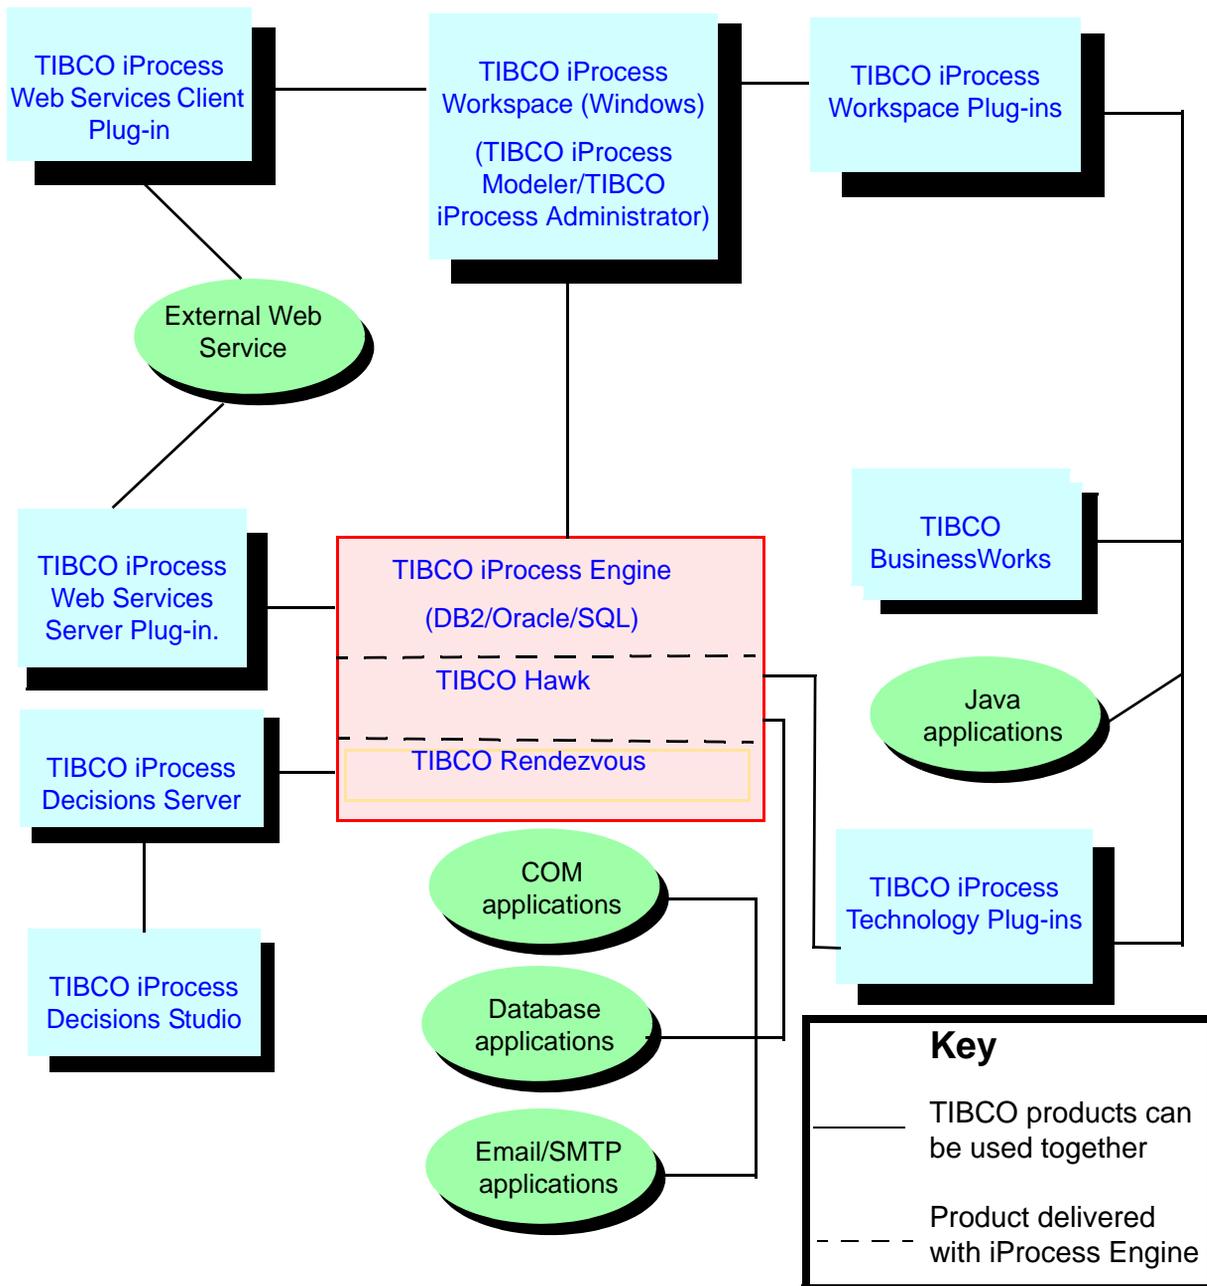
Topics

- [TIBCO Product Overview 1, page 3](#)
- [TIBCO Product Overview 2, page 4](#)
- [TIBCO Business Studio, page 5](#)
- [TIBCO BusinessWorks, page 6](#)
- [TIBCO Enterprise Message Service, page 7](#)
- [TIBCO Hawk, page 8](#)
- [TIBCO iProcess Analytics, page 9](#)
- [TIBCO iProcess Analytics Export, page 10](#)
- [TIBCO iProcess Decisions Server, page 11](#)
- [TIBCO iProcess Decisions Studio, page 12](#)
- [TIBCO iProcess Engine, page 13](#)
- [TIBCO iProcess Objects Server, page 15](#)
- [TIBCO iProcess Objects Director, page 16](#)
- [TIBCO iProcess Server Objects, page 17](#)
- [TIBCO iProcess Technology Plug-ins, page 18](#)
- [TIBCO iProcess Web Services Plug-in, page 20](#)
- [TIBCO iProcess Workspace \(Browser\), page 21](#)

- [TIBCO iProcess Workspace Lite, page 22](#)
- [TIBCO iProcess Workspace Plug-ins, page 24](#)
- [TIBCO iProcess Workspace \(Windows\), page 25](#)
- [TIBCO Rendezvous, page 27](#)

TIBCO Product Overview 1

Click on a product to find out more information.



TIBCO Business Studio

Description TIBCO Business Studio enables business analysts to implement business processes. It provides a standards-based modelling environment that supports both Business Process Modeling Notation (BPMN) and XML Process Definition Language (XPDL).

If you are using native service calls (database or email) or general service calls (such as web services), you can augment the process with execution details in TIBCO Business Studio and deploy it directly to the TIBCO iProcess Engine. However, if you need to make other types of service calls (for example, EAI Java), you must augment and implement the Process using another product such as TIBCO iProcess Modeler.

TIBCO Business Studio also enables business analysts to simulate processes that have been developed in Business Studio. This is useful to identify areas of the process that can be improved such as bottlenecks and areas of high cost or reduced service levels.

TIBCO Business Studio can be used instead of iProcess Workspace (Windows) to design business processes for use with the iProcess Engine.

For Information For information about TIBCO Business Studio, see following:

- *TIBCO Business Studio Process Modeling User's Guide*
- *TIBCO Business Studio Simulation Guide*

TIBCO BusinessWorks

Description TIBCO BusinessWorks is a scalable, extensible, and easy to use integration platform that allows you to develop integration projects. TIBCO BusinessWorks includes a graphical user interface (GUI) for defining business processes and an engine that executes the process.

TIBCO BusinessWorks also works with TIBCO Administrator, a web-based GUI for monitoring and managing run-time components.

The TIBCO iProcess BusinessWorks Connector enables the TIBCO iProcess Engine to interact with TIBCO BusinessWorks and vice versa. For example, BusinessWorks processes can be invoked from iProcess Engine procedures and iProcess Engine cases can be started from BusinessWorks.

For Information For information about TIBCO BusinessWorks, access <http://docs.tibco.com>.
For information about the TIBCO iProcess BusinessWorks Connector, see *TIBCO iProcess BusinessWorks Connector User's Guide*.

TIBCO Enterprise Message Service

- Description** TIBCO Enterprise Message Service implements JMS and integrates support for connecting other message services, such as TIBCO Rendezvous and TIBCO SmartSockets.
- Java Message Service 1.1 (JMS) is a Java framework specification for messaging between applications. Sun Microsystems developed this specification, in conjunction with TIBCO and others, to supply a uniform messaging interface among enterprise applications.
- Using a message service allows you to integrate the applications within an enterprise. For example, you may have several applications: one for customer relations, one for product inventory, and another for raw materials tracking. Each application is crucial to the operation of the enterprise, but even more crucial is communication between the applications to ensure the smooth flow of business processes. Message-oriented-middleware (MOM) creates a common communication protocol between these applications and allows you to easily integrate new and existing applications in your enterprise computing environment.
- TIBCO Enterprise Message Service is delivered with the TIBCO iProcess Engine to enable the iProcess Engine to communicate with other TIBCO products (for example, TIBCO BusinessWorks and TIBCO iProcess Insight) using a JMS protocol.
- For Information** For information about TIBCO Enterprise Message Service, access <http://docs.tibco.com>.

TIBCO Hawk

Description TIBCO Hawk is a monitoring system that monitors systems and applications on a specific computer. It consists of the following components:

- **TIBCO Hawk Agent.** A TIBCO Hawk agent is an autonomous process that resides on each computer and monitors systems and applications on that computer. Agents run independently of the TIBCO Hawk Display. Agents operate autonomously and are active whenever the operating system they monitor is active. Agents use sets of rules, called rulebases, to configure system management, status, and automation tasks. A TIBCO Hawk agent must be installed on each computer you wish to monitor.
- **TIBCO Hawk Microagent.** A TIBCO Hawk Microagent (HMA) is a partner process to the TIBCO Hawk agent and provides the local agent with methods for monitoring the host operating system. Like the agent, a TIBCO Hawk Microagent is generally installed on each computer you wish to monitor.
- **TIBCO Hawk Display Program.** The TIBCO Hawk Display program is used by system administrators to view network health and to create rulebases (sets of rules that automate monitoring activities). A TIBCO Hawk Display should be installed on any computers you wish to use for monitoring the network or for building rulebases.
- **TIBCO Hawk Event Service.** The TIBCO Hawk Event Service is a process that records TIBCO Hawk alerts and changes in agent status. When communication with an agent is lost, the Event Service can invoke a user-provided script. Alerts and notifications can be recorded to log files or a database. Typically, the TIBCO Hawk Event Service is installed on a minimal number of computers in the network.

TIBCO Hawk is delivered with the TIBCO iProcess Engine to enable you to use it to monitor the TIBCO iProcess Engine server processes. For example, you can use TIBCO Hawk to check what iProcess Engine server processes are running or stop and start them.

For Information For information about TIBCO Hawk, access <http://docs.tibco.com>.

For information on configuring TIBCO Hawk for use with the TIBCO iProcess Engine, see *TIBCO iProcess Engine Administrator's Guide*.

TIBCO iProcess Analytics

Description TIBCO iProcess Analytics is a tool that enables organizations to analyze, evaluate and monitor business processes, providing insight into the performance of key processes and the knowledge needed to continuously improve them. iProcess Analytics combines current and historical process data and presents this data via intuitive dashboard views. An integrated early-warning system monitors all running instances and triggers alerts when deviations against planned values occur, thus enabling organizations to react quickly and take rapid corrective action.

iProcess Analytics uses a SQL RDBMS as the repository that saves all configurations and data. It is developed in Java as a client/server application.

TIBCO iProcess Analytics, in conjunction with [TIBCO iProcess Analytics Export](#), is used to analyze the performance of iProcess procedures. iProcess Analytics Export is used to export the data from the iProcess Engine to prepare it for import into iProcess Analytics. Data from multiple sources can be analyzed providing data from iProcess is part of that data.

For Information For information about iProcess Analytics, see the following:

- *TIBCO iProcess Analytics System Architecture Guide*
- *TIBCO iProcess Analytics LDAP User's Guide*
- *TIBCO iProcess Analytics Advanced User's Guide*
- *TIBCO iProcess Analytics ARIS Toolset Interface Guide*
- *TIBCO iProcess Analytics CTK Quick Start Guide*
- *TIBCO iProcess Analytics Delta Notes*
- *TIBCO iProcess Analytics Customization Guide*
- *TIBCO iProcess Analytics Data Import Guide*

TIBCO iProcess Analytics Export

Description TIBCO iProcess Analytics Export is a utility that is used to export cases from iProcess procedures. It generates an XML file that can be imported into [TIBCO iProcess Analytics](#).

TIBCO iProcess Analytics Export, in conjunction with [TIBCO iProcess Analytics](#), is used to analyze the performance of iProcess procedures.

For Information For information about iProcess Analytics Export, see the following:

- *TIBCO iProcess Analytics Export (Java) User's Guide*
- *TIBCO iProcess Analytics Export (Windows) User's Guide*

TIBCO iProcess Decisions Server

Description The TIBCO iProcess Decisions Server enables you to define and use TIBCO iProcess Decisions steps in your iProcess procedures. iProcess Decisions Service steps extend the functionality of your business process by enabling you to integrate with a powerful rules engine. See [TIBCO iProcess Decisions Studio](#).

iProcess data can be sent to the TIBCO iProcess Decisions Server to be processed by a set of business rules and the resulting data can then be passed back to iProcess.

For Information For information about the TIBCO iProcess Decisions Server, see *TIBCO iProcess Decisions Plug-in User's Guide*

TIBCO iProcess Decisions Studio

Description The TIBCO iProcess Decisions Studio is used to build Rule Sets. Once deployed to the [TIBCO iProcess Decisions Server](#), a rule set becomes a Decisions Service.

A Decisions Service automates a discrete decision-making task. It is implemented as a set of business rules and exposed as a Web Service (or Java Service). By definition, the rules within a Decisions Service are complete and unambiguous; for a given set of inputs, the Decisions Service addresses every logical possibility uniquely, ensuring “decision integrity”.

Using an iProcess Decisions Service step in your iProcess procedure enables you to interact with a decision service in the [TIBCO iProcess Decisions Server](#).

This means that you can:

- Send iProcess case data to the [TIBCO iProcess Decisions Server](#) where it can process the case data against the decision service rules.
- Capture the resulting data from the decision service back in iProcess.

For Information For information about the TIBCO iProcess Decisions Studio, see the following:

- *[TIBCO iProcess Decisions Start Here](#)*
- *[TIBCO iProcess Decisions DB Access Tutorial](#)*
- *[TIBCO iProcess Decisions Studio Quick Reference Guide](#)*
- *[TIBCO iProcess Decisions Rule Language Guide](#)*
- *[TIBCO iProcess Decisions Rule Modeling Guide](#)*
- *[TIBCO iProcess Decisions Rule Modeling Tutorial](#)*
- *[TIBCO iProcess Decisions Sample Application Guide](#)*
- *[TIBCO iProcess Decisions Server Integration & Deployment Guide](#)*
- *[TIBCO iProcess Decisions Server Deployment Tutorial](#)*

TIBCO iProcess Engine

Description The TIBCO iProcess Engine is the core of iProcess and is responsible for the execution of business processes. In response to various input, for example, case start, completion of steps, events being triggered, deadlines expiring and data being modified, coupled with knowledge of the business process, it makes the next activity or step in the process available for execution. This might be the dispatching of a work item to a user queue or the invocation of an application or task via an EAI step.

When a case of a procedure is started by an individual from a TIBCO iProcess Workspace (Windows) computer, a number of processes are used to process the information contained in one step before the next step can be performed.

The TIBCO iProcess Engine maintains the list of work items in a user's work queue for all the active cases to be processed. If the first step was opened, completed and released, the TIBCO iProcess Engine determines the subsequent actions for the step and updates the necessary case data in the database.

The iProcess Engine enables dynamic processing and knowledge-based interaction with iProcess procedures. This means iProcess procedures can be processed based on events as they occur and that may not be known about at procedure definition time. This feature is available using graft steps. A graft step enables an external application to graft (attach) one or more sub-procedures to a particular point in your procedure at run-time. Therefore, when a case of the main procedure is started, the external application can start a number of sub-processes which are attached to the main procedure via the graft step. For example, a financial application determines that a credit check and a transfer of funds are required as part of the main procedure. When another case is started, it determines that only a transfer of funds is required. This means that the procedure is dynamic and cannot be decided at procedure definition time. One of the processes is an iProcess sub-procedure and the other is a process run by the financial system.

iProcess procedures can be defined in TIBCO Business Studio or TIBCO iProcess Modeler.

All the TIBCO iProcess Engine case data such as fields and their values are stored in a TIBCO iProcess Engine database instance. An iProcess Engine database instance can be a SQL Server, Oracle or DB2 database.



The TIBCO iProcess Engine is a 32-bit application but can be used on both 32-bit and 64-bit operating systems, and against a 32-bit or 64-bit database.

The iProcess Engine (in conjunction with other iProcess products, for example, the iProcess Technology Plug-ins, iProcess Server Objects, and iProcess Web Services Plug-in) enables connectivity to external applications. This means that iProcess procedures can integrate with other systems in your business process. iProcess case data can be sent to external applications using whatever communication method is required. Data can also be passed back from the application to iProcess.

For Information For information about the TIBCO iProcess Engine, see the following:

- *TIBCO iProcess Engine: Architecture Guide*
- *TIBCO iProcess Engine Administrator's Guide*
- *TIBCO iProcess Engine (SQL) Administrator's Guide*
- *TIBCO iProcess Engine (Oracle) Administrator's Guide*
- *TIBCO iProcess Engine (DB2) Administrator's Guide*
- *TIBCO iProcess swutil and swbatch Reference Guide*
- *TIBCO iProcess Engine System Messages Guide*
- *TIBCO iProcess User Validation API: Reference Guide*

TIBCO iProcess Objects Server

Description The TIBCO iProcess Objects Server acts as a gateway between iProcess client applications developed with TIBCO iProcess Server Objects (Java or .NET), and the TIBCO iProcess Engine.

After communication is established between the client application and the TIBCO iProcess Objects Server, the TIBCO iProcess Objects Server waits for request messages from the client. When the TIBCO iProcess Objects Server receives a request message, it in turn makes calls to the TIBCO iProcess Engine to perform functions such as locking work items, moving work items to other work queues, writing data to the database, etc.



The TIBCO iProcess Workspace (Browser) is an example client application that sends requests to the iProcess Engine through the iProcess Objects Server. For more information about this client application, see [TIBCO iProcess Workspace \(Browser\)](#).

For Information For more information about the iProcess Objects Server, see *TIBCO iProcess Objects Server Administrator's Guide*.

TIBCO iProcess Objects Director

Description The TIBCO iProcess Objects Director is a standalone program that maintains a list of TIBCO iProcess Objects Servers that are configured in a node cluster. When a client application needs access to a TIBCO iProcess Objects Server, it first establishes a connection to the TIBCO iProcess Objects Director. The TIBCO iProcess Objects Director then decides, based on a “pick method,” which TIBCO iProcess Objects Server the client should connect to.

The list of known TIBCO iProcess Objects Servers is updated dynamically as TIBCO iProcess Objects Server instances are started and stopped. The TIBCO iProcess Objects Director maintains this list by checking the process_config table of the iProcess Engine to which it is associated.

For Information For information about using and configuring the TIBCO iProcess Objects Director, see *TIBCO iProcess Objects Director Administrator’s Guide*.

TIBCO iProcess Server Objects

Description TIBCO iProcess Server Objects is an Application Programming Interface for iProcess. It provides access to all of the information and functionality required to write either a BPM user oriented client application or a batch oriented broker application.

It comprises a set of objects that are used to build applications that automate business processes. TIBCO iProcess Server Objects consists of an object model that provides access to the information and functionality needed in these applications.

The objects in the TIBCO iProcess Server Objects object model can be used to start cases, present information on screens to users, manipulate work items, remind users when actions need to be taken, and monitor and control the flow through the business process.

TIBCO iProcess Server Objects is designed to be used in server-side application architectures.

Client applications make use of the objects in the TIBCO iProcess Server Objects by making method calls that either retrieve or modify data. These method calls cause messages to be sent to a TIBCO iProcess Objects Server. The TIBCO iProcess Objects Server acts as a gateway between the client application created with TIBCO iProcess Server Objects, and the TIBCO iProcess Engine, where the actual processing and storage of data occurs. The TIBCO iProcess Engine manages all data, routing work items and updating the appropriate work queues.

There are two types of implementation available for iProcess Server Objects:

- TIBCO iProcess Server Objects (.NET)
- TIBCO iProcess Server Objects (Java)



Both 32-bit and 64-bit variants of TIBCO iProcess Server Objects (.NET) and TIBCO iProcess Server Objects (Java) are available. These variants allow customers to write applications that can interface with TIBCO iProcess Suite on 32-bit and/or 64-bit Windows platforms as required.

For Information For information about TIBCO iProcess Server Objects, see the following:

- *TIBCO iProcess Server Objects (JAVA): Programmer's Guide*
- *TIBCO iProcess Server Objects (.NET): Programmer's Guide*

TIBCO iProcess Technology Plug-ins

The TIBCO iProcess Technology Plug-ins is a package that includes a number of individual products. These products have been grouped together to ease their installation.

The TIBCO iProcess Technology Plug-ins include the following individual plug-ins:



To improve the installation process, a new assembly-based combined installer was introduced in Version 10.6, giving a standard installation process for all the iProcess Technology Plug-ins. Previously, these products were installed separately.

- **TIBCO iProcess Java Server Plug-in** - This plug-in enables the iProcess Engine to process EAI Java steps that have been added to the iProcess procedure. An EAI Java step enables you to design an iProcess procedure so that you can call out to a custom Java object to perform some additional work. For example, you can create a custom Java object to call business methods in Enterprise Java Beans (EJBs), call a database via a JDBC connection, or post a message on a Java Message Service (JMS) queue.
- **TIBCO iProcess BusinessWorks Server Plug-in** - This plug-in provides a way of initiating a BusinessWorks process from the TIBCO iProcess Engine. It provides the communication mechanism that allows the iProcess Engine to make calls to TIBCO BusinessWorks. Enabling the iProcess Engine to integrate with BusinessWorks provides a highly performant and versatile way of integrating with other applications and technologies.
- **TIBCO BusinessWorks iProcess Plug-in** - This plug-in provides a set of resources that allow a TIBCO BusinessWorks process to communicate with the TIBCO iProcess Engine. For example, a BusinessWorks process can start or suspend a case of an iProcess Engine procedure. The BusinessWorks iProcess Plug-in includes facilities that enable you to create a form using TIBCO BusinessWorks FormBuilder, and then provide it to the iProcess Workspace (Browser).
- **TIBCO iProcess XML Transform Server Plug-in** - This plug-in provides the facility to transform XML to/from iProcess field data and/or a designated URL. For example, an EAI Transform step could be used to take XML data from an iProcess memo field, apply a transformation to the data, and pass the result to another memo field. Equally, XML data can be parsed and mapped onto discrete iProcess fields.

For Information For information about the TIBCO iProcess Technology Plug-ins, see the following:

- *TIBCO iProcess Java Plug-in User's Guide*
- *TIBCO iProcess BusinessWorks Connector User's Guide*
- *TIBCO Business Studio Process Modeling User's Guide*

TIBCO iProcess Web Services Plug-in

Description The main function of the iProcess Web Services Plug-in is to provide an interface for both inbound and outbound communication between the iProcess Engine and external applications:

- Outbound - iProcess procedures make calls to external applications to perform some operation.
- Inbound - External applications make calls to iProcess to perform operations such as starting cases, triggering events or suspending cases.

The iProcess Web Services Plug-in consists of three components:

- **TIBCO iProcess Web Services Server Plug-in.**

This consists of:

- Jetty - This is a Java HTTP Server and Servlet Container that runs on a Java Virtual Machine (JVM) and the components contained in Jetty handle communication for both inbound and outbound calls.
- iProcess Engine Interface - This consists of an EAI Plug-in. It allows the iProcess background processes to communicate with Jetty.

- **TIBCO iProcess Web Services Client Plug-in**

This plug-in needs to be installed on your client machine that hosts your iProcess Workspace (Windows) and iProcess Modeler. This plug-in enables you to define EAI Web Service steps in your iProcess procedures.

For Information For information about the TIBCO iProcess Web Services Plug-in, see *TIBCO iProcess Web Services Plug-in User's Guide*.

TIBCO iProcess Workspace (Browser)

Description The TIBCO iProcess Workspace (Browser) is a client application that enables users to participate in business processes. It enables users to view and perform the work that has been assigned to them. This means that the movement of data/information is automated through your business process, whatever your business may be. For example, it may automate the process of filing an insurance claim, a loan application, or any process that has multiple steps.

The TIBCO iProcess Workspace (Browser) application is run in a browser (e.g., Microsoft Internet Explorer). You connect to a TIBCO iProcess Engine over the internet or an intranet by entering a web address that is determined by the location at which your TIBCO software is installed.

The type of functions you can perform are:

- Open and view work queues
- Apply filter and/or sort parameters to a work queue so only the desired work items are listed in the desired order.
- Open and view work items
- Enter information into forms that are displayed when you open a work item.
- Either “keep” a work item, which closes it and places it back in the work queue, or “release” a work item, which causes the case to advance to the next step in the procedure.
- Forward work items to a different work queue (user or group).
- Start a case of a procedure.
- Suspend or reactivate a case.
- View the history of a case.
- Allow temporary access to your personal work queue to another user.
- Temporarily redirect work items from your work queue to another work queue.

For Information For information about the TIBCO iProcess Workspace (Browser), see the following:

- *TIBCO iProcess Workspace (Browser) User's Guide*
- *TIBCO iProcess Workspace (Browser) Configuration and Customization Guide*
- *TIBCO iProcess Workspace (Browser) Action Processor Reference Guide*

TIBCO iProcess Workspace Lite

Description TIBCO iProcess Workspace Lite is a Business Process Management (BPM) client application that runs over the Internet or an intranet. As a BPM application, it automates the movement of data or information through your business process. For example, it automates the process of filing an insurance claim, a loan application, or any process that has multiple tasks.

It allows people with disabilities to access iProcess BPM functionality using standard accessibility aids, such as screen readers and magnifiers. It also allows selection and entry of data through the keyboard, as some disabilities prevent use of a mouse.

The type of functions you can perform are as follows:

- Customize the appearance of the iProcess Workspace Lite screen so that it is easier to view.
- Open and view work queues, which list the work items you are authorized to work on.
- Process work items, which includes:
 - Opening work items.
 - Entering information into forms that are displayed when you open a work item.
 - "Closing" the work item, which saves any entered data, closes the work item, and places it back in the work queue so it can be worked on later.
 - "Submitting" the work item, which saves any entered data, closes the work item, and causes the case to advance to the next task in the process.
 - "Cancelling" the work item, which discards any entered data, closes the work item, and places it back in the work queue so it can be worked on later.
 - "Forwarding" the work item, which removes the work item from your work queue and adds it to the work queue of another user or group.
- List the processes that are available on the system.
- Start a case of a process.
- List the current cases for a given process.
- View the audit trail (case history) for a given case.
- Close a case, which stops the process flow in that case.

For Information For information about TIBCO iProcess Workspace Lite, see the following documents:

- *TIBCO iProcess Workspace Lite User's Guide*
- *TIBCO iProcess Workspace Lite Customization Guide*

TIBCO iProcess Workspace Plug-ins

TIBCO iProcess Workspace Plug-ins is a package that includes a number of individual products. These products have been grouped together to ease their installation.

The iProcess Workspace Plug-ins include the following individual plug-ins:



To improve the installation process, a new assembly-based combined installer was introduced in Version 10.6, giving a standard installation process for all the iProcess Workspace Plug-ins. Previously, these products were installed separately.

- **iProcess Java Client Plug-in** - This plug-in enables the iProcess Engine to define EAI Java steps to be added to an iProcess procedure. An EAI Java step enables you to design an iProcess procedure so that you can call out to a custom Java object to perform some additional work. For example, you can create a custom Java object to call business methods in Enterprise Java Beans (EJBs), call a database via a JDBC connection, or post a message on a Java Message Service (JMS) queue.
- **iProcess BusinessWorks Client Plug-in** - This plug-in enables you to define EAI BusinessWorks steps in an iProcess procedure. EAI BusinessWorks steps allow an iProcess Engine procedure to invoke a TIBCO BusinessWorks process definition.
- **iProcess Conductor Client Plug-ins (Order and Orchestration)** - The Orchestration plug-in enables you to define Orchestration steps in your iProcess procedures. The Order plug-in enables you to define Order steps in your iProcess procedures.
- **iProcess Conductor XML Transform Client Plug-in** - This plug-in enables you to define Transform steps in your iProcess procedures.

For Information: For information about the TIBCO iProcess Workspace Plug-ins, see the following:

- *TIBCO iProcess Java Plug-in User's Guide*
- *TIBCO iProcess BusinessWorks Connector User's Guide*
- *TIBCO iProcess Conductor Implementation Guide*

TIBCO iProcess Workspace (Windows)

TIBCO iProcess Workspace (Windows) provides the ability to define and manage procedures.

The TIBCO iProcess Workspace (Windows) consists of:

- The Work Queue Manager
- The Procedure Manager

The Work Queue Manager displays users' queues and work items and the Procedure Manager displays all of the procedures currently available. From the Procedure Manager, you can:

- Create and edit procedures. The TIBCO iProcess Modeler is started from TIBCO iProcess Workspace (Windows).
- Organize and manage your procedures as a hierarchical structure of procedure libraries, in the same way as, for example, you manage files and directories.

TIBCO iProcess Workspace (Windows) includes:

- **TIBCO iProcess Modeler** - gives you a visual representation of your business process that is easy to follow and that can be enhanced or amended at any time. The TIBCO iProcess Modeler builds on the familiar flowchart metaphor to show in an unambiguous manner, the flow of work for a particular business process. The rules that you define graphically are stored by the iProcess Engine and can then be deployed across a wide ranging hardware architecture.

The TIBCO iProcess Modeler is automatically started by TIBCO iProcess Workspace (Windows) when you want to create or edit procedures. From the TIBCO iProcess Modeler you can access the Step Definer, which enables you to design the forms for each step in your procedure. The forms are the part of the step seen by the person who receives the work item in their queue. The forms contain text and fields into which users can enter information for a particular case, or instance, of a procedure.



Business processes can also be modelled in TIBCO Business Studio. See [TIBCO Business Studio on page 5](#) for more information.

- **TIBCO iProcess Administrator** - The TIBCO iProcess Administrator enables you to perform the tasks involved to administer iProcess Workspace (Windows), such as managing users, cases and work queues. It is a graphical utility that can be used on any iProcess Workspace (Windows) to perform common administration functions on the TIBCO iProcess Engine.

For Information For information about TIBCO iProcess Workspace (Windows), see the following:

- *TIBCO iProcess Workspace (Windows) User's Guide*
- *TIBCO iProcess Workspace (Windows) Manager's Guide*
- *TIBCO iProcess Modeler - Getting Started*
- *TIBCO iProcess Modeler - Basic Design*
- *TIBCO iProcess Modeler - Advanced Design*
- *TIBCO iProcess Modeler - Procedure Management*
- *TIBCO iProcess Modeler - Integration Techniques*
- *TIBCO iProcess Expressions and Functions Reference Guide*

TIBCO Rendezvous

Description Rendezvous software makes it easy to create distributed applications that exchange data across a network. You get software support for network data transport and network data representation. Rendezvous software supports many hardware and software platforms, so programs running on many different kinds of computers on a network can communicate seamlessly.

From the programmer's perspective, the Rendezvous software suite includes two main components—a Rendezvous programming language interface (API) and the Rendezvous daemon.

Rendezvous software includes several programming language interfaces, which are efficient, easy to use, and compatible with most other libraries (including window systems).

The Rendezvous daemon runs on each participating computer on your network. All information that travels between program processes passes through the Rendezvous daemon as the information enters and exits host computers. The daemon also passes information between program processes running on the same host.

Rendezvous programs are programs that use Rendezvous software to communicate over a network.

A Rendezvous distributed application system is a set of Rendezvous programs that cooperate to fulfill a mission.

TIBCO Rendezvous is delivered with the TIBCO iProcess Engine to enable the iProcess Engine to communicate with other TIBCO products, for example, TIBCO BusinessWorks and TIBCO iProcess Insight.

For More Information For more information about TIBCO Rendezvous, access <http://docs.tibco.com>.

For information on configuring TIBCO Rendezvous for use with the TIBCO iProcess Engine, see *TIBCO iProcess Engine Administrator's Guide*.

Chapter 2

Introduction to the TIBCO iProcess Engine

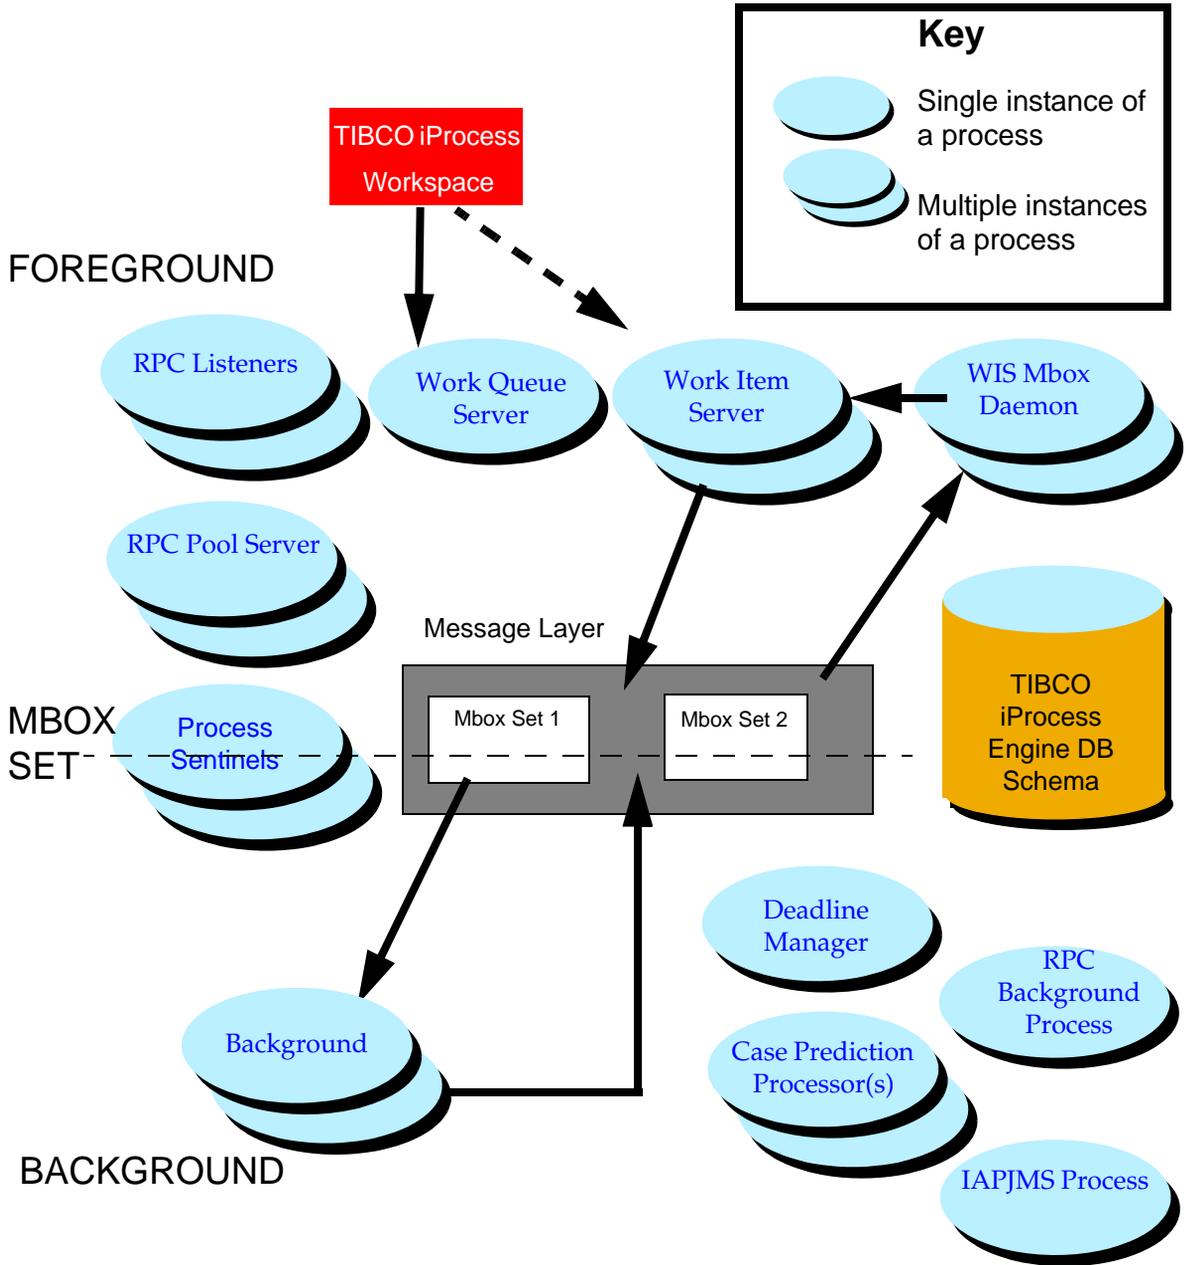
This chapter provides an overview of the TIBCO iProcess Engine architecture and the server processes that are involved to process business process information. It also contains some important TIBCO iProcess Engine concepts, which you need to be familiar with when administering the TIBCO iProcess Engine.

Topics

- [TIBCO iProcess Engine Architecture, page 30](#)
- [The Role of the TIBCO iProcess Engine, page 31](#)
- [iProcess Physical Architecture, page 32](#)
- [TIBCO iProcess Engine Process Structure, page 35](#)
- [Where is TIBCO iProcess Engine Case Data Stored?, page 39](#)
- [How Do Work Items Appear in Work Queues?, page 40](#)
- [24*7 TIBCO iProcess Engine Operation, page 44](#)

TIBCO iProcess Engine Architecture

Click on a process to find out more information



The Role of the TIBCO iProcess Engine

The following example demonstrates the role of the TIBCO iProcess Engine in the business process automation environment. In a typical business process where an item is purchased from an office supplies company you would go through a number of steps:

1. Obtain authorization.
2. Allocate requisition number.
3. Submit expense claim.

When a case of the procedure is started by an individual from a computer running TIBCO iProcess Workspace, a number of processes are used to process the information contained in one step before the next step can be performed.

The TIBCO iProcess Engine maintains the list of work items in a user's work queue for all the active cases to be processed. If the first step was opened, completed and released, the TIBCO iProcess Engine determines the subsequent actions for the step and updates the necessary case data in the database.

In this example, after the Authorization step is released, the TIBCO iProcess Engine reads the procedure definition and determines that the next step is the Allocate Requisition Number. The addressee of this step is determined and a work item with the relevant form is sent back to the appropriate work queue for the TIBCO iProcess Workspace to display.

iProcess Physical Architecture

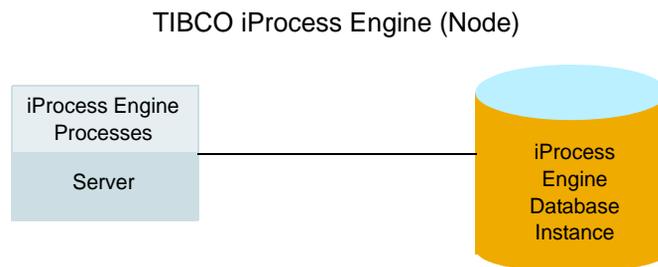
The following section describes:

- installing the TIBCO iProcess Engine on a single server
- installing the TIBCO iProcess Engine on multiple servers

Installing the iProcess Engine on a Single Server

A TIBCO iProcess Engine installation on one server is known as a **TIBCO iProcess Engine node**. All the iProcess Engine processes are run on a single server.

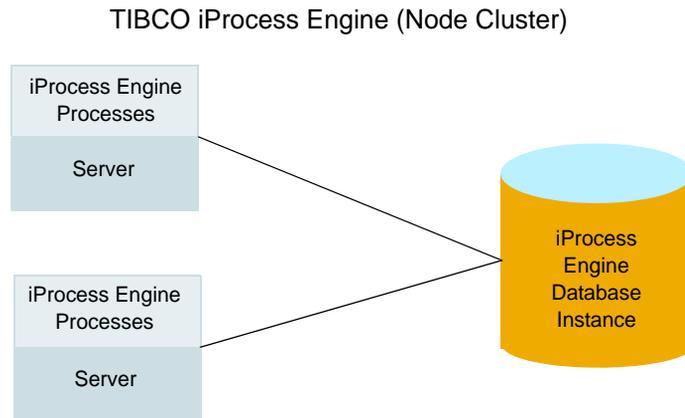
The following diagram shows a TIBCO iProcess Engine (node) that comprises of one server.



Installing the iProcess Engine on a Node Cluster

Installing the TIBCO iProcess Engine on multiple servers that all use the same database instance is known as a **node cluster** architecture. You can convert from a single-server to a node cluster at any time simply by adding another server to the installation. In the node cluster architecture, you can have a number of TIBCO iProcess Engine processes running on different servers, but they act as though they are a single iProcess Engine. Node clusters improve load balancing and performance. However, all the TIBCO iProcess Engine case data such as fields and their values in forms will be stored in one TIBCO iProcess Engine database instance.

The following diagram illustrates a TIBCO iProcess Engine (node cluster) that uses two servers to run TIBCO iProcess Engine processes. They can all access iProcess case data from the same TIBCO iProcess Engine database instance.



Note that:

- A complete installation of the iProcess Engine needs to be performed on each machine in the node cluster.
- All the iProcess Engines point to a single database instance. Within the database there is a single engine configuration that specifies which executables are configured to run on each server.
- The iProcess sentinels are configured to run on each machine. One server in the cluster is designated as the master server and the iProcess Sentinels on this server co-ordinate the startup, shutdown and management of the processes on all machines in the cluster.
- The iProcess Sentinels on the slave servers manage the executables designated to run on that machine.
- All servers that form part of the node cluster are required to provide full operation. See [The iProcess Engine and Hardware Clustering on page 44](#) for information about failover capability for the iProcess Engine.

TIBCO iProcess Workspace and TIBCO iProcess Engine Communication

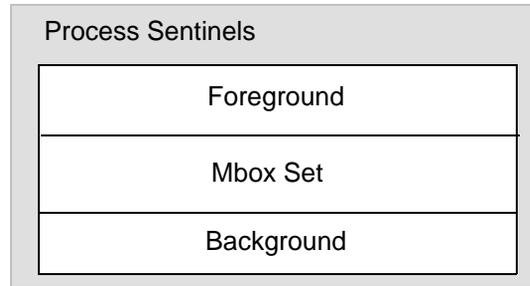
The TIBCO iProcess Engine uses a client/server model for communication where there is a two way communication path between each TIBCO iProcess Workspace client and TIBCO iProcess Engine server. The communication protocol used for this communication is called **Remote Procedure Call (RPC)**.

Information keyed in at the client needs to be passed to the TIBCO iProcess Engine for processing, for example when completing an insurance claim form, the new or updated data needs to be processed, stored, or updated so that subsequent steps can use the information. Information also needs to be passed from the TIBCO iProcess Engine to one of the TIBCO iProcess Workspaces, for example to tell the client what the next step in the business process is or show previous details about a work item that has already been entered.

Refer to [Network Communication on page 125](#) for more information about TIBCO iProcess Workspace and TIBCO iProcess Engine communication.

TIBCO iProcess Engine Process Structure

In order to understand the flow of data through the TIBCO iProcess Engine and the relationship of TIBCO iProcess Engine processes, you can think of the TIBCO iProcess Engine as being split into the following four areas.



A TIBCO iProcess Engine process operates in either the foreground or background layer. The Mbox set provides the communication link between the foreground and background processes. This is the Messaging layer where instructions are sent from the foreground to the background and vice versa.

The foreground area deals with the user interaction, list of queues and their contents. The background processes perform the case processing.

The Process Sentinels keep control of all the processes to make sure they are always running.

Process Sentinels

The Process Sentinels (*SWDIR\etc\procmgr*) are responsible for controlling all the TIBCO iProcess Engine processes. If a node cluster architecture is used, then the Process Sentinels will exist on each server to manage the processes running on that server. This process operates in different ways depending on whether you are using a third party transaction processing monitor.

In a node cluster, the Process Sentinels have a master process on one of the servers, which manages and controls slave Process Sentinels on all the other servers in the node cluster.

The Process Sentinels control the safe start up and shut down of processes on one or multiple servers i.e. making sure processes are started in the correct sequence.

The Process Sentinels can also monitor processes to check that they are running and can automatically restart them if they fail. It can also detect if `SWDIR\logs\sw_error` and `SWDIR\logs\sw_warn` files have been created and sends a work item to the administrator to inform them that these files have been created.

Refer to [Process Management on page 111](#) for more information about the Process Sentinels.

Foreground Processes

These processes are responsible for communicating with the TIBCO iProcess Workspaces and for passing any TIBCO iProcess Workspace requests such as released work items to the background area for processing.

The following table shows the processes that operate in the foreground:

Process Description	Logical Process Name
Work Queue Server	WQS
Work Item Server	WIS
WIS Mbox Daemon	WISMBD
RPC listeners	RPC_TCP_LI and RPC_UDP_LI
RPC pool server	RPC_POOL

Refer to [Foreground Processes on page 62](#) for more information about what these processes do.



All of the foreground processes *must* operate on the master server. See [Determining Where Processes Run on page 118](#) for more information.

Mbox Sets

The Mbox Sets are responsible for communicating information between the foreground and background processes. It is essentially the messaging system that stores and processes messages from TIBCO iProcess Workspace or TIBCO iProcess Engine requests. Mbox sets comprise of one or more message queues.

The messaging system used is:

- *in the UNIX Oracle and Windows Oracle versions*, Oracle AQ. This is a transactional messaging system provided by the Oracle database.

- *in the UNIX DB2 and Windows SQL Server versions, the iProcess Suite's own transactional messaging system, which uses database tables provided by the DB2 or Microsoft® SQL Server database.*

The messaging system provides the ability to store messages in tables. This means that message queues are highly scalable and have a guaranteed delivery mechanism.

See [iProcess Mbox Sets on page 89](#) for more information about the Mbox sets.

Background Processes

The background processes are responsible for processing message instructions received from the clients such as releasing a step or forwarding a step. They also monitor and process any deadlines that have been set up in the procedure and manage case prediction.

The following processes operate in the background:

Process Description	Logical Process Name
Background	BG
Case Prediction	BGPREDICT
Database Queue Daemon	DBQD ^a
Deadline Manager	DLMGR
IAPJMS	IAPJMS
RPC Background	RPCBG

a. This process is only present on the DB2 version of the TIBCO iProcess Engine.

Refer to [Background Processes on page 72](#) for more information about what these processes do.

Event Handling

The TIBCO iProcess Engine uses a publish/subscribe event mechanism to handle the following inter-process tasks:

- notifying processes to update caches.
- synchronization of process startup and shutdown.

The event handling mechanism used depends on the TIBCO iProcess Engine variant used, as follows:

- *UNIX Oracle version.* Events are handled using Oracle AQ's publish/subscribe interface.
- *Windows Oracle and Windows SQL Server versions.* Events are handled using the **iProcess Events** COM+ application, which is installed with the TIBCO iProcess Engine. All processes that want to subscribe to events register with the **iProcess Events** COM+ application.



If you are using a node cluster architecture, the **iProcess Events** COM+ application runs by default on the machine on which you installed the master server. If performance is or becomes an issue, TIBCO recommend that you dedicate a separate server, which is not running any TIBCO iProcess Engine processes, as the *event server*. This will reduce the load on the master server.

The name of the event server is stored in the Windows registry, so you will need to change this if you move the COM+ application to a new server. (See the *TIBCO iProcess Engine Administrator's Guide* for more information.)

- *UNIX DB2 version.* Events are handled using iProcess event/notification daemons.

The event daemon is a single thread created by the Process Sentinel watcher process. The event daemon:

- tracks which processes have subscribed to which events. (It maintains this information both in memory and in the **sw_subscription_list** database table.)
- receives events when they are published.
- forwards notification of an event to the notification daemon in each subscribed process.

A notification daemon is a thread created by any TIBCO iProcess Engine process that wishes to receive event notifications.

Where is TIBCO iProcess Engine Case Data Stored?

All iProcess case data (the information provided by users in forms such as customer details, orders, etc.) is stored in the TIBCO iProcess Engine database instance. This means that all important iProcess data can be backed up and restored as part of your regular database backup policy.

How Do Work Items Appear in Work Queues?

For any work item to appear in a user's work queue, a work item record has to exist on the server for the TIBCO iProcess Workspace to know about it. The record contains information about the work item such as its case number, where in the database the fields and values are stored, the addressee and procedure, etc.

The WIS process contains work item records that define what work items are in the queue. The WIS reads the **staffo** database table which contains all the work item records where one row in the table represents one work item record. Each work item entry in the table contains specific information about the work item such as its case number, unique identifier, case description, name and node of case starter, etc. The TIBCO iProcess Workspace also calls the RPC server which makes requests to the database tables to retrieve fields and their values and any memos and attachments.

Sending Instructions From the TIBCO iProcess Workspace to the TIBCO iProcess Engine

When a user processes a work item by either Keeping or Releasing it, the TIBCO iProcess Workspace creates a package of instructions that are sent to the TIBCO iProcess Engine. This package contains all the information about the work item, such as what fields and data are shown on the form, what fields may have changed, and the instruction to be performed.

The message is sent from the TIBCO iProcess Workspace to the TIBCO iProcess Engine in packets via RPC calls. The package is split up into its various components so that an instruction message is sent to the background and any case data updates are made by updating the database. The instruction message is enqueued to a Mbox message queue where it waits to be processed by one of the background processes.

The Background process reads the message queues for new messages so it dequeues the message from the queue and then processes the case instructions. The Background performs the necessary actions such as update the changed fields in the database and determine what the next step is and who to send it to.

Mbox Sets

The message system is able to distribute messages across physical Mbox queues, which means that high volumes of messages can be handled. The Mbox queues are grouped into logical sets (Mbox sets) with each set configured for a specific purpose. Different processes require access to different Mbox sets, for example the instruction processor requires write access to a background Mbox, and WIS Mbox.

Message Queues

All TIBCO iProcess messages are posted to a central queue repository (Mbox). This is hosted by:

- *in the UNIX Oracle or Windows Oracle versions*, a set of tables using Oracle Advanced Queues (Oracle AQ).
- *in the UNIX DB2 or Windows SQL Server versions*, a set of iProcess queue tables in the database.



The SQL Server database queue tables can be created in the same database as iProcess, or a different database on the same SQL Server. Refer to the TIBCO iProcess Engine Installation Guide for more information about the location of database queue tables.

The Mbox can be configured to be a number of queues in a number of database tables, and even split across physical disks. Therefore, even though the Mbox appears as a single queue, it could consist of multiple queues in multiple tables across multiple disks.

The messaging system provides an assured delivery mechanism for instruction and request messages between the foreground and background.

User Access to iProcess Engine Work Queues

The total number of work queues a single iProcess user can have access to is **32,767**.

If the number of work queues a iProcess user has access to exceeds 32,767, the iProcess user is not able to login to the iProcess Engine and the following error message is displayed:

Sal Queues Interface 'sal_queues_count()' returned fatal error -n. Called by wqsess_lcreat_folder.

Abort WorkQs?

where n is the total number of work queues.

This only affects the single user who has access to too many work queues. All other iProcess users will be able to log in to the iProcess Engine without any problems.

An iProcess user can have access to the following work queues:

- Their **Personal Queue** which only they can access.
- Any **Group Queues** that they are a member of.
- Any **Test User Queues** or **Test Group Queues**, if the user is allowed to define procedures.

An iProcess user can also be given access to work queues in the following ways:

- By being a supervisor of a work queue.
- By being a participant of a work queue.

For example, the following table illustrates the total number of work queues a user called **swusr001** has access to:

Work Queues that swusr001 has access to:	
Personal Queue	swusr001
Group Queue	group1
Test User Queue	swusr001
Test Group Queue	group1
Supervisor of	swusr002, swusr003
Participant of	swusr002
Total number of work queues that swusr001 has access to:	7

From the example above, you can see that the number of work queues a user can have access to can quickly increase if that user is a member of several groups and is a supervisor or participant of several queues.

This means it is important to plan in advance the number of work queues you want to create and how many supervisors and participants you are going to need. For example, if you want to create 17,000 work queues on your TIBCO iProcess Engine, you could not make a single user both a supervisor and a participant of all 17,000 work queues. This is because that user would then have access to 34,000 work queues which would exceed the limit.

If you do want to create a large number of work queues and you want to have supervisors and participants for them, you would have to split them between more than one user. For example, if you had 17,000 work queues you could make:

- **swusr001** a participant of all 17,000 work queues
- **swusr002** a supervisor of all 17,000 work queues.

24*7 TIBCO iProcess Engine Operation

The TIBCO iProcess Engine is designed to continue running non-stop even when administration tasks need to be performed. This is achieved by:

- all case data being stored in a database. Online backup and restore is performed via your database tools.
- automatically restarting TIBCO iProcess Engine processes if they fail at any time. This is achieved by the TIBCO iProcess Engine Process Sentinels monitoring all the processes that are running to make sure they are running and if not, restart them.
- TIBCO Hawk, supplied with the iProcess Engine, a monitoring system that enables you to monitor the TIBCO iProcess Engine server processes. For example, you can use TIBCO Hawk to check what iProcess Engine server processes are running or stop and start them.
- New or updated procedures and processes can be deployed without affecting work in progress.

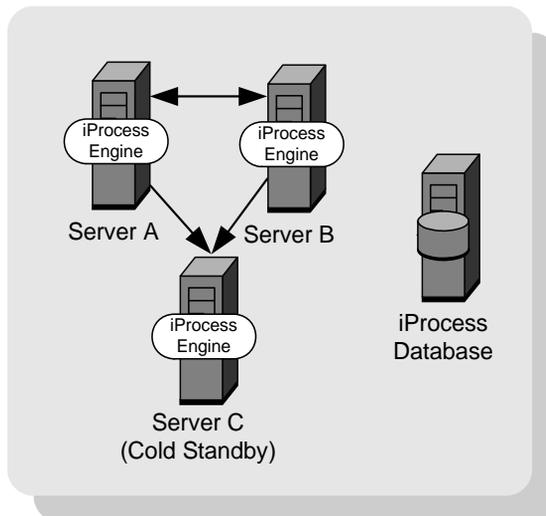
The iProcess Engine and Hardware Clustering

To cater for the loss of a physical server, the iProcess Engine can be deployed with a hardware cluster. The following section describes:

- Deploying the iProcess Engine in a hardware cluster for UNIX
- Deploying the iProcess Engine in a hardware cluster for Windows

Deploying the iProcess Engine in a Hardware Cluster for UNIX

There are two methods of deploying a hardware cluster, active-active and cold standby. The iProcess Engine is not supported in an active-active hardware cluster. However, it is supported in a cold standby hardware cluster. An example of a cold standby hardware cluster is shown below:



Each server is built in exactly the same way. Servers A and B are up and running and server C is a cold standby. If server A or B fails, server C is started and takes the place of the failed server.

To deploy the iProcess Engine in a cold standby hardware cluster:

- An iProcess Engine should be installed on each server in the cluster.
- If Server A or B fails, start the cold standby server and use `SWDIR\swadm move_server` to move the iProcess Engine processes from the failed to server to the cold standby server. See "Administering Servers" in the [TIBCO iProcess Engine Administrator's Guide](#) for more information about the `SWDIR\swadm move_server` command.



This can be done automatically by writing a script that uses the `SWDIR\swadm move_server` command.

Deploying the iProcess Engine in a Hardware Cluster for Windows

The iProcess Engine can be used with Microsoft Server Clustering when deployed in a Windows environment. Microsoft define a cluster as "a group of independent computers that work together collectively to provide a common set of services". When combining an iProcess node cluster with a Microsoft Server Cluster, each iProcess Engine node is contained within its own virtual server. Each virtual server is run on its own physical server. The iProcess Engine is installed on a shared disk that is accessible to all servers in the cluster. It is configured to work within the virtual server at installation by specifying a Cluster Network Name. If a physical server fails, the Microsoft Server Cluster automatically relocates the virtual server onto a secondary physical server and restarts the iProcess Engine.

To install the iProcess Engine in a Microsoft Windows Cluster, see either the *TIBCO iProcess Engine (Oracle) for Windows Installation* or the *TIBCO iProcess Engine (SQL) for Windows Installation* guide, depending on the version of the database you are using.

Chapter 3

Using the TIBCO iProcess Suite in a Multilingual Environment

This chapter describes how globalization is supported by TIBCO iProcess Suite. It also describes the steps you should follow to implement TIBCO iProcess Suite in an international environment and the advantages and costs of the different ways in which TIBCO iProcess Suite can support globalization.

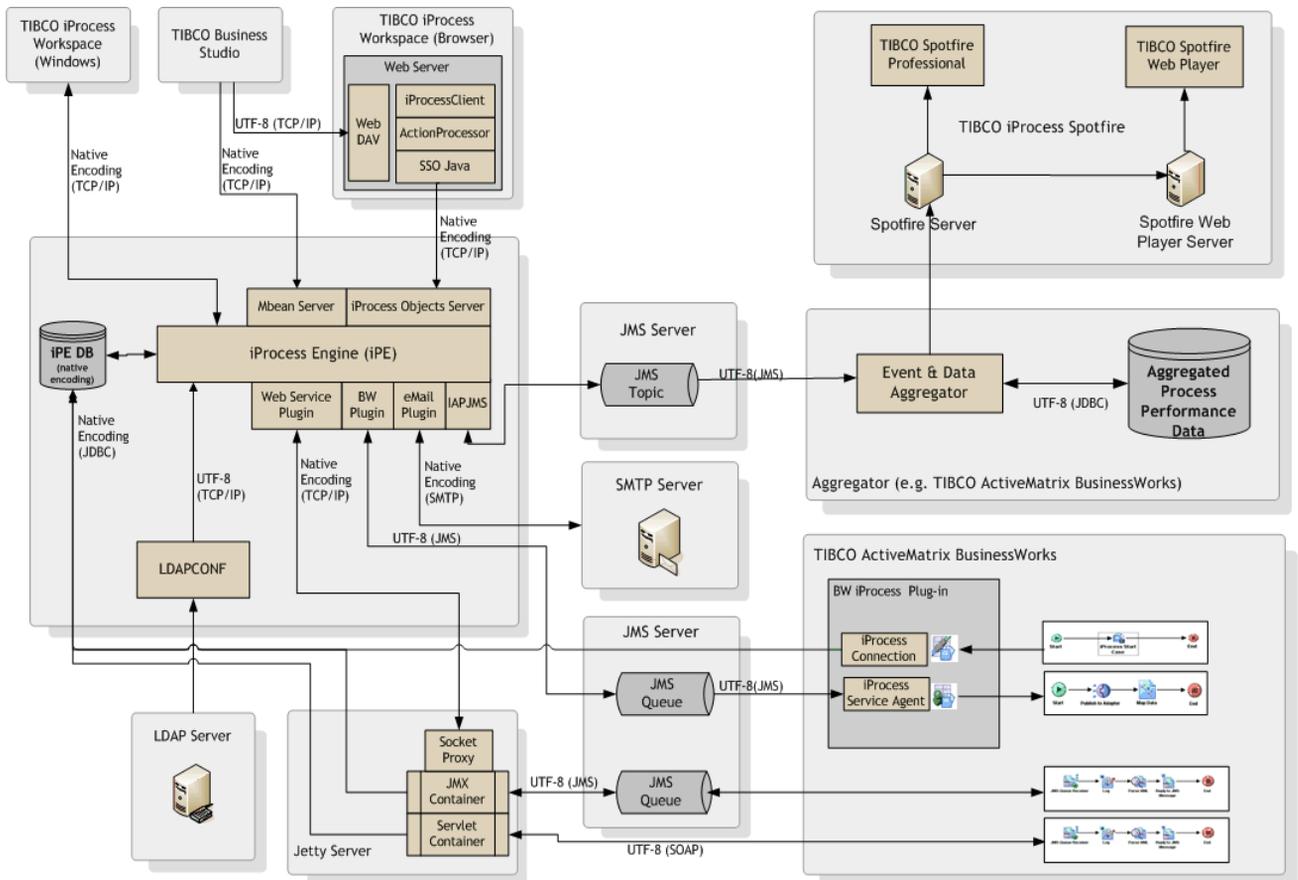
Topics

- [Overview, page 48](#)
- [Globalization Options in the iProcess Suite, page 50](#)
- [Recommendations, page 53](#)
- [Configuring the iProcess Suite Using UTF-8, page 54](#)
- [Globalization Support Using Native Encoding, page 58](#)

Overview

TIBCO iProcess Suite consists of many different products. A typical environment could be made up of any number or combination of iProcess products.

The diagram below illustrates how data is transmitted between products in the iProcess Suite.



Note that:

- The diagram indicates that some components, including the iProcess Engine can use either native encoding or UTF-8. Depending on the encoding used by other components in the iProcess Suite, the iProcess Engine performs the necessary data conversions. For information about setting the encoding attributes for iProcess components, refer to the documentation supplied with the components.

- The iProcess database can use either UTF-8 or native encoding. This is a choice that you can specify when you install or upgrade the iProcess Engine.
 - If you choose to use UTF-8 for the iProcess database, the iProcess Engine also operates in UTF-8.
 - If you choose not to use UTF-8 for the iProcess database, the iProcess Engine uses native encoding.

See the *TIBCO iProcess Installation* guide for your operating system for setting this option. Before deciding on an encoding, make sure you understand the issues discussed in this chapter.

- Each component that communicates with an external application must use the same character set as the external application. See [Implementing iProcess Suite in an International Environment with Native Encoding](#) on page 60.

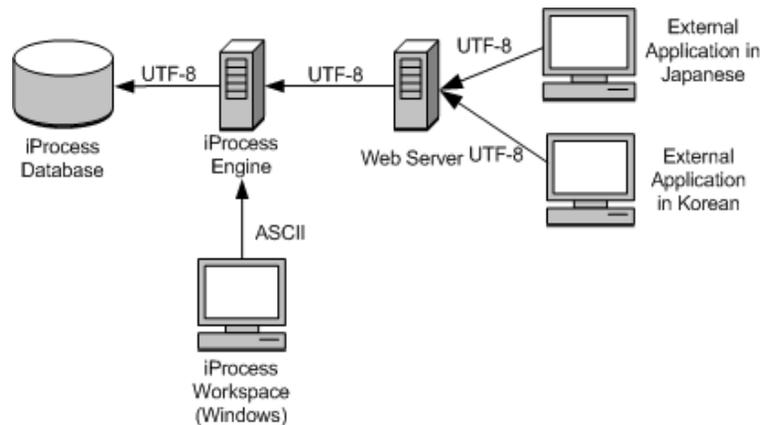
Globalization Options in the iProcess Suite

From version 11.1, the TIBCO iProcess Suite supports Unicode (UTF-8) character encoding. At installation time, you can enable UTF-8 support, or continue to use native encoding. See the TIBCO iProcess Engine Installation guide for your operating system for information on setting this option.

Advantages of UTF-8 Encoding

TIBCO iProcess Suite's support for UTF-8 (or UCS-2 in the case of SQL Server) enables you to work not simply in an international, but in a multi-national environment, using more than one multi-byte character set simultaneously.

This means that one TIBCO iProcess site can host multiple regions. For example, an enterprise operating globally could support data in Chinese, Japanese, Russian and English character sets in the same database. Work can be partitioned so that an end user working in this installation will normally see only work in his or her language:



In this example, you must ensure that users of the Japanese and Korean application do not receive each other's work items because the characters will not display correctly.

This is not possible without UTF-8 support because an iProcess Engine installation can support only one native character set. This option also simplifies using external UTF-8 applications. The conversions described in [Globalization Support Using Native Encoding on page 58](#) are not necessary if the iProcess Engine supports UTF-8.

An additional advantage of using UTF-8 is that WIS searches using wildcard characters are also more consistent. For example, the expected results for filtering work items with the expression `a?c` would be items such as `aac`, `abc`, `acc`, and so on. However, searches are performed by a comparison of bytes, not characters, so actual results could include `a`, followed by a Japanese character that includes `c` in its byte representation, followed by `c`. This situation is avoided when the iProcess Engine uses UTF-8.

Issues with UTF-8 Encoding

While UTF-8 encoding is more versatile than alternatives, there are some issues of which you should be aware when deciding whether to use this encoding:

- Some other multi-byte character encoding systems are more efficient than UTF-8, in terms of the number of bytes they use per character. This means that converting from such a system, for example from Big5 encoding of Chinese characters to UTF-8, will result in a larger database. Perhaps more importantly, some items may overflow the space designed for them. For example, an increase from using 2 bytes to represent a character to 3 bytes means that a case description limited to 24 bytes can only contain 8 characters rather than 12 characters.
- The TIBCO iProcess Workspace (Windows) client does not itself use UTF-8 encoding. It converts data from the native character encoding for its locale to UTF-8 to send to the iProcess Engine, and similarly converts data it receives from the iProcess Engine from UTF-8. It is therefore possible that the same issue of space mentioned in the previous point could arise. Data entered in a native character encoding could end up being stored in the iProcess database as a larger number of bytes.

For these reasons, if you are using UTF-8 encoding, TIBCO recommends using TIBCO Business Studio for modeling processes, and either TIBCO iProcess Workspace (Browser), or an application that utilizes SSO for processing work items.

- With multiple TIBCO iProcess Workspace (Windows) clients that use different locales, you must ensure that you do not address work items to users in different locales as unpredictable results can occur.

iProcess Names

Some internal names are supported in iProcess Engine only as 7-bit ASCII. These names are not normally seen by end-users, and generally do not exceed eight bytes in length. This affects the following items:

- Procedure name

- Field name
- Node name
- List name
- Table name
- Step name
- Attribute name
- Role name

If you access these internal names, do not use them to store data that exceeds 7-bit ASCII.



UserName is not affected by this limitation.

Recommendations

You should consider carefully whether you want to support UTF-8 encoding. While circumstances differ between individual cases, TIBCO recommends that:

- New iProcess Engine installations should always opt to use UTF-8 encoding. This enables your installation to convert to supporting multiple multi-byte character sets even if you don't need it now, at little cost.
- Existing iProcess Engine installations should upgrade based on the type of character encoding currently in use:
 - Existing iProcess Engine installations that use a single-byte character set should in most cases opt to use UTF-8 encoding. This ensures that you can convert to supporting multiple languages should it become necessary at a later date.
 - Existing iProcess Engine installations that already use a multi-byte native character set *should not* use UTF-8, because of the likely increase in database size and in the length of individual items, as well as the risk of data truncation in the conversion process (see [Issues with UTF-8 Encoding on page 51](#)).

Configuring the iProcess Suite Using UTF-8

When implementing an iProcess Suite installation in an multi-national environment using UTF-8 encoding, you need to consider the flow of data between the individual products that make up the iProcess Suite. Some products need particular settings configured in order to communicate with the UTF-8 character set.

Identify which parts of the iProcess Suite you are using and the physical locations of those products. The components that require consideration are:

- iProcess clients
- IAPJMS process
- LDAPCONF Utility
- iProcess Plug-ins
- TIBCO Business Studio

TIBCO iProcess Engine

The following configuration needs to be done on the TIBCO iProcess Engine:

- The IAPJMS process on the TIBCO iProcess Engine needs to be configured to use UTF-8 by setting the IAPJMS_LANGUAGE process attribute. Use the `SWDIR\util\swadm` utility as follows:
 - swadm set_attribute 0 ALL 0 IAPJMS_LANGUAGE UTF-8**
- On Oracle and DB2, the following environment variable must be set to specify the character set of the database:
 - **Oracle** If iProcess Engine is used with a database supporting UTF-8, set the variable `NLS_LANG` to `AL32UTF8`.
 - **DB2** If iProcess Engine is used with a database supporting UTF-8, set the variable `DB2CODEPAGE` to `1208`.
- Non-XML files (such as XFR or Abox files) are converted from the value specified in the environment variable `SW_FILE_ENCODING` or, if that is blank, from the encoding specified by the system locale. In most cases the encoding will conform to that specified by the system locale, so this environment variable can be safely left blank. However, in a situation where multiple iProcess Engine installations exist on the same machine and serve users with different language requirements, you may want override the locale

setting using the environment variable. For example, depending on your platform, you can set **SW_FILE_ENCODING** to UTF-8 as follows:

- In Windows


```
C:\> set SW_FILE_ENCODING=UTF-8
```
- In UNIX Korn Shell


```
$ export SW_FILE_ENCODING=UTF-8
```
- In UNIX C Shell


```
$ setenv SW_FILE_ENCODING UTF-8
```
- In UNIX Bourne Shell


```
$ SW_FILE_ENCODING=UTF-8;export SW_FILE_ENCODING
```

iProcess Clients

Any iProcess Clients, both iProcess Workspace (Windows) and iProcess Workspace (Browser), send to and receive from iProcess Engine data encoded in UTF-8. Any necessary conversions are carried out by the client.

TIBCO iProcess Workspace (Windows)

Using iProcess Workspace (Windows), no configuration is necessary. Conversions take place automatically, as follows:

- Sending data to iProcess Engine, iProcess Workspace (Windows) converts it from native encoding to UTF-8.
- Receiving data from iProcess Engine, iProcess Workspace (Windows) converts it from UTF-8 to native encoding.

TIBCO iProcess Workspace (Windows) determines the native encoding to use from the system locale to which it is set.

TIBCO iProcess Workspace (Browser)

If you are using iProcess Workspace (Browser), you must:

- On a UNIX system, set the **TISOUnicodeConverterName** environment variable to **UTF-8**.
- On a Windows system, set the **TISOUnicodeConverterName** registry entry (**HKEY_LOCAL_MACHINE\SOFTWARE\Staffware plc\Staffware SSO Client**) to **UTF-8**.

TIBCO iProcess Workspace (Browser) uses the **TISOUnicodeConverterName** value to determine what encoding to use for communications with iProcess Engine.

The LDAPCONF Utility

If you use the LDAPCONF Utility and the system default encoding is not UTF-8, ensure that option [10] **Enable Attribute Value Translation from UTF-8** option is deselected. If selected, this enables the conversion between the UTF-8 format used by LDAP and the system locale of the iProcess Engine. If iProcess Engine is using UTF-8, this conversion is redundant.

See "Setting up the Connection" in the *LDAPCONF Utility User's Guide* for details of this menu option.

TIBCO Business Studio

If you use TIBCO Business Studio, data sent from TIBCO Business Studio to iProcess Engine is converted to the encoding specified in the process attribute **DEPLOY_XSL_OUT_ENCODING**. Therefore you must:

- Set the iProcess Engine process attribute **DEPLOY_XSL_OUT_ENCODING** to **UTF-8** by using the `SWDIR\util\swadm` utility as follows:

```
swadm set_attribute 0 ALL 0 DEPLOY_XSL_OUT_ENCODING UTF-8
```

(The default is **ISO-8859-1**.)

iProcess Plug-ins

All iProcess Plug-ins send data to and receive it from the iProcess Engine in UTF-8 encoding. Any necessary conversion is carried out by the Plug-in.



Some iProcess Plug-ins do not require configuration for UTF-8 support; only those that require configuration are listed in this section.

iProcess BusinessWorks Plug-in

The **defaultEncoding** attribute in the `SWDIR/eaijava/encoding.properties` file must be set to the encoding used by the iProcess Engine. Therefore:

- Set the **defaultEncoding** attribute in the `SWDIR/eaijava/encoding.properties` file to **UTF-8**.

iProcess Web Services plug-in

The **WSDocumentHandler.Encoding** attribute in the `wsconfig.properties` file must be set to the encoding used by the iProcess Engine. Therefore:

- Set the attribute **WSDocumentHandler.Encoding** in the *WebServiceHome/jetty-6.1.1/staffware/wsconfig.properties* file to **UTF-8**.

iProcess Email Plug-in

The **Charset** attribute in the **eai_mail.cfg** file must be set to the encoding used by the iProcess Engine. The iProcess Email plug-in embeds this attribute in the encoded-words streams for the MIME header fields **subject**, **to address**, **cc address**, **from address** and others. Therefore:

- Set the attribute **Charset** in the *SWDIR/lib/eai_mail.cfg* file to **UTF-8** for the Windows platform.
- Set the attribute **Charset** in the *SWDIR/libs/eai_mail.cfg* file to **UTF-8** for UNIX platform

Globalization Support Using Native Encoding

The following sections describe how an iProcess Suite installation that uses native encoding rather than UTF-8 encoding. This type of environment supports either of the following:

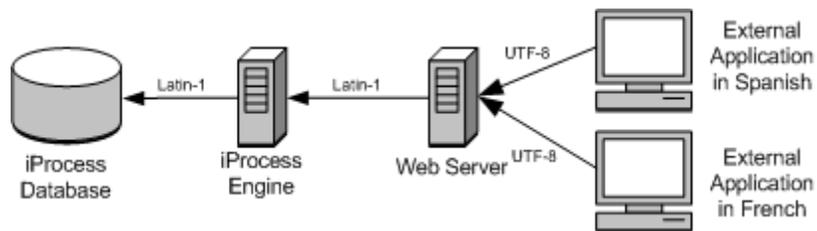
- a single-byte native character encoding environment
- a multi-byte native character encoding environment



Some native multi-byte character sets are referred to as double-byte character sets (such as BIG-5 or Shift-JIS).

Using iProcess Suite in a Single-Byte Native Encoding Environment

The following diagram shows iProcess Suite in an international environment that uses a single-byte native character encoding environment (in this example, Latin-1).



In this example:

- When text is entered, in either Spanish or French, into the external applications, the data is encoded using UTF-8. UTF-8 is used to send the data from the external applications to the web server.
- The web server converts the data from UTF-8 to Latin-1 and sends the data to iProcess Engine.
- iProcess Engine inserts the data into the iProcess Database.
- The iProcess Database stores the data as Latin-1.

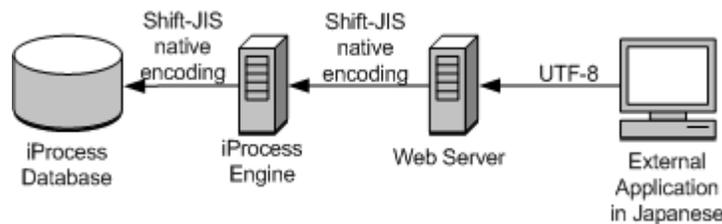
Note that:

- iProcess Engine supports a variety of character sets as its native character encoding. However, iProcess Engine still needs to handle data in other encodings. Therefore, iProcess Engine provides the functionality to convert different character encodings into its native character encoding.

- It does not matter if the machines hosting the iProcess Database, iProcess Engine, and iProcess Web Services Server and Client Plug-ins are all in separate locations to the machines hosting the external applications.

Using iProcess Suite With a Multi-Byte Character Encoding Environment

The following diagram shows the iProcess Suite in an environment that uses one multi-byte native character set.



In this example:

- When text is entered in Japanese into the external application the data is encoded using UTF-8. UTF-8 is used to send the data from the external application to the web server.
- The web server converts the data from UTF-8 to Shift-JIS (native encoding) and sends the data to iProcess Engine.
- The iProcess Engine inserts the data into the iProcess Database.
- The iProcess Database stores the data as Shift-JIS (native encoding).

Note that:

- iProcess Engine supports a variety of character sets as its native character encoding. However, iProcess Engine still needs to handle data in other encodings. Therefore, iProcess Engine provides the functionality to convert different character encodings into its native character encoding.
- It does not matter if the machines hosting the iProcess Database, iProcess Engine, and iProcess Web Services Server and Client Plug-ins are all in separate locations to the machines hosting the external applications.

Implementing iProcess Suite in an International Environment with Native Encoding

When implementing the iProcess Suite in an international environment, you need to consider the flow of data between the individual iProcess products that make up the iProcess Suite. Each iProcess product that needs to communicate with an external application needs to be configured to use the same character set as the external application.

An overview of steps you need to follow is:

1. Identify which iProcess products you are using and the physical locations of those products. The components that require configuration for encoding are:
 - your host system locale
 - iProcess Engine database encoding
 - iProcess Email Plug-in
 - iProcess Workspace (Browser)
 - iProcess Technology Plug-ins
 - iProcess Web Services Plug-in
 - LDAPCONF Utility
2. On the machines that are hosting the iProcess product components, check that host system locales are using the correct character sets.
3. Set the encoding attributes for each component. Refer to the Installation or user guides supplied with each component for information on how to do this.
4. 4. Non-XML files (such as XFR or Abox files) are always converted from the encoding specified by the system locale. The environment variable `SW_FILE_ENCODING` does not work in native mode.

TIBCO iProcess Engine Processes

This chapter provides detailed information about the TIBCO iProcess Engine processes and how they interact with each other. The iProcess Engine Process Sentinels are responsible for starting and stopping these processes and making sure that they keep running. The Process Sentinels are described in [Chapter 9 on page 111](#).

Topics

- [Foreground Processes, page 62](#)
- [Mbox Sets and Message Queues, page 70](#)
- [Background Processes, page 72](#)

Foreground Processes

The following table summarizes the function of each foreground process.

Process Description	Logical Process Name	Process Executable	Function
Work Queue Server	WQS	wqsrpc	The WQS is responsible for providing a complete list of work queues on the system, along with the RPC addresses (RPC number and machine name) of the WIS that handles each queue. It contains the list of all the users and groups and controls who has access to each queue. Refer to page 63 .
Work Item Server	WIS	wisrpc	<p>The WIS processes cache the contents of the work queues and provide lists of work items to TIBCO iProcess Workspace.</p> <p>The WIS processes concurrently produce messages (under transaction control) to be placed in the background Mbox set. Each WIS process receives messages from the WISMBD. These will be stored temporarily in an in-memory buffer (per physical work queue) for later retrieval when the WIS performs an update from the in-memory Mbox. Refer to page 67.</p>
WIS MBOX Daemon	WISMBD	wismbd	These are responsible for dequeuing/reading the messages from the WIS Mbox sets and passing them to the appropriate WIS that is handling the item. Refer to page 73 .
RPC Pool Servers	RPC_POOL	swrpcsvr	These manage and allocate the RPC connections. For each user that logs into iProcess using the TIBCO iProcess Workspace, a connection to a swrpcsvr pool process is allocated by the listener. Each pool process maintains a pool of connections.

Process Description	Logical Process Name	Process Executable	Function
RPC Listener (UDP)	RPC_UDP_LI	swrpcudp	For each user that logs into iProcess using the TIBCO iProcess Workspace, a connection to a swrpcsvr pool process is allocated by the listener.
RPC Listener (TCP)	RPC_TCP_LI	swrpcsvr	



All of the foreground processes *must* operate on the master server. See [Determining Where Processes Run on page 118](#) for more information.

Work Queue Server

The iProcess work queues, which contain all the user's work items, are managed by the following processes:

- **Work Queue Server (WQS)**, which handles the listing of queues. This process is run by `SWDIR\etc\wqsrpc`. There is only a single `wqsrpc` process running at any time.
- **Work Item Server (WIS)**, which handles the listing of work items in the queues. This process is run by `SWDIR\etc\wisrpc`. The number of `wisrpc` processes running is configured by the Process Sentinels (`process_attribute` table).



The WQS process handles what is displayed in the left hand pane of the Work Queue Manager (the queue list) and the WIS process handles the contents of the right hand pane (the work items list).

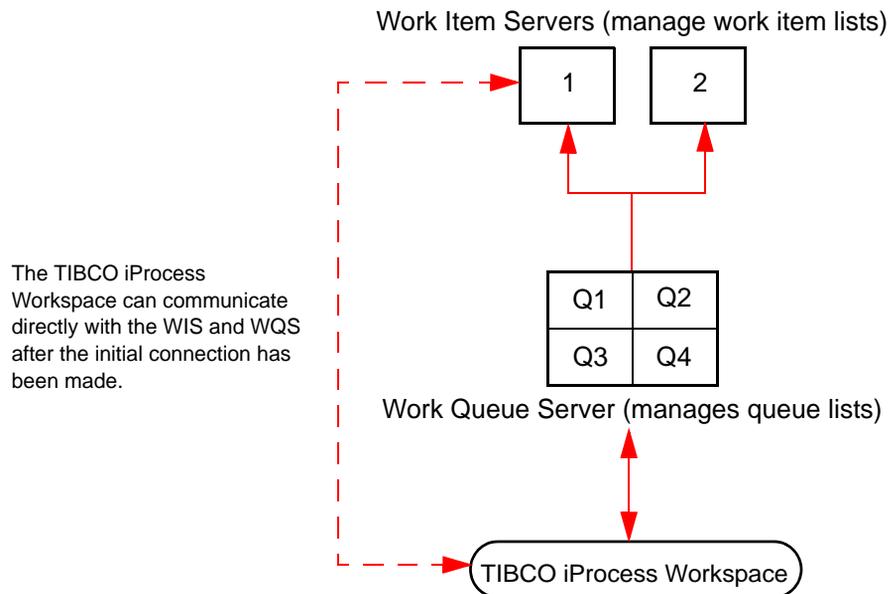
The work queue processes are started automatically when the other TIBCO iProcess Engine processes (such as the RPC pool servers) are started, and stopped when the other TIBCO iProcess Engine processes are stopped. The Process Sentinels start (and stop) processes in a specific sequence to make sure that processes have all their dependent processes running.

When a TIBCO iProcess Workspace logs in to the TIBCO iProcess Engine, the following sequence of events occur:

1. The TIBCO iProcess Workspace communicates with the RPC listener process (`RPC_TCP_LI` or `RPC_UDP_LI`).
2. The RPC listener process provides the RPC number of a pool server (it can start a new pool server if required).
3. The TIBCO iProcess Workspace connects to the pool server.

4. The RPC listener provides the RPC number of the Work Queue Server that the TIBCO iProcess Workspace then connects to.
5. The TIBCO iProcess Workspace communicates with the WQS.
6. The WQS provides the RPC numbers of the Work Items Servers that serve the work queues in the TIBCO iProcess Workspace.

The TIBCO iProcess Workspace communicates with the WQS and the WIS via RPC.



Refer to [TIBCO iProcess Workspace and TIBCO iProcess Engine Network Communication on page 126](#) for more information about RPC calls and how the TIBCO iProcess Workspace and TIBCO iProcess Engine communicate over the network.

The Work Queue Server controls which user and group queues a user will see on their TIBCO iProcess Workspace and which WIS process is handling the queue. The WQS provides a mapping to the appropriate work queues that are held in the WIS processes. Each TIBCO iProcess Engine instance runs with one WQS process.

Shared memory is used for caching this information. This memory is released when the instance has finished.

The WQS allocates work queues to the WIS processes that are running using either *round robin* or *on-demand* allocation. You can configure which allocation method is used by modifying the WQS_ROUND_ROBIN parameter in `SWDIR\etc\staffcfg`.

Allocation of Work Queues to WIS Processes

The WQS process performs the work queue allocation. The WQS reads the list of users and groups from the database and sorts them alphabetically and allocates them to a particular WIS. The WQS allocates which WIS a work queue is sent to in one of two ways.

Round Robin Queue Allocation

The WQS allocates a work queue to each WIS alphabetically, cycling round until all the work queues are allocated. For example, if a system has 5 WIS processes and 15 work queues (A-O) then the following allocation is performed:

- Queues A,F,K are allocated to WIS 1
- Queues B,G,L are allocated to WIS 2
- Queues C,H,M are allocated to WIS 3
- Queues D,I,N are allocated to WIS 4
- Queues E,J,O are allocated to WIS 5

This method of allocation takes no account of queue size so it is best used when queues are evenly distributed with work items and user access is evenly spread.

On-Demand Queue Allocation

This method allocates work queues alphabetically but only to the first available WIS. Therefore if a WIS is allocated a large work queue, it will take some time before it is ready to accept another queue. This means that other WIS processes that have smaller queues can accept more queues.

The effect of using on-demand allocation is that a more even distribution of work is achieved. However, the initial allocation is based upon the initial loading size of each queue so it may not be representative of the amount of requests to that work queue.



As both methods allocate work queues alphabetically, it is possible to have some control over queue allocation by careful naming of your queues. For example, if all the larger queues are first alphabetically, the early WIS processes will start getting these and the later WIS processes can pick up the smaller queues.

Controlling the Assignment of Queues to WIS Processes

There are two additional methods you can use to customize the assignment process to better reflect your system requirements, and so optimize performance.

- Use different WIS processes to handle user queues and group queues.

User queues and group queues frequently have different characteristics, in terms of the amount of load they carry. For example, if group queues are far more active than user queues on your system, you may want to give them higher priority for WIS allocation.

- Assign a queue explicitly to a WIS process

If you have certain queues that are very large or very busy, you may find it useful to dedicate specific WIS processes to handling *only* those queues (leaving the remaining queues to be dynamically assigned to the remaining WIS processes).

Refer to “Administering the Work Queue Server and Work Items Server” in the *TIBCO iProcess Engine: Administrator’s Guide* for more information.

RPC Pool Server

This process is responsible for handling RPC requests from a TIBCO iProcess Workspace to access and update data in the iProcess Engine instance.

A number of RPC pool servers can be created by the RPC Listener when the TIBCO iProcess Engine is started and each RPC server will be responsible for a configured pool of TIBCO iProcess Workspace connections. You can set up a number of pool servers that are pre-loaded if you have lots of TIBCO iProcess Workspaces logging in quickly. You do this by defining the `PRE_LOAD_POOL_SERVERS` parameter in the `SWDIR\etc\staffcfg` file. TIBCO iProcess Workspaces can be allocated to pool servers using either a round robin or load balanced method.

The RPC Pool servers are started by the RPC TCP Listener. The number of users that each pool server can support is configured using the `MAX_USERS_PER_PROCESS` parameter in the `SWDIR\etc\staffcfg` file.

RPC Listeners

The RPC Listeners are started by the Process Sentinels and are the first TIBCO iProcess Engine foreground server processes to be started. A listener is started for both TCP and UDP protocols. The RPC number for the Listener process is the same for TCP and UDP and is a start-up configuration parameter for the TIBCO iProcess Engine. Line 11 of the `SWDIR\swdefs` file defines this RPC number.

This RPC number is the published initial connection port for a TIBCO iProcess Engine for use by any client applications built using a TIBCO iProcess Workspace interface (e.g. the iProcess Applications Layer).

Work Item Server

The WIS handles the listing of work items in the queues. The process executable is `SWDIR\etc\wisrpc`. A number of WIS processes can be run and this is controlled by the Process Sentinels.

A WIS process is one instance of a Work Item Server and caches work queues that the WQS has allocated to it. Every iProcess work queue is hosted by a WIS process and each WIS can process more than one work queue. The iProcess Engine runs two WIS processes by default but you can increase or decrease this number using the `SWDIR\util\swadm` utility.

The WIS processes are multi-threaded processes. Different threads are used to perform different tasks - for example, responding to RPC requests, caching queue information, filtering queues or updating CDQP information.



Version 10.4 of the TIBCO iProcess Engine introduced multi-threaded WIS processes. This improved WIS performance significantly, making it possible to reduce the default number of WIS processes from 6 (the pre-Version 10.4 default) to 2.

There are many work queue performance issues related to the number of WIS processes you have, how many work queues they process, how many threads they use for different tasks and so on. Refer to “Administering the Work Queue Server and Work Item Servers” in the *TIBCO iProcess Engine: Administrator’s Guide* for more information.

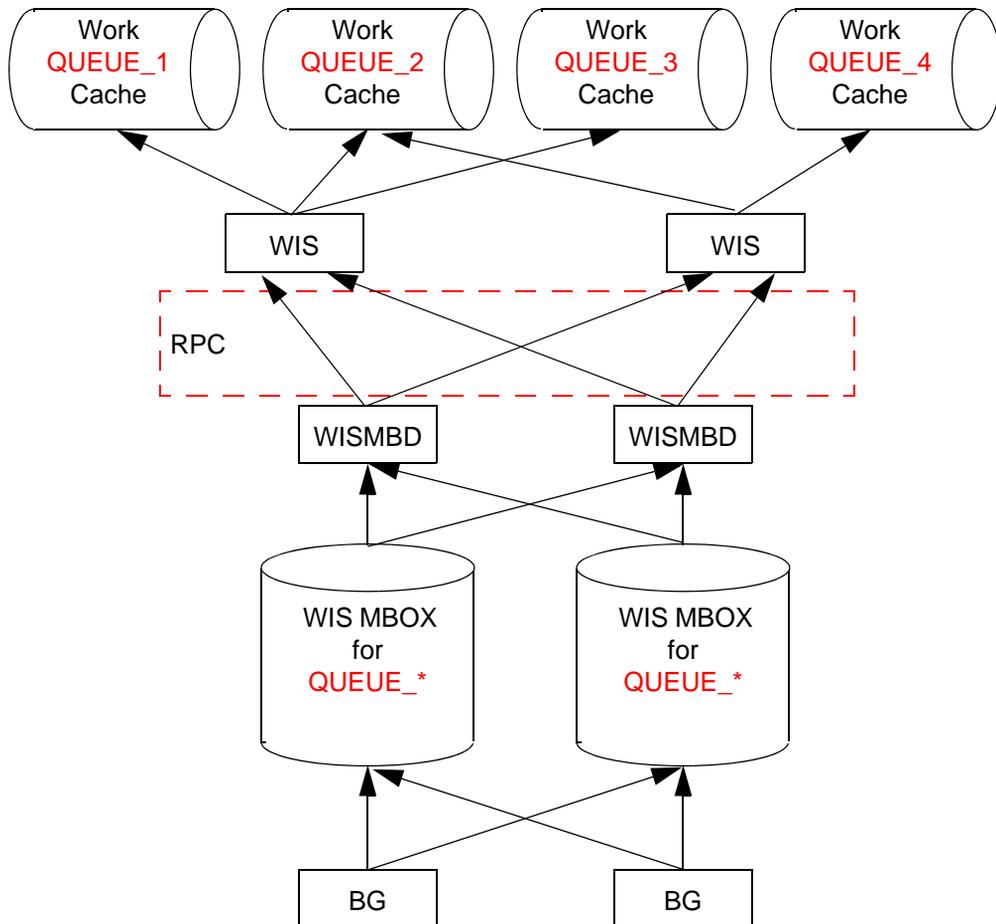
The WIS processes maintain a cache of the information they contain (which is the user’s work queue). This cache is synchronized with the same information stored in the user or group’s work queue (`staffo` database table). You can view the information in this table using `SWDIR\util\plist -m`.

WIS Mbox Daemon

This process (**WISMBD**) operates between the WIS Mbox set and the **WISRPC** processes forwarding messages from one to the other. The executable is `SWDIR\etc\wismbd`.



The WISRPC processes do not write outgoing messages via the WISMBD, they go straight to the Background Mbox set.



The **WISMBD** process is configured to read from a configurable number of physical WIS Mbox sets in a round-robin manner and it will deliver the messages to the appropriate WIS process.

When the WIS Mbox sets are empty, the **WISMBD** will sleep for a configurable amount of time. This is defined by the `EMPTYMBOXSLEEP`, `EMPTYMBOXSLEEP_INC` and `EMPTYMBOXSLEEP_MAX` process attributes. Refer to “Administering Process Attributes” in the *TIBCO iProcess Engine: Administrator’s Guide*.



The WQS and WIS processes are started up before the WISMBD processes so that no work queues will receive messages until ALL the WIS processes have started.

The **WISMBD** sends synchronous RPC requests to the WIS that maintains the work queue to which the message is addressed.

The **WISMBD** is initially set to read from the following Mbox sets:

- `MBSET_READ_WIS` - This is used to read Mbox messages
- `MBSET_WRITE_BG` - This is used to forward undelivered work items back to the local background for locally hosted procedures.

Mbox Sets and Message Queues

The foreground and background processes communicate with each other using messages. Messages contain information about the iProcess case and instructions about what to do with the case (such as release, keep or forward it.) A message is processed by the case instruction processors. **Mbox set** is the generic name used for a container in which these messages are stored. Processes can dequeue messages from a Mbox set as and when resources are available.

The messages are stored and managed in the following way:

- *in the UNIX Oracle and Windows Oracle versions*, using Oracle AQ. Oracle AQ uses message queues, which are defined as Oracle AQ tables.
- *in the UNIX DB2 and Windows SQL Server version*, using iProcess database tables that are managed by the database server.



Please see the appropriate TIBCO iProcess Engine Database Administrator guide for more information about the format of the AQ or database tables that are used to hold the message queues.

The TIBCO iProcess Engine uses a logical and physical grouping for queues to help improve message throughput. The physical message queues are grouped together as logical Mbox sets.

An Mbox set can be used for different purposes, for example two Mbox sets can be created for the Backgrounds. The Deadline Manager (**DLMGR**) can write to one set while other processes could write to the second set. The relative priority of processing the work can be changed by assigning different numbers of background processes to dequeue the messages in the different Mbox sets.

Processes can be configured to enqueue or dequeue messages to/from a specific Mbox set rather than to/from a single Mbox queue. The sets can be changed as the performance demands change in your system environment. Multiple queues in an Mbox set enable processes to distribute messages in a round robin manner resulting in an even distribution of messages on the queues.

Refer to [Chapter 6 on page 89](#) for more information about Mbox sets.

Transaction Control of Messages

The Mbox provides a message repository so that processes can post messages and continue with their own processing without having to wait for a reply. The messages are stored persistently and are processed only once. Messages have to be processed exactly once to preserve the data integrity of a transaction. For example, if a message instruction is to debit an amount from a bank account then this has to happen once only even if the systems fail. This is achieved using transaction control.

The messages in the database tables provides a reliable messaging system because the message queues are under transaction control. For example, if a message is delivered to the background case instruction processor but the server goes offline for a while, the message instruction is still persisted in the queue and is retried at a later time.

Background Processes

The following table summarizes the background processes:

Process Description	Logical Process Name	Process Executable	Function
Background	BG	swbgmd	<p>This is the core background process that interprets the business rules that have been defined in the iProcess Modeler. It performs the case instructions such as working out what the next step is in the procedure, updating the database with new information and checking for deadlines.</p> <p>This process is also responsible for dequeuing messages from the Mbox sets. There can be multiple background processes concurrently dequeuing messages. This process can be configured to read a specific Mbox set as defined by the process_attribute table.</p>
Case Prediction	BGPREDICT	swbgmd	This process is responsible for updating prediction data (stored in the database).
Database Queue Daemon	DBQD ^a	n/a	<p>This process:</p> <ul style="list-style-type: none"> • caches a configurable number of messages from the database for each available queue. • processes RPC requests to dequeue messages for the BG and WISMBD processes.
Deadline Manager	DLMGR	dlngr	This process is responsible for monitoring deadlines in cases of procedures.

Process Description	Logical Process Name	Process Executable	Function
IAPJMS	IAPJMS	iapjms	This process is responsible for receiving messages containing activity monitoring information from the BG process, and Work Queue Delta publication messages from the WIS process, and routing these to the specified JMS topic.
RPC Background	RPCBG	staffrpcbg	This process handles the Jump To and Case Suspend features in the iProcess Suite.

a. This process is only present in the DB2 version of the TIBCO iProcess Engine.

Background

This process retrieves messages (containing case instructions) from the Mbox sets and then processes the case instructions in the messages. There can be multiple Background processes all concurrently dequeuing messages from an Mbox set and processing case instructions

The number of Background processes is controlled by the Process Sentinels. The Process Sentinels read the **process_config** table to see how many instances of the process to run and on which computers.

Each message contains the case instructions from which the Background can determine what actions to take. The Background interprets the business rules defined in the procedure (such as the addressee of the next step) and routes work items to the necessary work queues or external applications. The process makes decisions based upon the iProcess data and procedure definition instructions as to what happens in the business process next.

Case Prediction Processor(s)

This process (*SWDIR\etc\bgpredict*) receives messages from the iProcess Background processes. When a case instruction that results in a change to a case has been processed, the iProcess Background processes notify the Case Prediction processes so that the prediction data (stored in the database) can be updated. There can be multiple case prediction processes running concurrently.

Each message contains information about the case that has changed and the procedure and instruction that caused the change. The process will read messages from the queue(s) and update the **predict** table so that it contains the latest prediction information about the case that the queued message was for.

The process attempts to determine a valid addressee for each step, so that `SW_USER:attribute` (where *attribute* is a user attribute, for example, `DESCRIPTION`) resolves to something valid. If there are multiple addressees on a step then the first one is taken, which is resolved in the following order:

1. First the process checks the user and group addressees and uses the first one in the list. If it finds a group it gets the first member listed in the group.
2. If no user or group was found it next checks the field addressees.
3. Finally it checks the roles.

If no addressee can be determined, if a group has no members, or a field or role contains an invalid user, then the process defaults to using the system administrator user (by default, **swadmin**).

Database Queue Daemon



This process is only currently used on the DB2 version of the TIBCO iProcess Engine.

This process (`SWDIR\etc\swdbqd`) processes RPC requests from the **BG** and **WISMBD** processes to dequeue messages from the database queue tables.

It caches a configurable number of messages from the database for each available queue. When a request to dequeue a message arrives, the process returns a message from the cache. If the cache is empty, the process first refills the cache from the database queue tables.

The number of messages to be cached is determined by the `DBQD_MAX_CACHED_MESSAGES` process attribute.

Deadline Manager

This single process manages the deadlines that have been defined in a procedure using the iProcess Modeler.

At defined intervals (deadline processing interval), the Deadline Manager checks the **outstanding_addressee** table for expired deadlines. If deadlines have expired, the Deadline Manager sends an Mbox instruction to the background Mbox set so that the case instruction process can process the deadlines for the case.

You can define a limit to the number of Mbox instructions the Deadline Manager sends. This is to avoid the Deadline Manager sending out duplicate Mbox instructions for the same unprocessed, expired deadlines. There are two process attributes that enable you to configure this:

- The `UNPROCESSED_DL_POST_LIMIT` process attribute sets a limit on the number of messages for expired deadlines that the Deadline Manager allows in the mbox queue at any one time.
- The `MAX_AGE_BEFORE_RESETPOST` specifies the time period before the Deadline Manager resets its internal marked of the last deadline it has processed to 0 (beginning of time).

See “Administering Process Attributes” in the *TIBCO iProcess Engine: Administrator’s Guide* for more information about process attributes.

You can define the priorities for the case deadlines. When the Deadline Manager processes the expired deadlines, two process attributes can be used to define the deadline message priority:

- The `FORCE_DEADLINE_PRIORITY` process attribute defines a priority value for the case deadlines that override all other priorities set for the deadline.
- The `USE_PRIORITY_DEADLINES` process attribute defines whether to apply the value of the `SW_CP_VALUE` field as the priority for the case deadlines if the value of the `FORCE_DEADLINE_PRIORITY` attribute is set to the default value, 0.

For more information about deadline message priority, see “Priority Setting for Case Deadline” in *TIBCO iProcess Modeler Advanced Design*.

When all the deadlines have been processed, the Deadline Manager “sleeps” until the deadline processing interval has expired. This interval can be set as an absolute or repeating value, for example:

- Absolute intervals are used to process deadlines once a day at a specific time.
- Repeating intervals are used to process deadlines at regular intervals and at set times throughout the day relative to midnight each day. The time is defined on the local computer on which the deadline manager is running.
- For example, an interval of five hours would set deadlines to be processed at 5.00am, 10.00am, 3.00pm and 8.00pm on Day 1 and 5.00 am, 10.00 am etc. on Day 2.

The deadline processing interval is set in the `process_attribute` database table. If the table does not contain an initial value, the Deadline Manager defaults to one minute.

IAPJMS Process

If activity monitoring is enabled on your TIBCO iProcess Engine, the **BG** process sends out a message when any of the TIBCO iProcess Engine activities that you have configured to monitor occur. This process (`SWDIR\etc\iapjms`) is responsible for receiving messages from the **BG** process and routing these to the JMS topic or queue.

If Work Queue Delta publication is in use, the **WIS** process similarly publishes messages to the **IAPJMS** process giving details of changes in monitored work queues.

See [Monitoring Activities on page 93](#) for more information and see "Activity Monitoring and Work Queue Delta Configuration" in the *TIBCO iProcess Engine: Administrator's Guide* for more information about configuring activity monitoring and Work Queue Delta publication.

RPC Background Process

This process (`SWDIR\etc\rpcbg`) handles synchronous RPC calls from the Jump To and Case Suspend features in the TIBCO iProcess Workspace.

Chapter 5

Introduction to Transactional Business Process Automation

This chapter provides an overview of the transaction management capabilities provided when using the iProcess Engine.

Topics

- [Overview, page 78](#)
- [What is a Local Transaction?, page 79](#)
- [What is a Distributed Transaction?, page 81](#)
- [Transaction Scope, page 82](#)
- [Using Enterprise Application Integration Steps in Procedures, page 84](#)
- [What is MSDTC?, page 85](#)
- [Examples of Transaction Control, page 86](#)
- [Transaction Failures and Rollbacks, page 88](#)

Overview

Typically, a business process involves a number of transactions such as looking up information from a customer database, taking an order, adding a new order to a database, updating prices, etc. To guarantee that all parts of the business process are completed, a resource manager is used to keep track of each step (or transaction) to make sure it is completed.

If all the steps can be completed successfully, then the complete process will be committed. This means that any updates to external systems and iProcess case data will be performed and committed in a single transaction.

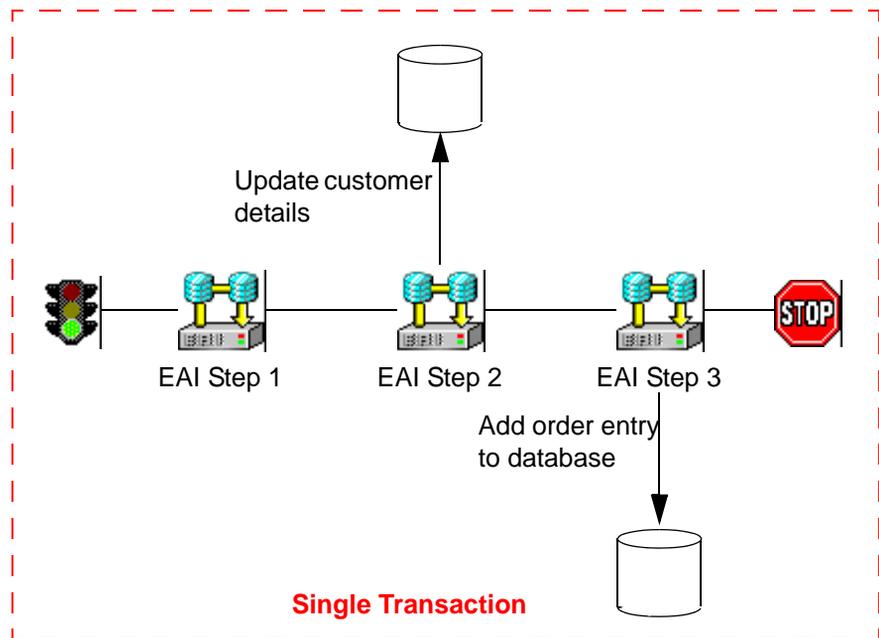
However, if one or more parts of the process cannot be completed for any reason (such as a database being offline), no parts of the process are committed. This means that all the external systems such as databases, document management systems and legacy systems are left in the same state as when the process was started i.e. no data updates are performed.

There are two types of transaction scope: local and distributed. The difference between them is the number of resource managers used to control the transaction. A local transaction uses just one resource manager and a distributed transaction involves using more than one resource manager. Both of these are described in the following sections.

What is a Local Transaction?

A local transaction is where a number of business operations are under the control of a single resource manager (typically the database resource manager). The TIBCO iProcess Engine process will be under the control of the resource manager so that the entire process will either be committed to the database or rolled back if one of the operations fail. The following diagram illustrates a procedure using EAI Database steps to update data in the local database.

Figure 1 A Local Transaction



Example of a Local Transaction

The following is an example of how a business procedure is designed so that all the steps are either committed or not committed using the local resource manager.

In a banking environment you want to make sure that all the accounts are updated correctly so that the balances are totalled. The ideal scenario is where all the processes are committed in one go so that money is deducted from one account and added to another. For example, if money is to be transferred from a savings account to a current account, two accounts will need updating. Both accounts need to be changed accordingly. It would cause serious problems if only one account was updated.

If the update to the savings account was successful but the update to the current account failed, there would be a discrepancy in the totals. However, because the step is part of the local transaction, the whole transaction will fail and the accounts are rolled back to their original state.

What is a Distributed Transaction?

A distributed transaction is where a number of related business operations (individual transactions) are grouped together so that they are put under the control of a transaction manager program (such as MSDTC). The integrity of the business data is maintained because all external systems being used (databases, document management systems, etc.) are either updated with information or left in their previous state if one or more of the steps fails.

If all the business operations (transactions) are successful, they can all be committed in one go. The information is stored in memory until the final transaction is reached when all the updates can be committed.

Transaction Scope

The following sections describe the scope of transactions on the different versions of the TIBCO iProcess Engine.

Oracle Server Transaction Scope

The Oracle version of the TIBCO iProcess Engine enables you to use one local resource manager (provided by the Oracle resource manager).

Because the TIBCO iProcess Engine uses Oracle Advanced Queues to store its internal instruction messages and Oracle tables to store case data, transaction management can be controlled by the internal Oracle database transaction manager.

This means that if updates to the Oracle database fail, the resource manager can roll back the updates and restore the database and case data to its original state.

DB2 Transaction Scope

The DB2 version of the TIBCO iProcess Engine enables you to use the local DB2 resource manager.

This means that if updates to the DB2 database fail, the resource manager can roll back the updates and restore the database and case data to its original state.

SQL Server Transaction Scope

When using the iProcess Suite with the Microsoft SQL Server, all transactions are managed by the SQL Server resource manager. SQL Server provides the following transaction functionality:

- Locking facilities that ensure transactional integrity.
- Logging facilities that ensure transaction durability. For example, if the hardware, operating system, or SQL Server fails, SQL Server uses the transaction logs, upon restart, to automatically roll back any uncompleted transactions to the point of the system failure.
- Transaction management features that enforce transaction atomicity and consistency. After a transaction has started, it must be successfully completed, or SQL Server undoes all of the data modifications made since the transaction started. Refer to the SQL Server documentation for more information about how SQL Server transactions work.

When all the processes use the local resource manager on the SQL Server, this is called a local transaction. If some processes use an external resource manager such as the EAI COM step, then this is known as a distributed transaction because more than one resource manager is used. A slightly different architecture is used for distributed transactions because a transaction manager program is used to control the resource managers. Refer to [Using Distributed Transactions with MSDTC on page 83](#) for more information.

Using Distributed Transactions with MSDTC

Windows has a built in transaction manager component called Microsoft Distributed Transaction Coordinator (MSDTC). The MSDTC controls resource managers from distributed sources such as SQL Server and COM.

The transaction scope provided with this architecture allows distributed transactions involving many resource managers to be controlled by the MSDTC. Some server EAI plug-ins such as EAI COM+ may require a distributed transaction and therefore need to use MSDTC. You can set the TIBCO iProcess Engine to use MSDTC by setting the EAI_NEEDS_MSDTC process attribute - refer to "Administering Process Attributes" in the *TIBCO iProcess Engine: Administrator's Guide* for more information about setting process attributes.

Using Enterprise Application Integration Steps in Procedures

To keep control of external transactions with third party applications and to keep the transactions under the control of the transaction manager, you can use EAI steps in your procedure definition. There are different types of EAI step used to communicate with different types of external system. For example:

- The EAI Database step is used to communicate with database stored procedures so you can use existing business logic in your procedures. Refer to the *TIBCO iProcess™ Modeler - Integration Techniques Guide* for more information about installing and using EAI Database steps.
- The COM+ EAI step is used so that iProcess can communicate with a COM+ application. Refer to the *TIBCO iProcess™ COM Plug-in: User's Guide* for more information about installing and using the COM EAI step.

An EAI step causes the iProcess background to call-out to an application defined sub-system in order to perform some work. In the MSDTC version, this work may be included within the current distributed transaction, hence combining data updates performed by that work with the data updates performed as part of the workflow case processing.

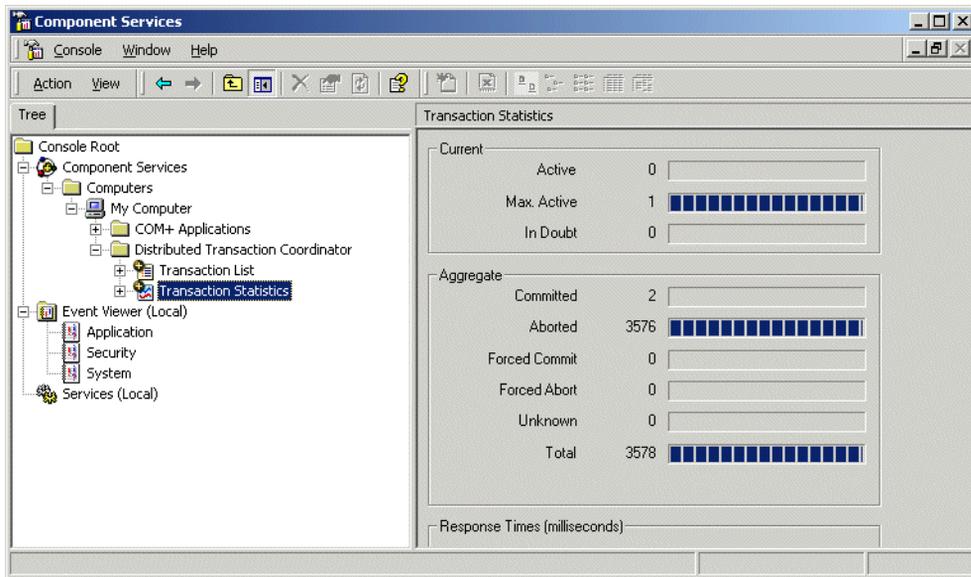
Refer to “Using Enterprise Application Integration Steps” in the *TIBCO iProcess Modeler - Integration Techniques Guide* for more information about using EAI steps.

What is MSDTC?

The Microsoft Distributed Transaction Coordinator is the transaction manager program built into the Windows 2003 Server. This runs as a service on Windows and coordinates transactions between distributed data stores.

The iProcess Suite integrates with the MSDTC by creating and controlling distributed transactions. This means that iProcess can start a transaction by sending a request to the MSDTC and then other resources (such as databases) can enlist in the transaction as required by communicating with the MSDTC.

The MSDTC can be administered using the **Component Services Console** (as shown in the following example).



You can configure iProcess to use MSDTC by setting the EAI_NEEDS_MSDTC process attribute - refer to “Administering Process Attributes” in the *TIBCO iProcess Engine: Administrator’s Guide* for more information about setting process attributes. This means that instead of the transaction being controlled by the local resource manager, the MSDTC is notified of the transaction and other distributed resource managers can enlist in the transaction.

Refer to the Microsoft documentation for more information about administering and monitoring MSDTC transactions.

Examples of Transaction Control

The following section provides some examples of iProcess transaction control.

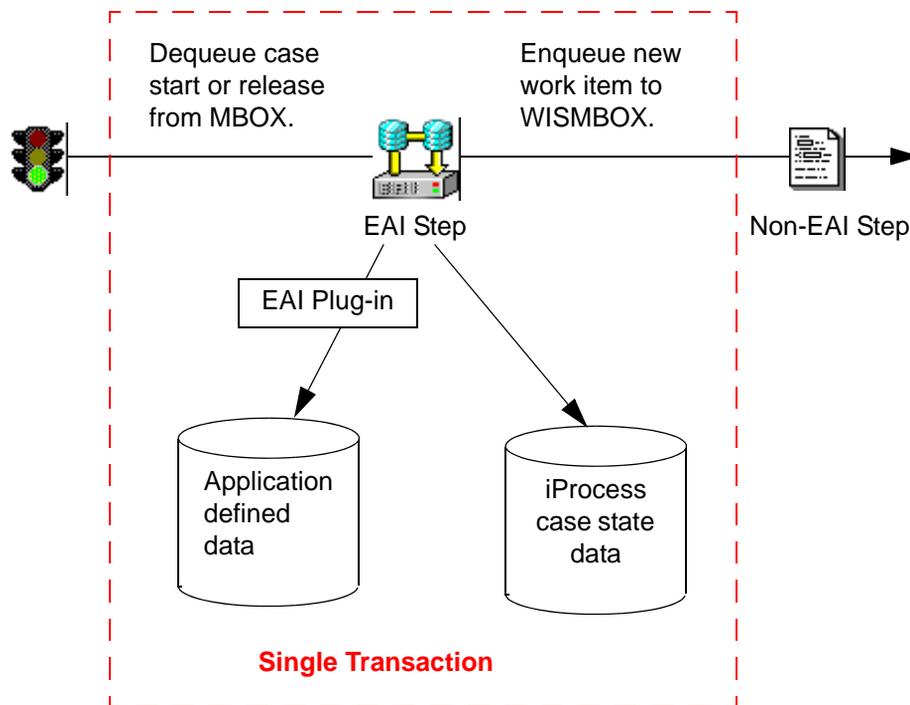
Case Data Updates to the SQL Server using MSDTC

In a simple process where iProcess case data is being updated in the SQL database and an external database using an EAI plug-in, the transaction process involves the following:

1. The iProcess Engine starts a transaction by notifying the MSDTC.
2. The SQL Server resource manager enlists itself as part of the transaction with the MSDTC.
3. The external database resource manager also enlists itself as part of the transaction with the MSDTC.
4. If all the case updates are successful, the MSDTC performs a two phase commit operation with all resources that are enlisted as part of the transaction.

External Updates Using EAI Steps

The Enterprise Application Integration (EAI) steps can be used in procedures to control updates to third party applications and iProcess case data under transaction control. If the external applications operate in the same TPM environment as iProcess, all business operations in a single procedure can be completed as one global transaction.



Transaction Failures and Rollbacks

In the transaction example in [Figure 1 on page 79](#), all the external data updates and iProcess case data updates are saved to memory until the case has been processed at which time all the updates are written to the database. This is known as the Commit process.

However, if one or more steps in the transaction are not possible (such as the database is not available), none of the updates are committed and the data is left in the same state as when the transaction started. This is known as the rollback process. Similarly, if the connection to the database is broken, any outstanding transactions are rolled back.

If the instruction is received by the Background but one of the steps it processes subsequently fails (such as a COM+ EAI step calling an application that is not running), then the resource manager returns an error to the Mbox daemon. The Mbox daemon aborts the transaction while the resource manager ensures that all of the steps involved in the transaction are rolled back to their initial state.

Poison Transactions

If a transaction is continually failing and being retried, this can have a serious impact on system performance. It is possible to restrict the number of times a transaction will be retried and to be notified of this failed transaction.

The iProcess installation process sets the message queues to have:

- a retry count of twelve (defined using the IQL_RETRY_COUNT process attribute).
- a retry delay of five minutes (defined using the IQL_RETRY_DELAY process attribute).

See “Administering Process Attributes” in the *TIBCO iProcess Engine: Administrator’s Guide* for more information about setting process attributes.

Chapter 6 **iProcess Mbox Sets**

This chapter describes how:

- TIBCO iProcess Workspaces communicate with the Background process.
- Mbox sets are used to help in the distribution of messages on the system.
- the middleware message queuing system is implemented for the TIBCO iProcess Engine.

Topics

- [Overview, page 90](#)
- [Transaction Management of Messages, page 92](#)

Overview

One of the core components of the iProcess Engine is the messaging system. This is used to deliver instructions from:

- all the TIBCO iProcess Workspaces to the Background process
- the Background process to the foreground processes

Because of the potentially high volumes of transactions in a TIBCO iProcess Engine system, there can be many thousands of instructions trying to be passed from TIBCO iProcess Workspaces to the Background to be processed. To prevent messages being lost or messages being held up for long periods of time, a message queuing system is used so that all messages sent from TIBCO iProcess Workspaces or the Background can be sent to a message queue. Once it is in the queue, the message has an assured delivery mechanism where it will be dequeued as system resources permit.

The message queues also makes sure that a message instruction is only processed once (unless a transaction rollback occurs - see [Transaction Management of Messages on page 92](#)) and is processed even if system resources fail because it is retried when the system is running again. For example, in the case for banking transactions, a message to update a bank account must only be performed once and must be performed at some point even if the system has failed.

The iProcess Suite currently supports two message queuing systems:

- iProcess database queues (for the UNIX DB2 and Windows SQL Server versions)
- Oracle AQ (for the UNIX Oracle and Windows Oracle versions)

Both control the delivery and storage of iProcess messages.

What are iProcess Messages?

A message consists of the business data captured by the information entered into iProcess forms such as what fields and values exist. A message also contains instructions so that the iProcess Engine knows what to do with the message such as release or forward the message. It can also contain control information used to manage the message. A message is delivered (enqueued) to the message queue and consumed (dequeued) by recipients.

The message queues and their messages are stored in database tables that operate as message queues.

The iProcess Engine is set up to use the Local message type, which is used for communication between the foreground and background processes.

Definition of Mbox Sets

The iProcess Engine uses a number of message queues to enable high volumes of transactions to be processed. iProcess uses a logical grouping of physical queues called Mbox sets in which processes read or write messages. An Mbox set can be defined as having one or more queues in which messages are posted in a round robin method. For example, an Mbox set can comprise a number of physical database tables operating as message queues.

There are many performance factors related to this design such as spreading queues over multiple disks. Instead of a process writing or reading messages from a specific queue, it can be configured to use an Mbox set. The Mbox set enables you to dynamically configure the queues used in a set as system resources change.

For example, two Mbox sets can be configured as follows:

```
MBOX_SET = queue_1, queue_2
MBOX_SET2 = queue_1
```



A physical queue can be included in multiple Mbox sets but can only be included once in the same Mbox set.

Each iProcess Engine process is configured to use a particular Mbox set and has a process attribute type set for it. This information is configured in the **process_attribute** table.

Configuring Mbox Sets

You can dynamically configure the Mbox sets using the `SWDIR\util\swadm` utility. For example, you can create Mbox sets, add queues to Mbox sets and remove queues.

Different processes require access to different types of Mbox set, for example, the Background process needs to have write access to a background Mbox set and a WIS Mbox set.

The **process_attribute** table specifies the access type for each process and which Mbox set is used.

The Mbox set tables contain the necessary information for each Mbox set such as the logical names of all the queues in the set. Multiple queues in a set enable iProcess to post messages on a round-robin basis through the queues.

Transaction Management of Messages

For a queuing technology to be transactional it must be under the control of a Resource Manager. The resource manager ensures that its resources, the queues, are always in a consistent state. Any updates, i.e. the sending or receiving of messages, made during a transaction will either all succeed or all fail.

UNIX Oracle Transaction Implementation

When using the UNIX Oracle version of an iProcess Engine, all resources are updated in an iProcess transaction as one unit of work. This is because the message queues are maintained in the database and can therefore be controlled by the Oracle Resource Manager. Transactions are limited to the scope of this database as the resource manager only controls this database.

Refer to the Oracle Database documentation for more information about how Oracle handles resource management.

Windows SQL Server Transaction Implementation

The TIBCO iProcess Engine uses the Microsoft SQL Server database to store iProcess messages. The SQL Server makes sure the queues are in a consistent state and will make sure that all resources updated during a transaction are committed as one unit of work.

UNIX DB2 Transaction Implementation

When using the UNIX DB2 version of an iProcess Engine, all resources are updated in an iProcess transaction as one unit of work. This is because the message queues are maintained in the database and can therefore be controlled by the DB2 Resource Manager. Transactions are limited to the scope of this database as the resource manager only controls this database.

Monitoring Activities

TIBCO iProcess Engine can be enabled to publish TIBCO iProcess Engine activity information and TIBCO iProcess Engine Work Queue Deltas to external applications. This chapter describes how messages are processed between the TIBCO iProcess Engine and external applications.

Topics

- [Overview, page 94](#)
- [How Messages are Processed From the BG Process to the IAPJMS Process, page 96](#)
- [How Activity Messages are Processed From the IAPJMS Process to the External Application, page 97](#)
- [How Messages are Processed From the WIS Process to the IAPJMS Process, page 98](#)
- [How Work Queue Delta Messages are Processed From the IAPJMS Process to the External Application, page 99](#)
- [Understanding the Message Types, page 100](#)

Overview

The TIBCO iProcess Engine can be enabled to publish both iProcess Engine activity information and iProcess Engine Work Queue Deltas to external applications, using JMS topics.

Both types of publication use the **IAPJMS** (Introspection Activity Publication JMS) process, but they differ in detail.

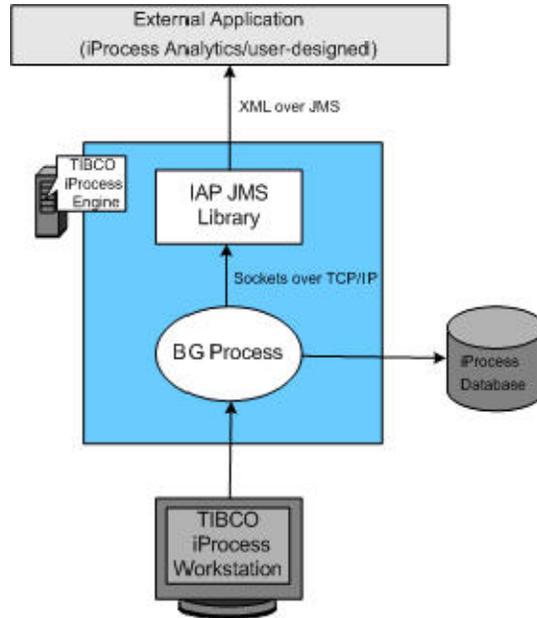
Activity Publishing

An activity is any instruction in the TIBCO iProcess Engine that creates an audit trail entry, for example, **Case started** or **Event issued**. You can configure any combination of step and/or activity to be monitored. This enables an external application to monitor important business events during the processing of cases.

The **BG** process can identify if a step is being processed and if activity monitoring has been configured for it. The **BG** process then sends details of the configured activities, in XML format to the [IAPJMS Process](#).

The **IAPJMS** process sends the XML message to a specified JMS topic, from which an external application can receive the XML format message.

The following diagram demonstrates Activity Publishing message processing between the TIBCO iProcess Engine and the IAPJMS process.



Work Queue Delta Publication

Work Queue Delta Publication enables an external application to monitor an iProcess work queue and to retrieve only those work items in a given work queue that have changed since monitoring began.

If any data in a monitored work queue changes - for example if an existing work item is modified, an existing work item is removed from the work queue or a new item is added to the work queue - the **Work Item Server (WIS)** process sends the modified data in XML format to the **IAPJMS** process. The **IAPJMS** process then sends the XML message on to a specified JMS topic (either a default topic, or one specified by the subscribing application), from which an external application can receive the XML message.

Work Queue Delta Publication functionality can only be accessed by using iProcess Server Objects (Java) or iProcess Server Objects (.NET).

How Messages are Processed From the BG Process to the IAPJMS Process

Once the **BG** process has been configured to enable audited activities to be published (defined in the IAPJMS_PUBLISH process attribute), it sends out messages whenever an audited activity occurs.

The messages are published to the **IAPJMS** process using a configurable port number (defined in the IAPJMS_PORTNO process attribute). The **BG** process has two delivery methods for publishing messages, synchronous and asynchronous (defined in the IAPJMS_SYNCHRONOUS process attribute).

- **SYNCHRONOUS** - When the message is sent, a receipt is requested. The **BG** process waits until the **IAPJMS** process has confirmed the message has been published. If the message is not published, an error is written to the *SWDIR/logs/sw_error* file. With the SYNCHRONOUS method the **BG** process has to block until a response/timeout occurs which could impact throughput. This means the SYNCHRONOUS method is slower than the ASYNCHRONOUS method.
- **ASYNCHRONOUS** - The message is assumed to have been processed correctly if the message was sent successfully to the **IAPJMS** process. The ASYNCHRONOUS method is faster than the SYNCHRONOUS method because it does not have to wait for a receipt.

How Activity Messages are Processed From the IAPJMS Process to the External Application

The **IAPJMS** process listens for activity messages from the **BG** process. When it receives an activity message from the **BG** process it routes this to a configurable JMS topic (defined in the **IAPJMS_TOPIC** process attribute). The **IAPJMS** process does not need to know anything about the contents of the message. It simply passes it to the JMS topic.

The external application is configured to listen to the same JMS topic that is defined in the **IAPJMS_TOPIC** process attribute.

For information on how to configure and administer activity monitoring see "Monitoring Activities" in the *TIBCO iProcess Engine Administrator's Guide*.

How Messages are Processed From the WIS Process to the IAPJMS Process

A user application - either using iProcess Server Objects (Java) or iProcess Server Objects (.NET) - subscribes to a work queue, creating an initial connection to the published queue. The application specifies a JMS topic name and subscribes to that topic (or, if no topic is specified at subscription time, the WQDJMS_TOPICNAME process attribute is used; see the *TIBCO iProcess Engine Administrator's Guide* for details of this attribute and of other process attributes involved), and then can start receiving JMS messages.

The first message sent is always a [Work Queue Synchronization \(DSY\) Message](#). This informs the listening application that it now needs to begin processing on other messages received from this Work Queue.



The data passed in JMS messages may be restricted by other components of the system. For example, the maximum length of work queue parameter fields that can be published is limited by the **WIS** process: see the section "Using Work Queue Parameter Fields in Procedures" in the *TIBCO iProcess Modeler - Advanced Design* guide.

Messages are published to the **IAPJMS** process using a configurable port number (defined in the WQDJMS_PORTNO process attribute).

How Work Queue Delta Messages are Processed From the IAPJMS Process to the External Application

The **IAPJMS** process listens for Work Queue Delta messages from the **WIS** processes. When it receives a message from a **WIS** process it sends this to the appropriate JMS topic. By default, this is the topic defined in the `WQDJMS_TOPICNAME` process attribute; but this default can be overridden if the subscribing application specifies a different topic. The **IAPJMS** process does not need to know anything about the contents of the message, it simply passes it to the JMS topic.

The external application must be set up and configured to listen to the same JMS topic to which the **IAPJMS** process sends the messages.

The first time that a subscribing application opens a work queue, it uses `iProcess Server Objects` interfaces to open the queue and to retrieve all the existing work items for that queue.

Whenever any change is made to the work queue after that - a new work item added, an existing work item modified or removed - the change is published to the appropriate JMS topic. If more than one application is subscribed to a given work queue, and they do not use the default topic specified by `WQDJMS_TOPICNAME`, the JMS message must be published to each topic used by a subscribing application.

Understanding the Message Types

Activity Monitoring (IAP) uses two message types, and Work Queue Delta publication (WQD) uses three types.

IAP Message Types

There are two types of message:

- The Monitor Event Detail message (MED). A **BG** process sends this message. The message is addressed to the JMS topic configured in the `IAPJMS_TOPICNAME` process attribute.
- The Monitor Event Request message (MER). The MER message is sent to the `iProcess` database to update the activity monitoring configuration information. It can be sent by SSO or by using the `SWDIR\bin\swutil IMPMONITOR` command.

The Monitor Event Detail Message (MED)

Every MED message sent from the **IAPJMS** process to the external application consists of:

- a message header
- message properties
- message body. The message body contains the detailed XML for the events that are being monitored. The MED XML format is defined by the **SWAuditMessage.xsd** schema.

These messages can be sent in either a basic format or an extended format. You can specify which format is used by setting the `IAPSCHEMA` parameter in the `SWDIR\etc\staffcfg` configuration file. See the section "Tuning the `iProcess` Engine Using `SWDIR\etc\staffcfg` Parameters" in the *TIBCO iProcess Engine: Administrator's Guide* for details of this parameter.

The extended format provides the following additional information in the message body:

- Description and type of the audit user. The basic format identifies the audit user; in the extended format the description and type of the audit user can be specified.
- Addressee. The basic format contains no addressee information. The extended format identifies the name, description and type of the addressee of the step for OPEN, KEEP, RELEASE and FORWARD actions.

- Main procedure of a sub-case. The basic format contains no information about the main procedure of a sub-case. The extended format identifies the procedure name and description.
- Parent procedure of a sub-case. The basic format identifies the parent procedure. The extended format adds a procedure description.

For examples of these IAPJMS messages, see the section "Examples of Configuring Activity Monitoring Information" in the *TIBCO iProcess™ Modeler Integration Techniques guide*.

Filtering Message Event Request (MER) Messages

Every MER message sent to the iProcess database to update the activity monitoring configuration information consists of XML requesting the events to monitor. The MER XML format is defined by the **SWMonitorList.xsd** schema.

WQD Message Types

There are three types of message:

- The Work Queue Delta message (WQD). The **WIS** process sends this message.
- The Work Queue Synchronization message (DSY). This contains the unique synchronization ID for the queue.
- The Work Queue Invalid message (DER). This informs the subscribing application that the work queue should be closed and re-opened.

Work Queue Delta (WQD) Message

Every WQD message sent from the **WIS** process consists of:

- a message header
- message properties
- the message body containing details of the Work Queue Delta. The WQD XML format is defined by the **apSSOTypes.xsd** schema. See the section "Schemas" in the *TIBCO iProcess™ Client (Browser) Action Processor Reference guide* for more information on this schema.

Whenever an item is added, modified or deleted from a monitored work queue, a WQD message is published by the **WIS** to the **IAPJMS** process.

Work Queue Synchronization (DSY) Message

Every DSY message consists of.

- a message header
- message properties
- the message body containing the unique synchronization ID, which will be published when the work queue has been fully loaded and JMS publication is to begin.

A DSY message is sent once the initial queue has been built. It is published to subscribing applications so that the subscriber knows it needs to begin processing any Work Queue Deltas for the specified Work Queue Delta ID.

Work Queue Invalid (DER) Message

Every DER message consists of.

- a message header
- message properties
- an empty message body.

The message type is an exception message sent to all applications subscribing to the work queue on which an error occurs. It tells the subscribing applications that the Work Queue Deltas for this queue are now invalid and that this work queue must be closed and re-opened.

See the *TIBCO iProcess Server Objects (.NET) Programmer's Guide* and the *TIBCO iProcess Server Objects (Java) Programmer's Guide* for more detail on Work Queue Delta publication via JMS and how it works.

Chapter 8

Database Failure and Failover

This chapter describes how the TIBCO iProcess Engine handles database failure or failover.

Topics

- [Overview, page 104](#)
- [TIBCO iProcess Engine Behavior, page 105](#)
- [TIBCO iProcess Workspace Behavior, page 107](#)
- [iProcess Objects and iProcess Server Objects Behavior, page 108](#)
- [TIBCO iProcess Engine Configuration Requirements, page 109](#)

Overview

The TIBCO iProcess Engine stores all iProcess data in its database (Oracle, SQL Server or DB2). It therefore relies on the availability of that database to be able to process work. If the database fails, the TIBCO iProcess Engine can do nothing until the database returns.

Database **failover capability** exists when the database is configured in such a way that if a failure occurs, a standby instance of the database takes over immediately with no or minimal impact on service. Failover capability can be provided in many ways, depending on the hardware, operating system and database systems you are using. Failover capability is therefore an essential component of any high-availability, production-level system involving the TIBCO iProcess Engine.

If a database failure or failover occurs, the TIBCO iProcess Engine is able to:

- detect that the failure has occurred.
- guarantee the integrity of any in-progress transactions.
- seamlessly continue working when the database connection returns (whether that connection is to the original database or a failover database).

TIBCO iProcess Engine Behavior

The following table shows how different TIBCO iProcess Engine components behave in the event of a database failure or failover.

TIBCO iProcess Engine Component	Behavior When Failure or Failover Occurs
Server processes	<p>On SQL and DB2: Each TIBCO iProcess Engine process clears down any database connections it already has (as they are used) and attempts to re-establish a new connection. If necessary, the process retries the connection attempt until the database becomes available again, when the attempt succeeds.</p> <p>On Oracle: The server processes use the Oracle Client to establish and maintain database connections. The Oracle Client therefore attempts to re-establish the database connection. If necessary, it retries the connection attempt until the database becomes available again, when the attempt succeeds.</p>
Database transactions	<p>Any in-progress Write transaction is rolled back. When the database becomes available again:</p> <ul style="list-style-type: none"> • Read-only transactions continue and complete as normal. • Write transactions are retried. <p>This means that transactional integrity is maintained; no transaction or workflow operation is left in an indeterminate state. Operations either succeed and are committed or fail and are rolled back.</p>
Process Sentinels	<p>Any processing that requires a connection to the database hangs until the database connection is restored.</p> <p>For example, if you use <i>SWDIR\util\swsvrmgr</i> to issue a Restart event for a process when the database has failed, the process is not restarted until the database connection is restored.</p> <p>Note - On the Oracle (UNIX/Linux) TIBCO iProcess Engine the event system used by the Process Sentinels is provided by the database. This means that the use of Oracle's Transparent Application Failover (TAF) feature is required to provide failure or failover support. See page 109 for more information.</p>

**TIBCO iProcess
Engine Component**
Behavior When Failure or Failover Occurs
System messages

Any TIBCO iProcess Engine process that attempts to contact the database will fail (until the connection is restored), causing it to write a database connection failure error to the *SWDIR\logs\sw_error* file.

For example:

```
2005/04/25 15:41(BG:1:2496:pro::7.1:4352): 0-lcase_add()
Unable to execute statement (Statement ID - 49) (S1002 -
[Microsoft][ODBC SQL Server Driver]Invalid Descriptor Index)
lcase_add() Unable to execute statement (37000 - Statement
failed due to database FAILOVER )
```

**Command line
utilities**

Any command that requires a connection to the database hangs until the database connection is restored.

For example, if you have a *SWDIR\util\swadm* session open when the database fails, and you then try to add a **BG** process, the command hangs until the database connection is restored. As soon as the database is available again, the command completes and the new **BG** process is added.

TIBCO iProcess Workspace Behavior

Existing TIBCO iProcess Workspace sessions may be unaffected by a database failure or failover, as many TIBCO iProcess Workspace operations - such as retrieving information about a work item from the WIS cache - do not require access to the database.

TIBCO iProcess Workspace behavior is only affected when an operation that requires access to the database is performed - for example, opening a procedure. The TIBCO iProcess Workspace then passes a request to the TIBCO iProcess Engine, which in turn tries to access the database. The request therefore hangs until the database connection returns. If this period exceeds the TIBCO iProcess Workspace's timeout period (by default, 25 seconds), a timeout dialog is displayed to the user. The user has the option to abort or retry the operation.

iProcess Objects and iProcess Server Objects Behavior

The following table shows how different TIBCO iProcess Objects and TIBCO iProcess Server Objects components behave in the event of a database failure or failover.

iProcess Objects/iProcess Server Objects Components	Behavior When Failure or Failover Occurs
iProcess Objects Server	From the viewpoint of a database failure or failover, the iProcess Objects Server is a TIBCO iProcess Engine server process. As such, its behavior is the same as any other TIBCO iProcess Engine server process - see Server processes on page 105 .
iProcess Objects Clients and iProcess Server Objects Clients	<p>Existing iProcess Objects/iProcess Server Objects client sessions may be unaffected by a database failure or failover, as many client operations - such as retrieving information about a work item from the WIS cache - do not require access to the database.</p> <p>iProcess Objects/iProcess Server Objects client behavior is only affected when an operation that requires access to the database is performed - for example, opening a procedure. The client passes a request to the iProcess Objects Server, which passes it to the TIBCO iProcess Engine, which in turn tries to access the database. The request therefore hangs until the database connection returns. If this period exceeds the iProcess Objects Server's RPC timeout period (by default, 25 seconds), a timeout exception is returned to the iProcess Objects/iProcess Server Objects client. It is the iProcess Objects/iProcess Server Objects client's responsibility to handle that exception.</p>

TIBCO iProcess Engine Configuration Requirements

The following sections describe what you need to do to provide failure/failover support on the TIBCO iProcess Engine.

Oracle (UNIX or Linux)



To provide failure/failover support on the Oracle (UNIX/Linux) variant of the TIBCO iProcess Engine, you **MUST** configure the use of Oracle's **Transparent Application Failover (TAF) feature**. Oracle TAF enables an application user (such as the iProcess Engine) to automatically reconnect to a database if the connection fails.

The TIBCO iProcess Engine Process Sentinels are event driven. On the Oracle (UNIX/Linux) variant of the TIBCO iProcess Engine, the event system is provided using Oracle AQ's publish/subscribe mechanism. This means that if the database connection fails, the Process Sentinels cannot receive events, and so appear to hang. To avoid this problem you need to configure TAF so that the connection can be restored. Once the connection is restored, all connections are recovered and normal operation is resumed.



If you **are** running parallel servers, using TAF allows iProcess to switch to an alternative instance if the instance that it is currently using fails.

If you **are not** running parallel servers, using TAF still means that although iProcess will not function while the database is down, it can recover immediately and automatically when the database is recovered.

To enable the use of TAF with the iProcess Engine, you need to configure TAF support for the service name that you intend to use to connect to the Oracle database. See the *TIBCO iProcess Engine (Oracle) for UNIX or Linux Installation Guide* for more information about how to do this.

TAF involves manual configuration of a net service name that includes the `FAILOVER_MODE` parameter included in the `CONNECT_DATA` section of the connect descriptor. For more information about TAF, and how to set up and configure it, please see your Oracle documentation.

Oracle (Windows), SQL Server and DB2 (UNIX or Linux)

Database **failure** support is provided automatically. You do not need to configure the TIBCO iProcess Engine in any way - it recovers immediately and automatically when the database connection is restored.

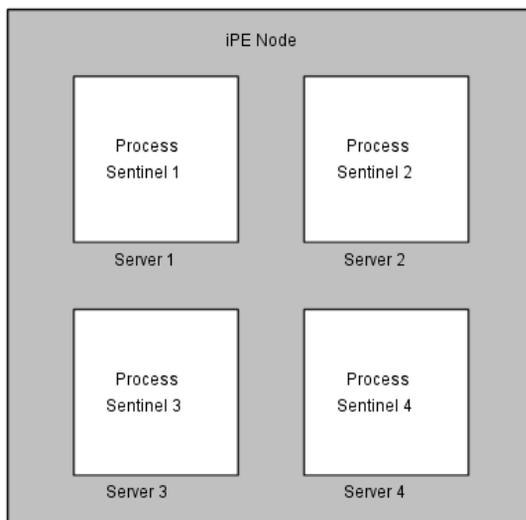
Database **failover** support is provided automatically **provided that you have configured your database system to provide failover capability**. For example, by implementing failover clustering in SQL Server, or by using DB2's high availability disaster recovery (HADR) database replication feature. (Failover capability can be provided in many ways, depending on the hardware, operating system and database systems you are using. Please see your SQL Server or DB2 documentation for more information about how to do this.) You do not need to configure the TIBCO iProcess Engine in any way - it recovers immediately and automatically when the failover completes and the database connection is restored.



Event handling is provided by the **iProcess Events** COM+ application on Windows variants, and by iProcess event/notification daemons on the UNIX DB2 variant. Consequently, the event system has no dependency on the database as it does in the UNIX/Linux Oracle variant.

Chapter 9 Process Management

This chapter explains the concepts of how the TIBCO iProcess Engine processes are managed by the Process Sentinels to make sure that the TIBCO iProcess Engine can operate on a 24*7 basis. If the TIBCO iProcess Engine is installed across multiple servers (a node cluster), processes can operate on the various servers within the node cluster. This means that each server in the cluster requires its own Process Sentinels to control the processes that are running on that server - see below.



Responsibilities of the Process Sentinels

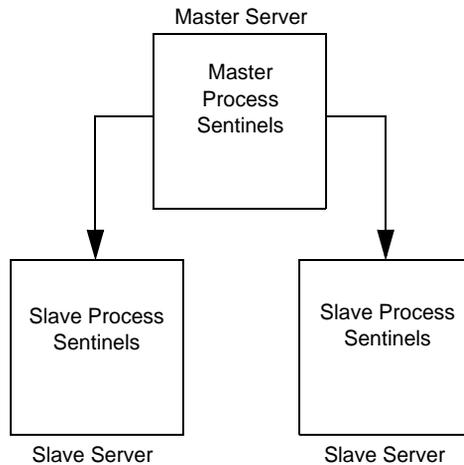
The following list describes the responsibilities of the Process Sentinels:

- Start processes during a server start up or upon a system administrator request. It will control the order in which the processes are started.
- Detect failed processes and restart them automatically (or manually using the command line `SWDIR\util\swsvrmgr` utility or the iProcess Server Manager).
- Shut down processes when the system is shut down or when the administrator requests a sub-set of the system to be stopped.
- Pause the entire system, a group of processes or a single process.
- Detect `SWDIR\logs\sw_error` and `SWDIR\logs\sw_warn` files and send a work item to the system administrator (**swadmin**) informing them that the file has been created and on which server.
- Check that there is enough disk space using the `SWDIR\fspart` file (you can configure which partitions to check). If there is not enough space, a work item is sent to the administrator informing them of the situation.
- Monitor other Process Sentinels and restart them if they fail - see [Restarting Failed Process Sentinels on page 119](#).
- Provide interfaces for integration with TIBCO Hawk® - See the *TIBCO iProcess Engine: Administrator's Guide* for more information.
- Maintain the list of all active user logins.

Distribution and Hierarchy of Process Sentinels

In a distributed TIBCO iProcess Engine node cluster where there are a number of servers running, it is necessary to have Process Sentinels running on each server. To keep the system processes synchronized, the Process Sentinels are designed so that there are master Process Sentinels running on one server and slave Process Sentinels running on all the other servers. See [TIBCO iProcess Engine to TIBCO iProcess Engine Network Communication in a Node Cluster on page 129](#) for more information about how process sentinels communicate between servers.

The following diagram shows that the Master process controls the slave processes in a hierarchical model. The master and slave processes have different responsibilities, which are described in [Master and Slave Responsibilities on page 113](#).



Master and Slave Responsibilities

The master and slave processes have different responsibilities for tasks that they can perform. There are also some tasks that are common across master and slave processes.

Common Process Sentinel Tasks

The following operations are performed by both the master and slave Process Sentinels:

- Starting and stopping processes on the local computer
- Detecting the failure of processes on the local computer and restarting them
- Maintaining the **process_config** table to store the status of processes.
- Check that there is enough disk space using the *SWDIR\fspart* file (you can configure which partitions to check). If there is not enough space, a work item is sent to the administrator informing them of the situation.
- Control the sequence of processes during a system start-up or shutdown.
- Check for *SWDIR\logs\sw_error* and *sw_warn* files. If any are detected, a work item is sent to the system administrator.

Master Process Sentinel Tasks

The master Process Sentinels can perform all of the above tasks but the following task is only performed by the master Process Sentinels.

- Report the status of the node/node cluster by publishing an event, for example, if the server processes are running or have been stopped.



The startup and shutdown requests for server processes are broadcast by the `SWDIR\util\swsvrmgr` utility. The master Process Sentinels will record the new system status when the shutdown or startup is complete.

How Processes are Controlled by the Process Sentinels

The following table lists all the processes that are controlled by the Process Sentinels and show how they operate.

Process Description	Startup	Shutdown	Restart	Can Be Paused Manually	Paused Automatically By Process Sentinels
Background	x	x	x	x	
Database Queue Daemon	x	x	x	x	
Deadline Manager	x	x	x	x	
IAPJMS ^a	x	x	x	x	
Prediction	x	x	x	x	
RPC Background	x	x	x		
RPC Pool Server	x	x	x		
iProcess Objects Server	x	x	x	x	
WIS Mbox Daemon	x	x	x	x	x
Work Item Server	x	x	x	x	x
Work Queue Server	x	x	x		x

a. This process is disabled unless you have chosen to enable it when installing the TIBCO iProcess Engine.

Starting the TIBCO iProcess Engine Processes

The Process Sentinels are responsible for starting all the TIBCO iProcess Engine processes. This can either be starting the complete system or starting individual processes as requested. The Process Sentinels make sure that the processes are started in the correct sequence.

Many processes require that dependent processes are already running, for example, the WQS needs to be started before the WIS processes can be started. The Process Sentinels make sure that all process dependencies are satisfied before starting a process.

All Process Sentinels (master and slaves) maintain a complete list of all processes running in the TIBCO iProcess Engine node cluster. Each server hosts their own Process Sentinels, which are only responsible for processes running on its computer. When the Process Sentinels attempt to start processes, they know if all the process dependencies are satisfied across all the TIBCO iProcess Engines in the node cluster.

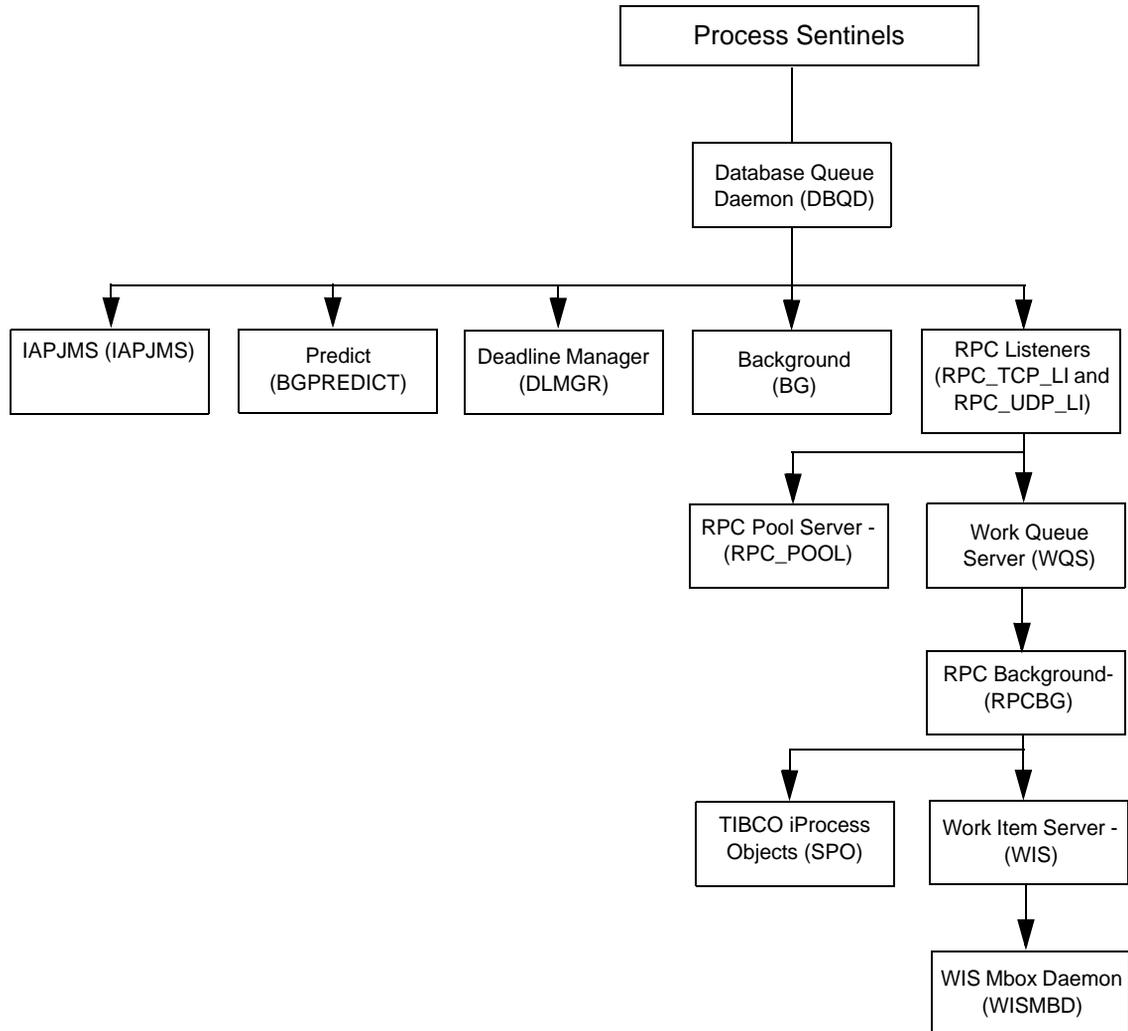
For example, if a system is configured to have 1 WQS, 5 WIS, 5 WISMBD, the processes would be started in the following order:

- Start the WQS processes
- Start all the WIS processes
- Start all the WISMBD processes.

The following diagram illustrates the order in which the Process Sentinels start the TIBCO iProcess Engine processes. All processes are started by the Process Sentinels apart from the RPC pool servers, which are started by the RPC TCP Listener process



The DBQD process is currently only present on the DB2 version of the TIBCO iProcess Engine.



Determining Where Processes Run

The `process_config` table determines on which computers individual processes will run and how many instances of those processes are started. There are a number of rules about where processes can be run:

- All of the foreground processes must operate on the master server. See [Foreground Processes on page 36](#) for a list of these processes.
- The background processes can operate on any server and can be distributed across several servers, for example you might run a background server instance on each of your servers if you are using a node cluster architecture. See [Background Processes on page 37](#).

Restarting Failed Processes

To make sure that the TIBCO iProcess Engine is always running, the Process Sentinels are constantly monitoring for any processes that fail. Each of the Process Sentinels operating in the node will monitor all processes running on the same server. If a process fails, the Process Sentinels will start a new copy of the process. In many cases, there can be many instances of a process running such as 5 Background processes and the Process Sentinels will make sure that the correct instance of the process is restarted.

The **process_attribute** table is responsible for the configuration of how processes are restarted. You can specify whether you want the process to be manually restarted, restarted a set number of times, and the minimum time a process must run before it can be restarted. Settings can be applied to the node, a single server in the node, a type of process, or an individual instance of a process.

Restarting Failed Process Sentinels

Process Sentinels consist of two processes: a worker process and a watcher process. On each server in an TIBCO iProcess Engine cluster, these two processes are started. This architecture is designed so that each process can monitor the other one and restart it if it fails. This makes sure that the Process Sentinels are always running. The watcher and worker processes perform different tasks:

Watcher Process

After starting the worker process and establishing a two way communication channel with the worker process, the watcher process monitors the worker process and can restart it upon a failure. If any errors occur, an error message is logged in the `SWDIR\logs\sw_error` file.

On the DB2 version, the watcher process also runs the iProcess event daemon (see [Event Handling on page 37](#)).

Worker Process

The worker process performs the process monitoring of the TIBCO iProcess Engine processes and monitors and restarts them if they fail. The worker process also monitors the watcher process to make sure that it is always running and will restart it if it fails. The following is the start-up routine for the worker process:

1. Initialize process by processing its command line arguments.
2. Establish a communication channel with its watcher process.

3. Connect to the database and read the **node_cluster** table to determine if the server is configured to be part of the TIBCO iProcess Engine. It can also determine if it needs to operate as a master or slave.

If there is no entry for this server in the **node_cluster** table, an error message is logged to the operating system event log (**syslogd**). The worker process will shut down and also shutdown the watcher process.

4. Read the **process_config** table to build an in memory process hierarchy model.
5. Start process monitoring.

Shutting Down Processes

Processes are shut down by the Process Sentinels in the opposite sequence to when they were started.

Configuring the Process Sentinels

There are several areas of configuration information related to the Process Sentinels:

- The Process Sentinels command line utility (*SWDIR\util\swsvrmgr*)

This utility enables system administrators to control the functions of the Process Sentinels such as setting which computer runs the master process.

Refer to “Using swsvrmgr to Administer Server Processes” in the *TIBCO iProcess Engine: Administrator’s Guide* for more information.

- The iProcess Server Manager

This graphical tool allows you to control the functions of the Process Sentinels in much the same way as with *SWDIR\util\swsvrmgr*. Refer to “Using the iProcess Server Manager to Administer Server Processes” in the *TIBCO iProcess Engine: Administrator’s Guide* for more information.

- *SWDIR\etc\staffpms* file

When the Process Sentinels start, it uses the database connection information in this file to establish the database connection.

Refer to “Using the TIBCO iProcess Engine Configuration Files” in the *TIBCO iProcess Engine: Administrator’s Guide*.

- **node_cluster** table

After the Process Sentinels have connected to the database, it reads this table to see if the computer on which it has been started is part of the TIBCO iProcess Engine node cluster. If it is not, the Process Sentinels will shut down.

Refer to the appropriate *TIBCO iProcess Engine Database Administrator’s Guide*.

- **process_config** table

The Process Sentinels read this table to determine which processes should be running on which computer. A process hierarchy model is created in memory, which is used to start and stop the processes in the correct sequence.

Refer to the appropriate *TIBCO iProcess Engine Database Administrator’s Guide*.

- **process_attribute** table

The Process Sentinels read this table to find out how processes should be restarted, for example if they should be restarted automatically or manually. This table also defines how long the Process Sentinels are idle for before checking for *SWDIR\logs\sw_error* and *sw_warn* files and disk space.

Refer to the appropriate *TIBCO iProcess Engine Database Administrator’s Guide*.

- You can use the TIBCO iProcess Engine Configuration utility `SWDIR\util\swadm` to administer the Process Sentinels and server processes. You can perform tasks such as setting the server on which to run the master Process Sentinels, adding a new server to the node cluster or enabling and disabling processes. This utility directly reads and updates the iProcess database tables.

Refer to “Administering Servers” in the *TIBCO iProcess Engine: Administrator’s Guide*.

Chapter 10 Network Communication

This chapter describes how:

- TIBCO iProcess Workspace communicates with the iProcess Engine. See [TIBCO iProcess Workspace and TIBCO iProcess Engine Network Communication on page 126](#).
- TIBCO iProcess Engines communicate when they are in a node cluster. See [TIBCO iProcess Engine to TIBCO iProcess Engine Network Communication in a Node Cluster on page 129](#).
- The iProcess Engine works when a firewall is implemented on your network. See [Using the TIBCO iProcess Engine in a Firewalled Environment on page 130](#).

TIBCO iProcess Workspace and TIBCO iProcess Engine Network Communication

This section provides a brief overview of how the TIBCO iProcess Workspace and TIBCO iProcess Engine communicate with each other via the network.

The TIBCO iProcess Workspace makes *Remote Procedure Calls* (RPC) to the server using ports. This means the TIBCO iProcess Workspace makes calls to the TIBCO iProcess Engine to find out information such as what work queues and work items it needs to display. It also means that procedures can be started from one computer but actually be running on another computer on the network. There are a number of reasons why this might happen:

- The local computer is not able to provide the functionality required.
- The requested service is shared amongst many clients but centralized arbitration, synchronization and communication is required.
- You want to use the superior performance of a remote computer.
- You want to balance the processing load across the network.

Any program that offers functions that can be remotely accessed via RPC must have a unique RPC number. The RPC number for a service is either pre-defined (fixed) or allocated dynamically.

All RPC numbers need to be bound to network ports before remote clients can communicate with the RPC servers. This is because the port is the fundamental method of communication between computers.

The allocation of ports, similar to the allocation of RPC numbers, can be fixed or dynamic and is determined by the program providing the RPC service and the operating system. Therefore, it is important to know the distinction between RPC ports and numbers.

Function of a Portmapper

The connection between the RPC client computers and RPC servers is serviced by a program known as the **portmapper** (or **rpcbind**). The portmapper provides a directory for the RPC services on the computer on which it runs.

A client application can request a connection to an RPC service by passing the portmapper the RPC number. The portmapper responds by sending back the actual port number that the service is bound to. The client can then connect directly to the service using the service's port.

The TIBCO iProcess Engine RPC Service

The process that monitors the RPC calls that request a connection to the TIBCO iProcess Engine is known as the RPC listener. This has a pre-defined RPC number - by default it is 391875. If a TIBCO iProcess Engine is communicating with another TIBCO iProcess Engine, the default port is 391870. The *SWDIR\swdefs* file contains the RPC numbers.

When the iProcess Engine starts:

1. The RPC TCP listener process starts and requests the operating system to allocate it a port.
2. It binds a TCP socket to this port.
3. It then contacts the portmapper and requests that its RPC number (391875) is registered on the port that the operating system has allocated it.
4. The portmapper makes an entry in an internal table of the RPC number and port.
5. The iProcess Engine starts a number of other RPC servers but with dynamically allocated RPC numbers. These servers bind to a port dynamically allocated by the operating system and then registers the relationship between the RPC number and port with the portmapper.

The following describes what happens when an TIBCO iProcess Workspace connects to an iProcess Engine:

1. The TIBCO iProcess Workspace connects to the portmapper program.
2. The TIBCO iProcess Workspace requests details of the port to which the client RPC listener is bound. The request specifies the RPC number (the default is 391875).
3. The TIBCO iProcess Workspace connects to the client RPC listener using the port details and then requests details of the RPC numbers of the other RPC servers the TIBCO iProcess Workspace needs.
4. For each of the RPC numbers obtained, the TIBCO iProcess Workspace:
 - a. Contacts the portmapper.
 - b. Requests details of the port to which the RPC number is bound.

TCP/IP

The RPC listeners and TIBCO iProcess Workspace use the TCP/IP protocol to communicate. In order for the TIBCO iProcess Workspace to communicate with the RPC listener processes, the TIBCO iProcess Workspace must be able to perform HOST name resolution. This means that data packets are sent to a host

name such as HERCULES rather than to an IP address such as 10.10.56.202 and the TCP/IP stack looks up the correct IP address. A typical way of setting this up is to use a HOSTS file on each computer which provides a simple look up table for a host name and its IP address.

Internet Protocol Version

TIBCO iProcess Engine to TIBCO iProcess Engine Network Communication in a Node Cluster

This section provides a brief overview of how TIBCO iProcess Engines communicate with each other when configured as part of a node cluster.

The only processes that communicate with each other between nodes in a cluster are the Process Managers and the `SWDIR\util\swsvrmgr` process. The `swsvrmgr` process tells the process managers what actions to perform. They communicate via the TIBCO iProcess Engine event mechanism by publishing and subscribing to events.

For example when you start the system by running `SWDIR\bin\swstart`, the `swsvrmgr` START utility publishes an event to say that the system needs to start up something. All the process managers on all the nodes in the cluster should receive that event and start up any processes that are configured to run on that node. Similarly on shutdown, a shutdown event is published and each machine's process manager should receive the event and shutdown any appropriate processes.

Using the TIBCO iProcess Engine in a Firewalled Environment

In many enterprise network models, a firewall is used to link logical networks together to provide access security into the protected network. The following section describes how the iProcess Engine can be configured to work in a firewall environment.

What is a Firewall?

A firewall is a computer that links two logical networks together and re-routes data between the two networks as required. The firewall computer also contains a filter. This filter only allows data to pass through it that is requesting a particular service (using a variety of filtering methods defined by the firewall administrator).

A typical use of a firewall is for Web servers. A Web server needs to be accessed by remote computers outside of the logical network so they can access the web service. However, these computers should not be able to access other services on that server that are more likely to be a security risk.

A firewall can vary in the way that they restrict access to the network, they can:

- only allow access to and from certain computers.
- analyze the port number the client is requesting and compare it to a list of port numbers allowed.
- analyze the data that is being sent and only allow it through if it conforms to some pre-determined rules that have been set up.

Within the data being sent (known as *packets*), many firewalls can obtain the RPC number requested for RPC calls and only allow data through if it is requesting a particular RPC number and therefore a particular RPC service.

iProcess RPC and Firewall Access

When an iProcess Workspace and the iProcess Engine are separated by a firewall, the iProcess Suite can fail because its communication method (remote procedure calls - RPC) is stopped by the firewall filter. Because iProcess Engine RPC services are allocated dynamically, the firewall filter is not set up to open all ports that the iProcess Engine is using. Not all the ports will be open because the firewall administrator has set up certain restrictions to enable security on the network.

The RPC numbers are allocated dynamically so there is no fixed set of RPC numbers for a firewall administrator to add to the filter. If the ports used are not opened up on the firewall, the iProcess Workspace and iProcess Engine cannot communicate because data requests are denied by the firewall. In order for the iProcess Engine to operate in this environment, the firewall administrator needs to know what ports the iProcess Engine is using so that iProcess RPC calls can be filtered through.

You can set up the iProcess Engine to use a specific range of ports and/or RPC numbers so that the firewall administrator has a range of port numbers to add to the firewall filter. You can use one or both of the following methods to do this:

- Port range filtering
- RPC number filtering.

You use the `SWDIR\util\swadm` utility to configure port range and/or RPC number filtering. Refer to “Administering Firewalls” in the *TIBCO iProcess Engine: Administrator’s Guide* for more information.

Port/RPC Number Resource Logging

When Port/RPC numbering is enabled, a log file containing resource allocation and release operations is stored in `SWDIR\logs\rpcport.log`. This is a text file containing entries for the following events:

- Start-up of the Port/RPC resource allocation service
- Shutdown of the Port/RPC resource allocation service
- Allocation of a Port/RPC number
- Release of a Port/RPC number
- Failure to re-bind a released port
- Successful re-bind of a previously failed port
- Errors in allocations/release of port/RPC number.

Using Oracle Events Through a Firewall

You can specify a range of AQ port numbers to be used in iProcess for communication through the firewall with Oracle Events.

You must ensure that you specify enough AQ ports for iProcess system processes. On a typical system, you should specify a minimum of 7 ports for iProcess (sentinel and utility processes), plus an additional port for each process defined in the `process_config` database table.

In addition to configuring the AQ port range in iProcess, you must also ensure that your firewall is configured to allow access for the iProcess AQ port range as well as the iProcess RPC port range, and any required Oracle ports (such as the default port 1521).

The following two tables are provided:

`aq_port_range_conf`: this records port ranges including the start port number and the count of the ports.

`aq_port_range`: this records the state of every port in every range.

You use the `SWDIR\util\swadm` utility to configure the aq port range. Refer to “Administering Firewalls” in the *TIBCO iProcess Engine Administrator’s Guide* for more information.

Using JMX Through a Firewall

JMX relies on a JAVA technology called RMI, which uses dynamic ports to be able to communicate between a client and a server. Firewalls cannot handle dynamic ports as they need to know the port number. iProcess overcomes this problem by statically assigning a listening port for the RMI server.

Refer to the Post Installation tasks in *TIBCO iProcess Installation* for more information.

Index

A

Activity Publishing [94, 96](#)

C

customer support [xix](#)

I

IAP Message Types [100](#)

IAPJMS process [94](#)

S

support, contacting [xix](#)

T

technical support [xix](#)

TIBCO_HOME [xvi](#)

W

Work Queue Delta Publicatio [95](#)

WQD Message Types [101](#)