HP Business Service Management

for the Windows operating system

Software Version: 9.00

Business Process Insight Server Administration Guide

Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Copyright Notices

Trademark Notices

Acknowledgements

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.

- Document Release Date, which changes each time the document is updated.

- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates, or to verify that you are using the most recent edition of a document, go to:

http://h20230.www2.hp.com/selfsolve/manuals

This site requires that you register for an HP Passport and sign-in. To register for an HP Passport ID, go to:

http://h20229.www2.hp.com/passport-registration.html

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

Visit the HP Software Support web site at:

http://www.hp.com/go/hpsoftwaresupport

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

http://h20229.www2.hp.com/passport-registration.html

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

Table of Contents

PART IV: APPENDIXES

Welcome to This Guide

This guide describes how to install the HP Business Process Insight (BPI) Server, which is an HP Business Service Management (BSM) component. Before installing the BPI Server component, you must already have installed HP Business Service Management.

The guide provides information on how the BPI Server installation fits within your BSM implementation and what you need to consider when installing the BPI Server for the first time.

This chapter includes:

How This Guide Is Organized

The guide contains the following parts:

Part I   **Business Process Insight Server Introduction and System Requirements**

Provides an overview of the BPI Server concepts.

Part II   **Installing the Business Process Insight Server**

Provides installation information.

Part III   **Configuring and Administering the Business Process Insight Server**

Provides configuration and administration information.

Part IV   **Appendixes**

Provides information on the database tables defined for BPI, migrating the BPI Server install to a 64-bit system, and reporting problems to Hewlett-Packard.

Who Should Read This Guide

This guide is intended for the following users of HP Business Process Insight:

➤ HP Business Process Insight administrators

➤ HP Business Process Insight application administrators

➤ HP Business Process Insight end users

Readers of this guide should be knowledgeable about navigating and using enterprise applications, and be familiar with HP Business Service Management and enterprise monitoring and management concepts.

Additional Online Resources

Troubleshooting & Knowledge Base accesses the Troubleshooting page on the HP Software Support Web site where you can search the Self-solve knowledge base. Choose **Help** > **Troubleshooting & Knowledge Base**. The URL for this Web site is http://h20230.www2.hp.com/troubleshooting.jsp.

HP Software Support accesses the HP Software Support Web site. This site enables you to browse the Self-solve knowledge base. You can also post to and search user discussion forums, submit support requests, download patches and updated documentation, and more. Choose **Help** > **HP Software Support**. The URL for this Web site is www.hp.com/go/hpsoftwaresupport.

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract.

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

To register for an HP Passport user ID, go to:

http://h20229.www2.hp.com/passport-registration.html

HP Software Web site accesses the HP Software Web site. This site provides you with the most up-to-date information on HP Software products. This includes new software releases, seminars and trade shows, customer support, and more. Choose **Help** > **HP Software Web site**. The URL for this Web site is www.hp.com/go/software.

Documentation Updates

HP Software is continually updating its product documentation with new information.

To check for recent updates, or to verify that you are using the most recent edition of a document, go to the HP Software Product Manuals Web site (http://h20230.www2.hp.com/selfsolve/manuals).

Part 1

**Business Process Insight Server
Introduction and System Requirements**

1

Introduction to the BPI Server

Business Process Insight (BPI) is an HP Business Service Management (BSM) application. Some components of BPI are installed as part of your BSM implementation, and some components need to be installed separately.

This chapter includes:

➤ BPI Server Features on page 17

➤ BPI Instance Database Schemas on page 18

➤ BPI Components on page 20

➤ Licensing on page 23

➤ High Availability Using Microsoft Server Clusters on page 25

➤ Summary of BPI Server Installation Steps on page 25

➤ What to Do Next on page 26

BPI Server Features

The BPI Server enables you to understand the business impact of system failures, and degraded service performance, on your operations and customers. It also enables you to monitor and report on the status of your business.

The BPI components installed as part of Business Service Management enable you to:

➤ model Business Processes using the BPI Modeler.

➤ link your Business Process Activities to IT operational resources, which are represented as CIs within the ODB.

➤ set KPI objectives on your Business Process and Business Process Activities in order to monitor the progress of your Business Processes against the defined objectives.

➤ View the health of your Business Processes within the Business Process Diagram in Service Health.

The following is a list of some of the additional features that are available when you have installed and configured the BPI Server:

➤ BPI Modeler, which is enhanced to enable you to model and link individual Business Process Activities to business events.

➤ Monitor Definer, which you use to define instance thresholds. You can configure the system to alert you when individual process instances violate these instance thresholds.

➤ Business Process Insight Application, which you use to monitor the health of the individual instances of your processes. This is accessed from the following menu option:

Application > **Business Process Insight** > **Health**

➤ Reporting capabilities, which enable you to create reports from the data that is created and stored by the Business Process Insight Application.

If you are interested only in monitoring the effect of IT operational resources on your processes, you do not need to install the BPI Server, and you do not need to read this document.

BPI Instance Database Schemas

The BPI Server uses the following two database schemas: BPI Server schema and Process Repository schema.

The information in these schemas is used by the BPI Server components for monitoring and progressing process instances, and storing the statistics relating to the ongoing status of the deployed Business Processes. It is also used as a source of information for the BPI Application to report on the health of your Business Processes and includes process impact data, instance threshold violation data and Business Process monitor data.

The Process Repository Schema and the BPI Server schema communicate through a deployer component, which is a component of the BIE. This means that when a change is made within the Process Repository and the Process Repository schema, the Process Repository notifies the deployer component within the BIE and the BPI Server schema is then also updated.

For more details on the database schemas, see "Database Schemas" on page 41.

BPI Components

The following is a high level diagram of the BPI components that you are installing and their connectivity to other components.
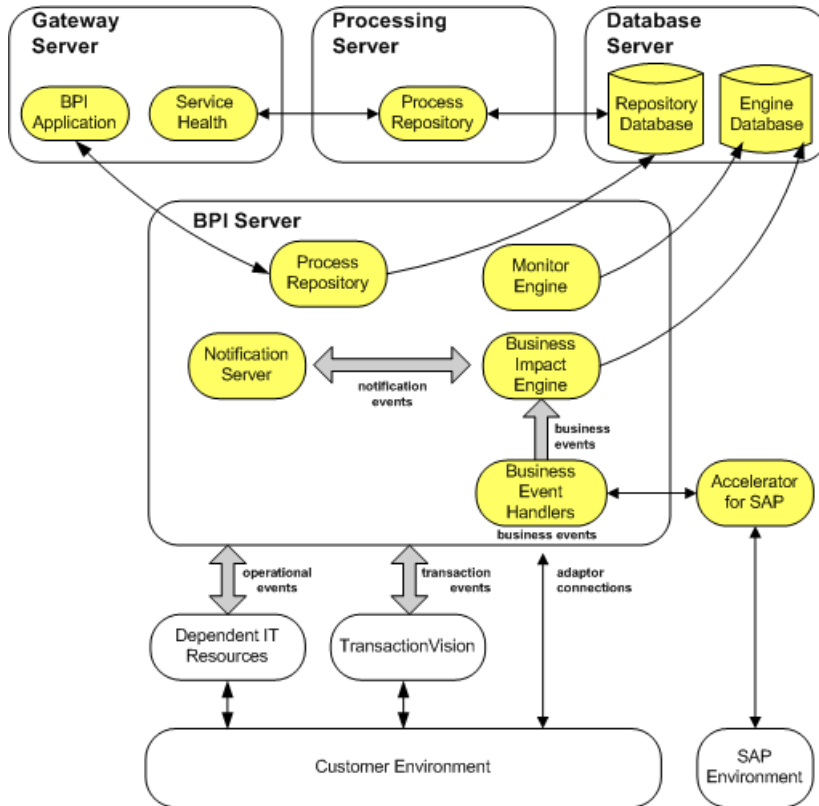
In addition to the BPI components, some third-party products are installed with BPI, one of which is the Tomcat Servlet Engine. Tomcat is used to manage the Java Server Pages (JSPs), which provide the dynamic web content for some of the BPI Web pages.

BPI Installation Programs

BPI has the following installation programs:

➤ BPI Server

➤ BPI Accelerator for SAP

There is an installation file for the BPI Accelerator for SAP, which contains the files and information required to install the Accelerator. For details of installing and using the Accelerator, see the *HP Business Process Insight Accelerator for SAP Guide*.

Note: Only one copy of the BPI Server can be run on any one machine. You cannot install and run multiple copies of the BPI Server component, or any install option that contains the Server component.

In the case of the Accelerator installation, it can be installed at any time, before or after the BPI Server; however, you need to have installed BSM and the BPI Server needs to be running before you can use the Accelerator.

For details on requirements before the installation and the installation procedures, see Chapter 4, "Installing the BPI Server".

Installation Options and Component Distribution for the BPI Server

You can install the BPI Server on the same machine as the BSM Processing Server, or on a separate machine. Where you install BPI depends on your requirements, the purpose of the installation and the number of business events that you anticipate the BPI Server will be processing. BPI uses a significant amount of database resource to process data, so unless you are using BPI for development or for a pilot, you need to make sure that the BPI Server has a dedicated database instance, or dedicated database server installation.

The following are the recommended options for installing the BPI Server:

➤ On the machine where the Data Processing Server is installed and using the same database instance as BSM.

This option is suitable only for development or pilot implementation purposes. You should not run an enterprise system using this option. Using this option means you can rely on the existing configuration of the database and you can just install the BPI Server without any further configuration of the database.

➤ On a different, dedicated machine using a separate database instance within the same database server.

This option is suitable for an enterprise system, but relies on the fact that it is possible to configure another database, or database instance on the existing database server; see also option 3.

➤ On a different, dedicated machine using a separate database server.

This option is required for an enterprise system, in cases where it is not possible to create an additional database or database instance within the existing database server.

The last two options (installing the BPI Server on a different, dedicated machine) require you to set up and configure a database server, or an instance within an existing database server for the BPI Server. Section "BPI Instance Database Requirements" on page 72 provides details of the database configuration requirements for the BPI Server.

Licensing

To use the full functionality of the BPI application on Business Service Management, you need the following:

➤ Installed BPI application license as described in this section.

➤ Installed BPI Server as described in Part II, "Installing the Business Process Insight Server."

➤ Enabled BPI option on the BSM Server Deployment page on the BSM user interface. For more information, see "Using BPI Within Business Service Management" in *Using Business Process Insight*.

Initially you are assigned a time-based trial license. This license covers both the BPI application and an unlimited number of BPI processes and the BPI Accelerator for SAP applications. Within this evaluation period, you must obtain either a trial license extension (see *Platform Administration)* or a permanent license key password (see "Permanent License Key Password" on page 24).

Note: Licensing is centrally managed by Business Service Management. For 9.00 and beyond, you no longer need to install individual licenses on each server.

Following is important information about licensing:

➤ You install the BPI application license from the Business Service Management License Management page at **Admin** > **Platform** > **Setup and Maintenance** > **License Management**. For licensing instructions, see *Platform Administration*.

➤ On the Business Service Management License Management page, you can:

   ➤ see whether the BPI licenses are valid.

   ➤ view the BPI consumed license count (deployed process count).

   ➤ view the BPI license capacity (maximum number of licenses).

➤ If BPI is not licensed and BPI is enabled on the BSM Server Deployment page, you can only access the BPI Modeler from **Admin** > **Business Process Insight** > **Modeling**. In this case, you can create and model a Business Process, but the ability to create or modify Data and Event definitions is disabled. For more details, see the "Modeler Management" chapter in *Using Business Process Insight.*

If BPI is disabled on the BSM Server Deployment page, you cannot access the BPI application, including the BPI Modeler. For more information, see "Using BPI Within Business Service Management" in *Using Business Process Insight.*

If the BPI engines had been installed, they remain running, but without the BPI license, you cannot modify elements of the deployed processes.

Note: The license keys are not backward compatible. Therefore, you need new license keys for BPI 9.00 and beyond.

Permanent License Key Password

When you purchased BPI you received one or more entitlement certificates. An entitlement certificate is issued for one or more of the licensed products that you purchased, for example:

TB145AA - HP BAC BPI 1 Process SW LTU

Use the information on your entitlement certificates and the Autopass License Management utility to retrieve and install the software license keys for BPI. You can access the Autopass license configuration utility from the Business Service Management License Management page. For details, see *Platform Administration.*

High Availability Using Microsoft Server Clusters

You can install BPI in a Microsoft Server Cluster environment in order to provide a high-availability solution for any of the BPI Server components that can be installed on a Microsoft Windows platform. When using the BPI Server within a Microsoft Cluster, the BPI Server must be deployed in Active/ Passive mode.

By supporting Microsoft Server Clusters in an Active/Passive mode, BPI Server components are automatically restarted on an alternative node in the cluster, if the node they are currently running on fails for any reason.

Information about how to configure your BPI Server to operate within a Microsoft Server Cluster environment is provided in Chapter 12, "High Availability Using Microsoft Clusters." If you intend to operate BPI within a Microsoft Server Cluster environment, you must read the information before starting the BPI installation. This guide makes no reference to the cluster environment when installing the BPI components for the first time.

Summary of BPI Server Installation Steps

This section provides a summary of the BPI Server installation steps. The summary is for a new installation and is not appropriate for upgrade or reinstallation. If you are reinstalling or upgrading, follow the instructions appropriate to these tasks; see Chapter 5, "Re-installing BPI Server Components," or Chapter 7, "Upgrading to BPI Version 9.00," respectively.

This summary provides pointers to sections later in the guide where the full instructions for the specified installation type are described.

To install the BPI Server, perform the following main steps:

1 Make sure that you have installed the correct version of Business Service Management; see "BPI Server Installation Requirements" on page 66.

2 Check the BPI Server section of the BSM Readme for any required pre-installation tasks.

3 Make sure that you have reviewed the database requirements for the BPI Server and that you have the database configuration appropriate to your implementation.

4 Verify all the hardware and software requirements for the BPI Server installation; see "BPI Server Installation Requirements" on page 66.

5 Ensure that the database you are planning to use for the BPI Server is correctly configured and is running and accessible to the machine where the BPI Server is to be installed; see "BPI Instance Database Requirements" on page 72.

6 Collect the data that you need for the installation; see "Information Required for the BPI Server Installation" on page 84

7 Install the BPI Server; see "Installing the BPI Server Components" on page 97.

8 Start the BPI Server; see "Starting BPI Server Components" on page 107.

9 Install a valid license key for BPI; see "Licensing" on page 23.

10 Run the verification step to check the installation as described in section "Verifying Your Installation Using the Integrity Checker" on page 110.

What to Do Next

Read the chapter appropriate to the task that you want to complete:

➤ Chapter 2, "BPI Server Architecture," describes the architecture of the BPI Server system and how it fits into the rest of your BSM implementation.

➤ Chapter 3, "Before Starting the Installation" describes the tasks that you need to complete before starting a new, or an upgrade installation.

➤ Chapter 4, "Installing the BPI Server" provides details of the tasks to install the Business Process Insight Server for the first time.

➤ Chapter 5, "Re-installing BPI Server Components" describes the tasks that you need to complete to re-install Business Process Insight Server components.

➤ Chapter 6, "Uninstalling BPI Server Components" describes the tasks that you need to complete to remove the Business Process Insight Server components from your machines.

➤ Chapter 7, "Upgrading to BPI Version 9.00" provides an upgrading caveat for 9.00.

➤ Chapter 8, "Component Administration," describes configuring and managing the BPI components using the Administration Console on the BPI Server.

➤ Chapter 9, "Deleting Instances from the Database," describes how to delete Process, Data and Monitor instances using the database instance cleaner options.

➤ Chapter 10, "Servlet Engine Authentication," describes how to use the Memory Realm mechanism within the Servlet Engine.

➤ Chapter 11, "Connecting BPI With TransactionVision," describes how to connect BPI with TransactionVision, a product within the Transaction Management application, which enables BPI to show the status of Business Process instances.

➤ Chapter 12, "High Availability Using Microsoft Clusters," describes how use a Microsoft Server Cluster environment to provide a high-availability solution for BPI.

➤ Chapter 13, "Backup and Recovery," describes the backup and recovery procedures for the BPI Server in cases where you are not using the Business Service Management high-availability solution.

➤ Chapter 14, "Log Files," describes the log files for the BPI components.

➤ Appendix A, "Database Tables Defined for BPI" provides lists of all the database tables and views that are created for BPI, to enable you to check that all the tables are included in any backup and recovery steps, as well as check for any potential database table name clashes.

➤ Appendix B, "Reporting Problems to Hewlett-Packard" describes how to report problems to Hewlett-Packard.

2

BPI Server Architecture

This chapter provides a high-level description of the architecture of the BPI Server system and how it fits into the rest of your BSM implementation. If you want an overall introduction to BPI as a BSM application, see Chapter 1, "Introduction to the BPI Server".

This chapter includes:

➤ BPI Architecture Overview on page 29

➤ BPI as a Component of Business Service Management on page 30

➤ HP Business Process Insight Application on page 35

➤ BPI Administration on page 40

➤ Database Schemas on page 41

➤ BPI Server on page 45

➤ BPI Integration Points on page 59

➤ BPI Security on page 61

BPI Architecture Overview

The architecture is presented to provide an understanding of the components and concepts that are described elsewhere within the BPI documentation.

➤ A high-level representation of the BPI component architecture; see "BPI as a Component of Business Service Management" on page 30.

➤ The BPI Application, which you use to track and monitor the impact information recorded on the BPI Server system; see "Architecture of the Business Process Insight Application" on page 36.

➤ The BPI administration clients; see "BPI Administration" on page 40.

➤ The BPI databases and how they are used by the BPI components; see "Database Schemas" on page 41

➤ The BPI Server components; see "BPI Server" on page 45.

➤ Integration points for BPI; see "BPI Integration Points" on page 59.

➤ "BPI Security" on page 61.

BPI as a Component of Business Service Management

When you first install Business Service Management, some BPI components are also installed as shown in the following diagram:
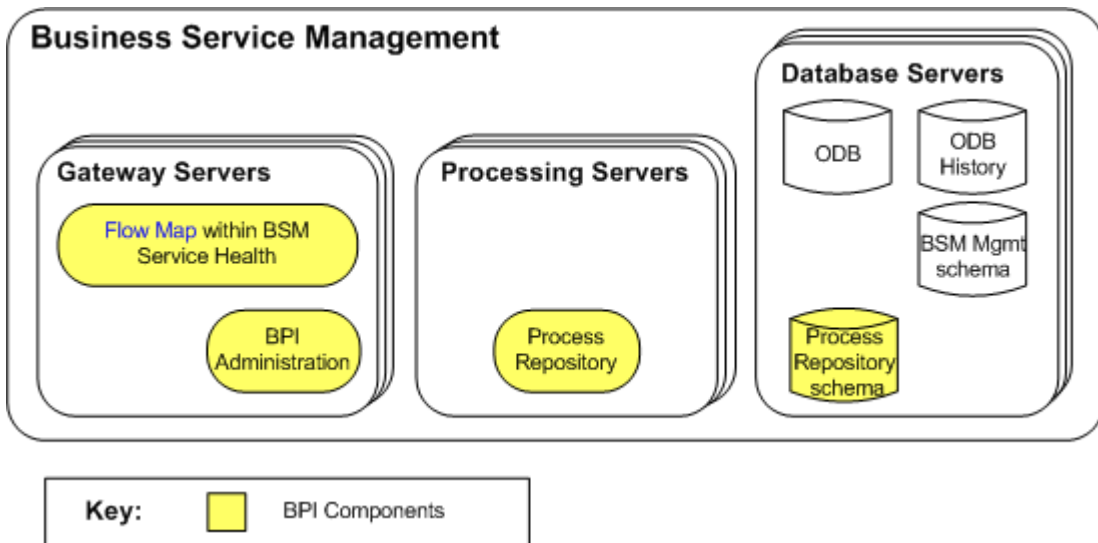
In this configuration, you can model a Business Process and map individual process Activities to the CIs representing the underlying IT infrastructure (IT operational resources).

The BPI Process Repository is installed as part of the Business Service Management Processing Server. The BPI Process Repository manages Process Definitions, which you create using the BPI Modeler, to monitor IT operational resources defined within the ODB. You create a database schema for the Process Repository as part of your BSM installation, or subsequent database configuration.

When you deploy the Business Process that you have created, or modeled, to the Process Repository, you can then use the Service Health Application to view the status of the health of the IT operational resources in the Business Process Diagram.

The BPI components that are installed with Business Service Management are:

➤ BPI Modeler, which you access from:

   Admin > **Business Process Insight**

➤ Process Repository and associated database schema.

➤ Business Process Diagram, which you access from:

   Applications > **Service Health** > **Service Health** and select **Process Diagram** from the **Select Page** drop-down list.

The following list describes the relevant components of your Business Service Management implementation:

➤ Gateway Server

   The Gateway Server is responsible for running BSM applications, producing reports, operating the Administration Console, receiving data samples from the data collectors and distributing this data to the relevant BSM components.

   In the case of BPI, the Gateway Server passes commands from the BPI Modeler to the Processing Server.

➤ Data Processing Server

The Data Processing Server is responsible for aggregating data, running the Business Logic Engine, and controlling the ODB service.

In the case of BPI, the Processing Server controls the IT operational resource CIs, which you can link to Activities within your Business Processes.

➤ Database Server

The Database Server manages the database schemas required for all BSM components, including BPI.

When you subsequently install the BPI Server, the BPI Application is installed.

The following diagram shows the layout of the BSM system following the BPI Server installation, including the location of the BPI components.
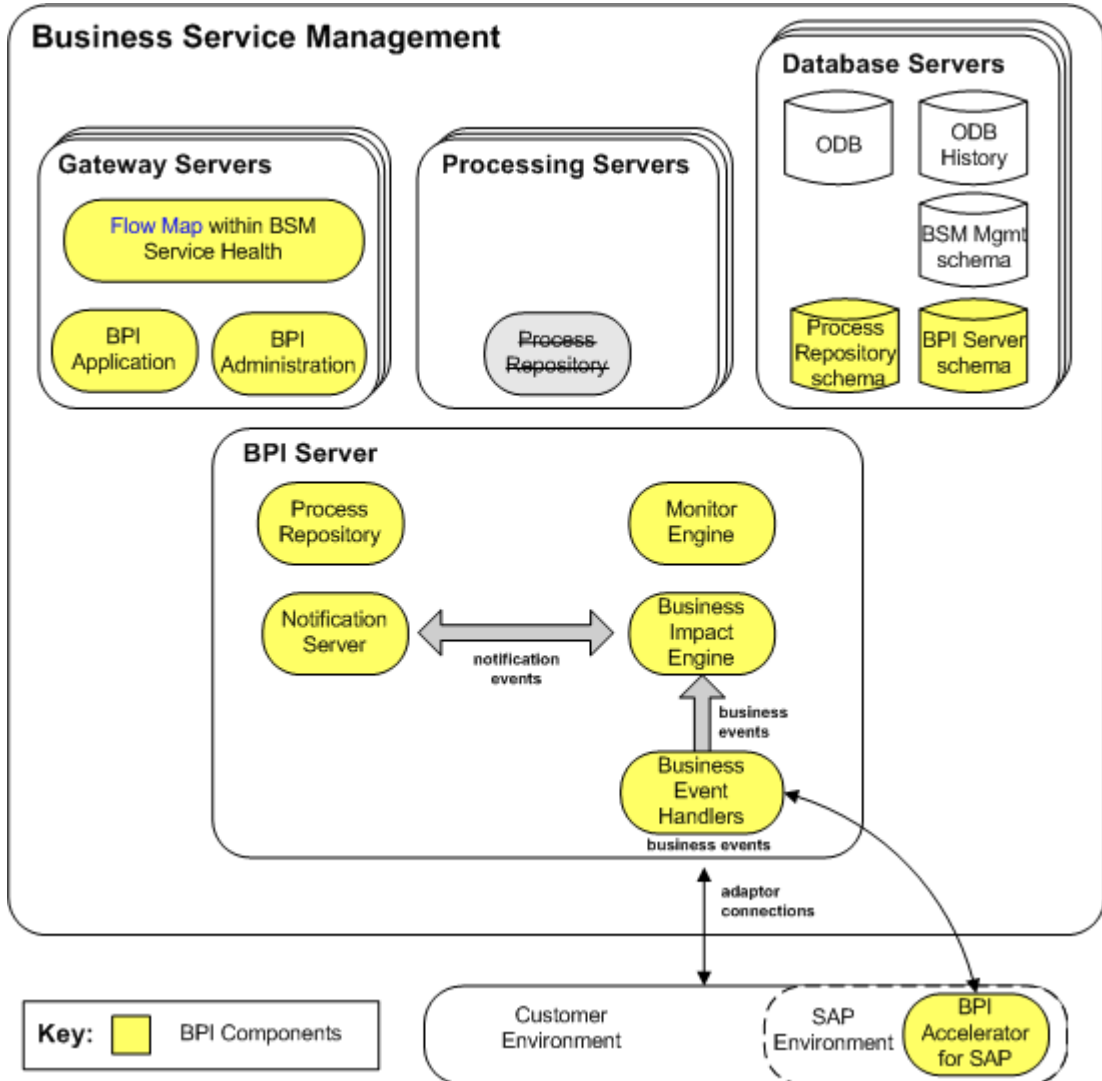
When you install the Server component of BPI, the existing Process Repository is replaced by a revised version, which is installed on the same system as the BPI Server. The revised Process Repository manages all Process Definitions from this point on, for processes that you create with instance tracking enabled and disabled.

The revised version of the Process Repository is installed on the same system as the BPI Server and automatically disables the Process Repository on the Processing Server.

In this configuration, you can model a Business Process and map individual process Activities to the CIs representing the underlying IT operational resources as before. In addition, you can model a Business Process and link business data and business events and to transaction events to the Process Activities in order to track and monitor the health of your Business Processes instances. You track and monitor your Business Process from the BPI Application as follows:

Application > **Business Process Insight**

When you have modeled and deployed your Business Processes the BPI automatically maintains basic statistics about the Business Process, such as:

➤ Overall state of the process.

➤ Number of active process instances.

➤ Specific process instance statistics such as:

   ➤ Start and stop times of each Activity

   ➤ Activity durations

These data are held in the BPI Server instance database.

In addition to this information, you can specify additional process measurements, known as Business Process monitors; for example, you can define measurements such as:

➤ Time taken to process an order.

➤ Current backlog of flights waiting to get airborne.

➤ Average value of an order.

Having defined Monitors, you can then configure instance Thresholds. An instance Threshold provides the data that BPI requires to raise a violation when a business process monitor does not meet its expected parameters; for example, you might want to configure an instance Threshold that issues a warning when the time taken to process an order exceeds four hours.

Business Process Monitors and instance Thresholds violations are reported using the BPI Application.

Defining Monitors and instance Thresholds is described in more detail in the *HP Business Process Insight Integration Training Guide - Defining Business Process Monitors* and the *Online Help* for the Monitor Definer.

HP Business Process Insight Application

When you install the BPI Server, and you purchase a suitable BPI license, all of the Business Process Insight options are accessible from the BSM **Application** and **Admin** menus.

The Business Process Insight Application functions enable you to monitor and track Business Process instances, in real time, using a number of different views and graphical displays. For example BPI enables you to:

➤ present the operational health of your Business Processes, based on changes in the status of the IT operational resources, which are mapped to your Business Processes and Business Process Activities.

➤ monitor the business health of your Business Processes and Business Process Activities based on the statistics derived from process instance-based information.

➤ display the business performance and business health, based on Monitor instance Threshold violation data for the process instances.

➤ display the business KPI values against objectives that you have defined.

➤ show historical business performance data in the form of reports.

These functions are described in more detail in "Business Process Insight Application Features" on page 38.

The following section ("Architecture of the Business Process Insight Application" on page 36) describes the architecture of the BPI Application.

Architecture of the Business Process Insight Application

The Business Process Insight Application is a BSM interface for viewing the business and operational Health and of the Business Processes that you have modeled and deployed in the BPI Modeler.

The data for the displays on the Business Process Application pages is taken from a number of different sources:

➤ Process Repository schema

This is data related to the Business Processes that you have defined.

➤ BPI Instance database

This is statistical data, which is collected as the BPI Server progresses the Business Processes and processes business event data.

➤ ODB

This is data related to IT operational resources and business objective KPIs.

All the sources of data are used to show the overall health of your Business Process, and include the Process Instance and Violation pages, which present the details of the statistics collected by the BIE and the Monitor Engine.

The Process Instance and Violation pages are delivered from the BPI Server and comprise a set of HTML pages, which are created using Java Server Pages (JSP). JSP is a simple way to create dynamic Web pages, which are both platform independent and server independent. A Servlet Engine is used to manage these pages; this Servlet Engine is Tomcat and is installed as a component of the BPI Server.

The following diagram shows how the Process Instance and Violation Pages are managed within BSM.
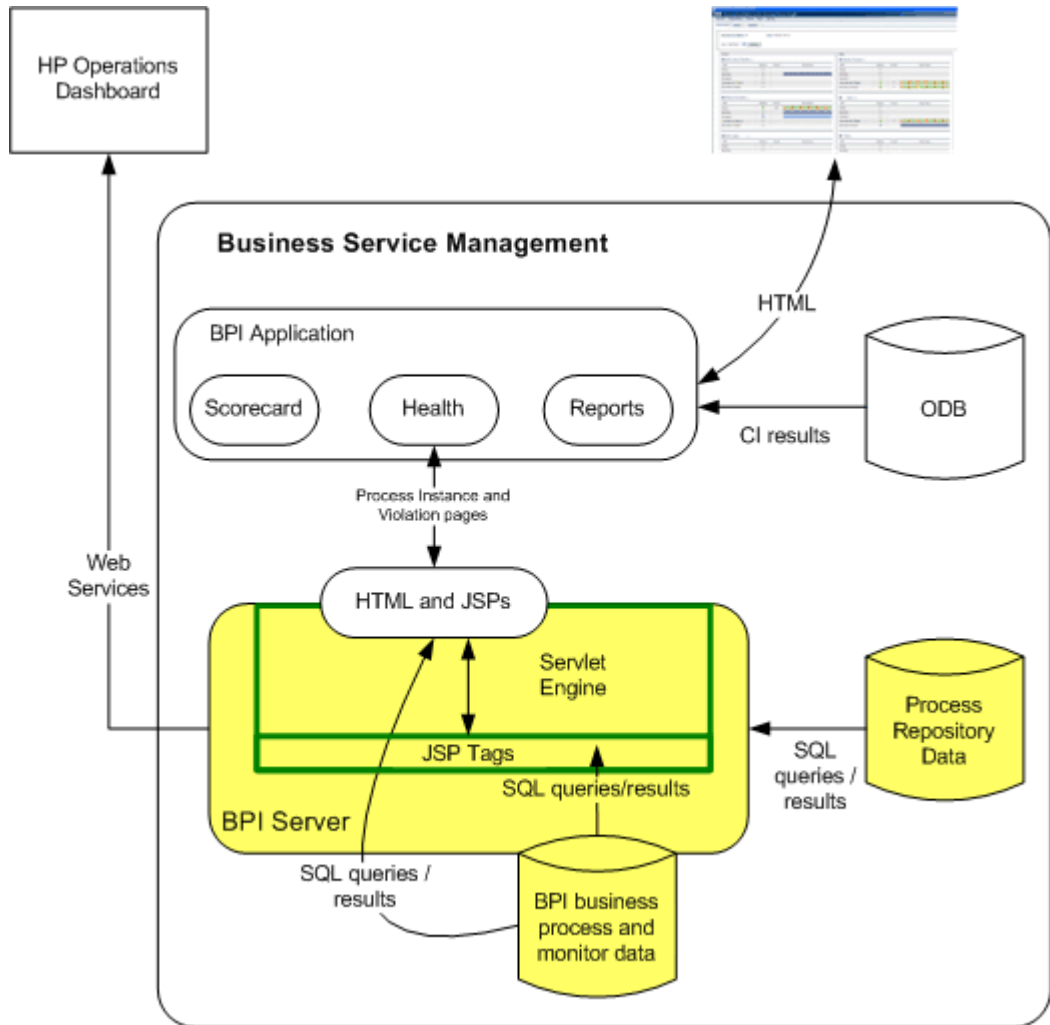
Business Process Insight Application Features

BPI provides you with a number of different options that you can use to graphically monitor and track your Business Processes and view the status of your Business Process Monitors and Monitor instance Thresholds.

Note: If you do not have a BPI license, these options are not available to you from BSM.

The BPI Application includes:

➤ Scorecard

This enables you to view all the deployed processes according to the current values of the Backlog, Duration and Value KPIs. The process that is least meeting its KPIs is shown first in the list and the process that is best meeting its KPIs is shown last in the list.

➤ Health

The Health page provides a number of options:

➤ Process Diagram

This is the same diagram as is shown on the **Application** > **Service Health** page. The Process Diagram provides a graphical image of the selected Business Process and enables you to view Value and Instance Counts for the each Activity in the process; the diagram also shows IT operation resource status for each Activity in the process.

➤ Summary

This main purpose of the Summary page is to show an overall (summary) view of the health of your Business Process, from the business and the operational standpoint.

Specifically the Summary displays Business Process or Activity summary data in tabular form, according to the context of the selection; for example, numbers of Healthy, At Risk or Blocked instances for the Business Process or Activity are shown.

You can navigate to pages which show more detail of individual process instances if required, including:

➤ Instances

   Displays Business Process or Activity instance lists in tabular form, according to the context of the selection. The tables shows information such as Value, Process or Activity Start and End Times.

   You can navigate to show more detail of individual Process and Activity instances if required.

➤ Instance Violations

   Displays lists of Monitor or instance Threshold violations for the selected Process or Activity, in tabular form.

   You can navigate to show more detail of individual Process instances if required.

➤ Business Health

   A series of dials showing the status of all the KPIs objectives defined for the selected Business Process or Business Process Activity.

➤ Transactions

   The information on this option is available only when you have configured the BPI Server integration with TransactionVision.

   The option displays the list and status of Transaction CIs and KPIs, which are linked to the selected Business Process or Activity, in tabular form.

   From this page, you can navigate directly to the BPI report, Business Process Over Time, to show a report based in the transaction CIs or to a selected Transaction. If you select a specific Transaction, you can navigate to the Aggregated Topology View within TransactionVision to obtain more details of the Transaction.

➤ Reports

Three reports are provided for the BPI Application:

➤ Business Process Summary

➤ BPI Monitors Over Time

➤ Business Process Activity Summary

The BPI Application is described in detail in *Using Business Process Insight*.

BPI Administration

Managing the BPI components can be divided into two logical areas:

➤ BPI Server Administration Console, which is accessed on the BPI Server system.

You use the Administration Console to manage the parameters of the BPI Server; for example, configuring the parameters required for the BPI connection to the BSM Gateway Server and starting and stopping the BPI Server Components.

For details on how to use the BPI Server Administration Console to manage the Server components, see Chapter 8, "Component Administration."

➤ Administration option from the Business Service Management user interface:

Admin > **Business Process Insight**

The BPI Administration option provides the ability to:

➤ model Business Processes using the BPI Modeler.

➤ link your Business Process Activities to IT operational resources, which are represented as CIs within the ODB.

➤ set KPI objectives on your Business Process and Business Process Activities in order to monitor the progress of your Business Processes against the defined objectives.

➤ configure your BPI Server to send email notifications relating to your process instances through the Notification Server. These email notifications contain information relating to the health of your Business Process instances.

For details on the BPI Administration functions, see *Using Business Process Insight*.

Database Schemas

This section describes the two database schemas used by the BPI Server. One is the BPI Server schema and the other is the Process Repository schema.

The information in these schemas is used by the BPI Server components for monitoring and progressing process instances, and storing the statistics relating to the ongoing status of the deployed Business Processes. It is also used as a source of information for the BPI Application to report on the health of your Business Processes and includes process impact data, instance threshold violation data and Business Process monitor data.

The Process Repository Schema and the BPI Server schema communicate through a deployer component, which is a component of the BIE. This means that when a change is made within the Process Repository and the Process Repository schema, the Process Repository notifies the deployer component within the BIE and the BPI Server schema is then also updated.

The following diagram shows the relationship between the database schemas and the BPI components. For details of BPI's database requirements, see "BPI Instance Database Requirements" on page 72.
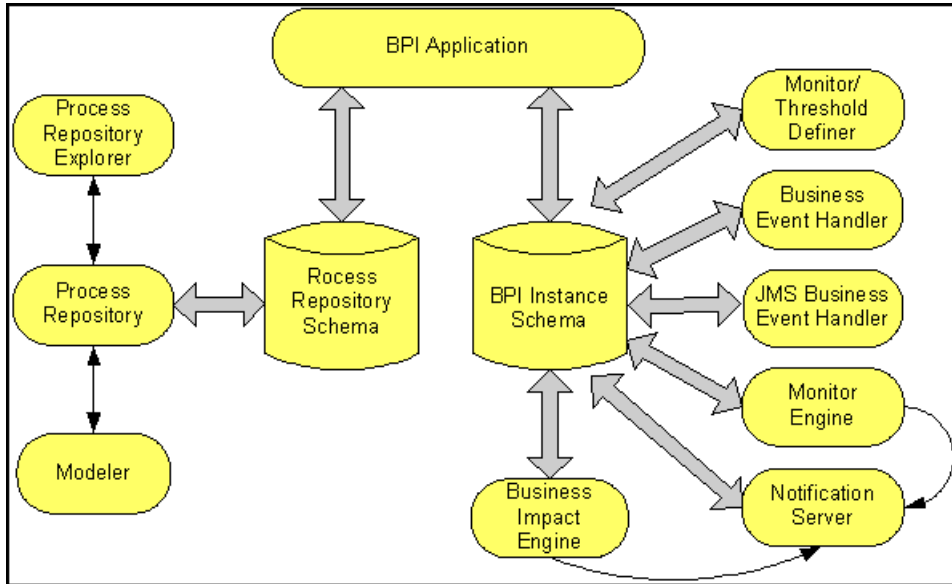
The BPI components' use of these schemas is described in the following sections:

➤ "BPI Server Schema" on page 42

➤ "Process Repository Schema" on page 44

BPI Server Schema

The BPI Server schema holds Monitor definition data, notification definition data and all the runtime statistics gathered as your Business Processes are progressed as described in the following sections:

➤ "Business Process Monitor Data" on page 43

➤ "Monitor Instance Threshold Definitions" on page 43

➤ "Event Hospital Data" on page 44

➤ "Notification Data" on page 44

Business Process Monitor Data

The BPI Instance database holds the data collected through the monitor tables as a result of the progress of Business Processes and the status of the processes relative to these defined monitors. Business Process monitors are evaluated as part of the process data as Business Process Activities are progressed by the Business Impact Engine. A Business Process monitor database table is populated using database triggers from the Process data table. The Monitor Engine then processes the data from this monitors table and uses the results to calculate statistics and populate the remaining monitors tables. The results of these calculations are then presented to you through the BPI Application.

You can also use reporting applications, such as HP Performance Insight, to access the monitor and statistical information in the database tables and generate reports from the data that is collected.

You can read more about the Monitor Definer and the monitor data in the *Business Process Insight Integration Training Guide - Defining Business Process Monitors*.

Monitor Instance Threshold Definitions

You can optionally define instance thresholds for your Business Process monitors. If you do this you might also want to be notified when these instance thresholds fall outside acceptable values. BPI enables you to create monitor instance threshold definitions for each Business Process monitor so you can then be notified when these monitor instance threshold values are violated.

Monitor instance threshold violations are shown:

➤ using the Business Process Insight Application

➤ as notifications using your email system, or as HP Operations notifications, or as scripts.

You can read more about monitor instance thresholds and creating them using the Monitor Definer in the *Business Process Insight Integration Training Guide - Defining Business Monitors*.

Event Hospital Data

Details of the events that are rejected by the Business Impact Engine because it does not recognize them, or they are out of sequence are written to the Event Hospital table in the BPI Instance database. From here they can be accessed, modified and resubmitted to the Business Impact Engine at a later date, if required.

Notification Data

These are the database tables relating to the information about the Notification Server subscriptions.

In addition to data relating to subscriptions, these tables contain data for the Notification Server retry mechanism.

Process Repository Schema

The Process Repository schema holds all the design-time information entered using the Modeler. Specifically, the Repository schema maintains the state of the Engine database objects (or definitions), for example, the business Data definitions and Business Process definitions, within the BPI system. These definitions comprise status information required by the Business Impact Engine for:

➤ Processes

➤ Data

➤ Events

These Process and Data components, plus the Event information, are all defined through the BPI Modeler. The data required to populate the Event definitions is obtained from your business applications through the Business Event Handlers. The data in the Process Repository can be viewed using the Repository Explorer.

The following are examples of the information stored in the Process Repository schema:

➤ Processes definition data

➤ Process Activity data

➤ Data definition data

➤ Event definition data

➤ progression rules

BPI Server

The BPI Server is a set of BPI components that receive and process business events and IT operational resource events and evaluate their impact on deployed Business Processes.

When installed, the BPI Server provides the components that you can use to monitor real-time information relating to the health of the Business Processes you have defined and deployed within Business Service Management.

The BPI Server is responsible for:

➤ Maintaining the Process, Data, Event and Business Process monitor definitions through the BPI database schemas.

➤ Enabling management for the Process Repository data, including the ability to browse and print the content of the Repository using the Repository Explorer.

➤ Monitoring business events, through adapters, in order to maintain the process context for the business.

➤ Defining Business Process monitors and monitor instance thresholds.

➤ Monitoring monitor instance thresholds and reporting on the instance threshold violations.

➤ Validating the status of processes according to progression rules that you have specified.

➤ Sending email notifications for process impact violations that you have subscribed to.

➤ Sending email notifications for monitor instance threshold violations that have been subscribed to.

The components of the BPI Server that provide these functions are the:

➤ Business Impact Engine

➤ Monitor Engine

➤ Business Process Monitor Definer

➤ Repository Explorer

➤ Notification Server

➤ Business Event Handler

➤ JMS Business Event Handler

The individual BPI Server components are shown in the following diagram and described, in more detail, in the following sections.
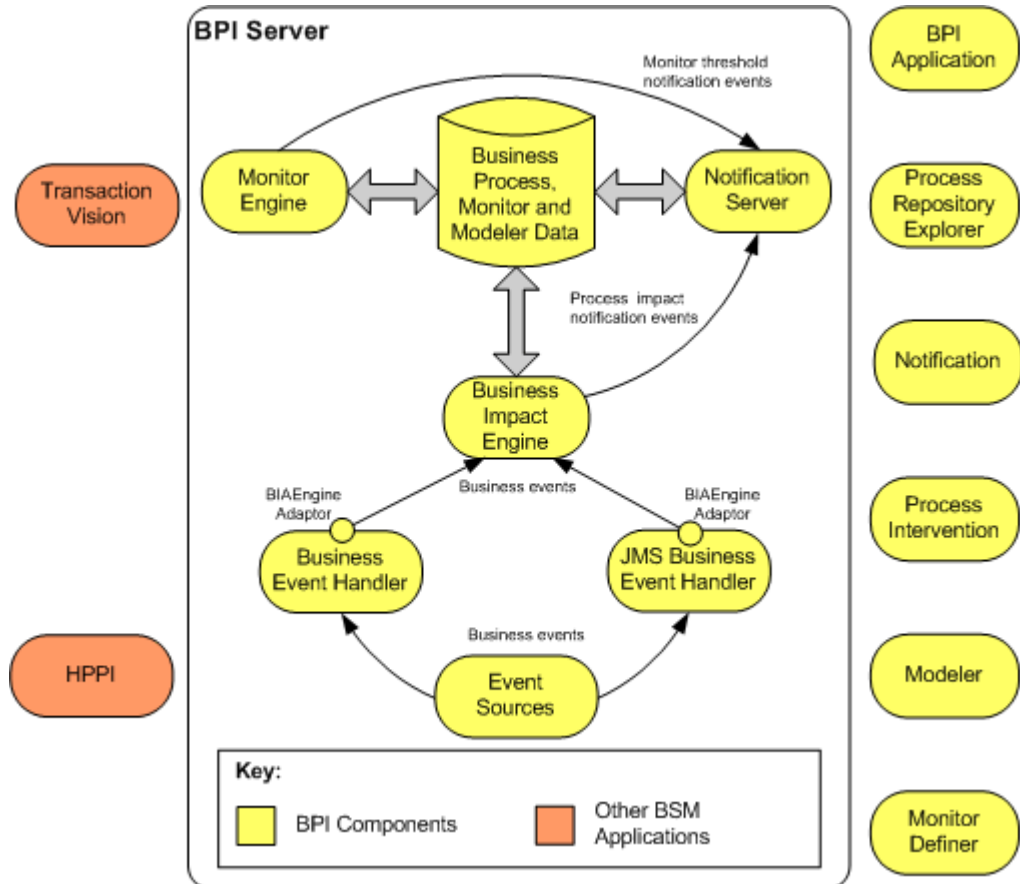
Business Impact Engine

The purpose of the Business Impact Engine is to process operational and business events and derive business impact information from these events. Using the information defined in Business Process definitions, the Business Impact Engine raises process impact violations when an underlying IT operational resource, which is relied on by the process, falls outside defined instance thresholds.

Note: Business Monitor Instance Threshold violations are managed by the Monitor Engine, which is described in the *<CiteTitle>Using Business Process Insight Guide*.

The following diagram shows the Business Impact Engine architecture.
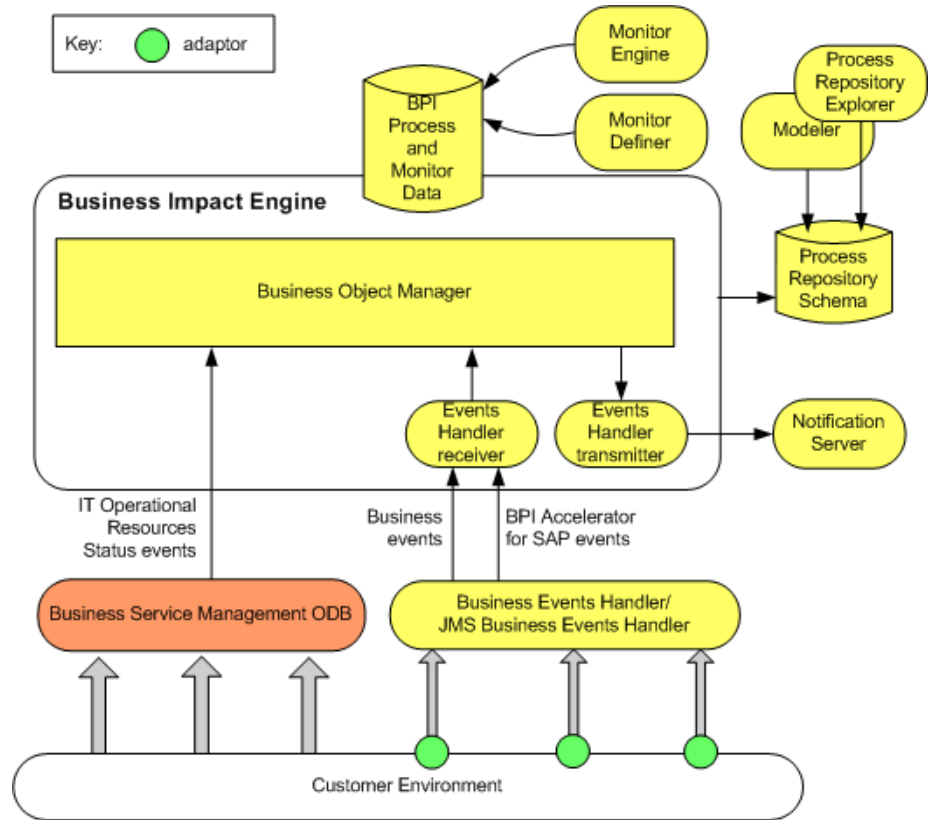
When an instance threshold violation is identified, the Business Impact Engine sends an impact violation to the Notification Server and writes data to the database, which can then be propagated to the BPI Application Health pages. This enables you either to be notified of the status of your Business Processes, or monitor the status of your processes according to your preferences and requirements.

The Business Impact Engine is responsible for:

➤ Managing the deployment and undeployment of Process and Data definitions.

➤ Receiving business events and IT operational resource events and handling them according to the deployed definitions, including updating the database as required.

➤ Generating notifications when necessary.

➤ Managing housekeeping tasks for the database through the instance cleaners.

You can also configure a business event queuing mechanism for the Business Impact Engine. In this case, all new business events are placed on a queue by both the Business Event Handlers. This means that the Business Event Handlers do not need to wait for the Business Impact Engine to process the business event and can immediately return to monitoring the event source and process the next incoming business event.

If you choose not to use a queue for business events, the Business Event Handlers wait until the business event has been processed by the Business Impact Engine. This includes waiting for all the necessary information from the event to be committed to the database. If you are not using a queue for business events, there is no multithreading capability and therefore there is no benefit of adding more Event Receivers.

The Event Transmitter accepts violations from the Business Impact Engine, converts them into an RMI message and sends them to the Notification Server.

Notification Server

This component of the BPI Server is responsible for notifying you of the process-impact, out-of-sequence and monitor-threshold violations that you have configured through the Notification pages. The Notification pages are described in the *<CiteTitle>Using Business Process Insight Guide*.

Notifications can be sent either through an SMTP server to an email client, or as an HPOM message to an HPOM Server, or as a script (for example SMS). There is a retry mechanism for both types of notifications if either the SMTP server or HP Operations Manager is not available for any reason. Notifications are queued and retried after a configurable interval.

Business violations are created by the Notification Server as a result of receiving:

➤ process-impact and out-of-sequence notification events from the business object manager through the Business Event Handler Transmitter.

➤ monitor instance threshold violation notification events from the Monitor Engine.

The Notification Server sends these events as business violations to you, based on a set of filters. These are filters that you create and that specify the impact events for which you want notification. You can filter email notifications based on event name, event type and the name of the Business Process.

In addition to sending violations when an event is received, the Notification Server can also run any script (.bat file) that you have created. Creating scripts is described in *Using Business Process Insight*.

Business Event Handler

The Business Event Handler provides the Business Impact Engine with normalized business events. Normalized events are events that have a standard (known) set of data attributes that enable it to be recognized by the Business Impact Engine. These events indicate the status changes in the underlying business applications.

Business events provide the data required to progress the process and also assess the business impact of a business failure; for example, the value of order or details of a new order that has been created.

The Business Event Handler is built using the **openadaptor** adaptor framework. **openadaptor** is an open source adaptor integration toolkit, written in Java, and provides a framework for adaptors. It is a message-based integration toolkit, based on the concept of adaptors providing one-way connections between one or more origins and one or more destinations, including publish/subscribe connections. It can be used to connect systems to systems or systems to middleware.

The following diagram shows how the Business Event Handler works.
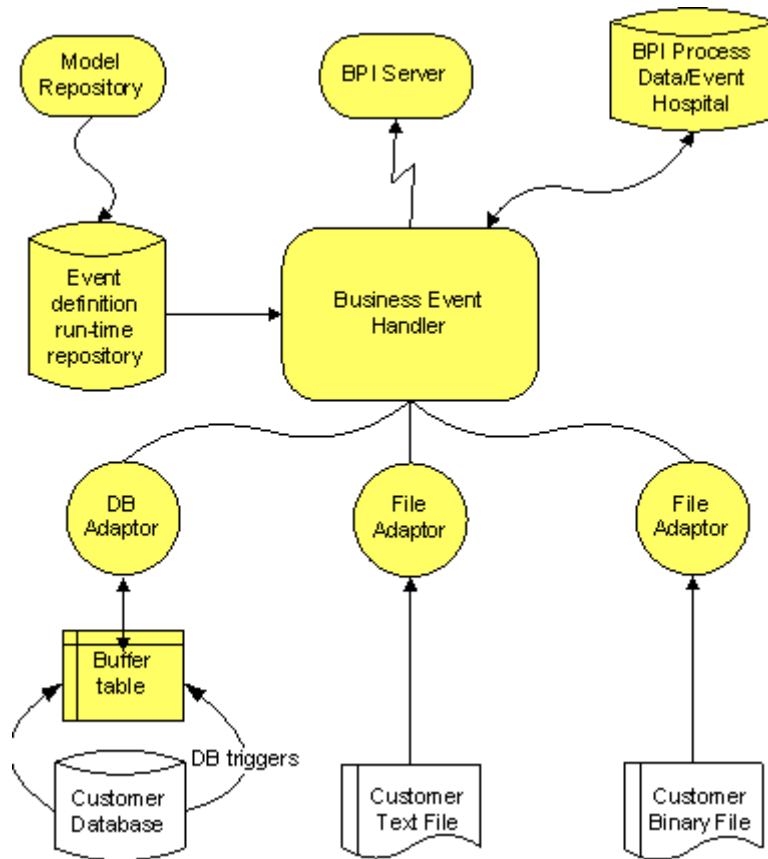
In the case of BPI, **openadaptor** is the framework that provides the following BPI event adaptors:

➤ File adaptors

The flat file adaptor enables you to access information in a binary or ASCII file.

➤ Database adaptors

The database adaptor allows you to access information in any JDBC-compliant database. The adaptor uses a polling mechanism to access buffer tables created by database triggers.

➤ Sockets adaptor

This is a multi-threaded socket adaptor that by default connects to BPI using a specific port.

The Business Event Handler accepts business events received through the adaptors and converts these events into a format that can be accepted by the Business Impact Engine. The Business Event Handler also manages the business events that cannot be immediately delivered to the Business Impact Engine; for example, when the Business Impact Engine has stopped for some reason.

The BPI Process Repository holds the design-time Event definitions that you create using the BPI Modeler. When these Event definitions are deployed, the Business Events Manager uses the definitions to convert the incoming events into the format required, and to make sure that the correct information is added to the event, so it can be used by the Business Impact Engine.

You can add additional data to events that are received from the individual applications. This allows you to monitor one key application and then assemble the complete set of data required for the event from other applications, rather than monitor all the applications waiting for all the data to become available.

If the Business Impact Engine cannot receive an event from the Business Event Handler for any reason, the Business Event Handler rolls back the event transaction so that it can be retried at a later time. Refer to the *Business Process Insight Integration Training Guide - Business Events* for more details of event adaptors.

In specific cases, the Business Event Handler places a business event in the Event Hospital. It does this if it receives out-of-sequence events or events that contain errors; an event error might be an event that had missing or corrupt data.

Event Propagation

This section gives a brief overview of how events that are received through the Business Event Handler are used to progress Business Processes. The *Business Process Insight Integration Training Guide - Business Events* provides more detail about the Business Event Handler and how to configure it to receive specific business events.

Each time an underlying business system is updated, the Business Event Handler is notified of the change, through a business event, and adds the event details onto the event as defined in the BPI Modeler.

The Business Event Handler includes an openadaptor-to-BPI adaptor (BIA EngineAdaptor), and business events for BPI are sent to this adaptor. How the adaptor detects the change, depends on how it is configured to communicate with the underlying systems. The adaptor might receive events from a messaging bus, or it could be notified through a database trigger; see the *Business Process Insight Integration Training Guide - Business Events* for more information about building and configuring the file and database adaptors.

Where the event mapping is successful, the BPI event details are updated and the event is then sent to the Business Impact Engine. If there is no BPI event defined that matches the data from the event, the event is moved to the Event Hospital by the adaptor. The business event then must be manually discharged from the Event Hospital so it can be resubmitted to the Business Impact Engine. In the case of an out-of-sequence event, the Business Event Handler places the event in the Event Hospital and marks it for automatic discharge. For other cases where the event cannot be delivered to the Business Impact Engine, the Business Event Handler rolls back the event transaction and returns the event to the source adaptor.

When the Business Impact Engine receives a business event, it determines if the event contains details that apply to any of the data instances that currently exist within the Engine. If there is a Data definition subscribed to the Event, and there are associated data instances, the Business Impact Engine updates the data details using the information from the event and the actions specified through the BPI Modeler for the appropriate Event and Data definitions. Alternatively, a Business Process can be progressed directly from the properties of an Event, or the arrival of a recognized Event type.

Any changes to the data cause the start and complete conditions for all relevant processes to be re-evaluated, and where appropriate the processes are progressed.

The Business Impact Engine also determines if there is an impact that needs to be reported for the Business Processes that use this data, for example, those processes with out-of-sequence Activities.

At the same time, the BPI Application is able to poll the BPI Process Data and refreshing its display to show the impacts of the changes.

You need to configure adaptors for each data source (application) from which you want to receive business events. These adaptors provide data for the process, or processes, that you are monitoring. There are different ways that you can create and configure adaptors, which are described in the *Business Process Insight Integration Training Guide - Business Events.* You do not need to understand how to create adaptors at this stage, although you do need to understand them when you come to develop your solution.

JMS Business Event Handler

The JMS Business Event Handler is an extension to the Business Events Handler described in section "Business Event Handler" on page 52.

The JMS Business Event Handler provides the Business Impact Engine with business events, in the form of XML documents, from a JMS-enabled system.

As described above the Business Event Handler is built using the openadaptor adaptor framework. An additional JMS Source has been added to the Business Event Handler in order to provide the additional capability of receiving business events from a JMS-enabled system.

Note: BPI supports the JMS message type javax.jms.TextMessage, where the content of the text is an XML file in a particular format. This means you must be integrating with an application that sends and receives this message type if you want to use the JMS Business Event Handler. One example of such an application is TransactionVision.

For details on configuring the JMS Business Event Handler, see "JMS Business Event Handler" on page 56.

Following is a diagram of the JMS Business Event Handler.
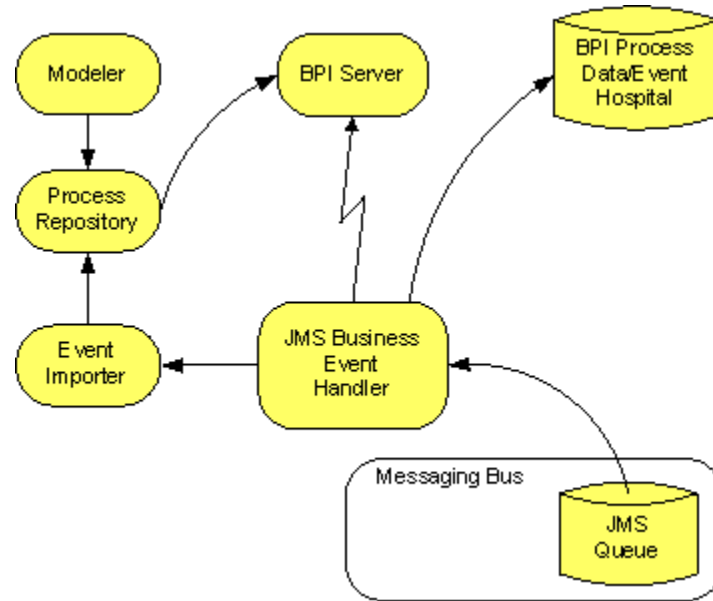
page 57 shows how Event definitions are imported from the JMS-enabled system through an Event Importer into the Process Repository. You can then access these Event Definitions from the Modeler, and add Data Definitions to reference these events. When the process, including the Event and Data definitions are deployed, the Business Events Manager uses the data to convert the incoming events into the format required. The Business Events Manager also makes sure that the correct information is added to the event, so it can be used by the Business Impact Engine.

If, for any reason, the Business Impact Engine cannot receive an event from the JMS Business Event Handler, the JMS Business Event Handler rolls back the event transaction in exactly the same way as the Business Events Handler.

Process Repository

The Process Repository is the server-based components that manages Process, Event and Data definitions as they are entered through the Modeler and then stores the results in the Process Repository Schema.

The Process Repository, stores the model definitions in the Process Repository schema and ensures that the details that you enter for the models are consistent and deployable. If the Process Repository finds inconsistencies, it reports them using a ToDoList within the Modeler.

If you have not installed the BPI Server, there is a lightweight Process Repository installed on the Gateway Server. This Process Repository manages the Business Process models that you create using the Modeler that link to IT operational resources.

When you subsequently install the BPI Server, the existing Process Repository is replaced by a revised version, which is installed on the same system as the BPI Server. The revised Process Repository manages the Process Definitions for processes that you create either with instance tracking enabled or disabled. This is in addition to the IT operational resource links that you have already created.

The deployer component of the Process Repository deploys the different definitions to their destinations as follows:

➤ Process definitions, including any imported BPEL definitions, are deployed to the Business Impact Engine as compiled Java code.

➤ Data definitions are deployed to the Business Impact Engine as compiled Java code.

➤ Event definitions are deployed to the Business Event Handler to be mapped to events from underlying business applications.

The Repository Explorer is a Web-based interface that you can use to view and manage the data stored in the Process Repository; see the <CiteTitle>Using Business Process Insight Guide for more information about the features of the Repository Explorer.

The BPI Modeler and its features are described in more detail in the <CiteTitle>Using Business Process Insight Guide*s*.

BPI Integration Points

BPI can integrate with the following:

➤ TransactionVision within the Transaction Management application; see "Connecting with HP TransactionVision" on page 60

➤ SAP applications; see "Integrating with SAP Applications" on page 60

The following are the integration points within BPI, some of which are enablers for the application integrations described later:

➤ Import service definitions from other HP BTO Software components.

This is achieved by linking IT operational resource CIs to your Business Process Activities using the BPI Modeler.

➤ Configure data sources through the Business Event Handler and import business events from underlying business applications.

Using the Business Event Handler, and openadaptor technology, you can import business events and use these events to progress your Business Processes. You can also use these business events to report the business impact where processes are not progressing as they should be. For details see "Component Configuration - JMS Business Event Handler" on page 191.

➤ Configure JMS data sources through the JMS Business Event Handler in order to import JMS-based Business Events. You connect BPI with TransactionVision using the JMS Business Event Handler. For details see "Component Configuration - JMS Business Event Handler" on page 191

➤ Configure notifications through the Notification Server by:

  ➤ Configuring email notifications that can be sent when a notification event is received by the Notification Server.

  ➤ Creating scripts that are executed when a notification event is received by the Notification Server.

  This enables you to integrate the data from notifications that are generated by BPI into your own spreadsheet applications. Alternatively, you can create a script that generates an SMS notification message.

➤ Configuring the Notification Server as described in "Component Configuration - Notification Server" on page 143.

➤ Importing BPEL definitions into the BPI Process Repository using the BPI Modeler.

When you have imported the BPEL definitions, you can use them as the basis, or starting point for your process definitions. For more information on the options for importing BPEL definitions, see the *Business Process Insight Integration Training Guide - Importing BPEL*.

Connecting with HP TransactionVision

You can configure BPI to connect with TransactionVision, which is a product that delivers a transaction tracing solution to the Transaction Management application. This enables you to view the life cycle and associated data for specific business transactions within your organization.

Once connected, you can link your Business Process Activities to event definitions that have been created, or discovered, within TransactionVision. For details about the advantages of connecting with TransactionVision and how to configure BPI to connect with TransactionVision, see Chapter 11, "Connecting BPI With TransactionVision."

Integrating with SAP Applications

In addition to the integrations with HP BTO Software products, BPI provides an Accelerator for SAP applications. The BPI Accelerator for SAP applications is specifically designed to integrate with IDocs within SAP; its design is based on the openadaptor framework. The Accelerator uses the Business Event Handler to access the status information within an IDoc header and is a separately licensed product.

The following diagram shows the integration between the BPI Accelerator and SAP.
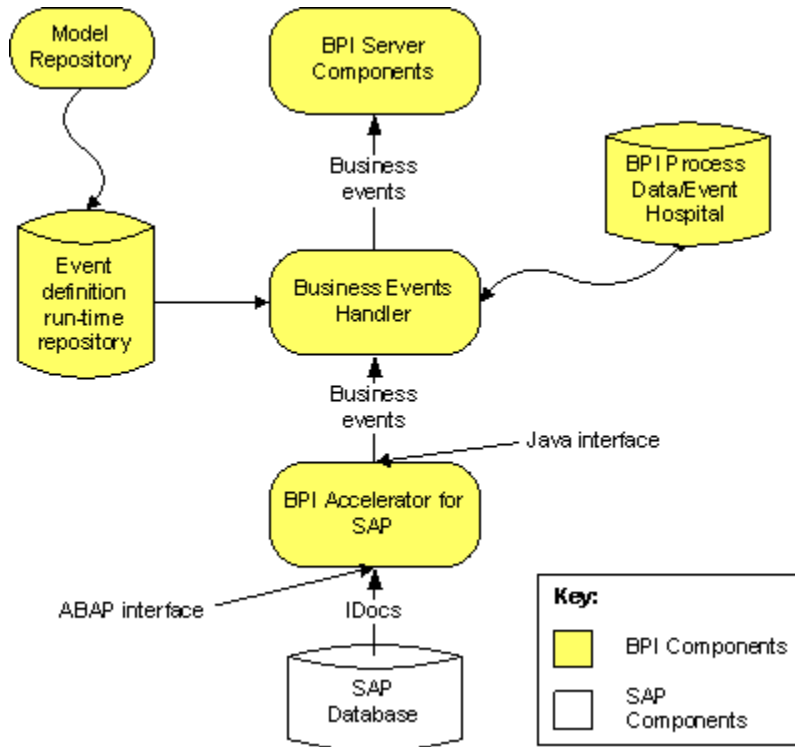
The *Business Process Insight Accelerator for SAP Guide* contains details of installing and setting up the Accelerator.

BPI Security

You can choose to implement access control for your BPI interfaces using Tomcat (Servlet Engine). When first installed, the BPI interfaces are configured to use the BSM lightweight single-signon mechanism for authentication. If you want to modify this and use Tomcat authentication mechanism refer to the Security options within the Administration Console as described in "Security" on page 200.

Part II

Installing the Business Process Insight Server

3

Before Starting the Installation

This chapter describes the tasks that you must complete in order to install the BPI Server. In addition, the requirements that BPI has on the database installation are described in section "BPI Instance Database Requirements" on page 72.

Before completing any of the tasks described in this chapter, read the *Business Service Management Readme* to check if there are additional requirements for your installation.

If you are re-installing the BPI components, follow the instructions in Chapter 5, "Re-installing BPI Server Components".

If you are upgrading from BPI version 8.0x, follow the instructions in Chapter 7, "Upgrading to BPI Version 9.00".

This chapter includes:

➤ BPI Server Installation Requirements on page 66

➤ Business Service Management on page 70

➤ TransactionVision on page 71

➤ BPI Instance Database Requirements on page 72

➤ What to Do Next on page 82

BPI Server Installation Requirements

The following sections describe the system, database and software requirements for a BPI Server Installation.

Note: The exact system requirements depend on all of the following:

➤ the number of processes you need to define

➤ the number of business events you intend to process

➤ the Business Process Monitors you define

➤ the load on your system

This section provides the minimum system requirements. You need to fine tune your system at regular intervals to take account of any increased load.

BPI Server System Requirements

The following table lists the minimum system requirements for the BPI Server installation.

Requirement	Value	How to Check
System Processor	Intel Pentium (at least 2Ghz)	Use: **Programs** > **Accessories** > **System Tools** > **System Information**
Windows version	➤ Windows 2008 Enterprise Server ➤ Windows 2003 Enterprise Server, plus SP1 ➤ Windows 2003 Enterprise Server, plus SP2 ➤ Windows 2003 64-bit	Use: **Programs** > **Accessories** > **System Tools** > **System Information**
Network	IPv4 or Dual-stack IPv4/IPv6	Run ipconfig. IPv4 addresses are of the form nnn.nnn.nnn.nnn where nnn is up to 3 digits. IPv6 addresses are of the form xxxx:xxxx:xxxx: xxxx:xxxx:xxxx:xxxx:xxxx where xxxx is up to 4 hexadecimal digits. **NOTE**: When referring to other servers on the IPv4 or Dual-stack IPv4/IPv6 network, you need to enter server names rather than IP addresses. Otherwise, the IP addresses may not function correctly in all areas of the product.

Requirement	Value	How to Check
Physical memory	1024MB	Use: **Programs > Accessories > System Tools > System Information**
Virtual memory	1536MB	Use: **Programs > Accessories > System Tools > System Information**

Disk Space Requirements

The following table describes the disk space requirements.

Free Disk Space	How to Check
1GB for static data, plus an allocation for dynamic data; see the note below.	For Windows: Use **My Computer** to check disk space. Click on the drive name where you intend to install BPI and check the available capacity.

Note: In addition to the disk space requirements for the BPI Server components, you also need to consider the requirements for a local database installation and take into account the disk space required for the database files and log files. 6GB is the initial recommendation for the dynamic data.

Software Requirements

The BPI Server installation requires the following software versions to be installed before the installation is started.

Software Component	Version
HP Business Service Management	HP Business Service Management 9.00 must be installed before you can install the BPI Server.
Databases[a]:	
➤ Microsoft SQL Server	➤ Microsoft SQL Server 2008 Enterprise Edition ➤ Microsoft SQL Server 2005 Enterprise You must configure Microsoft SQL Server to use static port numbers with BPI and not dynamic port numbers.
➤ Oracle Server	➤ Oracle 11g Enterprise Edition ➤ Oracle Database 10g Enterprise Edition Release 10.2.0.x ➤ Oracle Real Application Cluster (RAC) 11g (11.2.0.x) ➤ Oracle Real Application Cluster (RAC) 10g (10.2.0.x)

a. The database can be on the local system where you intend to install the BPI Server, or it can be on a remote system.

In addition to the above requirements, you might also need to:

➤ install HP SOA Policy Enforcer 3.x if you want to use HP SOA Policy Enforcer as a source of Business Events and to progress your BPI Business Processes.

➤ install HP Performance Insight, if you want to use Performance Insight to maintain historical reports for BPI using the process and monitor data in the BPI database.

➤ access an SMTP Mail Server if you want to be notified of BPI violations through email. This mail server is required for the Notification Server component of BPI. This component is responsible for sending business violations as email messages to the users who are configured to receive them.

Business Service Management

BPI requires that Business Service Management is installed. You cannot install the BPI Server unless you have already installed Business Service Management.

Note: BPI 9.00 is not compatible with earlier versions of Business Service Management (called Business Availability Center in previous versions) or TransactionVision. You must upgrade Business Service Management to 9.00.

If you intend to use a Load Balancer in your BSM configuration, do not install the BPI Server on the same system as the Gateway Server. This is because the Load Balance does not accept connections from Gateway Server and, as a result, if you install the BPI Server on the Gateway Server, it is unable to synchronize, send data samples or access infrastructure settings.

If you intend to use a Load Balancer, you must install the BPI Server on a different system.

Caution: If you have previously installed Business Service Management and BPI and subsequently removed the BSM installation, including all the BSM data, you must also remove the BPI data before attempting to reconfigure BPI to connect to BSM.

If you do not, some of the BPI data will have no parent CI data within the BSM database and there could be inconsistencies in your Business Process models.

TransactionVision

BPI can connect with TransactionVision, which is a product that delivers a transaction tracing solution to the Transaction Management application. With TransactionVision you can view the life cycle and associated data for specific business transactions within your organization. TransactionVision discovers transaction flows within your business and analyses the steps within these flows. The Transaction Management application then displays reports based on the results of the analysis. For more details, see Chapter 11, "Connecting BPI With TransactionVision."

Note: BPI 9.00 is not compatible with earlier versions of the Transaction Management application (called TransactionVision in previous versions). You must upgrade the Transaction Management application to 9.00.

BPI Instance Database Requirements

The BPI Server requires a database to store data for process instances, process monitors and for administration. This database information needs to be accessible to the BPI Server and can be configured within either a Microsoft SQL Server or Oracle Server.

This database is not the same as the BPI Process Repository database (BPI Schema) that was created as part of the BSM installation and which stores information relating to your Business Process Definitions and related CIs.

The following sections describe information that you need to be aware of, which relate to the BPI Server's use of the database. You must read these sections and use the information to tailor your database configuration where appropriate. Be aware that some of the tasks described must be completed before you start the BPI installation; this might impact whether or not you can use an existing database, or whether you need to install a database specifically for your BPI Server.

The following sections describe what you need to consider for the BPI Server's use of the database and includes:

➤ Configuration requirements for the Microsoft SQL Server database, which are specific to the BPI Server; see "Microsoft SQL Server Database Configuration" on page 73.

➤ Configuration requirements for the Oracle Server database, which are specific to the BPI Server; see "Oracle Database Configuration" on page 79.

It is strongly recommended that you allow the BPI Server installation to create a new database user (and a new database in the case of Microsoft SQL Server) rather than using an existing User or Database.

There are a number of reasons for this:

➤ it removes the possibility of name clashes for tables that the BPI Server needs to create in the database.

➤ it enables you to backup and recover the BPI Instance database tables in isolation from other data in the database, if required.

➤ you can copy (archive) the BPI Server data independently of other data in the database.

➤ if there is a problem with the installation, you can clean out the Microsoft SQL Server Database or Oracle User data and restart the step without impacting data for other database applications.

In the case of Microsoft SQL Server, the *Database* should be reserved for the exclusive use of the BPI Server. In the case of an Oracle Server, the *User* should be reserved for the exclusive use of the BPI Server.

The BPI Server installation assumes that the Microsoft SQL Server Database and Oracle User are exclusive to BSM and BPI, and therefore the installer can safely delete and recreate database schema objects, including tables, views, stored procedures. If this is not the case, there is a risk that existing database objects could be overwritten by BPI and their contents lost. Appendix A, "Database Tables Defined for BPI" lists all the instance database tables used by BPI. If your database is not for the exclusive use of BSM and the BPI Server, you need to check that the tables, views and stored procedure names are unique and not duplicates of existing ones.

Microsoft SQL Server Database Configuration

The following sections describe the BPI requirements on the Microsoft SQL Server database that you need to be aware of before starting the BPI installation. These requirements need to be addressed before installing any of the BPI components that use the BPI instance database.

If you plan to use Oracle as your database, refer to "Oracle Database Configuration" on page 79.

Static Port Number Required for Use with BPI

When installing and configuring Microsoft SQL Server for use with BPI, you must configure it to use a static and not a TCP dynamic port. If you configure a dynamic port for SQL Server, it can result in a different port number being assigned each time SQL Server is started. You cannot change the port number used by BPI following the initial installation, therefore, dynamic port numbers cannot be used.

Configure the port number used by Microsoft SQL Server using the SQL Server Configuration Manager as shown in the following figure.
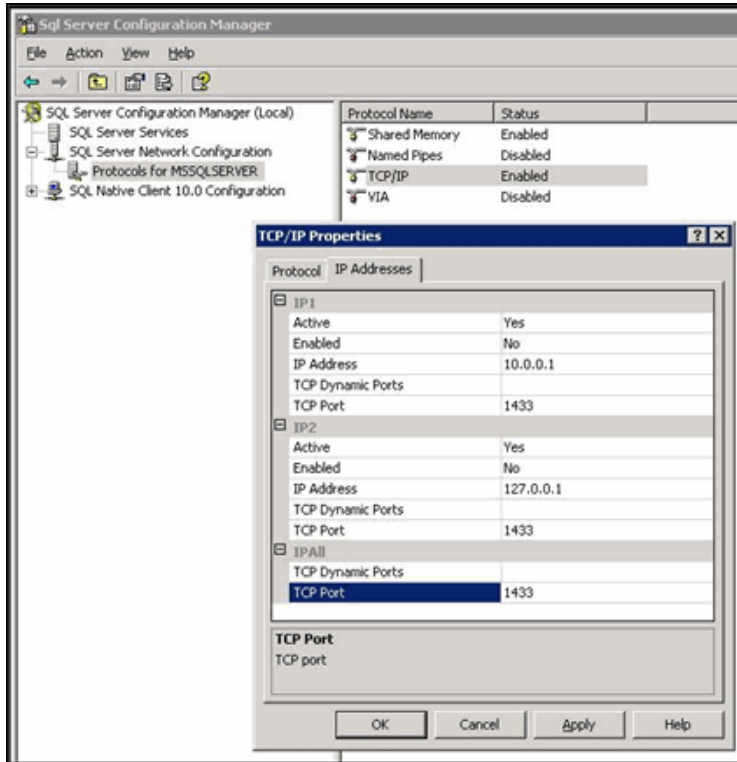
Restricting the Size of the Transaction Log File

You might also need to limit the size of the database transaction log files as these can get very large. By default, the transaction log files are not restricted to a particular size. To modify the settings for the transaction log file, select the Transaction Log tab from the database Properties option.

Note: The size of the transaction log files need to be large enough to accept the largest possible transaction that BPI can generate, in order to ensure that you can recover data following a failure. For details on considerations for BPI and how these impact the size of the transaction log, see Chapter 13, "Backup and Recovery."

Database Case Sensitivity

If you intend to allow the BPI installer to create the BPI database instance it requires, you do not need to read this section. When the BPI installer creates its database instance, it creates an instance with the characteristics that it requires, including the case sensitivity setting.

If you intend to use an existing instance of the SQL Server database and then provide the details of this database instance to the BPI installer, you need to be aware that the SQL Server database can be set up to be case insensitive or case sensitive. BPI does not operate with a database that is case sensitive.

If you are creating the database instance for BPI data, you must create the instance with the Collation Name set to Case Insensitive.

If your database instance has already been created, you can find out whether it is case insensitive or case sensitive by checking the properties of the database instance from within SQL Server Enterprise Manager.

hidden

The Collation Name property contains a string, which is one way of identifying the case sensitivity of the database instance, as follows:

➤ _CI, if the database instance is configured to be case insensitive.

➤ _CS, if the database instance is configured to be case sensitive.

➤ _BIN, if the database instance is configured for a binary collation, which is always case sensitive.

If there is no Collation value set for the database instance, the database instance inherits its Collation from the Database Server settings. If the Database Server Collation is set to be case insensitive, the BPI installation can proceed. If the Database Server Collation is set to be case sensitive, the BPI installation fails and reports the error shown above. The inherited collation is shown through the Collation Name parameter as described above.

Database Accent Sensitivity

If you configure the database to be accent insensitive, you must advise the BPI users not to create definitions that differ only by accent; for example:

Property: Propa

Property: Propä

If you do not, and users create processes that differ only by accent, the Process definition fails to deploy as the database does not consider the two data items to be unique.

Database Authentication Mode

When using Microsoft SQL Server with BPI, you need to ensure that when prompted (during the SQL Server installation) for the Authentication Mode to use within SQL Server that you select the option for SQL Server and Windows Authentication mode, and not Windows Authentication Mode. This is because the JDBC interface needs the SQL Server Authentication Mode activated.

You can modify the authentication mode from Windows Authentication to be SQL Server and Windows Authentication without affecting other applications that use the database.

Check the authentication mode for your installation of Microsoft SQL Server as follows:

1 Start SQL Server Management Studio and log on using your Windows account details. The Windows account that you use must have SQL Server Administrator rights.

2 Select the appropriate server, and then click **Properties**. (Make sure you select the server and not the database from the navigation pane.)

3 From the Server Properties dialog, select the Security page. There is a Server authentication option where you can check which authentication mode is being used.

4 If SQL Server and Windows Authentication mode is not set, you can change the selection in this dialog and click OK. If your configuration is correct, you can exit from the Management Studio without making any changes.

5 If you do change the configuration, click OK when prompted to restart SQL Server.

SQL Server Database User Properties

The BPI installation step creates a database user, with all the correct properties; however, if you have an existing database user that you want to use for the BPI database, you can do this. You are strongly advised to let the installation step create a database user that is specific to BPI.

If you allow the installation step to create the database User, you do not need to take any further action as the database User is created with all the required roles. If you intend to specify an already defined User for BPI within the installation step, you need to ensure that the following are completed before starting the BPI installation:

➤ the existing database User has access to the database where the BPI tables are created.

➤ the database used by BPI should ideally be the default database for machine; this is not mandatory, but preferred.

➤ the following database roles are permitted for the BPI database User:

  ➤ public

  ➤ db_ddladmin

  ➤ db_datareader

  ➤ db_datawriter

➤ the following database roles are not permitted for the BPI database User:

  ➤ db_denydatareader

  ➤ db_denydatawriter

When you run the BPI installation step and specify the database User, the installation step then configures the User to be a member of the following database roles:

  ➤ public

  ➤ db_ddladmin

  ➤ db_datareader

  ➤ db_datawriter

If you create your own database user, you are advised to check for the following log file after the BPI installation is complete to make sure that the SQL Server database tables have all been created correctly:

bpi-install-dir**HP_Business_Process_Insight_DbSetupLog.log**

This file is created when there are errors to report after BPI has made a successful connection to the database. If there are errors reported in the file, you need to investigate them and make any necessary corrections. You must then uninstall and reinstall BPI to reconfigure the database options.

If this log file does not exist, there are no errors.

Oracle Database Configuration

The following sections describe the BPI requirements on the Oracle database that you need to be aware of before starting the BPI installation. These requirements need to be understood before installing any of the BPI components that use the BPI database.

➤ "Configuring an Oracle Database Server on an HP-UX System" on page 79

➤ "Character Set for Oracle Database" on page 80

➤ "Oracle Server Database User Permissions" on page 80

If you plan to use a Microsoft SQL Server database, refer to section "Microsoft SQL Server Database Configuration" on page 73.

Configuring an Oracle Database Server on an HP-UX System

If you intend to use an Oracle Database Server that is installed on an HP-UX system, you must create an Oracle User for BPI before starting the installation. If you do not, the BPI installation fails.

The Oracle User that you create must be configured to have both a default and a temporary tablespace. You are strongly advised to assign a default tablespace that is specific to BPI and not shared by other applications.

In addition, you need to make sure that an INDX tablespace is assigned to the Oracle User that you create for BPI. Depending on the version of Oracle that you are using, you might have to create this INDX tablespace.

There are also permission requirements for the Oracle User that you create; these are described in section "Oracle Server Database User Permissions" on page 80.

Character Set for Oracle Database

BPI does not have any specific requirements for a database character set; however, if you want to support multiple languages simultaneously, from the same data, you are recommended to use a Unicode standard for the encoding form in the Oracle Database. UTF8 is the recommendation for BPI.

Using other database character sets can result in characters being lost during conversion between the user input and how the data is stored in the database; For example, using a database character set of ISO Latin1, does not enable you to store Korean language characters in the database. Using a character set such as UTF8 does enable you to store Korean characters in the database.

In addition:

➤ the BPI database clients use JDBC; therefore using UTF8 as the database character set has performance benefits as it reduces the amount of conversion between character sets.

➤ the Oracle database drivers always convert between character sets using UTF8.

If your database currently uses a form of encoding other than UTF8, then you need to make yourself aware of the restrictions of the encoding method and any potential loss of characters during conversion.

Oracle Server Database User Permissions

The BPI installation step creates a database user, with the correct profile, however, if you have an existing database user that you want to use for the BPI database, you can specify this User during the BPI installation.

Where possible, you are advised to let the installation step create the database user, unless there is a strong reason for using your own defined database user.

If you allow the installation step to create the Oracle User, you do not need to take any action as the database User is created with all the required privileges.

If you intend to specify an already defined User for BPI within the installation step, you need to ensure that the following are completed before starting the BPI installation:

➤ User has QUOTA UNLIMITED defined on its default tablespace (this is the tablespace where the BPI data is created).

➤ User has the CREATE SESSION privilege granted.

When you run the BPI installation, the installation step then completes the following for the User that you specify:

➤ Defines QUOTA UNLIMITED for the INDX tablespace. BPI uses the INDX tablespace to hold indexes for some of the tables in the BPI User's default tablespace. This is done for reasons of performance.

➤ Grants the following privileges to the user:

　➤ CREATE SESSION

　➤ CREATE SEQUENCE

　➤ CREATE TABLE

　➤ CREATE VIEW

　➤ CREATE PROCEDURE

　➤ CREATE TRIGGER

　➤ CREATE ANY INDEX

➤ creates a role called OA_HOSPITAL_ACCESS_ROLE and grants itself this role.

If you create your own database user, you are advised to check for the following log file following the BPI installation to make sure that the Oracle database tables have all been created correctly:

bpi-install-dir**HP_Business_Process_Insight_DbSetupLog.log**

This file is created when there are errors to report after BPI has made a successful connection to the database. If there are errors reported in the file, you need to investigate them and make any necessary corrections. You must then uninstall and reinstall BPI to reconfigure the database options.

If this file does not exist, there are no errors with the connection.

What to Do Next

When you have completed all the checks and tasks listed in this chapter, continue at Chapter 4, "Installing the BPI Server" to install the BPI components.

4

Installing the BPI Server

This chapter describes how to install the BPI Server.

This chapter includes:

➤ Collecting Required Information for the Installer on page 83
➤ Installation Log Files on page 96
➤ Installing the BPI Server Components on page 97
➤ Creating the BPI Instance Database Manually on page 103
➤ Starting BPI Server Components on page 107
➤ BPI Infrastructure Settings on page 108
➤ Enabling Oracle Database Table and Index Partitioning on page 109
➤ Verifying Your Installation Using the Integrity Checker on page 110
➤ Installation Troubleshooting on page 113

Collecting Required Information for the Installer

Before starting the installation be sure to do the following:

➤ Check the *Business Service Management Release Notes* for any additional tasks that you need to complete before and after the installation. The Release Notes are located on the distribution media.

➤ Work through the required preinstallation checks listed in Chapter 3, "Before Starting the Installation".

➤ Check that the Database Server you plan to use for BPI is configured and running (Microsoft SQL Server or Oracle Enterprise Server). The database can be on the same machine as the BPI Server or on a different machine.

The BPI Server uses a database to store process instance and violation information. This is a different database than the Process Repository database, which you configure later in the installation process (see "BPI Process Repository Database Type Dialog" on page 101). The Process Repository database holds data relating to process definitions and their data. For more information, see "Component Configuration - Process Repository" on page 185.

The BPI Instance Database is created and configured specifically for the BPI Server, as part of the BPI Server installation — it is not previously installed as part of the main installation of the BSM servers. For required settings, see "BPI Instance Database Required Information" on page 85. For details on creating the BPI Instance Database, see "BPI Database Type" on page 99.

➤ Be sure you have the required information ready as described in the next section.

Information Required for the BPI Server Installation

The following sections describe the required information needed for the BPI Server Installation:

➤ "Name of the BPI Server Installation Directory" on page 85

➤ "BPI Instance Database Required Information" on page 85

➤ "Business Service Management Information" on page 92

➤ "Process Repository Database Information" on page 93

Name of the BPI Server Installation Directory

The BPI Server installation requires the name of the BPI Server installation directory as follows:

➤ A default is offered by the installer, which you can override.

➤ You must not use a non-ASCII directory name for your BPI installation directory; the home directory for BPI must consist only of ASCII characters.

➤ If you do install BPI into a non-ASCII directory, the BPI installation step is successful, but the BPI components fail to start.

BPI Instance Database Required Information

Following is the required database information for either the Microsoft SQL Server Database or the Oracle Database Server, depending on which database you use.

Microsoft SQL Database Server Required Settings

The following table lists the Microsoft SQL Database Server information that is required for a Server installation. If you plan to configure BPI to use an Oracle Server, you can skip this table and continue at "Oracle Database Server Required Settings" on page 89.

Note: When using MS SQL server named instances, the hostname field must contain *hostname/instancename* and the port number must be the port number of the specific named instance (not the usual 1433).

Information	Notes
Database type	This defines the database that you plan to use for BPI process data. Enter Microsoft SQL Server.
Hostname of the machine where the database is running. This can be a local or remote machine.	The default of **localhost** is offered by the installer. Accept the default if you are using a local database, or enter the fully-qualified domain name of the machine where the database is running if you are using a remote database.
SQL Server Database Port Number	This is the static port number that Microsoft SQL Server uses to communicate with applications; by default this is port **1433**. You must configure Microsoft SQL Server to use static port numbers with BPI and not dynamic port numbers.

Login credentials for a database user that has permission to create schemas and users, plus assign privileges to users.	The BPI schema is created as part of the installation process, as is the database user. The installer needs a user with sufficient privileges to create the schema and the new database user.
	
	 Obtain the credentials for a suitable user from your database administrator.
	
	 For a BPI database user to create the BPI database on an MS SQL Server 2005/2008, the following permissions are needed:
	
	 ➤ Master database permissions:
	     ➤ **CONNECT**
	     ➤ **CREATE DATABASE**
	 ➤ Server permissions:
	     ➤ **ALTER ANY LOGIN**
	
	 By default the installer offers the sa account.

Information	Notes
Name for the SQL Server database	This is the name of the SQL Server database created for BPI. By default you are offered the name **hpbpischema**. If you want a different name, you can replace the default offered during installation. The database name must be between 1 and 16 characters.
	You are strongly recommended to assign a database name that is specific to BPI.
	CAUTION: The database name must be no larger than 16 characters. A larger name could cause errors.

Information	Notes
User name and password details of the BPI database user.	The BPI machine needs a database account that it can use to create and modify the BPI database tables. This might be a user that you have already created through SQL Server, in which case you need to enter the correct credentials. Alternatively, you can provide a new user name and the installation step creates the user.
	If you have created a user, you need to make sure that the user has the correct permissions set. See the specified permissions in the Login credentials for a database user on the previous page.
	The database user name must be between 1 and 16 characters and comprise letters, digits, and the underscore character (_); it must start with either a letter, or the underscore character.
	CAUTION: The user name must be no larger than 16 characters. A larger name could cause errors.
	Make sure that the password obeys the password policy for the database. As an example, Microsoft SQL Server 2005 requires a password to include some numbers. The install fails if the password policy is not adhered to.
Directory where the database data files are stored	The installation needs to know where to create the database files. A default is offered by the BPI installation. If the location of the database files has not been changed within SQL Server, accept the default.

Oracle Database Server Required Settings

The following table lists the Oracle Database Server information that is required for a Server installation. If you plan to configure BPI to use a Microsoft SQL Server, see "Microsoft SQL Database Server Required Settings" on page 85.

Information	Notes
Database type	This defines the database that you plan to use for BPI process data. Enter the Oracle Server.
Hostname of the machine where the database is running. This can be a local or remote machine.	The default of localhost is offered by the installer. Accept the default if you are using a local database, or enter the fully-qualified domain name of the machine where the database is running if you are using a remote database.
Oracle Server Database Port Number	This is the port number that the Oracle Server uses to communicate with applications. By default it uses port 1521.
Oracle database instance identifier	The installation requests a value for the Oracle database instance identifier (SID), which is where the BPI tables are created. Check with your database administrator for the correct value of the SID for BPI. There might be a SID that is specific to BPI.

Information	Notes
User name and password for a database user that has permission to create the tablespace and user for BPI.	The BPI tablespace is created as part of the installation process, as is the database user. The user name entered to create these needs the privileges to: create the BPI schema, create the BPI user and set up permission to read from the System tables. Obtain the credentials for a suitable user from your database administrator.
	For a BPI database user to create the BPI database on Oracle, the following permissions are needed:
	➤ **CREATE SESSION**
	➤ **CREATE PROCEDURE**
	➤ **CREATE TABLESPACE**
	➤ **CREATE USER**
	➤ **CREATE ROLE**
	➤ **DROP USER**
	➤ **ALTER USER**
	➤ **GRANT ANY PRIVILEGE**
	➤ **SELECT ON V_$DATAFILE**
	➤ **SELECT ON V_$TABLESPACE**

Information	Notes
User name and password details of the BPI database user.[a]	The BPI machine needs a database account that it can use to create and modify the BPI database tables. This might be a user that you have already created through Oracle Server, in which case you need to enter the correct credentials. Alternatively, you can provide a new user name and the installation step creates the user.
	If you have created a user, you need to make sure that the user has the correct permissions; see "Oracle Server Database User Permissions" on page 80.
	The database user name must be between 1 and 27 characters and can comprise letters, digits and the underscore (_) character; it must start with either a letter or the underscore character.
	The database password must be between 1 and 30 characters and can comprise letters, digits and any punctuation characters, with the exception of the space ( ) and the brace ({}) characters.
	You are strongly recommended to assign a User that is specific to BPI.
Directory where the database data files are stored	The installation needs to know where to create the database files. A default is offered by the BPI installation, but you need to check with your database administrator for confirmation about where the database files should be created.

a. If you intend to use an Oracle Database Server that is installed on an HP-UX machine, you must create a database user for BPI before starting the installation; if you do not, the installation cannot proceed. Refer to section "BPI Instance Database Requirements" on page 72 for more details.

Business Service Management Information

The following table lists the installation information required for the BPI Server installation, which is related to the location of Business Service Management. You can also enter and modify these details after the installation is complete, using the BPI Administration Console.

Information	Notes
Hostname for Gateway Server	The fully qualified DNS name of the machine where the Business Service Management Gateway Server is installed. This is the server that you want BPI to use as a source for Service Status information, as the destination for data samples and for ODB synchronization. If you have configured a Load Balancer, you can specify the name of the Load Balancer in place of the Gateway Server.
	If you plan to use the BPI Server in a domain that is different to the domain where the Gateway Server is installed, make sure that you add the BPI domain to the list of Protected Domains, on the BSM user interface under:
	Admin > **Platform** > **Users and Permissions** > **Authentication**
Use of SSL?	An indication of whether or not you are using of SSL when connecting to the Gateway Server.
HTTP(s) port of Gateway Server	The port number used by the Web Server on the Business Service Management Gateway Server. Usually this is port 80, or port 443 if you have configured SSL.

Information	Notes
Single sign-on init string	Single sign-on init string that was manually retrieved from the BSM user interface under **Admin** > **Platform** > **Users and Permissions** > **Authentication Management** > **Token Creation Key**, and copied to the BPI server.
	If you do not define the init string during installation or need to change it, you can do so in the **Servlet Engine** > **Security** > **Single sign-on token creation key (Init string)** setting in the Administration Console. For details see "Security" on page 200.
BSM user name and password for BPI Server.	The user name that you have defined to be used by BPI server components when connecting to Business Service Management. This is the account used by BPI to obtain the Service Status information from Business Service Management and for synchronization.

Process Repository Database Information

The BPI Server needs the location and login details of the system where the Process Repository database is installed. If you enter the correct details for the BSM Gateway Server, these details are determined automatically. If the BPI Server fails to contact the BSM Gateway Server for any reason, these details cannot be determined and you need to enter them as part of the installation.

Microsoft SQL Server Connection Required Settings for the Process Repository Database

The following table lists the information required in cases where the Process Repository schema is installed on a Microsoft SQL Server database.

Information	Notes
Database type	This defines the database that you plan to use for BPI process data. It is either, Oracle or Microsoft SQL Server.
Hostname of the machine where the Process Repository database is running.	The fully-qualified domain name of the machine where the Process Repository database is installed. You configured the Process Repository database as part of the BSM installation and configuration.
Database Port Number	The port number for the database where the Process Repository database is installed. By default this is port **1433.**
Process Repository database name	Name given to the Process Repository database when you installed and configured BSM.

Information	Notes
Authentication type	Type of authentication provided for the Process Repository database. This can be either SQL Server authentication or Windows authentication.
	NOTE: The Installer does not support connecting to the Process Repository database in the Windows authentication mode. However, this does not affect the installation process. A warning dialog will display, but you can continue the installation process by clicking **Ignore Connection Error**.
	The reason for this error is that the DLLs cannot be included on the Installer path. The database connection check is made before the bits are overlaid. Using the NT authentication mode requires that you provide the same NT account and domain information as configured in the SQL server for the **OVBPIModelRepository** windows service.
	For details, see the "Deploying MS SQL Server" section in the *HP Business Service Management Deployment Guide* PDF.
User name and password assigned to the Process Repository database.	The user name and password details as configured for the Process Repository database for BSM.

Oracle Connection Required Settings for the Process Repository Database

The following table lists the information required in cases where the Process Repository Schema is installed on an Oracle Server Database.

Information	Notes
Database type	This defines the database that you plan to use for BPI process data. It is either, Oracle or Microsoft SQL Server.
Hostname of the machine where the Process Repository database is running.	The fully-qualified domain name of the machine where the Process Repository database is installed. You configured the Process Repository database as part of the BSM installation and configuration.
Database Port Number	The port number for the database where the Process Repository database is installed. By default it uses port **1521**.
Process Repository SID	Oracle SID given to the Process Repository database when you installed and configured BSM.
Schema name and password assigned to the Process Repository database.	The schema name and password details as configured for the Process Repository database for BSM.

Installation Log Files

The installer creates the installation log files when the BPI components are installed, as described in the following sections.

BPI Server Installation Log Files

The log file for the BPI Server installation is located in the BPI root installation directory as follows:

bpi-install-dir**HP_Business_Process_Insight_InstallLog.log**

Database Log Files

The installation creates a log file when problems are encountered by the BPI installer when creating or migrating the database, database user and database tables. This file is called:

*bpi-install-dir***HP_Business_Process_Insight_DBSetupLog.log**

This file is created when there are errors to report after BPI has made a successful connection to the database. This log file contains errors relating to the SQL scripts that the BPI installer runs, and not database connection errors.

If there are no problems encountered during the installation, this log file is not created.

Installing the BPI Server Components

This section describes the steps for installing BPI components, using the BPI Server installer.

Before starting the installation:

➤ Read "Collecting Required Information for the Installer" on page 83.

➤ Set the system environment variable for **JAVA_HOME**.

If you intend to use an Oracle Database Server that is installed on an HP-UX machine, you must create a Database User for BPI before starting the installation; if you do not, the installation fails. For more details, see "BPI Instance Database Requirements" on page 72.

To install the BPI Server components for the first time using the automated BPI Server installer, complete the following steps (in order of the installer dialogs):

➤ "Start the Installation" on page 98

➤ "Choose Install Folder Dialog" on page 98

➤ "BPI Database Creation Dialog" on page 99

➤ "Business Service Management Details Dialog" on page 100

1 Start the Installation

a From an account with Administration privileges, download the installation step for Business Process Insight. Access the download from Business Service Management by selecting the **Admin** > **Platform** > **Setup and Maintenance** > **Downloads** page. Select the **Category** for Business Process Insight.

b Download the software to a location of your choice and run the setup file for the installation.

The installer **Introduction** page opens to introduce you to the installation.

c Click **Next** to start the installation and continue to the next installer screen.

Follow the instructions for each screen. You can click **Previous** if you need to go back.

2 Choose Install Folder Dialog

Click **Next** to accept the default installation directory, or click **Choose** to navigate to a directory of your choice. There might be a short delay before a navigation window is presented. Click **Next** to continue.

If you specify a non-default directory and wish to restore the default, click **Restore Default Folder**.

Caution: You can install only one BPI Server on each machine. You must not install and configure multiple BPI Servers on a single machine.

---

3 BPI Database Creation Dialog

You need to determine whether you want to have the installer create the BPI Instance Database to store the BPI process instance data, or if you want to manually create the BPI Instance Database after the installation is complete.

➤ Select **Yes** to have the installer create the BPI Instance Database automatically, and continue to "BPI Database Type" on page 99.

The installer requires login credentials for a user with sufficient privileges to create the BPI schema and user. The specific privileges needed are listed in "BPI Instance Database Required Information" on page 85.

➤ Select **No** to create the BPI Instance Database manually after installation. Skip to "Business Service Management Details Dialog" on page 100.

For instructions on manually creating the BPI Instance Database, see "Creating the BPI Instance Database Manually" on page 103.

➤ Click **Next** to continue.

4 BPI Database Type

If you selected **Yes** to have the installer create the BPI Instance Database, select one of the databases:

➤ Microsoft SQL Server. Select the **Microsoft SQL Server** option and click **Next** to provide the connection details. Click **Next** again to provide user details. For information on these settings, see "Microsoft SQL Database Server Required Settings" on page 85.

➤ Oracle Server. Select the **Oracle** option and click **Next** to provide the connection details. Click **Next** again to provide user details. For information on these settings, see "Oracle Database Server Required Settings" on page 89.

When you have entered all the required information, click **Next** to open "Business Service Management Details Dialog" on page 100.

Note: When you click **Next** after filling in the database details for the database you chose and you receive a warning dialog labeled **Database Tables Exist**, the installer has detected database objects in the database with the same names as those needed for BPI.

These database objects might be left from a previous installation, where you have uninstalled BPI and selected not to remove the BPI data. If the database is also used for other purposes, they could be objects from another application that have the same names as those used by BPI, or they could be user-defined.

For the database tables error, do one of the following:

➤ Click **Delete** to continue with the installation. The installer deletes and recreates all the database functions that BPI requires.

➤ Click **Back** and enter an alternative database name.

➤ Click **Back** and then **Cancel** to cancel the installation, remove or rename the BPI tables that already exist and then restart the installation.

5 Business Service Management Details Dialog

The BPI Server needs to know where the BSM Gateway Server (or Load Balancer) is located in order to send data samples to the Gateway Server.

Provide the details of the BSM Gateway Server. For information on these settings, see "Business Service Management Information" on page 92).

When you have entered all the required information, click **Next**.

Note: After clicking **Next**, if you receive a warning dialog indicating that the BPI Server failed to connect to the ODB, you can do one of the following from the warning dialog:

➤ Click **Modify Settings** to return to the Business Service Management Details dialog to modify your settings.

➤ Click **Continue** and allow the installation to continue. In this case, you can modify the settings for the BSM Gateway Server using the BPI Server Administration Console when the installation is complete.

6 BPI Process Repository Database Type Dialog

A Process Repository database was installed and configured as part of the BSM installation. The BPI Server needs to know the location of the Process Repository database; you therefore need to provide the details of the Process Repository that was configured on the BSM system. This change has no effect on your existing Business Process definitions, which will continue to work as before.

If the BPI Server successfully connected to the BSM Gateway Server in the previous step, these details are obtained by the installer from the BSM system.

Select the database type that was configured for the Process Repository. This is either Microsoft SQL Server or Oracle. Provide the details of these settings, see "Process Repository Database Information" on page 93.

When you have entered all the required information, click **Next**.

Note: After clicking **Next**, if you receive a warning dialog indicating that the BPI Server is not able to connect to the Process Repository database, you can do one of the following in the warning dialog:

➤ Click **Retry Connection** to return to the Microsoft SQL connection details dialog to modify your settings. Select this option if you checked **Windows authentication** for the MS SQL Server authentication mode. This error automatically appears for this mode, but the installation process is not affected. See the Microsoft SQL Server connection required settings in "Process Repository Database Information" on page 93.

➤ Click **Ignore Connection Error** and allow the installation to continue. In this case, you can modify the settings for the Process Repository using the BPI Administration Console when the installation is complete.

7 Automatic Start of Business Process Insight Dialog

Decide whether you want the BPI Server to be started automatically when the system reboots.

➤ Select **No** if you have a machine that takes some time for the BPI Server components to start, as the automatic restart is likely to time out. This might be the case if you are in the process of developing a BPI solution or running a pilot.

If you answer **No**, you need to manually start the BPI Server components, following a system restart. You do this using the BPI Server Administration Console.

➤ Select **Yes** if your machine starts up reasonably quickly.

8 Pre-Installation Summary Dialog

Check the details on the summary. If they are correct, click **Install**, if they are not, click **Previous** to correct them.

The installation step continues and installs the BPI Server directories and files. When the installation is complete, click **Done** to exit the installation step.

If you automatically created the BPI Instance Database during installation, skip the next section and go to "Starting BPI Server Components" on page 107.

Creating the BPI Instance Database Manually

If the BPI Instance Database was not created automatically by the installer, then it must be created manually before the BPI Server can be started.

The following sections describe the requirements used for each task and the tasks for creating the BPI Instance Database manually:

➤ "BPI Database Setup Command-Line Utility" on page 103

➤ "Tasks for Creating the BPI Instance Database Manually" on page 104

BPI Database Setup Command-Line Utility

The procedures in the next section require the use of the BPI database setup command-line utility — **bpidbsetup.bat** on Windows, located in *bpi-install-dir***bin**.

Each time the BPI database setup utility is executed, you need to specify the database option and its connection parameters for BPI:

➤ For the Microsoft SQL Server, use the **-mssq**l option with the following sequence of connection parameters (see "Microsoft SQL Database Server Required Settings" on page 85 for descriptions of each parameter):

-mssql <host> <port> <bpi_user> <bpi_pwd> <db_name>

➤ For the Oracle Server, use the **-oracle** option with the following sequence of connection parameters (see "Oracle Database Server Required Settings" on page 89 for descriptions of each parameter):

-oracle <host> <port> <bpi_user> <bpi_pwd> <db_sid>

Tasks for Creating the BPI Instance Database Manually

Perform the following tasks in order, to create the BPI Instance Database manually:

➤ "Create the BPI User and Tablespace or Schema" on page 104.

➤ "Create the Remainder of the BPI Instance Database Objects" on page 106.

➤ "Configure BPI to use the BPI Instance Database" on page 107.

Create the BPI User and Tablespace or Schema

The BPI user and tablespace or schema can be created manually using the command-line utility in one of two ways:

➤ Using the **-createScripts** option, which generates an SQL script that a database administrator can execute manually.

➤ Using the **-createDatabase** option, which creates the BPI user and tablespace or schema directly.

The user does not need an administrator role, but must have sufficient privileges. See the Microsoft SQL and Oracle database requirement settings in "BPI Instance Database Required Information" on page 85 for the necessary privileges.

To create the BPI user and tablespace or schema by executing an SQL script:

1 Invoke the BPI database setup command-line utility with the **-createScripts** option and either the **-mssql** or **-oracle** option (see "BPI Database Setup Command-Line Utility" on page 103).

Example: Using the **-createScripts** option on Windows with Oracle:

bpi-install-dir**bin>bpidbsetup.bat -createScripts -oracle localhost 1521 hpbpiuser hpbpiuserpassword oraclesid**

A set of SQL scripts are generated in *bpi-install-dir***misc****bia****DbSetupScripts.** These scripts are used in subsequent steps.

2 Create a copy of the **RDBMS_setup.sql** script and edit it. The following tokens in the file must be replaced:

Replace:	With:
{{BIA_ENGINE_ORACLE_TEXT_PASSWORD}} or **{{BIA_ENGINE_MSSQL_TEXT_PASSWORD}**	Password for the database login to be used by BPI.
{{BIA_ENGINE_ORACLE_DBF_DIR}} or **{{BIA_ENGINE_MSSQL_MDF_DIR}}**	Directory in which to create the data file for the tablespace or database. For example, **C:\ORACLE\PRODUCT\10.2.0\ORADATA\ORCL**
{{BIA_ENGINE_MSSQL_COLLATE}}	Default collation for the BPI database, which must be case-insensitive.

3 Save the file and have a database administrator use an interactive database query tool to connect to the database and execute it.

To create the BPI user and tablespace or schema by having the command-line utility do so directly:

Invoke the BPI database setup command-line utility with the **-createDatabase** option, either the **-mssql** or **-oracle** option (see "BPI Database Setup Command-Line Utility" on page 103), and the **-admin** option.

The **-admin** option requires the following sequence of connection parameters:

-admin *<admin_user> <admin_pwd>*

<admin_user> is the administrator user name for the database.

<admin_pwd> is the administrator password for the database.

The credentials supplied with the **-admin** option do not need to be for a user with an administrator role, but the user must have sufficient privileges. See the Microsoft SQL and Oracle database requirement settings in "BPI Instance Database Required Information" on page 85 for the necessary privileges.

Example: Using the **-createDatabase** option on Windows with Oracle:

bpi-install-dir\bin\bpidbsetup.bat -createDatabase -oracle localhost 1521 hpbpiuser hpbpiuserpassword oraclesid -admin adminuser adminpassword

Create the Remainder of the BPI Instance Database Objects

The remainder of the BPI Instance Database objects can be created manually using the command-line utility in one of two ways:

➤ Using the **-createScripts** option, which generates a set of SQL scripts that a database administrator can execute manually.

Invoke the BPI database setup command-line utility with the **-createScripts** option and either the **-mssql** or **-oracle** option (see "BPI Database Setup Command-Line Utility" on page 103).

The command line utility prints a list of SQL scripts that must be executed to create the BPI database objects. These scripts were generated in *bpi-install-dir*\misc\bia\DbSetupScripts. Have a database administrator use an interactive database query tool to connect to the database as the BPI database user and execute these scripts in the order specified.

➤ Using the **-createObjects** option, which creates the database objects directly.

Invoke the BPI database setup command-line utility with the **-createObjects** option and either the **-mssql** or **-oracle** option (see "BPI Database Setup Command-Line Utility" on page 103).

Example: Using the **-createObjects** option on Windows with Oracle:

bpi-install-dir\bin>bpidbsetup.bat -createObjects -oracle localhost 1521 hpbpiuser hpbpiuserpassword oraclesid

Configure BPI to use the BPI Instance Database

Invoke the BPI database setup command-line utility with the **-configureBpi** option and either the **-mssql** or **-oracle** option (see "BPI Database Setup Command-Line Utility" on page 103).

Example: Using the **-configureBpi** option on Windows with Oracle:

bpi-install-dir\bin>bpidbsetup.bat -configureBpi -oracle localhost 1521 hpbpiuser hpbpiuserpassword oraclesid

This option verifies the integrity of the database, making sure all the BPI database objects are present. Once verified, BPI is configured to connect to the database.

Starting BPI Server Components

This section describes how to start the BPI Server components.

If you are using an Oracle Server it can take a minute or so for BPI components to make a connection (on some systems). There is a retry mechanism for connecting to an Oracle Server when BPI is starting, but this retry interval eventually times out. In this case, you might need to have a second attempt at starting some of the BPI components that use the database.

You can start the BPI Server components in one of the following ways:

On the BPI Server Administration Console

1 Start the BPI components:

➤ On Windows, start the BPI Administration Console by selecting:

Start > **Programs** > **HP** > **Business Process Insight** > **Administration**

2 Click the **Start All** button to start all the BPI Server components.

On Windows from the Services Option from the Control Panel

1 Select **Start** > **Control Panel** > **Administrative Tools** > **Services**

The BPI Server components are listed as the following Windows services:

➤ **OVBPIAdaptorFramework**

➤ **OVBPIAdminServer**

➤ **OVBPIBacDataSamplesProvider**

➤ **OVBPIEngine**

➤ **OVBPIEventHandler**

➤ **OVBPIJmsEventHandler**

➤ **OVBPIMetricEngine**

➤ **OVBPIModelRepository**

➤ **OVBPINotificationServer**

➤ **OVBPIServiceAdapters**

➤ **OVBPIServletEngine**

➤ **OVBPIWebServicesProvider**

2 Start all the BPI Services as you would for any Microsoft Windows Service.

You can also configure further properties for these BPI Services, which are specific for your implementation. As an example, you can configure the Windows Service to restart automatically following a failure. As installed, the Services do not restart following a failure.

BPI Infrastructure Settings

When the BPI Server is started, it attempts to connect to the BSM Gateway Server. If this is successful, then the Business Process Insight settings are updated.

If you wish to override these settings, go to:

Admin > **Platform** > **Setup and Maintenance** > **Infrastructure settings** > **Business Process Insight**

Enabling Oracle Database Table and Index Partitioning

When you have completed the BPI installation, you can enable database table and index partitioning for the BPI database tables. This section is only relevant for Oracle database users.

Some of the BPI database tables and indexes are designed to enable you to take advantage of Oracle partitioning; however, you can do this only if you have purchased a suitable Oracle License.

If you have a license to use Oracle Database and Index Partitioning, and you consider that your database performance can be improved by using it with BPI, you can configure BPI to make use of this Oracle feature.

If you have the correct license terms, you can enable Oracle Database table and index partitioning for BPI as follows:

1 From the BPI Server Administration Console, stop all the BPI components.

2 Run the **OraclePartitioningSwitcher** command:

➤ On Windows, run:

   bpi-install-dir\bin\oraclepartitiningswitcher switch

The **switch** command can be either **true** or **false**.

true enables partitioning for some BPI database tables and indexes and **false** disables any partitioning for BPI.

Be aware that when you enable Oracle partitioning, there can be a significant delay while Oracle makes the required changes. The delay depends on the size of the database tables and indexes that are being partitioned and can be more than one hour.

3 When the partitioning is complete, restart the BPI components using the BPI Administration Console.

Verifying Your Installation Using the Integrity Checker

All the BPI Server components produce logging and tracing files that you can use to identify problems within your BPI system; however, these log and trace files do not detect changes to the characteristics and file permissions of the BPI files and directories, nor do they inform you of missing files. The Integrity Checker utility is designed to check for problems that cannot be detected by the logging and tracing mechanisms and is a useful tool for validating your system following an installation.

The Integrity Checker:

➤ Generates the files required as the benchmark for a future integrity check.

These files are created as part of the installation process and are tailored to each installation type. These files enable the Integrity Checker to be used to validate that all the installation files exist, are in the correct location, have the correct file permissions set and are complete.

➤ Uses the BPI configuration files to detail the current BPI configuration, when required. For example, the Integrity Checker can:

➤ Confirm port numbers that are currently configured for use

➤ Validate that all the database objects required by BPI have been created successfully; for example, functions, indexes and tables.

➤ Makes available the version numbers of all the BPI Server component JAR files (from the implementation version field from within the file). This does not include the third-party JAR files.

➤ Checks that BPI Server can connect to the BPI Instance database, and that all the database tables exist.

➤ Determines the version of J2SE (Java Virtual Machine) installed on the machine where BPI Server is running.

Running the Integrity Checker

If the log and tracing files seem to be providing misleading information, or you want to check the status of your BPI system for other reasons, you can run the Integrity Checker.

To run the Integrity Checker:

1 On the machine where you have installed the BPI Server, change the directory to:

bpi-install-dir**lbin****bia**

2 Execute the appropriate file:

➤ On Windows, execute the following file:

integritychecker.bat

3 Examine the contents of the report file generated by the Integrity Checker. This file is located at:

bpi-install-dir**data****log**

The file name is in the form **integrity_report_*systemtime*** and has a **.TXT** file extension. *systemtime* is the system time, which makes the file name unique.

Errors generated by the Integrity Checker, if it fails to run, are written to **stderr.**

Check the errors in the log file generated by the Integrity Checker and then find the problem statement in this chapter, and follow the instructions to recover from the error.

Integrity Report

The Integrity Checker creates a log file, which reports on the status of the BPI system and configuration. The Integrity Checker reports on:

➤ Files that have incorrect checksums

Following a new installation, the Integrity Checker calculates a checksum value for all the files installed. When you run the Integrity Checker, it re-evaluates these checksums and reports the files where the checksums are different.

There might be valid reasons why the checksums have changed, in which case you can ignore the problem, for example, you might have modified the content of one or more files.

➤ Java environment

The Integrity Checker checks the version numbers for the J2SE installed on your system to ensure that it is compatible with the BPI requirements.

If the Integrity Checker reports an error with the Java environment, you should correct the error before using BPI.

➤ Master configuration properties

The configuration options presented through the BPI Server Administration Console are held in a file called **masterconfig.properties**. The property values in this file are the properties currently being used by the BPI Server system. The Integrity Checker reports on the current values being used by BPI, so you can confirm that the values are as you expect.

If the values are not as you expect, you need to modify them using the BPI Administration Console.

If **masterconfig.properties** has been deleted, the Integrity Checker returns an error on the command line.

➤ Installed jar files

BPI has a number of jar files and the Integrity Checker reports on the implementation version of these files. This implementation version shows whether or not patches have been applied to BPI.

➤ Port numbers

All BPI port numbers need to be within the range of a number greater than zero (0) or a number less than 65535. The Integrity Checker reports on the port numbers and informs you if any are out of this range.

If the Integrity Checker reports that the port numbers are not valid, you need to change them as described in "Component Configuration - Port Numbers" on page 208.

Installation Troubleshooting

This section describes troubleshooting problems during the installation of the BPI Server.

➤ Installation Procedure Fails with Error Indicating a Service is Still Running

If the BPI installer notifies you that there is a service running that you cannot detect using the Administration Console, it is most likely because you have stopped the BPI components individually using the BPI Administration Console Status page.

To resolve this problem, select the **Stop All** option on the Status page and make sure that you close the Administration Console when the components have stopped.

5

Re-installing BPI Server Components

You can use the BPI Server installation procedure to re-install existing BPI Server components. This enables you to install new components and replace missing or corrupt BPI Server software files. You can use the Installation Integrity Checker to determine if there are missing installation files; see "Verifying Your Installation Using the Integrity Checker" on page 110.

If you want to replace BPI data files, you need to recover these files from a backup. You cannot use the reinstallation program to change details of the database configuration for BPI.

If you want to upgrade from a BPI Server version 8.0x system to a BPI Server 9.00 system, skip this chapter and see Chapter 7, "Upgrading to BPI Version 9.00."

This chapter includes:

➤ Before Starting the Re-installation Procedure on page 116

➤ Re-Installing the BPI Server on page 116

➤ Post Re-Installation Tasks on page 117

➤ What to Do Next on page 118

Note: If you want to create a new BPI Instance database and remove all the data associated with the current BPI Server installation, use the uninstall instructions provided in Chapter 6, "Uninstalling BPI Server Components." This involves removing all the installed BPI Server components and deleting the database used by the BPI Server. Be aware that by uninstalling BPI you are removing all your BPI Server files and data.

The schema for the Process Repository is not affected by the installation, re-installation or removal of the BPI Server components and data.

Before Starting the Re-installation Procedure

Before starting the re-installation process make sure that you have done a full backup of your machine, and make sure that it includes copies of any configuration files that you have modified for your current installation.

The reinstallation procedure overinstalls, makes copies of configuration files, or preserves relevant configuration changes as follows:

➤ Changes made to configuration parameters through the BPI Administration Console are preserved and are used to present the BPI configuration parameter values during the re-installation.

➤ Changes made to other configuration files must be copied and reapplied to the files after the re-installation is complete. Examples of files that you might have changed are the SQL script for archiving Active and Deleted instances, or the Tomcat configuration file (**server.xml**).

➤ Changes made to process diagrams using annotations must be restored manually after the re-installation is complete, as described in "Post Re-Installation Tasks" on page 117.

The following section describes how to re-install the BPI Server.

Re-Installing the BPI Server

the following procedure describes how to re-install the BPI Server.

To re-install any BPI Server components:

1 Make sure that you shut down the BPI Server components and close the BPI Administration Console.

If the installer detects any of the BPI Server components still running, the installation is aborted until the components are stopped.

2 Complete a full backup of the BPI Server installation directory and all the BPI database files.

You can now continue and start the installation.

3 From an account with Administrator privileges, download the installation procedure for Business Process Insight, which you access from the **Admin > Platforms > Setup and Maintenance > Downloads** page. Select the Category for Business Process Insight on the system Windows.

4 Download the software to a location of your choice and run the setup file for the installation.

You are presented with an installation wizard that leads you through the procedures to install your BPI components.

5 Click **Next** to start the installation and then follow the instructions presented by the wizard.

The installation procedure installs the BPI component directories and files. When the installation is complete, click **Done** to exit the installation procedure.

Post Re-Installation Tasks

You might need to perform the following procedures to complete the re-installation process:

1 Reapply any changes that you made to the BPI configuration files that you saved before starting the re-installation procedure.

2 Manually restore customized process diagrams using annotations:

The installer keeps a backup of the **annotations.xml** file and images in the <*BPIInstallDir*>**nonOV\jakarta-tomcat-5.0.19\webapps\hpbpi-bpd.ba k.0** directory.

Copy the **annotations.xml** file from the backup directory to the following directory to restore the customizations:

<*BPIInstallDir*>**nonOV\jakarta-tomcat-5.0.19\webapps\hpbpi-bpd**

3 Restart the BPI components that you stopped using the BPI Administration Console.

What to Do Next

When you have re-installed the BPI Server, and reapplied your changes, you have completed the re-installation tasks and you can continue using your BPI Server machine.

6

Uninstalling BPI Server Components

This chapter describes how to remove the BPI Server components.

You can choose to remove all the BPI Server schema objects, including the database tables, or you can preserve them. You might want to preserve them if they contain data that you want to use for other applications.

Note: The schema for the Process Repository is not affected by the installation, re-installation or removal of the BPI Server components and data.

This chapter includes:

➤ Uninstallation Overview on page 119
➤ Tasks to Complete Before Removing the BPI Server on page 120
➤ Removing BPI Server Components on page 122
➤ Problem Removing BPI Components on page 125

Uninstallation Overview

If you select to remove all the schema objects, the uninstall procedure deletes the table, but it does not delete:

➤ Tablespace (Oracle Server)
➤ Database (SQL Server)
➤ Database user

If you want to keep your data for reporting purposes, select the **Keep Data** option when you run the uninstall program. Keep the data and then move it to a suitable location when the uninstaller completes. When you have moved the data, you must delete the BPI schema objects and database User using your database management tools. You also need to delete the BPI files under the installation directory when you have made a copy of the data that you need from them.

The BPI uninstall procedure does not delete any CIs, or other related data that was created within Business Service Management. Data for BPI Processes, Activities, or Monitors needs to be manually deleted using the Business Service Management ODB management tools.

Tasks to Complete Before Removing the BPI Server

Before removing the BPI Server components, you must first remove the settings within Business Processes, which are used to track individual process instances and all the Data and Event definitions. If you do not remove these before uninstalling the BPI Server the following problems may occur:

➤ The Modeler might not be able to synchronize CIs with the ODB for any Business Processes, including those not configured to track individual process instances.

➤ Undeploying Business Processes configured to track individual process instances causes the Modeler to close as it cannot access the required information in the Process Repository.

You can remove the required processes by exporting and then deleting all your Business Processes and subsequently importing the Business Processes that you want to keep.

To remove Business Processes that are configured to track individual process instances, complete the following procedures:

1 Open the Process Repository Explorer:

Admin > **Business Process Insight** > **Process Repository Explorer**

2 Use the Process Repository Explorer option to export all definitions:

File > **Export All**

3 Select the **Download** button, to export and save the definitions to a .zip archive file. For more information about exporting and importing definition files, see "Exporting and Importing Monitor Definitions" in the *Using Business Process Insight Guide*.

You now have a complete backup of your definitions.

4 Move to the Modeler:

Admin > **Business Process Insight** > **Modeling**

5 Delete all your Process, Data and Event definitions,

6 Move back to the Process Repository Explorer.

7 Remove all the definitions listed under Recycled:

File > **Clean-up Recycled**

8 Open the BPI Administration Console on the system where the BPI Server is installed:

➤ On Windows, select:

Start > **Programs** > **HP** > **HP Business Process Insight** > **Administration**

9 From the Status pane, click the **Stop All** button to stop all the BPI Server components.

10 Modify the BPI- related configuration parameters in the BSM Infrastructure settings and remove the entries that list the details for the BPI Server and BPI database:

Admin > **Platform** > **Setup and Maintenance** > **Infrastructure Setting**

11 Move back to the Modeler. (It needs to be reopened as it shut down when the BPI Server components were stopped).

Admin > **Business Process Insight** > **Modeling**

12 Import the previously exported .zip archive file containing all the definitions into the Modeler:

File > **Import Definitions**

The Modeler imports the definitions in the correct context for the new configuration and none of the processes is now configured to track individual process instances.

13 Redeploy all the Process definitions.

You can now continue and remove (uninstall) the BPI Server components as described in "Removing BPI Server Components" on page 122.

Important: It is very important that you do not attempt to restart the BPI Server components after completing these procedures. If you do, the system reconfigures the Infrastructure settings, leaving the system in an indeterminate state. You must remove the BPI Server components when you have completed these procedures, or reconfigure your system so that the BPI Server components are not started automatically. On Windows, this requires reconfiguring the Windows Services.

Removing BPI Server Components

To remove the BPI components complete the following procedures:

1 Make sure that all the BPI Server components are shut down:

a Use the BPI Administration Console to shutdown all the BPI Server components.

b Close the BPI Administration Console.

2 On Windows, make sure that all the BPI Windows Services are stopped.

3 When all the BPI components are shut down and BPI Windows Services are stopped, you can continue and remove the BPI software as follows:

➤ On Windows:

➤ Select: **Start** > **Settings** > **Control Panel** > **Add/Remove Programs**.

➤ Select **HP Business Process Insight** from the currently installed programs.

➤ Click **Change/Remove** to start the uninstall program.

You are presented with an introduction screen labeled Uninstall HP Business Process Insight.

4 Click **Next** to continue.

You are asked to confirm that all the BPI components are shut down.

5 When all the BPI Server components are shut down, click **Uninstall** to continue.

6 You are presented with the Delete Data dialog where you are asked if you want to delete the contents of the BPI Instance database.

If you select this option, it affects only the data created by the BPI Server, that is the process instance and violation information. Data in the BPI Process Repository and any related CIs defined in BSM are not deleted. This means that the Business Processes that you have created remain and can be accessed using the Modeler.

You can view your Business Processes within the Business Process Diagram using **Application** > **Service Health**. However, all Monitor, Event and Data definitions associated with these processes are deleted.

7 Select the option to **Delete BPI data,** if you want to delete all the data from the database; otherwise, select the option to **Keep BPI data.**

If you choose to Keep the BPI Server data, be aware that if you want to install another copy of BPI Server in the future, the new installation cannot progress if it detects an existing BPI Instance database. To reinstall in the future, any existing BPI Instance database tables need to be removed, or copied to another location.

Select the option that you want and click **Uninstall** to continue with the uninstallation procedure.

If you have chosen to keep the BPI Server data, continue at step 9; otherwise continue at the next step.

8 If you have chosen to delete the BPI Server data, you are presented with the database user password dialog, where you are asked to enter the password for the specified BPI Instance database user. Enter the password and click **Uninstall** to complete the uninstall procedure.

9 The Uninstaller continues and removes the BPI Server components.

The uninstaller reports any files that it cannot delete as part of the uninstallation process and recommends that you restart your Windows machine to complete the installation.

Click **Done** to complete the uninstallation procedure.

If you have chosen to keep the BPI Server data, continue at step 11.

10 On Windows, restart your machine.

11 Make sure that all the BPI Server files and directories have been removed. Delete any remaining BPI files from the following directory, including the directory HP Business Process Insight:

➤ On Windows:

C:\Program Files\HP\HP Business Process Insight

You have now completed the tasks to remove the BPI Server components from your machine.

Post Uninstall Tasks

The following are tasks that you might need to complete after you have removed the BPI components from your machine.

1 Make a Copy of Data in Database

If you have chosen to keep the BPI Server data, you can make a copy of the data that you want to keep and then delete the BPI Server files and directories, BPI Instance database files and the BPI Instance database User.

If you are using Oracle 10g or 11g, you can also purge the BPI Database User's Recycle Bin using the Oracle **PURGE RECYCLEBIN** command. This Oracle database command purges all the database tables, and their dependent objects, for the specified user, plus any other indexes owned by the user.

2 Remove BPI Options from the Windows Start Menu

On Windows, if your BPI installation was previously an upgrade from an earlier version of BPI, you might have some residual options remaining on the **Start Menu**. You can delete all the **HP** > **HP Business Process Insight** menu options from the **Start Menu** when the uninstall is complete.

3 Remove Database User

The uninstall program removes all the BPI-related database tables; however, it does not remove the database user that owns these tables. This is because the database user could be used for other database objects.

If you know that the database user is required only for the BPI Instance database tables, you can delete it after the uninstall program has completed successfully.

Problem Removing BPI Components

You might experience the following problem when attempting to uninstall the BPI Server components:

Java error, could not find a valid Java virtual machine.

In this case, the uninstall fails due to a missing Java virtual machine. The installation program InstallAnywhere requires a JRE to be available for use in order to run and looks for the location of the JRE on the System Variable Path.

If you remove this JRE on your system for any reason, and subsequently attempt to use InstallAnywhere to uninstall the BPI Server, you see an error indicating that the installer could not find a valid virtual machine to load as follows:

Could not find a valid Java virtual machine to load.
You may need to reinstall a supported Java virtual machine.

If you see this error, you need to install a JRE and then run the uninstall program for BPI again.

7

Upgrading to BPI Version 9.00

Upgrading from a BPI 8.0x server installation to BPI Server version 9.0x is not supported in BPI 9.00. Upgrading will be supported in a future minor release.

Since BPI server 8.0x is not compatible with BSM server 9.0x, if you are using BPI server 8.0x do not upgrade the BSM server to 9.0x.

Part III

Configuring and Administering the Business Process Insight Server

8

Component Administration

This chapter describes managing the BPI components using the Administration Console on the BPI Server.

This chapter includes:

Component Administration Overview

Managing the BPI components can be divided into two logical areas:

➤ Administration Console, which is accessed on the BPI Server system. The Administration Console manages the parameters of the BPI Server — for example, configuring the parameters required for the BPI connection to the BSM Gateway Server and starting and stopping the BPI Server Components. For details, see "BPI Administration Console" on page 133.

➤ Administration menu option from BSM. To access:

Admin > **Business Process Insight**

For details, see the *Using Business Process Insight Guide*.

This chapter includes the following information:

➤ Parameters that can be modified using the Administrative Console.

➤ How to make sure that any modifications made through the Administrative Console are preserved.

Not all the parameters described in this chapter need to be modified. What you need to change depends on your particular configuration and requirements. You are advised not to change anything unless you are recommended to in the documentation, or by your support representative.

BPI Security

You can choose to implement access control for your BPI interfaces using Tomcat (Servlet Engine). When first installed, the BPI interfaces are configured to use the BSM lightweight single-signon mechanism for authentication. If you want to modify this and use the Tomcat authentication mechanism refer to the Security options within the Administration Console as described in "Security" on page 200.

BPI Administration Console

You use the BPI Administration Console to manage the BPI Server components. The BPI components that make up the BPI Server are described in "BPI Components" on page 20. The Administration Console is a graphical management interface that displays a list of the parameters that can be modified.

The Administration Console opens on the Status panel, which is where you start and stop the BPI Server components.

You navigate through the Administration Console configurations using the explorer-style menu in the left-hand pane. Select the option from the pane that you want to manage. For example, if you want to manage port numbers select the Port Numbers option.

When you select an option in the left-hand pane, the right hand pane is updated to contain the configuration details for the parameters that can be modified.

You do not necessarily need to complete all the tasks described in this chapter for your implementation. Many of the parameters that you can change are used to configure more than one BPI component, for example, changing the database password affects all components that access the database.

Note: Make sure that you follow the instructions in each section when making changes to the configuration for BPI Server components through the Administration Console. This is because the procedures for changing the parameters can vary according to the parameter that is being modified.

Starting the Administration Console

You start the Administration Console as follows:

Start > **Programs** > **HP** > **HP Business Process Insight** > **Administration**

The Administration Console opens at the **Component Status** pane as follows:
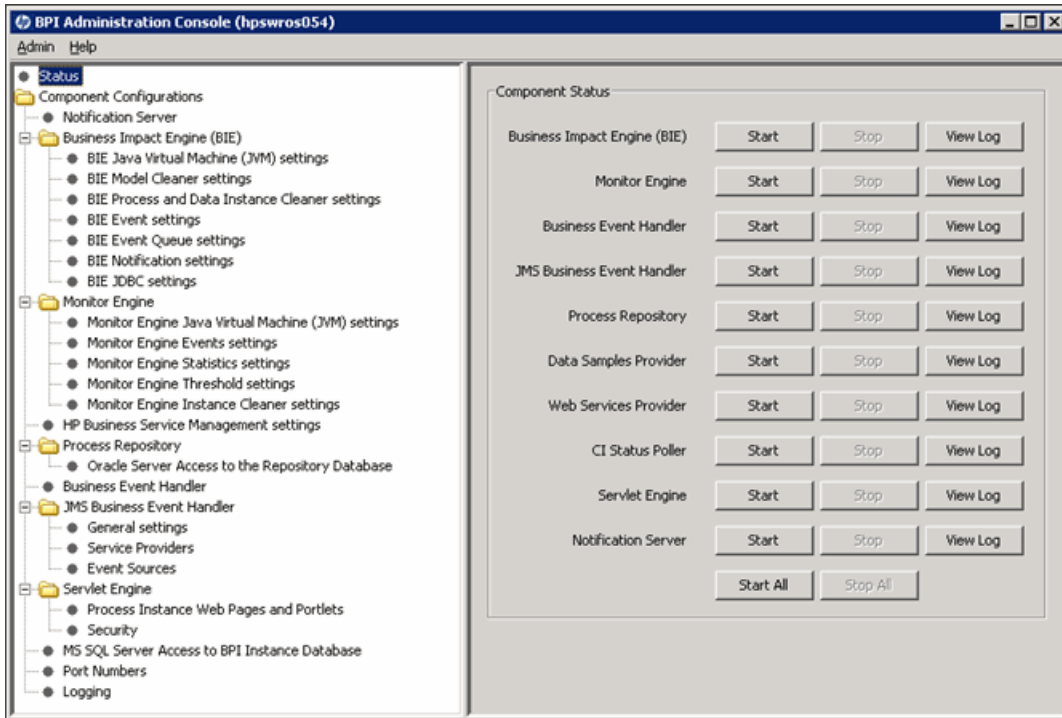
You can close the Administration Console at any time without applying outstanding changes and it does not impact the status of the system.

You might need to click the **Reset** button (in the other option parameter panes) to reset the page back to its state when you opened it. The next time that you start the Administration Console, it shows the current status of the BPI components.

Each time you start the BPI Server Administration Console, the BPI Server does the following:

➤ Sends details of the BPI instance database settings, Servlet Engine settings and installation directory details to BSM so that the BPI Application infrastructure settings can be set automatically within BSM; this means you do not need to add these settings manually.

If you change the settings using the Administration Console, the BPI Application infrastructure settings are also changed. If you change the BPI Application infrastructure settings directly from BSM (see "BPI Infrastructure Settings" on page 108), the settings are overwritten the next time the BPI Administration Console is restarted. This is with the exception of the Host name; the Host name is not affected because the same server can be known by different names to different systems.

➤ Polls for the BSM lightweight single-signon (LW-SSO) setting and uses the values to update the equivalent BPI settings under the Security option.

Functions Provided by the Administration Console

BPI comprises a number of individual components, each of which has its own configuration, or property file that define the parameters that can be used to control how the component behaves. Examples of parameter values that can be configured are:

➤ retry delays

➤ port numbers

➤ logging levels

Many of the parameters in these property files do not need to be changed. Therefore, in order to simplify the management of the BPI system, only those parameters that might need to be modified have been made available through the BPI Server Administration Console.

Caution: If you do make changes directly to the properties in the property files, you will need to reapply the changes each time you reinstall BPI, or add the same changes within the Administration Console, where they are preserved. This is because the files and the property values in these files are overwritten when the BPI Server is installed.

Specifically, the Administration Console provides access to the following BPI Server component areas:

➤ Server component status, which:

  ➤ shows whether or not the BPI Server components are running.

  ➤ enables you to start and stop components.

  ➤ enables you to view the content of the current log file.

  See "Status" on page 137 for details of the Status option.

➤ Notification Server parameters; see "Component Configuration - Notification Server" on page 143.

➤ Business Impact Engine parameters; see "Component Configuration - Business Impact Engine (BIE)" on page 147.

➤ Monitor Engine parameters; see "Component Configuration - Monitor Engine" on page 163.

➤ HP Business Service Management parameters; see "Component Configuration - HP Business Service Management Settings" on page 176.

➤ Process Repository parameters; see "Component Configuration - Process Repository" on page 185.

➤ Business Event Handler parameters; see "Component Configuration - Business Event Handler" on page 189.

➤ JMS Business Event Handler parameters; see "Component Configuration - JMS Business Event Handler" on page 191.

➤ Servlet Engine parameters; see "Component Configuration - Servlet Engine" on page 197.

➤ Microsoft SQL Server access.

  This option is available when you have configured Microsoft SQL Server as the database for the BPI Server; see "Component Configuration - MS SQL Server Access" on page 204.

➤ Oracle Server access.

  This option is available when you have configured an Oracle Server as the database for the BPI Server; see "Component Configuration - Oracle Server Access" on page 206.

➤ Port number settings; see "Component Configuration - Port Numbers" on page 208.

➤ Logging parameter settings; see "Component Configuration - Logging" on page 211.

The configuration parameters that you can modify for these component areas are described in the remaining sections of this chapter.

Note: The Modeler is not a BPI Server component. For details on configuring the Modeler, see the "Modeler" chapter in the *Using Business Process Instance Guide*.

Status

Use this option to start and stop the BPI Server components and also to view the component log files of these components.

The Status pane lists the BPI Server components that you can start and stop, and includes buttons labeled, Start, Stop and View Log. The components included on the Status pane, following a full installation on Windows, are:

➤ Business Impact Engine (BIE)

The component responsible for managing the progress of process instances, based on the business and operational events received from external sources. The BIE also generates impact notifications.

➤ Monitor Engine

The component that analyzes and provides the statistical results from Business Process Monitor, monitor instances and performance indicators. The Business Process monitors are created using the Business Process Monitors Definer.

➤ Business Event Handler

Receives business events from sources external to BPI and forwards them on to the Business Impact Engine in a standard format.

➤ JMS Business Event Handler

Received JMS-sourced business events from JMS-based business application sources external to BPI and forwards them on to the Business Impact Engine in a standard format. TransactionVision is an example of a JMS-based application.

➤ Process Repository

The design-time repository where the BPI Modeler stores its definitions. The Process Repository uses the information in the Process Repository schema when deploying the Business Processes.

➤ Data Samples Provider

The BPI component that is responsible for sending data samples to the BSM system.

➤ Web Services Provider

The BPI interface that exposes BPI Web Services to an external consumer.

➤ CI Status Poller

This is the service adapter that is used to manage the communication between the Business Impact Engine and the Gateway Server that BPI connects to for CI statuses.

➤ Servlet Engine

This is the Web Server component (Tomcat) that is installed as part of BPI. This component manages the server-side of all the BPI Web-based interfaces, for example: the Notification page and Intervention tool. It also manages the process instance and violation pages that are presented through the BPI Application pages.

➤ Notification Server

The component used to send email notifications of business events to configured recipients.

In addition, there are the Start All and Stop All buttons, which you can use to start and stop all BPI Server components.

You also access the log files for the BPI Server components from the Status pane. This enables you to identify status and error information for that component. The level of detail shown in the log files is controlled by the logging levels that you set. Setting log levels is described in section "Component Configuration - Logging" on page 211.

If the log files are not displaying correctly on your system, you can change the viewer used to display the log files as described in section "Log File Locations and Changing the Log Viewer" on page 141.

Starting and Stopping the BPI Server Components

You can use the BPI Administration Console to start and stop the BPI Server components that you have installed on the system where the Administration Console is running. The BPI Server components are also defined as Windows Services. For details on starting the Windows Services for the BPI Server components, see "Starting BPI Server Components" on page 107.

Starting the BPI Server Components Using the Administration Console

You start the BPI Server components from the Status panel. Click the Start All button to start all the BPI components installed on the system. You must make sure that the BPI Instance database is running and that BSM is installed and running before you attempt to start the BPI components as this command does not start the database.

You can start and stop the BPI components individually. To do this, click the Start button adjacent to the BPI component that you want to start.

Note: There are some dependencies between components, you are therefore advised to use the Start All option to start the components unless instructed otherwise.

Stopping the BPI Server Components Using the Administration Console

You stop the BPI Server components in the same way that you started them, from the Status panel of the Administration Console. Select the Stop All option to stop the BPI Server, or stop the components individually.

If you stop each BPI component individually, the Stop All button remains enabled. This is because the Administration Server components is still running. You are therefore advised to always use the Stop All command to shut down BPI Server before completing any installation tasks or tasks where it is important that all the BPI Server components are stopped.

BPI Clients

BPI has the following clients, which are accessed from the BSM Gateway Server where the BPI Application is available. These clients require both the BSM Gateway Server to be running and the BPI Servlet Engine to be enabled:

➤ Modeler

This is an application that you use to create and deploy your Business Processes. Using the Modeler to configure Business Processes and managing the Modeler are described in the "Modeler" chapter in the *Using Business Process Insight Guide*.

➤ Monitor Definer

You use the Monitor Definer to configure Business Process monitors and monitor instance threshold definitions. Examples and scenarios for using the Monitor Definer to configure Business Process monitors and thresholds are described in the "Monitor Definer" chapter of the *Using Business Process Insight Guide*.

➤ Process Repository Explorer

Use the Process Repository Explorer to view and manage entries in the Process Repository database. The Process Repository database holds the details of the Business Processes that you have created using the BPI Modeler. The Repository Explorer is described in the "Process Repository Explorer" chapter of the *Using Business Process Insight Guide*.

➤ The Process Intervention

Use the Process Intervention Tool to make changes to definitions and instances that are active or completed in the Business Impact Engine. The Process Intervention tool is described in the "Process Intervention" chapter of the *Using Business Process Insight Guide*.

➤ Notification Administration

Use the Notification Administration to:

➤ add and configure users who want to receive email notifications of instance violations from BPI.

➤ configure HPOM to receive HP Operations Manager messages.

➤ configure scripts.

Using the Notification page to configure the Notification Server is described in the "Notification" chapter of the *Using Business Process Insight Guide*.

In addition there are BPI Application Reports pages. These are the pages that enable you to monitor and track your Business Processes and are described in the "BPI Application Concepts."

You can also configure your own security options for accessing the BPI interfaces. You do this using the Servlet Engine (Tomcat); see Chapter 10, "Servlet Engine Authentication".

Log File Locations and Changing the Log Viewer

All the BPI components have log files; however the log files presented through the Status pane relate only to the BPI Server components, which are managed using the Administration Console.

The Modeler is not a BPI server component. Its log file is described in the "Modeler" chapter of the *Using Business Process Instance Guide*. The structure of all BPI log files is described in the *Business Process Insight Problem Solving Guide*.

All the BPI Server log files are located in the following directory:

*bpi-install-dir***\data\log**

Use the **View Log** button on the Status pane (adjacent to the Start and Stop buttons) to view the log files for the server components. By default the logs are displayed using notepad.

Section "Component Configuration - Logging" on page 211 describes the parameters for modifying the viewer application that is used to display the log files for BPI Server components.

Component Configuration - Notification Server

The following table lists the **Notification Server** parameters that you can modify to change the configuration using the Administration Console:

Descriptive Parameter Name	Description
SMTP hostname	The fully qualified domain name of the system where the SMTP email server is running; you can also enter localhost where appropriate.
	The BPI system needs access to a running SMTP email server in order to send impact violations through email to registered users.
	For information about registering users to receive email notifications see the "Notification" chapter in the *Using Business Process Instance Guide*.
SMTP port number	The port number that the SMTP server is communicating on.
	The BPI system communicates with the SMTP email server through a specific port number, which is usually port 25.
Sender email address	The email address that you want to appear in the Sender: field on email notifications sent from the BPI system. This is the email address that identifies where the message originated from.

Descriptive Parameter Name	Description
Retry interval to SMTP server (minutes)	The time (in minutes) that the Notification Server waits before attempting to send a notification to the SMTP server, if it fails on the first attempt. If the Notification Server is unable to send the message, it logs a warning message and temporarily stores the message in the database. The message is removed from the database when the notification is successfully delivered, or the retry limit is exceeded.
Maximum retries to SMTP server	Maximum number of retry attempts made by the Notification Server to send a notification message to the SMTP server. When the maximum number of retries is reached, the Notification Server deletes the message from the database, issues an error message in its log file and makes no further attempts to deliver the message to the email server.
Path to HP Operations Manager Opcmsg utility	The full path name for the **opcmsg** utility within HP Operations Manager. Used by the Notification Server to send notifications to HP Operations Manager.

Descriptive Parameter Name	Description
Retry interval to HP Operations Manager (minutes)	The time (in minutes) that the Notification Server waits before attempting to send an HP Operations Manager message to HP Operations Manager, if it fails on the first attempt. If the Notification Server is unable to send the message, it logs a warning message and temporarily stores the message in the database. The warning message is removed from the database when the notification is successfully delivered.
Maximum retries to HP Operations Manager	Maximum number of retry attempts made by the Notification Server to send an HP Operations Manager message to HP Operations Manager. When the maximum number of retries is reached the Notification Server deletes the message from the database, issues an error message in its log file and makes no further attempts to deliver it.
Retry interval for script execution (minutes)	The time (in minutes) that the Notification Server waits before attempting to execute a script that you have created, if the script fails to execute on the first attempt. If the Notification Server is unable to execute a script, it logs a warning message in the Notification Server log file.

Descriptive Parameter Name	Description
Maximum retries for script execution	Maximum number of retry attempts made by the Notification Server to execute a script that you have created. When the maximum number of retries is reached the Notification Server, issues an error message in its log file and makes no further attempts to execute the script.
Timeout for script execution (seconds)	The time (in seconds) that the Notification Server waits before aborting the execution of the script, after the script has started executing. Once aborted, the Notification Server waits for the retry interval before attempting to execute the script again.

Modifying the Notification Server Parameters

To change the **Notification Server** parameters, complete the following steps on the Windows system, where the BPI Server is installed:

1 Start the Administration Console by accessing **Start** > **Programs** > **HP** > **HP Business Process Insight** > **Administration**.

2 Select the **Component Configurations** > **Notification Server** option from the Navigator pane on the BPI Administration Console.

3 Enter the changes that you want to make to the Notification Server parameters in the right-hand pane as appropriate.

4 Click **Apply** to apply the changes to the BPI configuration.

5 Select the **Status** option to view the Component Status panel of the BPI components.

6 Stop and restart all the BPI Server components using the **Stop All** and **Start All** buttons.

Note: You can stop and restart only the Notification Server option and the Servlet Engine, but it is recommended to use **Stop All** and **Start All** to make sure that components are started and stopped in the correct order.

The new **Notification Server** parameter values are now applied to your BPI system.

If you want to change the login password for the **Notification Server** Administration Console, see Chapter 10, "Servlet Engine Authentication."

Component Configuration - Business Impact Engine (BIE)

The **Business Impact Engine (BIE)** configuration parameter settings are divided into the following logical sections:

➤ "BIE Java Virtual Machine (JVM) Settings" on page 148

➤ "BIE Model Cleaner Settings" on page 148

➤ "BIE Process and Data Instance Cleaner Settings" on page 150

➤ "BIE Event Settings" on page 155

➤ "BIE Event Queue Settings" on page 157

➤ "BIE Notification Settings" on page 160

➤ "BIE JDBC Settings" on page 161

The parameters relating to these settings appear on the right-hand pane of the Administration Console when you select one of these options in the console's navigation tree. If the options are not visible, you need to expand the entries under **Business Impact Engine (BIE)**.

BIE Java Virtual Machine (JVM) Settings

This section describes the parameters that you can modify through the **BIE Java Virtual Machine (JVM) settings** option. These parameters enable you to modify the amount of memory heap available to the JVM for the system. A heap is a storage management structure for tracking and allocating memory. In this case, the Java heap is used for allocating the Java objects used by the Engine.

Following are the **BIE Java Virtual Machine (JVM) settings** parameters:

Descriptive Parameter Name	Description
Initial size of the JVM heap (MB)	The initial size of the storage allocated by the JVM for Java objects.
Maximum size of the JVM heap (MB)	The maximum size of the storage allocated by the JVM for Java objects.

BIE Model Cleaner Settings

This section describes the parameters that you can modify through the **BIE Model Cleaner settings** option.

The Model Cleaner is a thread that runs at intervals and deletes Process and Data definitions that have been undeployed and that no longer have any associated instances, either in progress or completed.

The following table shows the Business Impact Engine parameters that enable you to remove undeployed Process and Data definitions from the BPI Instance database:.

Descriptive Parameter Name	Description
Delay between checking for all Models (minutes)	The time interval (in minutes) that the Business Impact Engine waits before reading the latest list of definitions. This list is used as input for the **Delay between checking each model** parameter.
Delay between checking each Model (seconds)	The time (in seconds) that the Business Impact Engine waits between checking each entry in the list of definitions created as a result of the Delay between checking for all models parameter. It is used to establish those definitions that have been undeployed and which are therefore candidates to be deleted. Definitions are not deleted if: ➤ the definition is deployed. ➤ there are associated instances in the database, either in progress or completed. ➤ the definition is referenced from another definition that is currently defined in the BPI Instance database; for example, a Data Definition is referenced from a Process definition.
Disable cleanup of files for Models not in database?	A check box indicating whether or not to remove the Java source, Java Class and script files for process and data models that have been cleaned from the Business Impact Engine. When checked, the source, classes and scripts are not removed.

BIE Process and Data Instance Cleaner Settings

The **BIE Process and Data Instance Cleaner settings** option determines how often the instance cleaner thread is run and the Completed and Active instances are deleted (or not deleted) from the database.

Instances are initially marked as being candidates to be deleted and the Business Impact Engine Instance Cleaner then runs at intervals that you specify to delete the instances that have been marked as candidates. For more information on the Business Impact Engine Instance Cleaner Threads and the parameters that control it, see Chapter 9, "Deleting Instances from the Database."

The **BIE Process and Data Instance Cleaner settings** are divided into three logical sections:

➤ "BIE Process and Data Instance Cleaner Settings" on page 151

➤ "Completed Process and Data Instance Cleaner Settings" on page 152

➤ "Active Process and Data Instances Cleaner Settings" on page 153

BIE Process and Data Instance Cleaner Settings

The **BIE Process and Data Instance Cleaner settings** parameters enable you to control how often the instance cleaner thread is executed. Following are the parameters:

Descriptive Parameter Name	Description
Mark instances for deletion:	The option for marking process and data instances as candidates to be deleted by the Instance Cleaner thread at a specified interval. You have the following options: ➤ Never - instances are never marked to be deleted. ➤ Periodically - instances are marked to be deleted at the interval that you specify. ➤ Once a day - instances are marked to be deleted once a day, at the time that you specify. ➤ Twice a day - instances are marked to be deleted twice a day, at the times that you specify.
Delete interval for marked instances	The time interval used to control when the instance cleaner thread runs to delete any instances that have been marked for deletion.
Delete batch size	The maximum number of active and complete process and data instances that the instance cleaner thread deletes in one Delete interval. There are four batches, one each for active and complete process instances, and active and complete data instances.

Completed Process and Data Instance Cleaner Settings

These settings control when Completed processes and data instances are deleted from the BPI Instance database. You also have the option to delete completed instances as soon as they are complete. There is no requirement for the Engine Instance Cleaner to run in this case, as the instances are deleted as soon as they are complete.

Following are the **Completed Process and Data Instance Cleaner settings** parameters:

Descriptive Parameter Name	Description
Delete completed instances from the database	The option for deleting Completed instances from the database. You have the following options for deleting Completed instances: ➤ Never - completed instances are never deleted and remain in the database ➤ As scheduled above - completed instances are deleted from the database at an interval that you specify in the BIE Engine Process and Data Instance Cleaner settings options. ➤ Immediately on completion - completed instances are deleted as soon as they reach the state COMPLETED and are not under the control of the settings for the Engine Instance Cleaner thread.
Age of the completed instances to be removed (minutes)	The age (in minutes) that the Completed instances must be before they can be deleted. The following are quick references to common time durations expressed as minutes: 1 day = 1440 minutes 1 week (7 days) = 10080 minutes

Active Process and Data Instances Cleaner Settings

These settings control how often Active process and data instances are deleted from the BPI Instance database. Following are the parameters:

Descriptive Parameter Name	**Description**
Delete active instances from the database	This is an indication of whether or not you want to delete active instances from the database. ➤ Never - active instances are never deleted and remain in the database ➤ As scheduled above - active instances are deleted from the database at an interval that you specify in the Engine Process and Data Instance Cleaner settings options.
Age of the active instances to be removed:	The age (in days, hours or minutes) that the Active instances must be before they can be deleted. If you enter a value of 60 minutes, the next time that you access this pane, the value is presented in terms of hours and not minutes.

For a new installation, the Engine Instance Cleaner Threads does not run and instances are not deleted when they complete. Completed and Active instances can therefore accumulate in the database and this, eventually, has an impact on the performance of your BPI system.

If performance, or space, is a consideration, you need to modify this default to remove these instances and also consider archiving data in order that you can view historical data. See the "Process Intervention" chapter in the *Using Business Process Instance Guide* for details of how you can specify your own SQL to archive process data before it is removed from the BPI database.

Note: If you delete the data for all Completed instances as soon as they are completed, data is no longer available to be viewed through the Process Instance and Violation Web pages for viewing purposes.

These parameters, and information about archiving instances before deleting them, are further described in Chapter 9, "Deleting Instances from the Database".

BIE Event Settings

This section describes the parameters that you can modify through the **BIE Event settings** option.

The parameter settings listed in the following table enable you to minimize the amount of time that the components spend in loops and possible database deadlock situations and provides settings for throughput rates.

Descriptive Parameter Name	Description
Maximum event generation number	A limit for the number of child events created from an initial incoming event. This is intended to interrupt potential infinite loops for subscriptions to data definitions.
Maximum retries for an event transaction deadlock	When processing an individual event, this is the maximum number of attempts by the Business Impact Engine to retry a database transaction before aborting it. It is possible for deadlocks to occur when the Engine and other applications are accessing the BPI database simultaneously. When the maximum number of retries is exceeded, the Business Impact Engine generates an error message and sends an exception to the Business Event Handler. The Business Event Handler then rolls back the event transaction and retries at a later time.

Descriptive Parameter Name	Description
Event sample count used for calculating rates	The number of Events used as a basis for calculating throughput rates. Set this parameter according to the incoming Event rate (to the Business Impact Engine). You want to use sufficient numbers of Events to give you a realistic throughput rate.
Decay period for rates when no events are received (seconds)	The time period after which the throughput rates are decremented, in stages, to zero. This is used for periods of inactivity, to maintain realistic results for throughput rates.
Maximum number of group values per process grouping	The maximum number of groups per process for which separate per group statistics are maintained. The default is 100. A large number of groups can severely impact the performance of the engine. If the number of groups exceed the specified limit then instances belonging to extra groups are bundled together under one group name called **Other**.

BIE Event Queue Settings

This section describes the parameters that you can modify through the **BIE Event Queue settings** option.

These settings enable you to configure a business event queuing mechanism within the Business Impact Engine (BIE). When you select the option to use a queue for incoming events, all new business events are placed on a queue by the BIE. This means that any source of business events, for example: the Business Event Handler or the JMS Business Event Handler does not need to wait for the BIE to process the business event and can immediately return to tracking the event source and process the next incoming business event. This allows the BIE to have multiple threads to process business events.

If you choose not to use a queue for business events, the Business Event Handler waits until the business event has been processed by the BIE; this includes waiting for all the necessary information from the event to be committed to the database. The effect of this is to make the BIE single-threaded for event processing; the options are all disabled and any setting for the number of event handling threads is ignored.

The settings listed in the following table control whether a queue is utilized for business events and the characteristics of the queuing mechanism.

Descriptive Parameter Name	Description
Use a queue for incoming events?	Select this option if you want BPI to use a queue for incoming business events. Clear the check box if you do not want to use a queue. When cleared, all the remaining options for the BIE Event Queue settings are unavailable.
Maximum number of events in the queue	The maximum number of business events that can be placed on the queue. When the number of business events on the queue reaches this maximum, no further business events can be added to the queue. In this case, the Business Event Handler has to wait until the queue is reduced before it can add further events.

Descriptive Parameter Name	Description
Save the queue to the database	Select this option if you want to write each business event to the database either as it arrives on the queue, or each time the Business Impact Engine is shut down (where shut down is stopping the Business Impact Engine from the Administration Console, or stopping the Windows Services.)
	Using this option enables you to rebuild the queue in the event of a system failure. If the business events are not written to the database on arrival, there is a risk that the events might be lost following a system failure. Be aware that there is a possible performance impact when you select the option to save each event as it arrives to the database.
Event queue cleaner interval (mins)	The time interval used to control the frequency that the event queue cleaner runs. Each time the event queue cleaner runs, it deletes, from the database, all business events that have been successfully processed by the BIE.
	This setting is available when the option to save the queue to the database for every event is selected.
Number of event handling threads	The number of individual processes available to process business events held on the event queue.
	A recommendation is to set this parameter to be twice the number of CPU cores that are available on your machine.

BIE Notification Settings

The **BIE Notification settings** option parameters control at what impact level notifications are sent, and the number of retry attempts and the interval between the retries for the notifications. If the Business Impact Engine initially fails to send an event notification to the Notification Server, it can try again.

Following are the **BIE Notification settings** parameters you can modify:

Descriptive Parameter Name	Description
Business activity impact notification severity threshold	The minimum instance threshold for the severity for a Business Process Activity's KPI that can trigger an impact of impeded or blocked for the process. When one (or more) Activity services reach the instance threshold specified a notification event is issued. Possible levels of severity are: Normal, Warning, Minor, Major and Critical. These map to other HP BTO Software products' levels of severity.
Maximum number of retries	The maximum number of attempts that the Business Impact Engine tries to send a notification event to the Notification Server, if the initial attempt fails. When the maximum number of retries is reached, the Business Impact Engine issues an error message in its log file and makes no further attempts to send the notification event.
Retry delay (seconds)	The delay (in seconds) between the Business Impact Engine's attempts to retry sending notification events to the Notification Server. Each time the Engine fails a retry attempt, it issues a warning message to the log file.

BIE JDBC Settings

This section describes the parameters that you can modify through the **BIE JDBC settings** option.

The following table lists the settings that enable you to tailor the JDBC connection between the BIE and the database.

Descriptive Parameter Name	Description
Maximum number of active JDBC Connections	The maximum number of JDBC connections that the Business Impact Engine can have active at any one time. If a connection does not become available in the required time, the connection is refused and the Engine is unable to commit the transaction. In this case, the Engine writes an error to its log file.
Maximum wait time for a JDBC Connection (seconds)	The maximum length of time (in seconds) that the Business Impact Engine waits for an available JDBC connection before reporting an error. If a connection does not become available in the required time, the connection is refused and the Engine is unable to commit the transaction. In this case, the Engine writes an error to its log file.
Maximum number of idle JDBC Connections[a]	The maximum number of JDBC connections that can be idle at any one time. Set this to at least the number of active connections.
Maximum number of active JDBC Prepared Statements[1]	The maximum number of database Prepared Statements that can be active at any one time.

161

Descriptive Parameter Name	Description
Maximum wait time for a JDBC Prepared Statement (seconds)	The maximum length of time (in seconds) that the Business Impact Engine waits for a Prepared Statement to be executed before reporting an error in the log file and aborting the transaction.
Maximum number of idle JDBC Prepared Statements[1]	The maximum number of Prepared Statements that can be idle at any one time.

a. These connections are managed by Apache Commons DBCP (Database Connection Pool).

Modifying the Business Impact Engine (BIE) Settings

To change any of the **BIE settings**, complete the following steps on the Windows system, where the BPI Server is installed:

1 Start the Administration Console by accessing **Start** > **Programs** > **HP** > **HP Business Process Insight** > **Administration**.

2 Select the **Component Configurations** > **Business Impact Engine (BIE)** > **<*BIE_settings*>** option from the Navigator pane on the BPI Administration Console.

3 Enter the changes that you want to make to the Engine settings in the right-hand pane as appropriate.

4 Click **Apply** to apply the changes to the BPI configuration.

5 Select the **Status** option to view the Component Status panel of the BPI components.

6 Stop and restart all the BPI Server components using the **Stop All** and **Start All** buttons.

The new Business Impact Engine settings are now applied to your BPI system.

Component Configuration - Monitor Engine

The **Monitor Engine** configuration parameter settings are divided into a the following logical sections:

➤ Monitor Engine Java Virtual Machine (JVM) settings

➤ Monitor Engine Events Settings

➤ Monitor Engine Statistics settings

➤ Monitor Engine Threshold settings

➤ Monitor Engine Instance Cleaner settings

The parameters relating to these settings appear on the right-hand pane of the Administration Console when you select one of the Monitor Engine options in the console's navigation tree. If the Monitor Engine options are not visible in the navigation tree, expand the entries under Monitor Engine using the usual Explorer-style navigation techniques.

You can modify these settings, see "Modifying the Notification Server Parameters" on page 146 for instructions.

Monitor Engine Java Virtual Machine (JVM) Settings

This section describes the parameters that you can modify through the **Monitor Engine Java Virtual Machine (JVM) settings** option. These are the Monitor Engine parameters that enable you to modify the amount of memory heap available to the JVM and are listed in the following table.

A heap is a storage management structure for tracking and allocating memory. In this case, the Java heap is used for allocating the Java objects used by the Monitor Engine. Following are the Monitor Engine Java Virtual Machine (JVM) settings:.

Descriptive Parameter Name	**Description**
Initial size of the JVM heap (MB)	The initial size of the storage allocated by the JVM for Java objects.
Maximum size of the JVM heap (MB)	The maximum size of the storage allocated by the JVM for Java objects.

Monitor Engine Events Settings

This section describes the parameters that you can modify for the **Monitor Engine Events settings** option.

These settings provide control for the rate at which the Monitor Engine processes monitor events and the maximum sizes of the resulting database transactions as monitor events are processed.

The values for the **Monitor Engine Events settings** have been set to be suitable for most implementations. You should not change them unless you are advised to by your support representative.

The **Monitor Engine Events settings** are divided into two logical sections:

➤ "Monitor Engine Event Processing Settings" on page 164

➤ "Monitor Engine Performance Checking Settings" on page 166

Monitor Engine Event Processing Settings

The Monitor Engine Event processing settings parameters are listed in the following table:

Descriptive Parameter Name	Description
The maximum number of pooled monitor event processors (1 to19)	The number of parallel threads available to the Monitor Engine for processing different monitor events.

Descriptive Parameter Name	Description
The maximum number of monitor events to process in each batch	The number of monitor events that can be processed in a single batch. This provides a balance between the monitor event processing rate and database transaction size. It also optimizes how the monitor events are processed across all available threads.
The monitor events idle wait time (seconds)	Each time the Monitor Engine determines that there are no more monitor events to process, it waits for this time period before checking again. This prevents the Monitor Engine continuously initiating searches that can return no results.

Monitor Engine Performance Checking Settings

The following settings provide some control to regulate the rate at which monitor events are generated by the Business Impact Engine. It is used to make sure that the Business Impact Engine does not create monitor events at a rate that the Monitor Engine is unable to keep up with,

The Monitor Engine performance checking settings parameters are listed in the following table:

Descriptive Parameter Name	Description
The monitor events queue length high-water mark	This is the number of unprocessed monitor events that signal that the Monitor Engine is not keeping up with the Business Impact Engine. When the high-water mark is exceeded, the Business Impact Engine temporarily stops processing Business events until the Monitor Engine has reduced the backlog to below the low-water mark.
The monitor events queue length low-water mark	This is the number of unprocessed monitor events that signal that the Monitor Engine has cleared its backlog and is now keeping up with the Business Impact Engine. As a result, the Business Impact Engine can resume its processing of Business events.
Interval for checking monitor events queue length (seconds)	Time interval (in seconds) at which the number of Monitor Events waiting to be processed is checked against the values of the high-water mark and low-water mark parameters.

Monitor Engine Statistics Settings

The **Monitor Engine Statistics settings** parameters control how the Monitor Engine processes the staged statistical data and stores them in the configured statistical time intervals (Data Samples setting).

Staged statistical data are produced and stored as monitor events are processed. At intervals that you configure, the Monitor Engine collects the staged statistics into the statistics table for the completed time intervals.

The following table lists the parameters that you can set to configure how often the Monitor Engine polls the BPI database.

Descriptive Parameter Name	Description
Statistics generation polling interval (seconds)	The time interval (in seconds) that the Statistics generator polls the BPI database tables in order to update the statistics that have been configured using the Monitor Definer.
Maximum age of generated statistics on startup (days)	The time interval (in days) that the Monitor Engine uses to calculate historical monitor statistics following a period of BPI shutdown.
The maximum number of statistics consolidators (1 to 19)	The number of parallel threads available to consolidate the staged statistical data, for different monitor events, into the user-defined statistical time intervals. This includes monitor events that have Groups defined.
The maximum number of staged statistics' rows per transaction	Controls the size of database transactions as the statistical data is being consolidated.
The maximum number of staging transactions each time	Controls the rate at which statistics are consolidated. This is a multiplier of The maximum number of staged statistics' rows per transaction and is applied each time the Monitor Engine processes statistics.

Descriptive Parameter Name	Description
The maximum number of statistics aggregators (1 to 19)	The number of parallel threads available to aggregate statistics that have Groups defined into their overall statistical time intervals.
The maximum number of statistics group values per monitor	Used to specify the maximum number of statistics group values permitted per monitor.

Most of the parameters control the rate at which statistics are generated and these settings should not be modified unless you are instructed to do so by your support representative. The exception is the settings that controls how far back in time monitor statistics are generated when the Monitor Engine is restarted.

When BPI is restarted following a system shutdown, the monitor statistics are calculated for all the monitor definitions defined for processes. If you have shut your BPI system down for a significant amount of time (several days), then these statistic calculations can take considerable time, and impact the performance of your system. You can configure how far back in time to calculate these statistics following a restart using the Maximum age of generated statistics on startup (days) setting.

The following settings are used to control how the staged statistics are processed into the statistical time intervals:

➤ The maximum number of statistics consolidators (1 to 19)

➤ The maximum number of staged statistics' rows per transaction

➤ The maximum number of staging transactions each time

➤ The maximum number of statistics aggregators (1 to 19)

In the description of the parameters, consolidation is the process of producing the data for specific time intervals directly from the staged data. If you have defined Business Monitors that have Groups, aggregation is then the process of taking all the individual groups' time interval data and producing the overall time interval data, where overall is the combined data for all the groups.

Monitor Engine Threshold Settings

This section describes the parameters that you can modify through the **Monitor Engine Threshold settings** option.

The Monitor Engine Threshold Settings are divided into two logical sections. These sections appear on the Monitor Engine Threshold settings pane within the following categories:

➤ "Monitor Engine Threshold Checking Settings" on page 169

➤ "Monitor Engine Threshold Violation Notification Settings" on page 170

Monitor Engine Threshold Checking Settings

The following table lists the Monitor Engine Threshold checking settings that enable you to control how often the Monitor Engine polls its data in order to identify instance threshold information for alarms, notifications.

Descriptive Parameter Name	Description
Delay before raising violations on active instances (seconds)	The default of 300 seconds (5 minutes) should be fine for most installations.
Threshold minimum sample count, used for calculating relative thresholds	The minimum number of monitor instances required for the calculation of the relative scope thresholds. If the number of monitor instances available for sampling is less than this minimum figure, the instance threshold is not calculated. As a result, there are no notifications or alarms generated based on the threshold.
The maximum number of pooled threshold monitors (1 to 19)	The number of parallel threads available to the Monitor Engine for processing different Business Monitor Thresholds.

Monitor Engine Threshold Violation Notification Settings

The following table lists the Monitor Engine Threshold Violation Notification settings that enable you to control the number of, and frequency at which, notifications (email, HP Operations Manager messages and script invocations) are sent out when monitor instance threshold violations occur:

Descriptive Parameter Name	Description
Threshold violation notification polling interval (seconds)	The time interval (in seconds) that the Monitor Engine waits between the time that is sends the last notification message at the end of one polling period, to the start of the next polling period.
Maximum number of individual violation notifications per threshold per polling interval	The maximum number of violation notifications sent to the Notification Server for each business monitor instance threshold. The actual number of threshold violations generated might be significantly more than this maximum; however, if you are using the default Velocity email templates, each email notification also has summary information, which includes all the instance threshold violations generated during the polling interval.

When an instance threshold is violated, the Monitor Engine database records the detail of the violation. You can then be notified of these instance threshold violations in a number ways; for example, email; this is done using the Notification Server.

Where there is a significant failure in your business, for example, an application becomes unavailable, you can have the situation where thousands of thresholds are violated, or a small number of thresholds are violated many times. This can result in large numbers (or a storm) of notifications being generated. In the case of email, this can quickly fill up your email Inbox and it is possible that your email Server is unable to keep up with the rate at which the notifications are generated. This in turn creates a backlog of notifications and can severely impact the performance of the BPI system. In the case of HP Operations Manager messages or scripts, the backlog created depends on how your system is configured to receive the notifications.

Use the parameters to control the frequency and numbers of notifications that you receive. This reduces the possibility of these notification storms occurring, and still notifies users of the monitor instance threshold violations.

The following is an example of the effect of setting these parameters:

> The parameter values set when you install BPI for the first time are:
>
> ➤ 300 seconds for the polling interval
>
> ➤ a maximum of 10 notifications for each instance threshold violated
>
> This means that you receive a maximum of 10 notifications for each instance threshold violation, every 5 minutes (300 seconds); this is for each instance threshold. If there are 20 thresholds violated, there are potentially 200 (10 x 20) notifications in each 5 minute polling interval.

If you know the number of instance threshold violations your business is likely to generate, you can set these parameters such that:

➤ you receive all notifications for the occasional problems.

➤ you have the maximum instance threshold set such that when there is a system failure, notification storms are prevented and you receive a summary instead.

This also has the effect of reducing the impact of notification storms on BPI performance.

Monitor Engine Instance Cleaner Settings

This section describes the parameters that you can modify through the Monitor Engine Instance Cleaner settings option.

The Monitor Engine instance cleaner settings section controls how often the Monitor Engine instance thread is run and therefore how often Active and Completed monitor instances, Monitor violation instances and statistics are deleted.

Note: If you set a low Collection interval for your Business Process monitor definitions, the monitor data can accumulate very quickly and consume a considerable amount of disk space. In addition, if you shutdown your BPI system and restart it after a significant amount of time, the monitor statistics are calculated for the period of the shutdown as soon as the BPI Server is restarted. As an example, if you set a Collection interval of five minutes and shut down your BPI system for a week. When you restart your BPI Server, the Monitor Engine calculates all the monitor statistics for each five-minute period since the last time the monitors were calculated. This can result in a significant amount of calculation time, and a significant amount of disk space for the results.

If these results are immediately deleted as a result of the instance cleaner settings, this can also have an impact on the overall operational performance of your machine.

The Collection interval is set as the value for the Data Samples period.

The configuration parameters are divided into a number of logical sections. These sections appear in the Monitor Engine Instance Cleaner settings pane within the following categories:

➤ "Monitor Engine Instance Cleaner Settings" on page 173

➤ "Active Monitor Instances Settings" on page 173

➤ "Completed Monitor Instances Settings" on page 174

➤ "Monitor Statistics Settings" on page 174

➤ "Monitor Violation Instances Settings" on page 175

The parameters relating to these settings are described in the following sections.

Monitor Engine Instance Cleaner Settings

The following setting controls how often the Monitor Engine instance cleaner thread is executed. This interval impacts all the other sections on the Monitor Instance Cleaner pane

➤ Instance cleaner execution interval (minutes)

The time period (in minutes) that you want the Monitor instance cleaner thread to run. When it runs, it deletes all eligible Active monitor instances, Completed monitor instances, Monitor Alarm instances and Statistics. This thread also deletes all monitor instances for monitors that have been marked as deleted in the Monitor Definer.

Active Monitor Instances Settings

The following table lists the settings that enable you to control the amount of Active monitors instance data held in the BPI database tables.

Descriptive Parameter Name	Description
Delete active monitor instances from the database?	Select this check box if you want to delete Active monitor instances from the BPI database tables. Clear the check box if you do not want to delete the Active monitor instances.
Age of active monitor instances to be removed (days)	This option is available when you select the check box for deleting Monitor violation instances. Enter the age (in days) of the Active monitor instances that you want to be deleted when the Monitor Engine instance cleaner thread is run.

Completed Monitor Instances Settings

The following table lists the settings that enables you to control the amount of Completed monitors instance data held in the BPI database tables.

Descriptive Parameter Name	Description
Delete completed monitor instances from the database?	Select this check box if you want to delete Completed monitor instances from the BPI database tables. Clear the check box if you do not want to delete the Completed monitor instances.
Age of completed monitor instances to be removed (days)	This option is available when you select the check box for deleting Completed monitor instances. Enter the age (in days) of the Completed monitor instances that you want to be deleted when the Monitor Engine instance cleaner thread is run

Monitor Statistics Settings

The following table lists the settings that enable you to control the amount of statistical data held in the BPI database tables.

Descriptive Parameter Name	Description
Delete monitor statistics from the database?	Select this check box if you want to delete statistics from the BPI database tables. Clear the check box if you do not want to delete the statistics data.
Age of monitor statistics to be removed (days)	This option is available when you select the check box for deleting statistics. Enter the age (in days) of the statistics data that you want to be deleted when the Monitor Engine instance cleaner thread is run

Monitor Violation Instances Settings

The following table lists the settings that enables you to control the amount of Monitor Violation instance data held in the BPI database tables.

Descriptive Parameter Name	Description
Delete monitor violation instances from the database?	Select this check box if you want to delete Monitor Violation instances from the BPI database tables. Clear the check box if you do not want to delete the Monitor Violation instances.
Age of monitor violation instances to be removed (days)	This option is available when you select the check box for deleting Completed monitor instances. Enter the age (in days) of the Monitor Violation instances that you want to be deleted when the Monitor Engine instance cleaner thread is run

Modifying the Monitor Engine Settings

To change any of the Monitor Engine settings, complete the following procedures on the Windows system, where the BPI Server is installed:

1 Start the Administration Console by accessing **Start** > **Programs** > **HP** > **HP Business Process Insight** > **Administration**.

2 Select the **Component Configuration** > **Monitor Engine** > <**Monitor Engine_Settings_option**> from the Navigator pane on the BPI Administration Console.

3 Enter the changes that you want to make to the **Monitor Engine** > <**Monitor Engine_Settings**> option as appropriate in the right-hand pane.

4 Click **Apply** to apply the changes to the property files.

5 Select the **Status** option to view the Component Status panel of the BPI components.

6 Stop and restart all the BPI Server components using the **Stop All** and **Start All** buttons.

The new settings are now applied to your BPI system.

Component Configuration - HP Business Service Management Settings

The **HP Business Service Management settings** options enable you to configure the details of the Gateway Server system (or Load Balancer) in order that the BPI Server can send data samples to the Gateway Server.

The **HP Business Service Management settings** page is divided into the following logical sections:

➤ "Business Service Management General Settings" on page 177.

These are the properties relating to the connection details for the Business Service Management server that you want BPI to use and are required for both the Data Samples settings and for CI Poller settings.

➤ "Data Samples Settings" on page 180.

These are the properties relating to the data that you want BPI to send to Business Service Management.

➤ "CI Poller Settings" on page 184.

These are the properties that enable BPI to receive status information relating to IT operational resources, which are configured for BPI CIs, within Business Service Management.

To configure the properties of the Business Service Management integration, make sure that the **HP Business Service Management settings** option is selected in the Administration Console navigator pane. The parameters are displayed in the right-hand pane.

You can modify these settings, see "Modifying the Notification Server Parameters" on page 146 for instructions.

Business Service Management General Settings

The following table describes the settings related to connection details.

Descriptive Parameter Name	Description
Hostname of Gateway server	The fully qualified DNS name of the machine where the Business Service Management Gateway Server that receives data from the BPI Data Collectors is installed. This is the Gateway Server that provides the status information for Business Process Activity KPIs or access to the ODB for BPI. This is the Gateway Server to which the BPI Data Samples are sent.
	Alternatively, this could be the name of a Load Balancer if there is one configured.
	This value can also be an IP address.
Use SSL?	Select the checkbox to enable HTTP over SSL for BPI:
	➤ Data Sample Provider
	➤ ODB synchronizer
	➤ component obtaining status information for process Activity KPIs
	Clear the checkbox if you do not want to use HTTP over SSL. This parameter also requires that Business Service Management is configured to use SSL; refer to the Business Service Management help documentation for details of how to set this up.

Descriptive Parameter Name	**Description**
HTTP(S) port of Gateway server	The port number used by the Web Server on the Business Service Management Gateway Server for Data Collectors. This is usually port 80 if SSL is not enabled, or port 443 is SSL is enabled.
User name	The user name for a Business Service Management account, which is for use by BPI to send data samples. This can be an existing user, or a new user, as long as it has permissions to access the View name as specified in "CI Poller Settings" on page 184. For a new installation the view name is Business Processes.
Password	The password set for User name. If you have created your own user, enter the appropriate password details for this user name.

Descriptive Parameter Name	Description
Basic authentication enabled?	Select the checkbox to enable basic authentication. Clear the checkbox to disable basic authentication.
	Select this option only when Business Service Management is configured to use basic authentication for Web Server and Web browser communications.
	Basic authentication is the BSM mechanism for enabling a Web browser (or Web client) to authenticate with the Web Server using a user name and a password; refer to the Business Service Management help documentation for details of basic authentication.
Basic authentication user	User name configured for basic authentication credentials.
Basic authentication password	Password configured for basic authentication credentials.
Maximum retries for synchronization deadlock	Maximum number of times that you want BPI to attempt to send data to the Gateway Server.
	This parameter ensures that, in the case of a deadlock, the synchronization attempt to the Gateway Server is aborted in order to break the deadlock.
	Failed synchronizations are automatically retried at regular intervals.

Data Samples Settings

The following table describes the settings related to the data that you want
to send to Business Service Management.

Descriptive Parameter Name	Description
Data samples send interval (minutes)	The time interval (in minutes) that you want BPI to use to send data samples to the named BSM Gateway Server or Load Balancer. Select a time from the drop-down list. The selection of time intervals is aligned with the time periods for the Business Service Management SLM calculations. in order to provide the most accurate results. This time interval is the measurement interval for all BPI CI types.
Connection timeout (seconds)	The period of time (in seconds) allowed for the BPI Server to request a connection with the BSM Gateway Server when sending data samples. If this parameter is set to zero (0), the time allowed for requesting a connection is unlimited, that is, no timeout occurs. If a timeout does occur, BPI attempts to resend the data sample after waiting for the interval configured for **Data samples send interval**.

Descriptive Parameter Name	Description
Guaranteed retention period (minutes)	The period of time (in minutes) that a Data Sample is guaranteed to be made available for transfer to the BSM Gateway Server and before it is discarded. In cases where it is not possible to make a connection to the BSM Gateway Server, or where the Gateway Server's buffer tables are full, BPI retains copies of each Data Sample for the period specified by this parameter.
Delay before sending samples (minutes)	Delay (in minutes) following the Data samples send interval to allow time for the Business Impact Engine and the Monitor Engine to collate data to be sent to Business Service Management. Setting a value for this parameter enables you to take account of the fact that there might be a time difference between when a business event is recorded by a business application and when BPI receives notification of the event from the application; this is because some business applications do not make data available immediately.

Descriptive Parameter Name	Description
Delay before sending samples (minutes) (cont'd)	Setting a delay enables BPI to collect complete and accurate statistics for each send interval before sending data samples to the Gateway Server. It also means that BPI data shown within BSM lags behind the Last Update time by the value of this delay; for example, if the Last Update time in BSM is 10:30 and you set a delay of 900 seconds (15 minutes), the values you see for the BPI data samples displayed within the BPI Application and elsewhere in BSM related to the collection period ending at 10:15.
	Note: If you configure a delay that is more than 90 seconds, you should also increase the Decay setting for BPI samples in BSM (see the BSM Service Health Help). If you do not, a No Data status is shown for the BPI Application instance data.

Descriptive Parameter Name	Description
Send updates to earlier samples?	If there are occasions when BPI events arrive even later than the interval specified by the **Delay before sending samples** setting, you can enable updates to be sent to earlier samples. Updates are only sent if the original sample data is found to be out-of-date when it is re-checked later. **Note**: Enabling this option can adversely affect the performance of both the BPI and BSM servers. Use the **Delay before sending samples** setting to incorporate delayed BPI events as far as possible and only enable this option if you need to capture exceptional late events.
Further delay before sending updates (minutes)	If you have enabled the **Send updates to earlier samples** option, then use this field to enter the delay before which earlier BPI data samples are re-calculated and updated if necessary. For example, if the delay before sending samples is 5 minutes and this further delay value is set to 20 minutes, the BPI samples for 10:00 are first calculated and sent at 10:05. They are then re-calculated and updated if necessary at 10:25. The value should be between 5 and 120 minutes.

CI Poller Settings

The following table describes the settings related to the status information that you want to receive from Business Service Management.

Descriptive Parameter Name	Description
View name	The name of the Business Service Management Service Health view that contains the Business Process information (IT operational resource information in BPI terms), required for display in the BPI Application. The default view for the Business Service Management Service Health is the Business Processes view.
Status event poll interval (minutes)	The time interval (in minutes) that you want to be used for the CI Poller to poll the Business Service Management Gateway Server (or Load Balancer) for status events.
Connection timeout (seconds)	The period of time allowed for the BPI Server to request Service status information from the BSM Gateway Server (or Load Balancer) when requesting information relating to services. If this parameter is set to zero (0), the time allowed for the connection is unlimited, that is, no timeout occurs. If a timeout does occur, the BPI Server waits for the interval configured for **Status event poll interval** before issuing another request for status events.

Modifying the Business Service Management Settings

To change the HP Business Service Management settings, complete the following procedures on the Windows system, where the BPI Server is installed:

1 Start the Administration Console by accessing **Start** > **Programs** > **HP** > **HP Business Process Insight** > **Administration**.

2 Select **Component Configurations** > **HP Business Service Management settings** option from the Navigator pane on the BPI Administration Console.

3 Enter the changes that you want to make to the settings as appropriate in the dialog presented.

4 Click **OK** to dismiss the dialog.

5 Click **Apply** to apply the changes to the BPI configuration.

You do not need to stop and start the components after changing these settings as the configuration changes are made automatically.

The new settings are now applied to your BPI system.

Component Configuration - Process Repository

The Process Repository configuration has one of the following sections. The section that appears depends on the database type for the Process Repository Database and is related to the option that you selected and configured when you installed the BPI Server:

➤ "MS SQL Server Access to the Process Repository Database" on page 186

➤ "Oracle Server Access to the Process Repository Database" on page 187

Note: Changing the database connection details relating to the Process Repository does not automatically migrate data that is stored in the Process Repository database. You should not reconfigure the database connection unless you have already used the database tools to migrate your data.

The parameters relating to these settings appear on the right-hand pane of the Administration Console when you select one of these options in the console's navigation tree.

To change the password for the Process Repository database user that you created during the BPI Server installation, follow the instructions on "Changing the Password Details" on page 188.

MS SQL Server Access to the Process Repository Database

The following table lists the parameters that you can modify for an MS SQL Server Database.

Descriptive Parameter Name	Description
MS SQL Server hostname	The hostname of the system where the Microsoft SQL Server Database for the Process Repository is installed.
MS SQL Server port number	The port number that the Microsoft SQL Server Database is expecting connections on.
MS SQL Server database name	The name of the database configured for the Process Repository database.
MS SQL Server authentication	Select the type of authentication that you want to use when BPI connects to the Process Repository database. You can select either SQL Server authentication, or Windows authentication.
	If you select SQL Server authentication, the MS SQL Server login user and MS SQL Server login password options are enabled, as you need to provide a user name and password for the SQL Server authentication.
	If you select Windows authentication, no additional information is required.

Descriptive Parameter Name	Description
MS SQL Server user name	The name of the user who has access privileges to the database schema for the Process Repository. This field is enabled only when you have selected SQL Server Authentication.
MS SQL Server password	The password for the user who has access privileges to the database schema for the Process Repository. This field can be used to re-synchronize the BPI Server version of the password with the Process Repository database password after the database password is changed using the database management tools. This field is enabled only when you have selected SQL Server Authentication.

To change the password for the database, follow the instructions in "Changing the Password Details" on page 188.

Oracle Server Access to the Process Repository Database

The following table lists the parameters that you can modify for an Oracle database

Descriptive Parameter Name	Description
Oracle Server hostname	The hostname of the system where the Oracle Server Database for the Process Repository Database is installed.
Oracle Server port number	The port number that the Oracle Server Database is expecting connections on.

Descriptive Parameter Name	Description
Oracle Server SID	The SID for the database configured for the Process Repository.
Oracle Server schema user	The name of the user who has access privileges to the database schema for the Process Repository.
Oracle Server schema password	The password for the user who has access privileges to the database schema for the Process Repository. This field can be used to re-synchronize the BPI Server version of the password with the Process Repository database password after the database password is changed using the database management tools.

Changing the Password Details

To make the changes to the database password, complete the following procedures:

1 Start the Administration Console by accessing **Start** > **Programs** > **HP** > **HP Business Process Insight** > **Administration**.

2 Stop all the BPI components using the **Stop All** button.

3 Select the **Component Configurations** > **Process Repository** > *<server_type>* **Server Access to the Repository Database** option on the BPI Administration Console and make the changes to the following field:

login password

4 Click **Apply** to apply the changes to the BPI configuration.

5 Return to the **Status** option and restart all the BPI components, using the **Start All** button.

The new settings are now applied to your BPI Server system.

Component Configuration - Business Event Handler

The **Business Event Handler** options enable you to configure the parameters for the Adaptor (based on **openadaptor**) that sends and receives events between the Business Impact Engine and the Business Event Handler.

These parameters are listed in the following table:

Descriptive Parameter Name	Description
Log package information?	A check box that when checked means that the Business Event Handler logs package information.
Log thread information?	A check box that when checked means that the Business Event Handler logs thread information.
Log time information?	A check box that when checked means that the Business Event Handler logs time information.
Maximum number of socket source threads	Maximum number of threads that can be created for the adapter socket source. When this instance threshold is exceeded, the request for a thread is refused and the Business Event Handler reports a **Connection Refused** error in its log file.
Maximum number of retries to deliver events into Business Impact Engine	The maximum number of times the Business Event Handler attempts to deliver an event to the Business Impact Engine (using RMI) before rolling back the transaction for the business event.
Business Impact Engine event retry delay (seconds)	Time (in seconds) between each attempt to retry an event from the Business Event Handler to the Business Impact Engine using RMI.

Descriptive Parameter Name	Description
Maximum number of retries to deliver events from Hospital	The maximum number of times the Business Event Handler attempts to deliver an event from the Event Hospital to the Business Impact Engine.
	The Business Impact Engine marks particular categories of event errors to result in the event being sent to the event Hospital and marked to be automatically discharged; for example, events that are received out of sequence before a data instance has been created for the event.
	If the Business Event Handler does not succeed in delivering the event to the Engine in within the specified number of retries, the event remains in the event Hospital.
Hospital event poll interval (seconds)	Time (in seconds) that the Business Event Handler polls the Event Hospital looking for events that have been marked as ready for discharge and that can be delivered to the Business Impact Engine.

Modifying the Business Event Handler Parameters

To change these settings, complete the following procedures on the Windows system, where the BPI Server is installed:

1 Start the Administration Console by accessing **Start** > **Programs** > **HP** > **HP Business Process Insight** > **Administration**.

2 Select the **Business Event Handler** option from the Navigator pane on the BPI Administration Console.

3 Enter the changes that you want to make to the settings as appropriate in the right-hand pane.

4 Click **Apply** to apply the changes to the BPI configuration.

5 Select the **Status** option to go to the Component Status panel of the BPI components.

6 Stop and restart all the BPI Server components using the **Stop All** and **Start All** options.

You can stop and restart the **Business Event Handler** component; however, you are advised to use **Stop All** and **Start All** to ensure that the BPI components are stopped and started in the correct order.

The new settings are now applied to your BPI system.

Component Configuration - JMS Business Event Handler

Use the **JMS Business Event Handler** options to configure the parameters for your JMS connections., which enables BPI to receive business events from a JMS-based system.

These options appear hierarchically in the Administration Console, under the JMS Business Event Handler navigation option as follows:

➤ "General Settings" on page 192

➤ "Service Providers" on page 193

➤ "Event Sources" on page 195

The parameters relating to these settings appear on the right-hand pane of the Administration Console when you select one of these options in the console's navigation tree. If these options are not visible, you need to expand the entries under **JMS Business Event Handler** using the usual Explorer style navigation techniques.

Note: BPI supports the JMS message type **javax.jms.TextMessage**, where the content of the text is an XML file in a particular format. This means you must be integrating with an application that sends and receives this message type if you want to use the JMS Business Event Handler. One example of such an application is HP Business Service Management TransactionVision.

General Settings

The following table describes the parameters that you can set to configure the behavior of your JMS connection:

Descriptive Parameter Name	Description
Log package information?	A check box that when checked means that the JMS Business Event Handler logs package information.
Log thread information	A check box that when checked means that the JMS Business Event Handler logs thread information.
Log time information?	A check box that when checked means that the JMS Business Event Handler logs time information.
Maximum number of retries to deliver events into Business Impact Engine	The maximum number of times the JMS Business Event Handler attempts to deliver an event to the Business Impact Engine (using RMI) before rolling back the transaction for the business event.
Business Impact Engine event retry delay (seconds)	Time (in seconds) between each attempt to retry an event from the JMS Business Event Handler to the Business Impact Engine using RMI.

Descriptive Parameter Name	Description
Maximum number of retries to deliver events from Hospital	The maximum number of times the JMS Business Event Handler attempts to deliver an event from the Event Hospital to the Business Impact Engine.

The Business Impact Engine marks particular categories of event errors to result in the event being sent to the event Hospital and marked to be automatically discharged; for example, events that are received out of sequence and where there is no process or data instance created for the event.

If the JMS Business Event Handler does not succeed in delivering the event to the Engine in within the specified number of retries, the event remains in the event Hospital. |
| Hospital event poll interval (seconds) | Time (in seconds) that the JMS Business Event Handler polls the Event Hospital looking for events that have been marked as ready for discharge and that can be delivered to the Business Impact Engine. |

Service Providers

This section describes the parameters that you can set to specify the details of the service providers that you are using for your JMS connections to BPI. You must add a Service Provider before you can configure a JMS Event Source. Note that there is a Service Provider already configured for a standard TransactionVision integration.

From the list of JMS Service Providers, you can:

➤ select HP TransactionVision from the list, if you want to configure an integration with TransactionVision. Click the Modify option to make changes to the TransactionVision JMS Service Provider parameters.

➤ click the **Add** button to add a new Service Provider to the list of available providers.

The JMS Service Provider Properties dialog opens and the properties for adding a new service provider are described in the following table:

Descriptive Parameter Name	Description
Service Provider Name	Unique name that identifies the service provider within the BPI interfaces for administrative and identification purposes. The name can be up to 40 characters. In the case of TransactionVision, this name is predefined as HP TransactionVision.
Description	An optional description of service provider.
Enabled?	Select the check box to enable this service provider to send JMS Business Events to BPI. Clear the check box to disable the service provider.
Service Provider Classpath	Location of the Service Provider client Java classes.
JNDI Provider Location	Location of the JNDI service for the named Service Provider, typically this is a URL. This parameter is predefined as tcp:// analyzerhost:21110 for the HP TransactionVision service provider. You need to replace analyzerhost with the host name of machine where the TransactionVision Analyzer is installed.

Descriptive Parameter Name	Description
JNDI Factory Class	The Java class name for the initial context factory for the JNDI service.
	This parameter is predefined as **com.sonicsw.jndi.mfcontext.MFCon textFactory** for the HP TransactionVision service provider. If this value is not correct for your implementation, you can modify it.
JMS Factory Class	The Java class name for the initial context factory for the JMS service provider.
	This parameter is predefined as **BPIQueueFactory** for the HP TransactionVision service provider. If this value is not correct for your implementation, you can modify it.

Event Sources

This section describes the parameters that you can set to specify the details of the event sources that you want to access from your service provider. These are the event sources that will supply the transaction events that you have configured for your processes.

Note: You cannot add a new Event Source until you have configured at least one Service Provider.

There is a preconfigured JMS Event Source for HP TransactionVision, or you can click the Add button to add a new Event Source to the list.

The JMS Event Source Properties dialog opens and the properties for adding a new Event Source are described in the following table:

Descriptive Parameter Name	Description
Service Provider Name	Select the name of the service provider that you have configured to source the events that you require.
JMS Queue Subject	The name of the JMS queue created for BPI at the selected Service Provider.
	This parameter is predefined as TV2BPI.EVENT.QUEUE for the HP TransactionVision Queue Subject. You cannot modify this parameter, you must create a new JMS Event Source with a new name and Queue Subject if required.
User Name	If required by the Service Provider, the name of the user with authority to access the queue identified by JMS Queue Subject.
	Not all JMS Service Providers require a user name to be specified.
Password	If required, the password for User Name.

Modifying the JMS Business Event Handler Settings

To change these settings, complete the following procedures on the Windows system, where the BPI Server is installed:

1 Start the Administration Console by accessing **Start** > **Programs** > **HP** > **HP Business Process Insight** > **Administration**.

2 Select the **JMS Business Event Handler** option from the Navigator pane on the BPI Administration Console.

3 Enter the changes that you want to make to the settings as appropriate in the right-hand pane.

4 Click **Apply** to apply the changes to the BPI configuration.

5 Select the **Status** option to go to the Component Status panel of the BPI components.

6 Stop and restart all the BPI Server components using the **Stop All** and **Start All** options.

You can stop and restart the **JMS Business Event Handler** component. However, you are advised to use **Stop All** and **Start All** to ensure that the BPI components are stopped and started in the correct order.

The new settings are now applied to your BPI system.

Component Configuration - Servlet Engine

Use the Servlet Engine options to configure the parameters that relate to how the process instance pages are displayed within the BPI Application and the security of the pages served by the Servlet Engine.

These options appear hierarchically in the Administration Console, under the **Servlet Engine** navigation option as follows:

➤ "Process Instance Web Pages and Portlets" on page 197

➤ "Security" on page 200

Process Instance Web Pages and Portlets

This option enables you to manage the parameters of the BPI process instance pages, instance violation pages and the BPI portlets. BPI portlets can be found under My BSM and on the **Applications** > **Business Process Insight** > **Health** > **Business Health** page.

The parameters relating to these settings appear on the right-hand pane of the Administration Console when you select the **Servlet Engine** > **Process Instance Web Pages and Portlets** in the console's navigation tree. If the **Process Instance Web Pages and Portlets** options are not visible in the navigation tree, expand the entries using the Explorer-style navigation techniques.

You can also control the security settings for the **Process Instance Web Pages and Portlets** as described in section "Security" on page 200.

The Process Instance Web Pages and Portlets option enables you to configure parameters related to the behavior of the Process Instance pages, Violation Pages and portlets within the Business Process Insight Application Health pages and My BSM.

The Process Instance page and Portlet settings parameters are listed in the following table:

Descriptive Parameter Name	Description
Process Instance Web Pages refresh delay (seconds)	The time interval (in seconds) that the Process Instance pages use to set the refresh interval for the Web browser.
Maximum number of retries on database deadlock	The maximum number of attempts by the Process Instance pages to retry a database transaction before aborting it. It is possible for deadlocks to occur when the Process Instance pages and other applications are accessing the BPI database simultaneously. This parameter ensures that, in the case of a deadlock, the Process Instance pages abort the transaction in order to break the deadlock and generates an error message. which is displayed on a Web page. If the Process Instance pages abort the transaction, use the Refresh button on your browser to reload the page.
Database deadlock retry delay (seconds)	The time (in seconds) between each attempt to retry a database transaction.

Descriptive Parameter Name	Description
Show service state for a completed activity?	A check box indicating whether or not to show the status of services for Activities that have already been completed in the Business Process.
Include data properties when exporting process lists as CSV?	A checkbox indicating whether or not to include the properties of a Data instance when exporting Process instances, using the Export All option, within the Process Instance pages. This is the option, on the Process Instance pages, to export the current process instance list to a file in a comma-separated value (CSV) format.

Modifying the Process Instance Web Pages and Portlet Settings

To change these settings, complete the following steps on the Windows system, where the BPI Server is installed:

1 Start the Administration Console by accessing **Start** > **Programs** > **HP** > **HP Business Process Insight** > **Administration**.

2 Select the **Servlet Engine** > **Process Instance Web Pages and Portlets** option from the Navigator pane on the BPI Administration Console.

3 Enter the changes that you want to make to the settings as appropriate in the right-hand pane.

4 Click **Apply** to apply the changes to the BPI configuration.

5 If you have modified the **Show service state for a completed activity?** parameter, you need to stop and restart the Servlet Engine as described in step 6. If you have not modified this parameter, continue at step 7.

6 Select the **Status** option to go to the Component Status panel and Stop and restart the Servlet Engine component using the **Stop All** and **Start All** options.

7 If you have modified the **Page refresh delay (seconds)** parameter, and the parameter was previously set to zero (0), you need to refresh the Web browser window where the Process Instance page is active.

The new settings are now applied to your BPI system.

Security

The **Security** option enables you to set the type of authorization used between BSM and the BPI Server.

Important: Each time a BPI Server component is started, the BPI Server polls BSM for the lightweight single-signon settings and sets the BPI Server parameters based on the values set for BSM. These settings overwrite any settings that you have entered, using the Administration Console; you are therefore advised not to modify the Authorization settings and to leave the BPI Server to set them, based on the BSM values, when you start components.

The parameters relating to security appear on the right-hand pane of the Administration Console when you select Security in the Administration Console's navigation tree.

The right-hand pane is divided into the following logical sections:

➤ "Authorization Settings" on page 200
➤ "HTTP Over SSL Settings" on page 202

Authorization Settings

The Authorization settings hold the values obtained for the lightweight single-signon properties within BSM.

If you have installed the BPI Server in a domain that is different to the domain where the Gateway Server is installed, make sure that you add the BPI domain to the list of Protected Domains, under:

Admin > **Platform** > **Users and Permissions** > **Authentication**

The following table lists the **Authorization** parameters.

Descriptive Parameter Name	Description
Authorization method	Use one of the following settings for this option to select which method, or product, that you want to use to authorize your BPI components: ➤ **Use BSM lightweight single sign-on** ➤ **Use Servlet Engine** This parameter is set at installation and reset each time you restart a component. Do not modify the parameter as it will be overwritten.
Domain name of local host	The domain name for the Business Service Management Gateway server. This is a fully qualified domain name for the Gateway Server host, named in the Business Service Management settings in section "Component Configuration - HP Business Service Management Settings" on page 176. This parameter is set at installation and can be modified if required.
Single sign-on token creation key (Init string)	This parameter must be manually retrieved from the BSM user interface (**Admin** > **Platform** > **Users and Permissions** > **Authentication Management** > **Token Creation Key**) and copied to the BPI server. This setting allows you to set or change this parameter after installing the BPI Server (see "Business Service Management Information" on page 92).
Bypass security for Process Instance Web Pages and Portlets?	This parameter is set at installation and reset each time you restart a component. Do not modify the parameter as it will be overwritten.

HTTP Over SSL Settings

HTTP over SSL (secure socket layer) is a secure version of HTTP. SSL enables the data from a client, such as a Web browser, to be encrypted prior to transmission to prevent unauthorized access to the data being transmitted.

See Chapter 10, "Servlet Engine Authentication" for more information about BPI and HTTPS.

The following table lists the properties that you can configure to encrypt the data that you send to the BPI Web-based clients.

Descriptive Parameter Name	Description
Enable HTTP over SSL for BPI Web Pages?	Select the checkbox to enable HTTP over SSL from the Web Server to the Web Browser, or Web Browsers, for the BPI: ➤ Process instance pages, instance violation pages and the BPI portlets. ➤ Process Repository Explorer ➤ Monitor Definer ➤ Process Intervention ➤ Notification The setting applies to all Web Browsers where these pages are running for a specific installation of the Servlet Engine. Clear the checkbox if you do not want to use HTTP over SSL.
SSL Keystore	This is the location of the signed certificate file required for the HTTPS connection.

Descriptive Parameter Name	Description
SSL Keystore Password	This is the password for the signed certificate that you supply.
Enable Certificate-based Client Authentication?	Select the checkbox to enable certificate-based client authentication (from the Web Browser to the Web Server). Clear the checkbox if you do not want to use this form of authentication.

Modifying the Security Parameters

To change these settings, complete the following procedures on the Windows system, where the BPI Server is installed:

1 Start the Administration Console by accessing **Start** > **Programs** > **HP** > **HP Business Process Insight** > **Administration**.

2 Select the Security option from the Navigator pane on the BPI Administration Console.

3 Enter the changes that you want to make to the settings as appropriate in the right-hand pane.

4 Click **Apply** to apply the changes to the BPI configuration.

5 Select the **Status** option to go to the Component Status panel where the BPI component status are shown.

6 Stop and restart all the BPI Server components using the **Stop All** and **Start All** options. If you do not start all the components, some of the security settings can be compromised.

The new settings are now applied to your BPI system.

Component Configuration - MS SQL Server Access

The **MS SQL Server Access** parameters are available only when you select MS SQL Server as the database from within the BPI installation procedure. If you select Oracle as your database, the MS SQL Server parameters are not available through the Administration Console.

The **MS SQL Server Access** option enables you to view the current configurations for the SQL Server access and modify the password for the database access.

The database is used by the following components:

➤ Business Impact Engine

➤ Monitor Engine

➤ Business Event Handler

➤ JMS Business Event Handler

➤ Notification Server

➤ Web pages hosted within the Servlet Engine, for example, the Process Intervention pages.

The MSSQL Access parameters are listed in the following table.

Descriptive Parameter Name	Description
MS SQL Server hostname	This is a non-modifiable field and shows the fully qualified domain name of the system where the Microsoft SQL Server Database is installed.
MS SQL Server port number	This is a non-modifiable field and shows the port number that the Microsoft SQL Server Database is expecting connections on.
MS SQL Server database	This is a non-modifiable field and shows the name of the database configured for BPI.

Descriptive Parameter Name	Description
MS SQL Server login user	This is a non-modifiable field and shows the name of the user who has permission to read and write the BPI database files.
MS SQL Server login password	The password for the user with permission to read and write the BPI database files. This field can be used to re-synchronize the BPI version of the password with the database password after the database password is changed using the database management tools.

To change the password for the database, follow the instructions on section "Changing the Password Details" on page 205.

Changing the Password Details

This section describes how to change the password for the BPI database user that you created during the installation.

To make the changes to the database password, complete the following procedures:

1 Start the Administration Console by accessing **Start** > **Programs** > **HP** > **HP Business Process Insight** > **Administration**.

2 Stop all the BPI components using the **Stop All** button.

3 Select the **MS SQL Server Access** option on the BPI Administration Console and make the changes to the following field:

   MSSQL login password

4 Click **Apply** to apply the changes to the BPI configuration.

5 Return to the **Status** option and restart all the BPI components, using the **Start All** button.

You have now completed the tasks to change the database password details used by BPI components.

Component Configuration - Oracle Server Access

The Oracle Server Access parameters are available only when you select Oracle as the database from within the BPI installation procedure. If you select MS SQL Server as your database, the Oracle Server parameters are not available through the Administration Console.

The Oracle Access option enables you to manage the properties of the connection to the Oracle database.

The database is used by the following components:

➤ Business Impact Engine

➤ Monitor Engine

➤ Business Event Handler

➤ JMS Business Event Handler

➤ Notification Server

➤ Web pages hosted within the Servlet Engine, for example, the Process Intervention pages.

The Oracle Server Access parameters are listed in the following table:

Descriptive Parameter Name	Description
Oracle Server hostname	This is a non-modifiable field and shows the fully qualified domain name of the system where the Oracle Database is installed.
Oracle Server port number	This is a non-modifiable field and shows the port number that the Oracle Database is expecting connections on.
Oracle Server SID	This is a non-modifiable field and shows the name of the database sid configured for BPI.

Descriptive Parameter Name	Description
Oracle Server username	This is a non-modifiable field and shows the name of the user who has permission to read and write the BPI database files.
Oracle Server password	The password for the user with permission to read and write the BPI database files. This field can be used to resynchronize the BPI version of the password with the database password after the database password is changed using the database management tools.

If you are changing the password for the database user with permission to read and write the BPI database tables, follow the instructions on section "Changing the Password Details" on page 207.

Changing the Password Details

To change the password for the BPI database user, complete the following procedures:

1 Start the Administration Console by accessing **Start** > **Programs** > **HP** > **HP Business Process Insight** > **Administration**.

2 Stop all the BPI components using the **Stop All** button on the **Status** option.

3 Select the Oracle Access tab on the BPI Administration Console and make the changes to the following fields:

   Oracle Server login password

4 Click **Apply** to apply the changes to the BPI configuration.

5 Return to the **Status** option and restart all the BPI components, using the **Start All** button.

You have now completed the tasks to change the database password details used by BPI components.

Component Configuration - Port Numbers

This section describes the port numbers used by the BPI components. You might need to modify these if there is a port number clash, or if you have specific requirements with personal firewalls or other security processes.

The port number parameters are listed in the following table:

Descriptive Parameter Name	Description
RMI Registry	The port number for SUN's implementation of the Java RMI Activation System Daemon (RMID), which is used by BPI as a reliable RMI registry. The RMI registry holds the directory of port numbers used by the BPI components to receive RMI requests.
	If this port number is in use by another application, the BPI Server components are unable to start.
Administration Console Server	The port number that the server component of the Administration Console uses to listen for incoming administration requests.
Business Impact Engine SOAP Receiver	The port number used by the Business Impact Engine to receive incoming SOAP requests.
Business Impact Engine XML Receiver	The port number used by the Business Impact Engine to receive XML requests.
Business Impact Engine RMI Receiver	The port number used by the Business Impact Engine to receive incoming RMI requests.
Monitor Engine RMI Services	The port number used by the Monitor Engine to receive connections to RMI services.

Descriptive Parameter Name	Description
Business Event Handler Source	The port number used by the Business Event Handler Engine Adapter to receive incoming Events into the Business Event Handler.
Business Event Handler Control	The port number used to stop and start the Business Event Handler Engine Adapter programmatically; see the *Business Process Insight Integration Training Guide - Business Events*, for more details.
JMS Business Event Handler Control	The port number used to stop and start the JMS Business Event Handler Engine Adapter programmatically; see the *Business Process Insight Integration Training Guide - Business Events*, for more details.
Process Repository	The port number used by the Process Repository to listen for incoming RMI requests.
Servlet Engine HTTP	The port number that the Tomcat Web Server uses to receive HTTP requests.
Servlet Engine HTTPS	The port number that the Tomcat Web Server uses to receive HTTP over SSL encrypted requests.
Servlet Engine Shutdown	The port number used by the Tomcat Servlet Engine to shutdown.
Servlet Engine AJP 1.3	The AJP 13 protocol is a packet-based protocol that allows a Web server to communicate with the Tomcat JSP/Servlet Container over a TCP connection.

Descriptive Parameter Name	Description
Servlet Engine Admin	The port number used by the Tomcat Servlet Engine administration tools.
Web Services Provider	Not used in BPI.
Data Samples Provider RMI Services	The port number defined for the Data Samples Provider and used by the Business Impact Engine and Monitor Engine to send data samples to Gateway Server (or Load Balancer).

If you want to modify a port number, see "Modifying BPI Port Numbers" on page 210.

Modifying BPI Port Numbers

The BPI components use a number of different port numbers for connections. There should not be any need to modify these port numbers unless you know that other applications are already using them, or if you are using personal firewall software that filters incoming connections based on port number. You must stop all the BPI components before you can modify the BPI port numbers.

To change the port number for the BPI Server components, complete the following procedures:

1 Start the Administration Console by accessing **Start** > **Programs** > **HP** > **HP Business Process Insight** > **Administration**.

2 Select the **Status** option and stop the BPI components using the **Stop All** button.

   You cannot modify any of the port numbers until you have stopped the BPI components.

3 Select the **Port Numbers** option in the Navigator pane.

4 Make the changes to the port numbers that you want to modify.

5 Click **Apply** to apply your modifications.

6 Restart the BPI components that you previously stopped using **Start All** in the Component Status panel.

The port number changes have now been applied.

Component Configuration - Logging

This section describes the parameters that you can set for the BPI component log levels. For details, see "Logging Parameters" on page 211.

Logging Parameters

BPI includes a number of third-party products, which use their own logging mechanisms, for example:

➤ Business Event Handler component, which is based on the openadaptor framework, uses Log4j as its logging mechanism.

➤ Hibernate, which is used by the Business Impact Engine and other BPI Server components, uses Jakarta Apache Commons logging.

Log4j and Java J2SE logging are two alternative implementations and Apache Commons is an layer above a specific logging implementation. In the case of BPI, Apache Logging is a layer over the Java J2SE logging implementation.

This Logging option enables you to set logging levels for the BPI components, and for the Business Event Handler; the Business Event Handler logging is based on openadaptor. There is no mechanism for setting the logging levels for Hibernate using the Administration Console, as the logging for the components using Hibernate is usually sufficient.

The logging parameters, which can be set through the Logging option, are listed in the following table. For your BPI implementation, you see only those logging parameters that relate to the components that you have installed.

Descriptive Parameter Name	Description
Change the log viewer application	
View log files with user-defined viewer?	If unchecked, BPI uses the system-defined application associated with *.log.
	If checked, BPI uses the application specified in the next line.
Path for the log viewer application	The full path name for the application that is being used to display the log files.
	If the application is defined as being on your PATH, you can enter the application name without its full path name.
Logging Levels - HP BTO Software Logging	
The Administration Console log level	The log level set for the Administration Console.
The Administration Console Server log level	The log level set for the Administration Console Server.
The Monitor Engine log level	The log level set for the Monitor Engine
The Business Impact Engine (BIE) log level	The log level set for the Business Impact Engine.
The Process Repository log level	The log level set for the Process Repository.
The Notification Server log level	The log level set for the Notification Server

Descriptive Parameter Name	Description
The Servlet Engine log level	The log level set for the Tomcat Servlet Engine.
The Web Services Provider log level	The log level for the Web Services interface.
The CI Status Poller log level	The log level for all the CI Poller, unless overridden by specific instances of the CI Poller. You override the log level on the configuration page for the individual CI Poller. Each entry in the log file is identified by the name that you give the CI Poller in the Administration Console.
The Data Samples Provider log level	The log level for all the Data Sample Destinations, unless overridden by specific instances of the Destination. You override the log level on the configuration page for the individual Data Sample Destination. Each entry is identified by the name that you give the Destination in the Administration Console.

Descriptive Parameter Name	Description
Log BIE events (to bie_event0_0.csv)?	Check the checkbox to log all business events to the file **bie_event0_0.csv** in the *bpi-install-dir*\data\log

The file is in a comma-separated format in order that you can order and filter the events within a tool such as Microsoft Excel.

When the maximum number of bytes that should be written to a file is reached, or the BPI Server is restarted, a new log file is opened and named **bie_event0_0.csv**. The previous log file is renamed, for example, **bie_event0_1.csv** |
Logging Levels - Log4j Logging	
The Business Event Handler log level	The logging level set for the Business Event Handler.
The JMS Business Event Handler log level	The logging level set for the JMS Business Event Handler.
Component Startup Settings	
Startup timeout (seconds)	The time, in seconds, that is allocated to start each BPI component. All components are allocated the same value for their startup time.

The following are the logging levels that you can set for the **HP Operations Manager Logging** in ascending level of log detail provided:

➤ **Info**, which reports runtime informational messages and is the default level of logging.

➤ **Fine**, which provides more detail and is used when debugging or diagnosing problems.

➤ **Finer**, which provides the greatest detail and is used for tracing and debugging low level problems.

The following are the logging levels that you can set for Log4j Logging in ascending level of detail:

➤ **Info**, which reports at the informational level and is the default level of logging.

➤ **Trace**, which should be used for debugging purposes.

Note: When setting logging levels to report more detail, make sure that these are temporary settings as the detail logged can consume a significant amount of disk space and impact the performance of the system where BPI is installed.

Modifying the Log File Parameter Values

To modify the log levels for the BPI server components complete the following procedures:

1 Start the Administration Console by accessing **Start** > **Programs** > **HP** > **HP Business Process Insight** > **Administration**.

2 Select the **Status** option and stop the BPI Server components using the **Stop All** button.

3 Select the **Logging** option in the Navigator pane.

4 Make the changes to the log levels that you want to modify.

5 Click **Apply** to apply your modifications.

6 Restart the BPI components that you previously stopped using the **Start All** button.

The logging level changes have now been applied.

Important: The log levels set when BPI is installed are suitable for most error reporting. If you need to increase the granularity of the logging information provided, make sure that you set the log levels for a temporary period only. When set the Fine and Finer logging levels, the log files become large very quickly and can have a detrimental effect on your system's performance.

9

Deleting Instances from the Database

This chapter describes how to delete Process, Data and Monitor instances using the database instance cleaner options.

Use the Business Impact Engine (BIE) Process and Data Instance Cleaner parameters for regular removal of process and data instances from the database.

The instance cleaner parameters are not designed to enable individual instances to be manipulated nor do they enable you to name the process definitions or data definitions whose instances are to be deleted; these functions are provided by the Intervention tool.

Use the Monitor Engine Instance Cleaner parameters for regular removal of monitor instances from the database.

This chapter includes:

➤ Business Impact Engine Process and Data Instance Cleaner Parameters on page 217

➤ Monitor Engine Instance Cleaner Parameters on page 227

Business Impact Engine Process and Data Instance Cleaner Parameters

The Business Impact Engine (BIE) Instance Cleaner parameters enable you to control the numbers of Completed and Active Process and Data instances that are stored in the database. There are separate parameters for deleting Completed and Active Monitor instances from the database; see "Monitor Engine Instance Cleaner Parameters" on page 227.

Using the BIE Process and Data Instance Cleaner parameters, you can control:

➤ how often the BIE Instance Cleaner thread marks Completed and Active process and data instances as candidates to be deleted.

➤ when the BIE Instance Cleaner clean-up thread is run

➤ how often active and completed process and data instances are deleted from the database (if at all)

➤ the age of the instances that you delete

Following a new BPI installation, the BIE Instance Cleaner thread does not run and Completed and Active instances therefore accumulate in the database. Active instances include **Healthy**, **At Risk** and **Blocked** instances.

This section describes the behavior of these parameters and the SQL procedures, which you can modify to enable you to archive the data, before it is deleted, to a location of your choice. By archiving the data you have it available for your reporting tools, but it is no longer stored in the BPI database tables and therefore does not impact the performance of your system.

In summary, the instance cleaner parameters are intended for regular removal of Completed and Active instances as they occur during normal operations. They are not intended to be used for the removal of individual completed instances, or removal of instances in states other than the **COMPLETED** or **ACTIVE** state. Use the Intervention tool to remove individual instances; see the "Process Intervention" chapter in the *Using Business Process Insight Guide*.

Configuring the BIE Process and Data Instance Cleaner Parameters

You modify the instance cleaner parameters using the BPI Server Administration Console as described in "BIE Process and Data Instance Cleaner Settings" on page 150.

These parameters enable you to mark instances as being candidates to be deleted. Instances are initially marked to be deleted but are not immediately deleted, to reduce the impact of the database activity on other Business Impact Engine activities. Once the instances have been marked as candidates to be deleted, they are subsequently deleted by the clean-up thread at intervals specified by the instance cleaner parameter values. You can specify the frequency that you want them deleted and the numbers of instances that you want deleted in one database Delete statement, using the parameters.

In addition to modifying the instance cleaner parameters, you can modify the stored procedures invoked by the instance cleaner thread and add SQL commands to archive the Business Impact Engine data to a location of your choice; see "Archiving Completed and Active Process Instances" on page 222. You should only archive BIE data when you understand the performance of the database.

The following are some examples of how you can set the instance cleaner parameters to achieve specific outcomes, or requirements.

Example for Deleting Completed Instances

In the following example. you are deleting completed process instances from the database that are more than one day old and with a candidacy of **once a day**:

1 Select the **Administration Console** > **Business Impact Engine (BIE)** > **BIE Process and Data Instance Cleaner settings** option.

2 From the **BIE Process and Data Instance Cleaner settings** section:

    a Set the **Mark instances for deletion** option to **Once a Day**.

        Enter 01:00 as the time for the instances to be marked. Pick a time when the database is not busy processing other BPI data.

b Set the **Delete interval for marked instances** option to 1 minute.

This is the interval at which the Process and Data instance cleaner thread runs and attempts to delete any Process and Data instances that are marked as being candidates to be deleted.

c Set the **Delete batch size** to 500

This is the number of instances that the instance cleaner thread attempts to delete each time it runs.

3 From the **Completed Process and Data Instances Cleaner settings**:

a Set the **Delete completed instances from the database** option to **As scheduled above.**

The Completed instances are then deleted on the basis of the times set for the BIE Instance Cleaner thread.

b Set the **Age of the completed instances to be removed (minutes)** to two days, which is expressed in minutes (2880 minutes).

Example for Deleting Active and Completed Instances

In the following example, you are deleting Completed process instances that have an age of more than 30 days, and Active process instances that have an age of more than 50 days, with a candidacy set to **once a day**:

1 From the **BIE Process and Data Instance Cleaner settings**:

a Set the **Mark instances for deletion** option to **Once a Day**.

Enter 01:00 as the time for the instances to be marked. Pick a time when the database is not busy processing other BPI data.

b Set the **Delete interval for marked instances** option to 1 minute.

This is the interval at which the Process and Data instance cleaner thread runs and attempts to delete any Process and Data instances that are marked as being candidates to be deleted.

c Set the **Delete batch size** to 1000

This is the number of instances that the instance cleaner thread attempts to delete each time it runs.

2 From the **Completed Process and Data Instance Cleaner settings**:

   a Set the **Delete completed instances from the database** option to **As scheduled above**.

   b Set the **Age of the completed instances to be removed (minutes)** to 30 days, which is expressed in minutes (43200 minutes).

3 From the **Active Process and Data Instance Cleaner settings**:

   a Check the **Delete active instances from the database** option to **As scheduled above** so that the active instances are deleted on the basis of the BIE Instance Cleaner thread settings.

   b Set the **Age of the active instances to be removed** to 50 days.

Example for Deleting Completed Instances as Soon as They Are Complete

In the following example, Completed instances are deleted as soon as they are complete:

1 From the **Completed Process and Data Instance Cleaner settings**:

   a Set the **Delete completed instances from the database** option to **Immediately on completion.**

   b **Age of the completed instances to be removed (minutes)** is not applicable and is not an editable field.

   Be aware that using this setting has the effect of removing data that might have a value for tracking purposes. If your process instances do not have long lives, then there is going to be very little data in the database to be reported on. Make sure that you do not need the data that you are deleting from the database.

2 Set the **Active Instances settings** as required.

3 Set the **BIE Process and Data Instance Cleaner settings** as required for the Active instances, as these settings are ignored in the case of setting the **Delete completed instances from the database** option to **Immediately on completion**.

Archiving Completed and Active Process Instances

You have the option to archive Completed and Active instances before they are deleted from the database by the instance cleaner thread. This enables you to remove data from the BPI database to improve performance, but still keep the data for reporting purposes.

Note: You cannot archive **Completed** instances if you choose the option to **Delete completed instances from the database Immediately on completion**. This is because this option deletes the instances as soon as the event that completes the instance is received.

To archive the instance data, you need to make changes to the stored procedures or functions provided with BPI to delete the process and data instances. There are stored procedures for Microsoft SQL Server and functions for the Oracle Server. Both the stored procedures and the functions contain the SQL commands required to delete data from the BPI database and also provide a return code to the calling program.

Following are the stored procedures/functions:

➤ **bia_DeleteDataInstances**

This procedure deletes the data instances in the **COMPLETED** state that meet the criteria specified in the configuration.

This stored procedure is available for the Microsoft SQL Server only.

You can add your custom SQL statements to the stored procedure or function to archive the Business Process data before it is deleted.

➤ **bia_DeleteFlowInstances**

This procedure deletes all process instances in the **COMPLETED** state that meet the criteria specified by the configuration. This stored procedure also deletes the associated Activity instances, Activity started and Activity completed times.

The **bia_DeleteFlowInstances** procedure cleans out the following tables as specified:

Flow_Instance
Node_Instance
Node_Instance_StartedTimes
Node_Instance_CompletedTimes

You can archive as much or as little of these tables as you like. To archive all records as is, you first need to create new data tables to hold the archive records. For instructions on creating new data archive tables, see "Creating Archive Tables" on page 225.

➤ **DeleteActiveDataInstances**

This procedure deletes the data instances in the **ACTIVE** state and that meet the criteria specified in the configuration. This stored procedure also provides the following SQL Cursor that you can use:

activeInstances_cursor

This stored procedure is available for the Microsoft SQL Server only.

As for the **bia_DeletedDataInstances** procedure, you can add your custom SQL statements to the stored procedure or function to archive the Business Process data before it is deleted.

➤ **DeleteActiveFlowInstances**

This procedure deletes all process instances in the **ACTIVE** state that meet the criteria specified by the configuration. This stored procedure also deletes the associated Activity instances, Activity started and Activity completed times.

These stored procedures or functions are contained in the following SQL file:

bpi-install-dir**misc\bia\EngineStoredProcedures_script.sql**

There are also examples of this file in the **examples** directory on the distribution media, located at:

bpi-install-dir**examples\bia\EngineScripts**

Following are the example files, according to the database:

➤ **EngineStoredProcedures_script.mssql**

➤ **EngineStoredProcedures_script.oracle**

Important: Be aware that changes to the SQL file **EngineStoredProcedures_script.sql** are lost if you reinstall or upgrade BPI. Make sure that you have made backup copies of your modifications, which you can then reapply following a reinstallation.

You need to make your changes to the stored procedures, which have been created in the database for BPI, to add the SQL to archive the Completed or Active instances.

You are also advised to make changes to the following SQL script and make a backup copy of the script to ensure that you have a copy of the changes:

bpi-install-dir**misc****bia****EngineStoredProcedures_script.sql**

The instance cleaner parameters are executed at the frequencies specified in the **BIE Instance Cleaner Settings**. The value of the **Age of the completed instances to be removed** and **Age of the active instances to be removed** parameters is passed to the stored procedure, or function, when it is called.

Note: It is possible for two data instances to be created, each representing the same Business Process: one in the BPI Instance database and one in the archive database. This occurs when an out-of-sequence event is received by the Engine after the Completed or Active instances are archived and deleted. An out-of-sequence event is where an event is received by the Business Impact Engine for an activity in a process instance that is earlier than the latest event received for the same process instance.

In summary, you need to ensure that your archive code takes into account that there might be more than one instance that represents the same Business Process when this scenario exists.

Creating Archive Tables

You create new archive tables, using the **misc\bia\FlowTables_script.sql** script, which was used to create the Engine data tables when the BPI Server was installed. You can copy the create statements for the instance tables above and use these as a template for creating the set of archive tables.

To create new archive tables:

1 Open the FlowTables_script.sql file and grab the lines that create the Flow_Instance table.

2 Paste these lines into a new text file.

3 Change the final '#' character for a ';' character.

4 Remove any lines that mention a foreign key.

5 Change the table name to be Archive_Flow_Instance.

Your create table command looks as follows:

CREATE TABLE Archive_Flow_Instance (
  FlowInstance_ID NVARCHAR(36) NOT NULL PRIMARY KEY,
  Flow_ID NVARCHAR(36) NOT NULL,
  Identifier NVARCHAR(40),
  Primary_entity NVARCHAR(36),
  Primary_entity_inst NVARCHAR(36),
  Weight FLOAT,
  Weight_type NVARCHAR(40),
  StartTime DATETIME,
  StartTimeLongMillis BIGINT,
  EndTime DATETIME,
  EndTimeLongMillis BIGINT,
  Status NVARCHAR(12),
  Subclass NVARCHAR(256),
  );

6 Run this command to create the archive table.

7 Repeat these procedures for each of the tables:

Node_Instance
Node_Instance_StartedTimes
Node_Instance_CompletedTimes

You now have the following set of tables:

Archive_Flow_Instance
Archive_Node_Instance
Archive_Node_Instance_StartedTimes
Archive_Node_Instance_CompletedTimes

The next procedure is to set up the appropriate indexing on these tables. This depends on the kind of reports you want to run.

Indexing

To get the same indexing as the original set of tables, open the **FlowTables_script.sql** file and edit the indexing lines to use the correct **Archive_ names**.

For example:

CREATE NONCLUSTERED INDEX Archive_Node_Instance11 ON Archive_Node_Instance (FlowInstance_ID ASC);

CREATE NONCLUSTERED INDEX Archive_Flow_Instance12 ON Archive_Flow_Instance (Flow_ID ASC, Primary_entity_inst ASC, FlowInstance_ID ASC, Identifier ASC, Primary_entity ASC, Weight ASC, Weight_type ASC, StartTime ASC, EndTime ASC, Status ASC);

CREATE NONCLUSTERED INDEX Archive_Flow_Instance13 ON Archive_Flow_Instance (Status ASC);

CREATE NONCLUSTERED INDEX Archive_Flow_Instance1 ON Archive_Flow_Instance (EndTime ASC);

CREATE CLUSTERED INDEX Archive_Node_Instance_StartedTimes3 ON Archive_Node_Instance_StartedTimes (NodeInstance_ID ASC);

CREATE CLUSTERED INDEX Archive_Node_Instance_CompletedTimes4 ON Archive_Node_Instance_CompletedTimes (NodeInstance_ID ASC);

Monitor Engine Instance Cleaner Parameters

The Monitor Engine Instance Cleaner parameters enable you to control the numbers of monitor instances that are stored in the database. These include:

➤ Active monitor instances

➤ Completed monitor instances

➤ Monitor statistics

➤ Monitor alarm instances

There are separate parameters for deleting Completed and Active business impact Engine instances from the database; see "Business Impact Engine Process and Data Instance Cleaner Parameters" on page 217.

Using the Monitor Engine Instance Cleaner parameters, you can control:

➤ how often the Monitor Engine Instance Cleaner thread is run.

➤ how often active and completed monitor instances are deleted from the database (if at all).

➤ how often monitor statistics are deleted from the database (if at all).

➤ how often monitor violation instances are deleted from the database (if at all).

➤ for all monitor instances, the time and the age of the instances that you delete.

Following a new BPI installation, the Monitor Engine Instance Cleaner thread does not run and the different monitor instances therefore accumulate in the database.

Configuring the Monitor Engine Instance Cleaner Parameters

You modify the Monitor Engine instance cleaner parameters using the BPI Administration Console as described in "Monitor Engine Instance Cleaner Settings" on page 172.

These parameters enable you to delete the Monitor Engine instances periodically, using the instance cleaner thread, as specified by the instance cleaner parameter values.

You can control how often the monitor instance cleaner thread is executed and the age of the specific monitor instances that you want to delete. Each type of monitor instance can be separately controlled.

The following are some examples of how you can set the Monitor Engine instance cleaner parameters to achieve specific outcomes, or requirements.

Example for Deleting Completed Monitor Instances

If you want to delete completed monitor instances from the database, which are more than one day old and once a day, set the Monitor Engine Instance Cleaner parameters.

To set the Monitor Engine Instance Cleaner parameters:

1 From the **Monitor Engine Instance Cleaner settings**, set the **Instance cleaner execution interval (minutes)** to 5.

2 From the **Completed Monitor Instances settings**, do the following:

  a Check the **Delete completed monitor instances from the database?** option in order to activate the **Age of completed monitor instances to be removed (days)** option.

  b Set the **Age of completed monitor instances to be removed (days)** option to 2.

3 Set the other option as appropriate for your implementation.

10

Servlet Engine Authentication

This chapter describes how to use the Memory Realm mechanism within the Servlet Engine (Tomcat).

This chapter includes:

➤ Servlet Engine Overview on page 229

➤ Tomcat Realm Mechanism on page 230

➤ Modifying Servlet Engine Files on page 234

➤ Tomcat and HTTP Over SSL on page 235

Servlet Engine Overview

You use the Memory Realm mechanism within the Servlet Engine (Tomcat) to:

➤ modify the login details for the BPI Web-based interfaces, which are the process instance and violation pages, which are accessible from the Business Process Insight Application pages.

➤ configure the use of HTTP over SSL for BPI

The chapter is structured as follows:

➤ "Tomcat Realm Mechanism" on page 230

➤ "Tomcat and HTTP Over SSL" on page 235

The BPI components that can be configured for authentication using the Servlet Engine are:

➤ Process Repository, which affects both the Repository Explorer and the Modeler.

Setting authentication for the Process Repository also impacts any Contributed components that use the repository.

➤ Intervention tool

➤ Notification pages

➤ Monitor Definer

➤ Process instance and violation pages

With the exception of the process instance and violation pages, each of these components requires you to provide a user name and password in order to access the Web screens. Following a new BPI installation, the user name and password for these components are predefined as:

➤ User name: **admin**

➤ Password: **hpbpi**

Tomcat Realm Mechanism

BPI uses Tomcat as its Servlet Engine to present JSPs to the user, through a Web browser; starting the Servlet Engine is described in section "Starting and Stopping the BPI Server Components" on page 139. For authentication, BPI uses the Tomcat Realm mechanism, specifically it uses the Memory Realm mechanism, which is more fully described in the Tomcat documentation at the following URL:

http://jakarta.apache.org/tomcat/tomcat-5.0-doc/realm-howto.html

Within the Tomcat documentation, a Realm is described as a database of usernames and passwords, which identify valid users of a Web application (or set of Web applications), plus a list of roles associated with each valid user.

Within BPI the following roles have been defined:

➤ **ovbpi-model-repository,** which is the role for the Repository Explorer and the Modeler (see the "Modeler" and "Process Repository Explorer" chapters in the *Using Business Process Insight Guide*).

➤ **ovbpi-notify-admin**, which is the role for the Notification pages (see the "Notification" chapter in the *Using Business Process Insight Guide*).

➤ **ovbpi-int-client,** which is the role for the Intervention tool (see the "Process Intervention" chapter in the *Using Business Process Insight Guide*).

➤ **ovbpi-metric-definer**, which is the role for the Web-based Monitor definition interface (see the "Monitor Definer" chapter in the *Using Business Process Insight Guide*).

➤ **ovbpi-event-injector**, which is for a contributed tool and is provided on the BPI media.

There is an XML file called tomcat-users.xml, which has the credentials for these roles. The file is located at:

bpi-install-dir/**nonOV/jakarta-tomcat-5.0.19/conf**

If you have upgraded from a previous version of BPI, there is another XML file located in the same directory as follows:

bpi-install-dir/**nonOV/jakarta-tomcat-5.0.19/conf/ tomcat-users.xml.***versionstring*

where `versionstring` is the version number of the product.

You need to compare the new file, which includes the version string, with your existing tomcat-users.xml configuration file and modify tomcat-users.xml to include any changes shown in the newly installed file.

The reason that the BPI installation procedure does not replace the file is because you might have made your own changes to the file, which you do not want to be overwritten.

Following a new installation, the file has the following content:

```
<?xml version='1.0' encoding='utf-8'?>
<tomcat-users>
<role rolename="ovbpi-model-repository"/>
```

```
<role rolename="ovbpi-int-client"/>
<role rolename="ovbpi-notify-admin"/>
<role rolename="ovbpi-metric-definer"/>
<role rolename="ovbpi-event-injector"/>
<user username="admin" password="hpbpi"
roles="ovbpi-notify-admin,ovbpi-int-client,ovbpi-event-injector,ovbpi-metric-defin
er,ovbpi-model-repository"/>
</tomcat-users>
```

You are strongly advised to assign different credentials to the different roles in order to increase the security of your system. This is particularly the case for the Intervention tool, from where you can delete active process instances and data instances.

Modifying Login Credentials

To assign different credentials, you need to edit the XML file that defines the roles and the login credentials. The following is an example of how you can do this:

```
<?xml version='1.0' encoding='utf-8'?>
<tomcat-users>
 <role rolename="ovbpi-model-repository"/>
  <role rolename="ovbpi-int-client"/>
  <role rolename="ovbpi-notify-admin"/>
  <role rolename="ovbpi-metric-definer"/>
  <role rolename="ovbpi-event-injector"/>

  <user username="inject" password="inject"
    roles="ovbpi-event-injector"/>
  <user username="interv" password="secret"
    roles="ovbpi-int-client"/>
  <user username="notif-admin" password="admin"
    roles="ovbpi-notify-admin"/>
  <user username="monitor" password="secret"
    roles="ovbpi-metric-definer"/>
  <user username="model-repository" password="admin"
    roles="ovbpi-model-repository"/>
</tomcat-users>
```

In this example, the following clients now have a separate user name and password:

➤ The **ovbpi-event-injector** role for the contributed tool has a user name and password of inject.

➤ The **ovbpi-int-client** role for the Intervention tool has a user name of interv and a password of secret.

➤ The **ovbpi-notify-admin** role for the Notification pages has a user name notif-admin with a password of admin.

➤ The **ovbpi-metric-definer** role for the Monitor Definer interface has a user name of monitor and a password of secret.

➤ The **ovbpi-model-repository** role for the Process Repository (used within the BPI Modeler) has a user name **model-repository** with a password of admin.

To make changes to the tomcat-users.xml file, complete the following procedures:

1 Make sure that you have logged out of the BPI Modeler, Notification pages and the Intervention tool.

2 If you are using the Event Injector contributed component, you must also close the Event Injector.

3 Open tomcat-users.xml in a text editor:

   bpi-install-dir/nonOV/jakarta-tomcat-5.0.19/conf/tomcat-users.xml

4 Make your required changes to the file.

5 Save and close the **tomcat-users.xml** file.

6 Make your changes to the Security settings in the Administration Console as required.

7 Stop and restart all the BPI components from the Status pane of the BPI Administration Console for the changes to take effect.

   You must stop and restart all the BPI components when you change the security options, if you do not, some of the security settings can be compromised.

Memory Realm is a simple implementation of the Tomcat 5 Realm interface. It is described in the Tomcat documents as a demonstration implementation and not designed for production use.

Modifying Servlet Engine Files

If you intend to modify the Servlet Engine configuration files, you need to be aware that changes made directly to the Servlet Engine files are overwritten by BPI configuration files. As an example, if you want to configure another form of Tomcat Realm security, you would expect to make the changes in the web.xml and server.xml files under the following directory:

   bpi-install-dir\nonOV\jakarta-tomcat-5.0.19\conf

However, before making any changes to these files, you need to be aware of how BPI manages configuration changes to the Servlet Engine files. BPI overwrites the Servlet Engine configuration files in order to add BPI-specific configuration details and as a result of your making changes to configuration parameters using the Administration Console. BPI overwrites the Servlet Engine configuration files whenever a change is made to any parameter using the Administration Console, it does not have to be a Servlet Engine specific change.

If you want to make changes directly to the Servlet Engine configuration files, you need to make the modifications in the BPI-specific configuration files and then make sure that BPI propagates these modifications to the Servlet Engine files.

BPI has a number of configuration files and the Servlet Engine files related to web.xml and server.xml are located under the following directory:

   bpi-install-dir\newconfig\DataDir\conf\bia\tomcat\conf

These files are database specific and you need to modify the files that are appropriate to your configuration.

If you are using Oracle, you need to modify the following files:

   bpi-install-dir\newconfig\DataDir\conf\bia\tomcat\conf\
   server.engine.oracle.xml

bpi-install-dir\newconfig\DataDir\conf\bia\tomcat\conf\
server.views.oracle.xml

If you are using SQL Server, you need to modify:

bpi-install-dir\newconfig\DataDir\conf\bia\tomcat\conf\
server.engine.mssql.xml

bpi-install-dir\newconfig\DataDir\conf\bia\tomcat\conf\
server.views.mssql.xml

When you have made the required changes to these files, you need to artificially make a change to any parameter using the Administration Console and apply the change. As a result of this change, the Administration Console updates all the relevant configuration files, including the Servlet Engine files, and your modifications are then propagated to these files.

Tomcat and HTTP Over SSL

HTTP over SSL (secure socket layer) is a secure version of HTTP. SSL enables the data from a client, such as a Web browser, to be encrypted prior to transmission to prevent unauthorized access to the data being transmitted.

Using SSL, the Web Server and the Web browser (client) establish a secure connection, which enables them to trust each other and send data knowing that the connection is secure. In order to do this, a level of authentication is required and this is provide through Certificates. You obtain Certificates from a Certificate Authority (CA).

You can configure HTTPS for the Tomcat Servlet Engine installed with BPI and then enable the use of HTTPS for your BPI Web applications using the Administration Console. Note that the use of SSL is between the BPI Servlet Engine and the BPI Web applications only. Other HP BTO Software products can also be configured to use SSL; however, BPI can not be configured to use SSL as a means of communicating with these other HP BTO Software products.

There are two aspects of configuring HTTPS for BPI, one aspect is related to the Servlet Engine configuration and one to the Web client configuration. These are described in the following sections.

Servlet Engine Configuration

In order to configure the Tomcat Servlet Engine to support HTTPS, you need to obtain an SSL certificate. SSL uses public key cryptography, which is based on key pairs. These key pairs contain one public key and one private key. These keys are used by the Servlet Engine to enable HTTPS on the machine where the BPI components are running. When installed the certificate encrypts the data that is sent from the Servlet Engine to the Web client, or Web browser.

There is more than one form of keystore, which is a repository of certificates used for identifying a client or a server. The Java Keytool generates a key based on the JKS format of keystore. There is also a PKCS12 format of keystore; however, this cannot be used with BPI. You can use only the JKS format as generated from the keytool command.

The following are ways that you can obtain, or generate, a certificate:

➤ Using the Java Keytool

Keytool is a key and certificate management utility that ships with your Java distribution. It enables users to administer their own public/private key pairs and associated certificates for use in self-authentication or data integrity and authentication services, using digital signatures.

Use the keytool command, from a Command Window, to generate a Java Key.

➤ From a Certificate Authority

You can obtain SSL certificates from Certificate Authorities (CA) such as Verisign.

In order to obtain and SSL certificate from a CA, you first need to generate a Certificate Signing Request (CSR) for the Certification Authority to sign. You can do this using the Java Keytool. When generating this certificate make sure that you supply the fully qualified domain name of the machine where the BPI is installed as the common name (cn) attribute.

The CSR certificate and the certificate chain received from the CA need to be stored in a JKS keystore file.

When you have obtained all the required information and configured the Servlet Engine to send HTTPS requests, you can use the Administration Console to enable HTTPS as the protocol to be used to send data to one or more of the BPI Web Clients. This is done through the Security options (see "Security" on page 200.

Web Client Certificate-Based Authentication

You can also use the Servlet Engine to provide certificate-based client authentication, based on HTTPS, for your Web applications:

➤ Process instance and violation Web pages

➤ Repository Explorer

➤ Monitor Definer

➤ Intervention tool

➤ Notification Server Administration

In order to use client authentication, you first need to configure the Servlet Engine to be able to send HTTPS requests to the Web Browser, as described in section "Servlet Engine Configuration" on page 236. You need to configure the Servlet Engine on each system where it installed. You then need to obtain a client certificate, which is signed by a CA, plus a CA root certificate.

In order for the Servlet Engine to be able to receive encrypted data from the Web Browser, you need to enable client authentication to ensure only Web browsers presenting a valid certificate can authenticate with the Servlet Engine.

Each CA provides a root certificate, which you install, or import, for your Web Server (in this case the Tomcat Servlet Engine). The Servlet Engine uses this certificate to determine whether or not it trusts the data being sent from the Web Browser. You obtain this root certificate from your CA. When you have the certificate, you import it into the keystore for your Java distribution. The root certificate needs to be imported so it can become a trusted certificate, enabling a chain of authorization to be established between the Web Server and the Web Browser.

The keystore location for the Servlet Engine is as follows:

java-home\jre\lib\security\cacerts

The Administration Console allows you to select a keystore file from any location; however, the Tomcat Servlet Engine requires the keystore to be located on a local hard drive. The keystore must not be located on a virtual or network drive.

You import the certificate using the keytool utility that comes with your Java distribution, using the following command:

> **keytool -import -keystore** *java-home***/jre/lib/security/cacerts -storepass changeit -alias rootca -file rootca.pem -trustcacerts**

where:

➤ *java-home* is the directory where your Java distribution is installed.

➤ **changeit** is the default password for the Java security keystore for a new Java installation.

Enter **yes** when prompted if you trust this certificate

When you have imported the root certificate on the Web Server, you then need to import both the root certificate and the client certificate on each Web Browser where you intend to use the BPI Web-based interfaces (or those for which you want to enable HTTPS). Follow the instructions in you Web Browser documentation for installing Certificates.

When you have completed the configuration to set up your Web Browser to Web Server HTTPS connection, you can use the Administration Console to enable certificate-based client authentication. This is done through the Security options; see "Security" on page 200. Client authentication is used only for those clients where you have also enabled HTTPS as described in section "Servlet Engine Configuration" on page 236.

11

Connecting BPI With TransactionVision

This chapter describes how to connect BPI with TransactionVision, a product within the Transaction Management application, which enables BPI to show the status of Business Process instances. The Business Process instances are linked to specific business transactions within TransactionVision.

This chapter includes:

Connecting With TransactionVision Overview

TransactionVision is a product that delivers a transaction tracing solution to the Transaction Management application. TransactionVision enables you to view the life cycle and associated data for specific business transactions within your organization. TransactionVision discovers transaction flows within your business, analyses the steps within these flows and then displays reports based on the results of the analysis. Using TransactionVision, you can better monitor and manage these business transactions. By connecting BPI with TransactionVision, you can link the business transactions into your Business Process Activities.

Connecting Business Process Insight with TransactionVision enables you to:

➤ Import TransactionVision business transaction event definitions into the Business Process Insight Model Repository. You can then use the Modeler to create processes, which have Data definitions that subscribe to these imported events.

➤ Receive events from TransactionVision and use the data within them to progress BPI Business Processes.

➤ Link from the process instance details page within the BPI Application Health pages to the details of the instances within TransactionVision.

➤ Provide the detail of TransactionVision transaction level problems to Business Process Insight.

➤ Drill down from BPI to Transaction Management reports and topologies.

This provides BPI users with more information relating to problems that might be occurring in a process at the transaction level.

The following diagram shows the high level architecture of the BPI and TransactionVision connection.
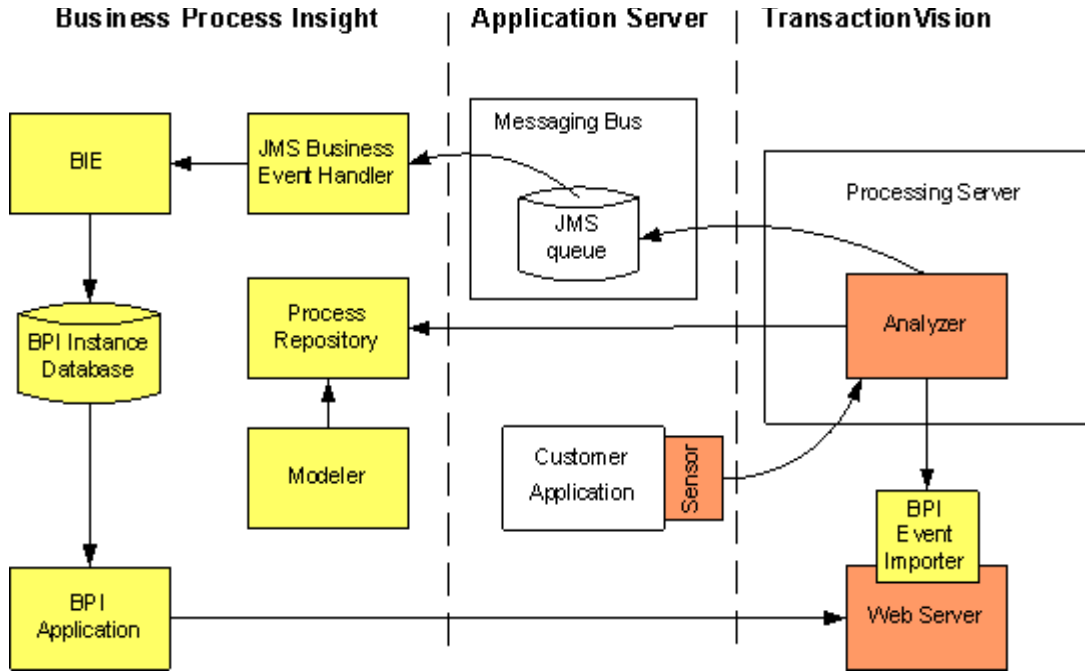
See the "Working with Other Applications in HP Business Service Management" chapter in *Using Transaction Management* for details on:

➤ Design-time integration and how data is moved between BPI and TransactionVision when you are designing your processes and configuring the BPI and TransactionVision systems.

➤ Run-time integration and how event data flows from each TransactionVision Processing Server to the BPI Server when TransactionVision and Business Process Insight are both deployed in the Business Service Management deployment environment.

Configuring BPI to Connect with TransactionVision

This section describes the tasks that you must complete to enable the BPI and TransactionVision connection.

The configuration tasks described in these sections assume that the necessary BPI components have been installed on the TransactionVision and BPI systems.

TransactionVision Configuration Tasks

The Transaction Management administrator sets up the connection configuration between TransactionVision and BPI. For details on configuring TransactionVision to connect with BPI, see the "Working with Other Applications in HP Business Service Management" chapter in *Using Transaction Management*.

BPI Configuration Tasks

The following tasks need to be completed on the BPI system. Most of these tasks are described in the sections that follow:

➤ Import the TransactionVision events on the TransactionVision Processing Server to create the required event definitions in the Model Repository. See "Export BPI Action Definitions into the Process Repository" on page 243.

➤ Using the BPI Modeler, create Data definitions that subscribe to the imported TransactionVision events; refer to the *Integration Training Guide - Modeling Processes* and the Modeler Online Help for more information.

Using the BPI Modeler, define progression rules based on the data being imported from TransactionVision events. The following progression rule enables you to link an Activity in your Business Process to a TransactionVision event:

Start and complete on TransactionVision events

See the *Integration Training Guide - Modeling Processes* and the Modeler Online Help for more information.

➤ Deploy the BPI process that you have created for your TransactionVision connection from within the BPI Modeler.

➤ Configure BPI to receive Business Events from the JMS queue that you have created for BPI; see "Configure a JMS Business Event Handler" on page 243.

➤ As the process progresses, monitor it using the BPI Application.

You can also link from the BPI Application to TransactionVision for more details of the transactions related to selected process instances. You do this from the individual process instance page, where there is a link from the Data Definition tab for an option to View Transaction Tracking Details. This link opens the Transaction Tracking Report for the selected process instance Data Definition.

Export BPI Action Definitions into the Process Repository

You need to export the Business Events defined in TransactionVision as BPI Actions, into BPI. This makes the Business Events available within your process definitions. When you have exported the Business Events, you can then subscribe to the events from the Data definitions that you have created for the Business Process.

Export the definitions from the BSM user interface as follows:

Admin > **Transaction Management** > **Administration** > **Monitoring**, click **Transaction Tracing Rules**, then click the **Export BPI action definition to BPI** button.

Configure a JMS Business Event Handler

You configure the integration with TransactionVision through the BPI Administration Console. BPI uses an instance of a JMS Business Event Handler to integrate with TransactionVision.

If BPI is on a different machine to TransactionVision, you need to configure a JNDI repository and add the JNDI entries for accessing the shared JMS queue on the messaging bus for the BPI and TransactionVision integration. Refer to the documentation for your JMS messaging bus for the details that you need for the JNDI repository configuration. Some values are preconfigured for the integration, assuming the use of the Sonic Bus.

You then need to make sure that the correct JMS Queue and JNDI entries are set up for each JMS Business Event Handler

You also need to supply details of a Service Provider and Event Source as described in "Service Provider" on page 244 and "Event Source" on page 246.

Service Provider

You can configure the JMS Business Event Handler to receive events from one, or more, JMS sources. You configure a Service Provider for each JMS service on the TransactionVision Processing Server from which you want to receive event data. You need to configure a Service Provider before you can configure a JMS Event Source.

There is a Service Provider preconfigured for the connection with TransactionVision. You need to modify some of the parameter values, but most of the details should be correct for your implementation.

If you want to create additional Service Providers, you need to create a new entry.

To configure the Service Providers settings using the BPI Server Administration Console:

1 Start the Administration Console as follows:

   Start > **Programs** > **HP** > **HP Business Process Insight** > **Administration**

2 Click the **Service Provider** under **JMS Business Event Handler**.

   A configuration page is opened in the right-hand pane, where you can add a new Service Provider, or modify an existing entry.

3 If this is the first connection with TransactionVision, select **HP TransactionVision**, then click the **Modify** button to modify the parameters for the preconfigured HP TransactionVision Service Provider entry.

   A **JMS Service Provider Properties** dialog is opened where you can modify values for the properties of the Service Provider. The properties are fully described in "Service Providers" on page 193:

   ➤ **Service Provider Name**

      Unique name used within BPI to identify the service provider. The name can be up to 40 characters.

      This is set to **HP TransactionVision**.

   ➤ **Description**

Optional description of the service provider.

➤ **Enabled?**

Check the box to enable this Service Provider.

➤ **Service Provider Classpath**

Location, or set of locations, relating to the service provider Classpath.

This is not required for the Sonic implementation as the Classpath is already configured to be on the BPI Server Path, and can therefore be left blank.

➤ **JNDI Provider Location**

Location of the JNDI service for the named service provider, typically this is a URL.

This parameter value includes the name of the machine where Transaction Management is installed. In the value presented, this is **analyzerhost**. Substitute the correct fully-qualified host name for **analyzerhost**; for example:

 tcp://<*fully-qualified-analyzerhostname*>:21110

The port number remains as 21110, unless you have specifically modified it.

➤ **JNDI Factory Class**

The Java class name for the initial context factory for the JNDI service. This is set for the integration and should not be changed.

➤ **JMS Factory Class**

The Java class name for the initial context factory for the specific JMS service provider.

This is set for the integration and should not be changed.

4 Click **OK** to close the dialog.

5 Click **Apply** to apply your configuration changes.

If you want to create subsequent entries, click **Add** to add a new Service Provider.

Event Source

When you have configured a service provider, you can then configure an event source.

Having the event source configuration separate from the service provider configuration enables you to have multiple event sources for one service provider. The event source parameters enable you to configure JMS queue and login details.

To configure the Event Sources settings using the Administration Console:

1 Start the Administration Console:

Start > **Programs** > **HP** > **HP Business Process Insight** > **Administration.**

2 Click the **Event Source** option under **JMS Business Event Handler**.

A configuration page is opened in the right-hand pane. If you are modifying the preconfigured HP TransactionVision Service Provider entry. select the **Modify** button to modify the existing Event Source.

You can also select the **Add** button to add a new **Event Source** where appropriate.

A JMS Event Sources Properties dialog is opened where you can enter values for the properties of the **Event Source**. The properties are fully described in "Event Sources" on page 195:

➤ **Service Provider Name**

If there is more than one Service Provider defined, select the name of a service provider that you have previously configured for this event source.

➤ **JMS Queue Subject**

The JNDI value that references the JMS queue at the selected service provider. JMS Queue Subject needs to reference the name of the JMS queue that has been configured for BPI to receive transaction events. For the BPI and Transaction Management integration this is predefined for the Sonic BUS as:

TV2BPI.EVENT.QUEUE

➤ **User Name**

The name of a user with authority to access the named **JMS Queue Subject**. Not all service providers require a user name to be specified.

➤ **Password**

Password for **User Name** where required.

3 Click **OK** to close the dialog.

4 Click **Apply** to apply your configuration changes.

When you have configured an event source, BPI generates the appropriate files to enable the JMS Business Event Handler to access the transaction events from the JMS queue. To complete, the configuration, you need to stop and restart the JMS Business Event Handler from the **Status** screen on the Administration Console.

If TransactionVision has been configured to push events to this shared queue, BPI should now be configured to receive them. Any errors are reported to the JMS Business Event Handler log file, which can be accessed from the Status screen of the BPI Administration Console.

12

High Availability Using Microsoft Clusters

This chapter includes:

➤ High Availability System Using Clusters Overview on page 249

➤ Modes of High Availability on page 250

➤ Configuring BPI Within a Microsoft Cluster on page 253

High Availability System Using Clusters Overview

A highly available system is one that enables applications and services to continue after hardware and software failures, thus reducing the time when the system is not available for use; also known as downtime. You can have planned downtime, where you have set aside time for maintenance tasks, such as upgrading, and you can have unplanned downtime. Unplanned downtime is where the system fails due to events such as hardware and software failures.

High availability in terms of BPI is where your BPI system is always available when it is planned to be available and fail over mechanisms are in place to minimize, or eliminate, unplanned downtime.

A Microsoft Server Cluster can provide some of the high availability capabilities, but it cannot provide a complete solution for high availability with no single point of failure. You need to make sure that your BPI deployment includes other protection mechanisms for disks, networks, power supplies, memory, controllers and systems; for example:

➤ dual power supplies to BPI systems

➤ dual paths from system to disks

➤ dual network controllers

➤ uninterruptable power supplies (UPS)

➤ disk mirroring and (or) the use of redundant arrays of inexpensive disks (RAID)

In other words, to achieve the maximum availability of your system, you need to build in redundancy at all levels.

You can install BPI in a Microsoft Server Cluster environment in order to provide a high-availability solution for BPI. By using the cluster technology when one node in a cluster fails or is taken offline another node in the cluster takes over the failed node's operations. When BPI functions are moved from one node in the cluster to another node, users can experience some temporary interruption in their use of BPI, but this interruption is minimal.

Modes of High Availability

There are two modes of high availability that are offered by Microsoft:

➤ Active / Passive

In this mode, one node in the cluster hosts an application at any given time; this can be likened to having a standby system in place.

➤ Active / Active

In this mode, different instances of the same application are running concurrently on different nodes in the cluster. In this mode, when a node fails the remaining nodes in the cluster need to take on additional load and this can impact performance.

BPI must be deployed in the Active/Passive mode within the cluster. You cannot use the Active/Active mode with BPI.

By supporting Microsoft Server Clusters in an Active/Passive mode, BPI Services are automatically restarted on an alternative machine, if the machine on which they are currently running fails for any reason.

A Microsoft Server Cluster can support up to eight nodes; however, in an Active/Passive mode, BPI requires only two nodes to be configured within the cluster. You can configure additional passive nodes for BPI to keep to a minimum the single points of failure, but it is not needed in order for BPI to operate within the Server Cluster.

Each node within the cluster can host one or more resource groups. A resource group is a collection of components that are managed and viewed as a single unit within the cluster; for example, a public file share, a Web server, and a database application can all be managed as resources. BPI is managed as a Resource Group within the cluster. The use of Resource Groups enables you to manage the order in which resources are started within the cluster, and dependencies between resources. This then enables you to specify specific dependencies for BPI; for example, the dependency between BPI and its database.

The nodes and the resources within the cluster are constantly monitored and in the event of a failure on the node where BPI is running, the workload of the failed node is re-started on another designated node. Note that there is an interruption to service while BPI is in the process of failing over and following the failover, the BPI Modeler needs to be restarted as it has lost its connection to the Process Repository.

In the following diagram, Node A and Node B are both configured to run BPI in Active/Passive mode. Node A is currently hosting BPI. If Node A fails for any reason, the cluster can start BPI on Node B.
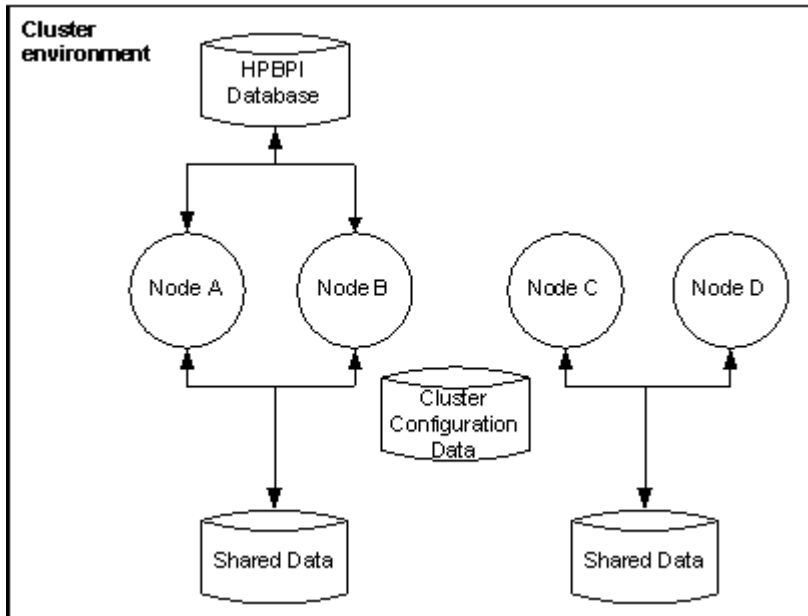
In order for the failover to be successful, both Node A and Node B must have access to the BPI data. This includes the installation data and configuration files on disk, plus all the Process, Monitor and Events data in the database. In this diagram it is only Nodes A and B that have this access. Nodes C and D are also operating in the cluster but are not available for BPI failover.

The BPI shared data device is a resource that needs to be managed by the cluster and must therefore reside within the bounds of the cluster. The BPI database might or might not be managed by the cluster, depending on how you have configured your system. In this case, the device where the database is located might or might not be managed as part of the cluster.

The documentation for the Microsoft Server Cluster provides details of the hardware configuration require, plus recommendations for using technology such as RAID for the shared storage to minimize the single points of failure.

There is also a cluster configuration database that must be accessible to all the nodes in the cluster and which contains information such as the membership of the cluster, the configuration data and quorum information.

You can have additional nodes in the cluster made available for BPI failover if you want to minimize the your single points of failure. In the following diagram, BPI has also been installed and configured on Node C, which means that either Node B or Node C can be used as a failover system if needed. Note that Node C is also available as a failover node for another application whose data is stored on a different device to that used by BPI.
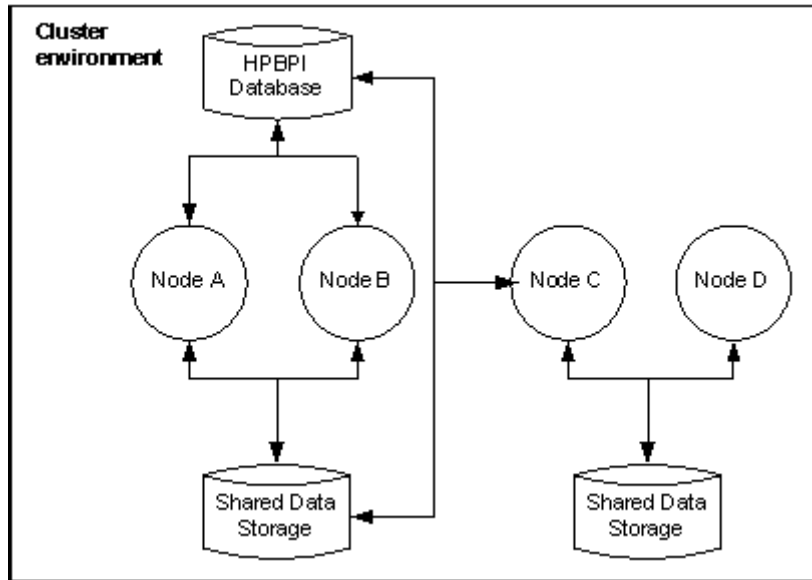
Configuring BPI Within a Microsoft Cluster

This section describes the BPI-specific tasks that you need to complete in order to have BPI operate within the Microsoft Server Cluster environment. The section assumes that the Server Cluster has been configured to meet your organizational requirements; it does not provide any information relating to the installation and setup of the cluster. Refer to the Microsoft Server Cluster documentation for this information.

In order to have BPI operate within a cluster environment, it is necessary to install BPI on each node in the cluster that you want BPI to run on. This does not mean you have to install BPI on every node in the cluster; see "Modes of High Availability" on page 250.

You need to install BPI multiple times, even though each node has access to the same shared install directory and database. This is to ensure that the BPI Windows Services are created on every node where BPI is to run.

Groups

BPI installs a number of Windows services, therefore one way to configure BPI within the cluster is to set up the services using the Generic Service resource type; for example, resources can be grouped as self-contained resource groups, called Groups. How you configure a resource within the cluster is outside the bounds of this document. You need to read the information provided by Microsoft to determine the best configuration for your requirements.

The following are suggested scenarios for grouping the BPI resources:

1 Configure a self-contained resource group for BPI resources and have the database outside the cluster environment. The means that if there is a catastrophic cluster failure, the BPI data is still accessible for browsing and for analysis.

This configuration enables a specific disk drive (or partition), IP address and virtual server name to be assigned to BPI and BPI can then be moved between nodes as an independent resource.

2 Configure a self-contained resource group for BPI resources and have the database installed as a cluster application. You then configure the BPI resources to be in the same resource group as the database.

In this case, BPI can share the disk, virtual server name and IP address of the database. This enables you to define specific dependencies between BPI and the database and the database and BPI can be treated as a single entity and operates as described in scenario 1. In addition, you need only allocate one disk and host name.

3 Configure a self-contained resource group for BPI resources and have the database installed as a cluster application. You then configure the BPI resources to be in a different resource group to the database.

By allocating different resource groups to BPI and the database, each can failover independently of the other, which minimizes the potential downtime of the BPI system. In addition, BPI can be moved between nodes as an independent resource as in scenario 1.

Note: Note that when a specific physical disk partition is assigned to BPI, it cannot be shared by another resource group.

You also need to assign the following resource types to BPI:

➤ An IP Address resource

➤ A Network Name resource

These resource types become the virtual server name and address that BPI clients must use to access the BPI Server within the cluster. This virtual server name is valid for whichever node in the cluster BPI is configured to run on and should be used as the BPI Server name when configuring the BPI Host Name property in the Infrastructure Settings:

Admin > **Platform** > **Infrastructure** > **Business Process Insight**

Having all the BPI resources in the same resource group enables you to define dependencies between the BPI services, and the disk and network resources. This means that the cluster manager does not try to start BPI unless all the required resources are available.

When you have identified how you want to manage the BPI resources within the cluster, use the Cluster Administrator to create the new resource groups.

"Summary of Configuration Requirements" on page 256 provides a summary of the nodes that you need to take to set up BPI within a Microsoft cluster environment.

Summary of Configuration Requirements

The following is a summary of the nodes that you need to complete in order to install and configure BPI on an existing Microsoft Server Cluster:

1 Make a decision on how you are going to configure the BPI Services as resource groups; that is, whether they have a dedicated resource group, or whether the BPI services are added to an existing resource group.

2 If BPI is to have a dedicated resource group, allocate a shared physical disk for BPI.

If BPI is to share an existing resource group, you can choose to allocate a different (and unique) shared physical disk for BPI, or use the shared physical disk already defined for the group. As an example, there might be a resource group already allocated for the database and you can use the same shared physical disk or allocate a different one.

3 Allocate an IP address and network name for BPI. If BPI is sharing a resource group, these details already exist. If you are creating a resource group specifically for BPI, you need to allocate a unique IP address and network name.

4 If you are configuring a new resource group for BPI, use the Cluster Administrator to create a new resource group. In this new group, create the resources for the allocated physical disk, the BPI IP address and network name.

5 Ensure that the version of the J2SE required by BPI is accessible to each node of the cluster where you want to install and run BPI.

To avoid multiple installations, you can install the J2SE on the shared disk, where it is accessible to all the BPI instances operating in the cluster. You can also choose to install it on a disk that is local to each node, in this case, the installation path must be the same on every node, for example c:\Program Files\java.

6 From a node in the cluster where the physical disk resource that you want to use for BPI is active, install BPI. You need to provide the following, cluster specific, information as part of the installation:

   a When the BPI installer prompts for the installation directory, enter the drive letter of the allocated physical disk for BPI.

   b When the BPI installer prompts for the database host name and the database is installed within the cluster environment, use the virtual network name allocated to the database; do not use localhost as the hostname.

7 For each of the other nodes where you want BPI to be able to run (or failover to), use the Cluster Administrator to make the disk for BPI active on that node. When the disk is active, you can install BPI on the node. These subsequent installations are identified and treated as reinstallations by the BPI installation procedure. The installer detects the previous installation on the shared disk and presents as defaults the installation details that it finds, including the database details. This means that any existing data in the database is preserved and it is therefore possible to add new nodes to the cluster at a later time.

8 Use Cluster Administrator to create the following new Generic Service resources in the BPI resource group for the following services:

   ➤ OVBPIAdminServer

   ➤ OVBPIBACDataSamplesProvider

   ➤ OVBPIEngine

   ➤ OVBPIEventHandler

   ➤ OVBPIJmsEventHandler

   ➤ OVBPIMetricEngine

   ➤ OVBPIModelRepository

   ➤ OVBPINotificationServer

   ➤ OVBPIServiceAdapters

   ➤ OVBPIServletEngine

   ➤ OVBPIWebServicesProvider

If you have BPI adapters that also have services, you can create generic services for these in the same way.

9 Set the OVBPIAdminServer resource to be dependent on the shared disk and virtual network name resources used by BPI. If the OVBPIAdminServer service is in the same resource group as the database, it must also be made dependent on the database server resource.

10 For each of the BPI services, set the resource name to be the same as the service name.

11 Set all the remaining BPI resources to be dependent on the OVBPIAdminServer resource.

12 Make sure that any BPI adapters access BPI using the virtual server name.

You can test the cluster installation and configuration by checking that the BPI group can be made online using the Cluster Administrator and that the BPI Modeler can access BPI using the configured virtual network name.

Reinstalling BPI In a Cluster

You can reinstall BPI in the cluster using the instructions in Chapter 5, "Re-installing BPI Server Components." You need only install BPI on one node in the cluster to reinstall, as this process updates the files and directories on the shared disk.

The only time you need to install BPI on all the nodes in the cluster is if there is a change to the local information, for example, BPI service definitions or shortcuts.

Removing BPI From the Cluster

Use the instructions in "Removing BPI Server Components" on page 122 to remove the BPI files from the shared disk.

Note that you can use the uninstall procedure only for the first node in the cluster where you want to remove BPI. This is because the all the files and directories on the shared disk have already been removed.

The following are still present on each of the other nodes in the cluster where BPI is installed:

➤ BPI component shortcuts

➤ BPI services

➤ personal configuration files for the BPI Modeler and the BPI
   Administration Console

These need to be removed as required; see the *Business Process Insight Problem
Solving Guide* for information about removing BPI Services.

Other Considerations

The following are other points that you need to be aware of when
configuring BPI to operate within a cluster environment:

➤ When you shut down the BPI components using the BPI Administration
   Console, the cluster server attempts to automatically restart them. You
   can modify this behavior from the Cluster Administrator tool; there is a
   property that specifies that specific resources are not restarted if they are
   shut down. However, if you use the Cluster Administrator and not the BPI
   Administration Console to start and stop the BPI components, you do not
   have this problem.

➤ If you have created a customized adaptor, using openadaptor, and the
   adaptor is running on a machine outside the cluster environment, there is
   additional information that you need to be aware of. This information is
   appropriate only for adaptors running outside the cluster and sending
   business events to the Business Event Handler installed within the cluster;
   the information covers configuring the adaptor to retry in order to
   re-establish a connection with the BPI cluster when there is a cluster
   failover. The information that you need is covered in the *Business Process
   Insight Integration Training Guide - Business Events* in the section describing
   Service Wrappers.

13

Backup and Recovery

This chapter describes the backup and recovery procedures for the BPI Server in cases where you are not using the Business Service Management high-availability solution (see the Business Service Management Deployment Guide for more details of the high-availability solution).

This chapter includes:

Backup and Recovery Overview

It is important to back up your BPI Server system to ensure that the BPI files and data can be recovered in the event of a media corruption. It is also important to back up your system to ensure that any processes being processed and their associated data and events can be recovered to a known state.

There are general high reliability configurations, which are recommended for all systems, for example, using mirror or Redundant Array of Inexpensive Disk (RAID) storage. Restoring or recovering from a backup should be a last resort. Your goal should be to make sure that the backups are never required.

This chapter does not cover how to set up your system for high availability - see Chapter 12, "High Availability Using Microsoft Clusters," for details. The section focuses on the recommended procedures for backing up and recovering BPI instance data.

Note: The Process Repository data is part of the BSM backup and recovery process and is not covered in this chapter.

The chapter also describes the behavior of the BPI system when specific components are shutdown, and provides recommendations for minimizing the effect of the shutdown.

If you have a system shutdown or failure, you want to keep your BPI system as synchronized as possible during the shutdown period. This chapter describes techniques for design that can help achieve this.

Key Considerations When Designing Your BPI System

The following list describes some key considerations that you need to be aware of when designing your BPI system for effective backup and subsequent recovery:

➤ For planned component shutdowns:

  ➤ Shut your BPI system down during quiet periods where possible.

  ➤ Use high reliability tools where appropriate.

➤ Design adapters to minimize the number of events lost during a system failure.

For most adapters, it is possible to design them such that they process events from the last time the adapter was executed. The *Business Process Insight Integration Training Guide - Business Adapters* describes how you do this.

➤ Minimize the number of Check Sequence arcs that you use in your processes.

When the BPI system is recovered following a shutdown, it is likely to receive a burst of out-of-sequence events. You are recommended to use Check Sequence arcs only where essential. Overuse can cause the BPI system to generate an excess of violations when it restarts. The *Business Process Insight Online Help for the Modeler* describes Check Sequence arcs in more detail.

Files and Directories to Backup

The following list describes the BPI data that is persisted within your BPI Server system and that therefore needs to be part of your backup policy:

➤ All BPI schema objects for the BPI Instance database schema.

You specified the database details when you installed BPI, you need these details in order to complete your backup.

The database for BPI should be backed up using the utilities and procedures provided as part of your database management system.

➤ Event definitions

BPI stores details about Event definitions that have been deployed to the Business Impact Engine in a series of files; these files are located as follows:

bpi-install-dir\data\repository\events

➤ Configuration information

BPI includes configuration information, which controls how BPI is setup and is relatively static. This configuration information is located in the following directory:

bpi-install-dir\data\conf\bia

➤ Template configuration files

The following directory contains the template definitions for most of the configuration files for the BPI Server and presented (managed) within the Administration Console. When you configure a component using the Administration Console, all the configuration files in data\conf\bia are regenerated, based on the content of the template files located at:

bpi-install-dir\newconfig

➤ Log files

These do not need to be backed up to restore a consistent system; however, you are advised to back up the log files, which are located at:

bpi-install-dir\data\log

➤ Data files

These are the files (Class files, Java files and so on) that contain the details of the deployment status of your definitions. BPI includes configuration information relating to the processes and data that are deployed within the database. This information changes whenever the components of a process are deployed to the Business Impact Engine. This information might be static in a production system, but variable in a development system where Business Processes are being designed, developed and tested and are therefore in a state of change.

These data files are located at:

bpi-install-dir\data\datafiles

➤ Business Event Handler

Events received by the Business Impact Engine from the Business Event Handler are not acknowledged by the Engine until they have been successfully received. Out-of-sequence events and events that contain errors (for example, missing data) are stored in the Business Event Handler hospital.

The Business Event Handler includes database tables for the hospital and event store. By default these tables are created in the same database as the rest of the BPI data, and are therefore backed up with the BPI data.

➤ Application adapter files and data

Application adapters are those developed using the Business Event Handler component to access business data from applications within your organization. The *Business Process Insight Integration Training Guide - Business Events* describes how to build application adapters.

It is usual for application adapters to be located close to (or co-located with) their data source, which means that they are likely to be running on a different system to the system where BPI is installed.

Any data used by that application adapter that needs to be synchronized with the application data (for example buffer tables populated using database triggers or adapter history files), should be backed up as part of the backup procedures for the application. Where possible the buffer table should exist within the same database as the application data that is being monitored. This makes it easier to backup the data together and keep it synchronized.

➤ TransactionVision

You need to back up the TransactionVision event data related to the BPI Server as described in the TransactionVision Help.

With the exception of the application adapter files and data, you must back up all of the data described above at the same time in order to maintain its consistency.

How often you back up the BPI data depends on how frequently the data is changing. The data can be backed up while the BPI system is operational; however, you need to ensure that you do not deploy or undeploy and definitions during the backup as the backup includes directories.

Important: As the BPI instance database cannot be synchronized with the Event definition files and data files, it is not possible to recover the BPI system to the point of failure from a backup. You must use the appropriate database management system tools to restore your system to the most recent backup and not to the point of failure.

As you cannot recover to the point of failure, if you have installed a database solely for BPI, you can consider not backing up all the transaction log files. This will reduce the amount of disk space required by the database for recovery purposes. In the case of Microsoft SQL Server, you could use the Simple Recovery model, rather than the Full Recovery model. The Simple Recovery model enables you to recover to the most recent backup. The Full Recovery model enables you to recover the database to the point of failure.

In the case of Oracle, you need to use ARCHIVELOG to enable online backups. If you do not want to use ARCHIVELOG, you can stop BPI and complete an offline backup.

How the BPI System Behaves When Components are Shutdown

This section describes the behavior of the BPI components when they are shutdown. It provides additional information that helps you understand the behavior of the system for backup and recovery purposes.

Engine Shutdown

If the Business Impact Engine component is shutdown, it is no longer able to receive and process events.

As a result of the Engine shut down, the Business Event Handler component stops accepting business events until the Business Impact Engine is restarted. When the Engine is restarted, any business events waiting to be processed are automatically submitted to the Business Impact Engine.

Database Shutdown

If the database is shut down for any reason, the Business Impact Engine can no longer communicate with it and also shuts down, as described in section "Engine Shutdown" on page 266. In addition, all other BPI components that use the database are shut down.

Notification Server Shutdown

When the BPI system determines that a notification is required, the Business Impact Engine passes a notification to the Notification Server. All notifications are held by the Business Impact Engine, in the database, until the Notification Server has confirmed that it has accepted responsibility for them.

If the Business Impact Engine shuts down before the notification is sent to the Notification Server, the notification is retried when the Business Impact Engine restarts. Note that the service status shown is the latest status. No notifications are generated for any interim changes that occur when the Business Impact Engine is shut down. For example, if the Engine shuts down when a service is in the state Critical, and the service subsequently changed from Critical to Normal and then back to Critical, when the Engine restarts, the service status is shown as Critical. In this case, the information about the state change to Normal and then back to Critical is not recorded.

Email Server Shutdown

If the Notification Server cannot make a connection to the email server to send the notification to the designated user, the Notification Server stores the notification in the database and resends it at a later time when a connection is established. The Notification Server continues the attempt to resend the notification until is reaches the value for the Maximum retries to SMTP server parameter. When this maximum is reached, the Notification Server stops retrying and logs an error in its log file.

Business Event Handler - BPI Adapter Shutdown

If the Business Event Handler shuts down for any reason, the effect of the shutdown depends on the design of the individual adapters that are connected to it.

For database adapters, where you are using buffer tables, make sure that the buffer tables are backed up and synchronized with the application data that is being tracked.

For file adapters, make sure that the current position in the file is backed up with the source application data.

This means that when the Business Event Handler BPI adapter is restarted, the application adapters can be also restarted and any events marked as **DISCHARGE_WAIT** can then be sent back into the BPI system.

Servlet Engine Shut Down

If the Servlet Engine is shut down for any reason, all the currently open BPI clients using Browser Windows also stop. You might need to restart the Web Browser Windows and the BPI clients when the Servlet Engine is restarted.

Completing an Online Backup for BPI

Before starting an online backup, you need to ensure that your database is configured such that it is capable of completing a backup online. For example, in the case of an Oracle Server, you need to make sure that the Oracle instance you have specified for BPI is running in ARCHIVELOG mode. If it is not, you cannot complete an online backup.

You should also align the BPI backup process with that of the overall BSM system backup processes in order to make sure that the data is as consistent as possible across your deployment after you have recovered the data.

To complete an online backup, follow these procedures:

1 Ensure no one is using the BPI Modeler.

2 Close the BPI Modeler.

3 Use the BPI Administration Console to shut down the Process Repository component. This prevents changes being added to the repository whilst the backup is taking place.

Do not make any further administration changes using the Administration Console until after the online backup is complete.

4 Backup the following BPI directories and their subdirectories on the system where the BPI Server is installed:

bpi-install-dir**data****repository**
bpi-install-dir**data****conf****bia**
bpi-install-dir**data****datafiles**
bpi-install-dir**data****log**
bpi-install-dir**newconfig**

If you prefer, you can backup the entire *bpi-install-dir***data** directory.

Note: Any log files currently locked by BPI processes cannot be backed up, so there are some error messages in the backup process referring to these files. Older log files are successfully backed up.

---

5 Back up the BPI data in the database using the database management system backup utilities.

6 Use the BPI Administration Console to restart the Process Repository component on the system where the BPI Server is installed.

7 Restart the BPI Modeler.

Your backup is now complete.

It is possible that the Engine Model Cleaner process is running at the same time as the online backup is completed. This is not a problem, but can result in inconsistencies between the data stored in the Engine's database and its configuration files. If you see warnings in the Business Impact Engine log file indicating that the Engine cannot find a definition in the database, it is as a result of this scenario.

Restoring your BPI Data from a Backup

This section describes how to restore your BPI system following:

➤ a complete system failure

➤ a BPI file system failure

➤ a database failure

In all cases, you are restoring both the database and the file system for the BPI components.

The following are points to note when restoring the database and the file system:

➤ Process instances created after the last database backup are lost.

➤ Process instances completed after the last database backup remain active and are unlikely to complete.

➤ Process, Data or Event definitions created since the last database backup need to be redeployed.

The Modeler detects that the current process is not the latest one and the Process is set to a deployment state of:

Not deployed, but can be deployed; an earlier revision is deployed

You can select to deploy the process and the later version that has been detected in the Process Repository is deployed.

➤ Modifications to Process, Data or Event definitions since the last database backup need to be redeployed.

The Modeler detects that the current process is not the latest one and the Process is set to a deployment state of:

Not deployed, but can be deployed; an earlier revision is deployed

You can select to deploy the process and the later version that has been detected in the Process Repository is deployed.

➤ Process, Data or Event definitions deployed since the last database backup need to be redeployed.

➤ Monitor definitions that are new or have been modified since the last database backup need to be redeployed.

➤ Email users and subscriptions configured through the Notification Server Administration Console after the last backup are lost and need to be re-added.

➤ Scripts and files created since the last backup need to be recreated; for example, any Velocity templates or Notification Server scripts that you have recently created might need to be redone.

➤ Email users and subscriptions deleted since the last database backup reappear and will need to be deleted again.

➤ Process, Data, Service and Event definitions that have been undeployed reappear and need to be undeployed again.

➤ Recent configuration changes made using the Administration Console need to be redone.

Where appropriate, you are advised to restore the database files, BPI Event definition files and BPI configuration files to maintain consistency. If you do this, you also need to re-enter any Event definitions created since the backup was completed.

Procedures to Restore your BPI System

This section describes the procedures for restoring your BPI system.

If you have had a severe system failure, you are recommended to reinstall BPI including all its dependencies, for example, the database management system (Microsoft SQL Server or Oracle Server) and the J2SE, before restoring any files.

To restore the BPI data, complete the following procedures:

1 Ensure no one is using the BPI Modeler.

2 Close the BPI Modeler to prevent new and modified definitions being created.

3 Use the BPI Administration Console to shut down all the BPI components.

Do not make any further administration changes using the Administration Console until after the recovery is complete.

4 Restore the files that you previously backed up under:

bpi-install-dir\data
bpi-install-dir\newconfig

If your BPI system has been running since the last backup you are likely to receive warnings for log files that already exist; this is expected. However, other error and warning messages should be fully investigated before continuing.

If you are using an Oracle database continue at step 6. If you are using a Microsoft SQL Server database, continue at step 5.

5 Use the Microsoft SQL Server database utilities to restore the BPI database to the last backup. Be aware that when restoring files following the reinstallation of the database, the system identifiers for the database users are not the same. As a result, the newly created BPI user is not able to access the data in the restored database.

For Microsoft SQL Server you use the Management Studio to execute an SQL command similar to the following (if you are not using the default schema name or database user defined within BPI, substitute the correct schema name or user for your database):

use hpbpiSchema;
exec sp_change_users_login 'Update_One','hpbpiuser', 'hpbpiuser';

You also need to ensure that the default database for your database user is set correctly; for example: hpbpiSchema.

Note: As the BPI database cannot be synchronized with the BPI file data relating to Event definitions and configuration information, it is not possible to recover to the point of failure from a backup. You must use a database management system command to restore your system to the most recent backup and not to the point of failure.

Continue at step 7.

6 Use the Oracle Server database tools to restore the BPI database to the last backup.

Note: As the BPI database cannot be synchronized with the BPI file data relating to Event definitions and configuration information, it is not possible to recover to the point of failure from a backup. You must use a database management system command to restore your system to the most recent backup and not to the point of failure.

---

7 Restart the Business Impact Engine, Monitor Engine and Process Repository from the BPI Administration Console.

8 Restart the BPI Modeler.

9 Open a Monitor Definer Window.

10 Re-apply any Business Process data that has been entered through the BPI Modeler, since the last backup.

11 Re-apply any monitor definition data that has been entered using the Monitor Definer, since the last backup.

You have now restored your BPI data.

12 Undeploy any Process, Data, Service or Event definitions that are deployed and that were undeployed before you restored your BPI system.

13 Use the BPI Administration Console to restart the remaining BPI components.

14 Delete any email users and subscriptions that have reappeared since you restored your BPI system.

15 Use the Process Intervention tool to re-synchronize with the ODB.

Now go to the next section "Completing the BPI Recovery Procedure" on page 273 to complete the recovery procedure.

Completing the BPI Recovery Procedure

This section describes points to note after you have completed a BPI recovery from backup.

After you have completed the procedures described in "Restoring your BPI Data from a Backup" on page 269, you have recovered to the latest backup. You have therefore potentially missed events that are relevant to your Business Processes. This means that BPI can raise erroneous events, particularly in cases where you have Check Sequence arcs defined.

You do also have the option to take a differential backup, which is faster. However, you are advised to adopt the backup and recover policies adopted for other applications within your organization, using the information provided in this section, which is specific to BPI.

14

Log Files

This chapter describes the log files for the BPI components.

This chapter includes:

➤ BPI Server Log Files on page 275

➤ BPI Modeler Log Files on page 276

➤ openadaptor Log Files on page 277

➤ BPI Service Wrapper Log Files on page 277

➤ Logging Properties Structure and Description on page 277

➤ Adjusting Logging Levels on page 279

BPI Server Log Files

The BPI Server log files are accessible through the BPI Administration Console. For information about the BPI Modeler log files, refer to section "BPI Modeler Log Files" on page 276.

The log files for the Administration Console are called bia_adminservern_n.log and bia_adminclientn_n.log, and are located in the following directory on the system where BPI components are installed:

bpi-install-dir\data\log

The BPI log files are text files that you can open using a text editor. The structure and naming of the BPI log files can be configures and is described in section "Logging Properties Structure and Description" on page 277.

There are a number of logging levels available that change the volume and level of logging information provided by each BPI component; for most components these levels are:

➤ Info

➤ Fine

➤ Finer

These logging levels listed are in descending order of level and volume of data reported. These log levels compliment the levels used by other HP BTO Software products, to provide a consistent level of logging across your HP Software product portfolio.

By default the log levels are set at the INFO level. This level includes WARNING and SEVERE log levels.

In the case of components using Log4J, for example, the Business Event Handler, the levels are:

➤ Info

➤ Trace

You can set the levels of logging through the BPI Administration Console using the Logging option.

The generated log files can be viewed through the BPI Administration Console, or from the following location:

➤ *bpi-install-dir*\data\log (Windows)

➤ *bpi-install-dir*/data/log (HP-UX)

In the case of a failure in the BPI Administration Console, you might need to open the log file directly as you cannot access it through the console.

BPI Modeler Log Files

The Modeler log files are not accessible using the BPI Administration Console.

The location of the Modeler log files is described in the *Business Process Insight System Administration Guide*. The properties of the BPI Modeler log files is described in section "Logging Properties Structure and Description" on page 277.

openadaptor Log Files

The openadaptor log files are not accessible using the BPI Administration Console.

The location of the log files for adaptors developed using openadaptor depend on how your adaptors are configured and are described in the *Business Process Insight Integration Training Guide - Business Events*.

BPI Service Wrapper Log Files

The Windows Service Wrapper log files are located at:

bpi-install-dir/bin/bia/wrapper.log

Logging Properties Structure and Description

Most of the BPI logging property files all have a similar structure; the following is an extract from the Business Impact Engine property file:

```
####################################################################

# File output is in BIA's log directory.
java.util.logging.FileHandler.pattern=C:\\Program Files\\HP \\HP Business
Process Insight\\data\\log\\bia_bce%u_%g.log
java.util.logging.FileHandler.limit=10000000
java.util.logging.FileHandler.count=100
java.util.logging.FileHandler.formatter=java.util.logging.SimpleFormatter

####################################################################
# Logging levels
####################################################################
```

```
# Available levels are OFF, SEVERE, WARNING, INFO, FINE, FINER, FINEST,
ALL

# Default global logging level.
.level=INFO

# BIA logging level.  Recommended setting is INFO.
# The handling of Events is logged at FINE.  The contents of Events are logged
at FINER.
# However these settings will produce a large amount of logging which may
consume disk space.
com.hp.ov.bia.level=INFO

# This is intended to turn off INFO and below in Hibernate.
net.sf.hibernate.level=WARNING

# This is intended to turn off RMI logging.
sun.rmi.level=OFF
```
The Business Event Handler log files have a different structure, which is
described in the *Business Process Insight Integration Training Guide - Business
Events*.

In the case of the Business Impact Engine example, the following parameters
affect how information for the Business Impact Engine component is
logged:

➤ java.util.logging.FileHandler.pattern

This parameter determines location and name of the log files. The
following tags are used within this parameter:

➤ %g

This is a file generation number and is used as the log files are rotated.
When the maximum number of bytes that should be written to a file is
reached, or the BPI component is restarted, a new log file is opened
and named bia_bce0_0.log. The previous log file is then renamed, for
example, bia_bce0_1.log.

➤ %u

This is a unique number that is used to resolve conflicts. This is the
second part of the number in the log file name, for example: 0_0, 0_1
or 0_2. %u generates the 0, 1 or 2 in these file names.

➤ java.util.logging.FileHandler.limit

This parameter specifies an approximate maximum amount to write (in bytes) to a log file. If this is zero, then there is no limit. If this parameter is not set, the BPI system default is zero, no limit.

➤ java.util.logging.FileHandler.count

This parameter specifies how many output log files to cycle through. If this parameter is not set, the BPI system default is one file.

➤ java.util.logging.FileHandler.formatter

This parameter specifies the name of a Formatter class to use for the log files. If this parameter is not set, the BPI system defaults to using the java.util.logging.XMLFormatter.

Adjusting Logging Levels

You can adjust the logging levels for the BPI components as described in *Using Business Process Insight*. However, if you set the logging levels to be FINE or FINER, a significant amount of logging information is generated that can quickly fill up disk space. Much of this logging information is RMI logging information.

You can selectively set the RMI logging information to the INFO level, and keep the BPI component logging at FINE or FINER, to significantly reduce the amount of information produced.

To do this, modify each of the BPI component logging property files (listed in *Using Business Process Insight* and add the following lines:

```
# This is intended to turn off RMI logging.  RMI only logs at
#FINE or FINER, so setting these to
# INFO effectively turn off RMI logging.
sun.rmi.server.level=INFO
sun.rmi.loader.level=INFO
sun.rmi.server.level=INFO
sun.rmi.client.level=INFO
sun.rmi.dgc.level=INFO
sun.rmi.loader.level=INFO
sun.rmi.transport.level=INFO
sun.rmi.transport.tcp.level=INFO
sun.rmi.transport.proxy.level=INFO
```

You also need to change the template file for the BPI logging property files. If you do not, the changes that you made are reset by the contents of the template file whenever you apply changes to BPI components through the BPI Administration Console.

The template logging property files are located as follows:

*bpi-install-dir***\newconfig\DataDir\conf\bia\...**

You now need to run the following file to make the changes to the configuration file, or files, active on your BPI system:

*bpi-install-dir***\bin\biaadmin.bat applyconfig**

This causes all template configuration files to replace the active configuration files with any new or revised settings.

Note that if you reinstall BPI for any reason, the template logging property files are overwritten and you will need to reapply your changes.

Part IV

Appendixes

A

Database Tables Defined for BPI

This appendix provides tables that list all the database tables and views that are created for BPI. Many of these tables are for internal use only and should not be modified.

This list enables you to make sure that all the tables are included in any backup and recovery procedures and to check for any potential database table name clashes.

The capitalization used in the tables are for ease of reference only. How the table names are represented in your database is database specific, for example, for an Oracle Server all table and view names are upper case.

This chapter includes:

Complete List of BPI Database Tables

The following table lists all the BPI database tables and "Complete List of BPI Database Views" on page 284 lists all the views.

Table Name	Internal Use Only (yes/no)
AdaptorConfig	yes
Arcs	no
BACDataSampleQueue	yes

Table Name	Internal Use Only (yes/no)
BACDataSampleDestinations	yes
BACSynchronizeStatus	yes
Data_Objects	no
DBusmh_Attribute	yes
DBusmh_EditableAttribute	yes
DBusmh_Patient	no
DBusmh_Role	yes
DBusmh_User	yes
DBusmh_UserRole	yes
EventQueue	yes
Event_Store	no
FlowDataFilters	yes
FlowNodeZones	yes
Flows	no
Flow_Group_Statistics	no
Flow_Instance_Table	yes
FlowZoneMaps	
ImpactThreshold	yes
Metric_CustomTypes	no
Metric_Definitions	yes
Metric_Dim_Dates	no
Metric_Dim_Flows	no
Metric_Dim_Flow_Instances	no
Metric_Dim_Groups	no
Metric_Dim_Metrics	no

Table Name	Internal Use Only (yes/no)
Metric_Dim_Thresholds	no
Metric_Event_Seq	yes
Metric_Events	yes
Metric_Events_Engine_Status	yes
Metric_Events_Seqs_To_Remark	
Metric_Events0 (to 18)	
Metric_Fact_Alerts0 (to 18)	no
Metric_Fact_Statistics0 (to 18)	no
Metric_Fact_Values0 (to18)	no
Metric_Generate_Errors	yes
Metric_Notification_Check_Time	yes
Metric_Staging_Statistics0 (to 18)	yes
Metric_Staging_Stats_Modified0 (to 18)	yes
Metric_Stored_Procedures[a]	yes
NodeIdAndResourceStatus	yes
NodeInstToBeDeleted	yes
Nodes	no
Nodes2Resources	no
Node_Group_Statistics	no
Node_Instance_Table	yes
Node_Instance_CompletedTimes	no
Node_Instance_StartedTimes	no
NS_Digesterstore	yes
NS_DigesterSubscriptions	yes
NS_EmailRetry	yes

Table Name	Internal Use Only (yes/no)
NS_EmailSubscriptions	yes
NS_EmailUsers	yes
NS_OVORetry	yes
NS_OVOSubscriptions	yes
NS_ScriptRetry	yes
NS_ScriptSubscriptions	yes
NS_SLAEmailSubscriptions	yes
NS_SLOEmailSubscriptions	yes
OVIS_AlarmStatus	yes
OVISTimeStamps	yes
Resources	no
RMIEventRetry	yes
SOAPEventRetry	yes
UcmdbIdToMonitorId	yes
UcmdbIdToOProcessId	yes
UcmdbToStepId	yes
Version	yes

a. this table is available only with Microsoft SQL Server.

Complete List of BPI Database Views

The following table lists all the BPI database views.

View Name	Internal Use Only (yes/no)
Dbusmh_View	yes
Flow_Instance	no

View Name	Internal Use Only (yes/no)
Metric_Events	yes
Metric_Fact_Alerts	yes
Metric_Fact_Statistics	yes
Metric_Fact_Values	yes
Metric_Staging_Statistics	yes
Metric_Staging_Stats_Modified	yes
Metrics	no
Metric_Values	no
Node_Instance	no

Complete List of BPI Database Indexes

The following table lists the BPI database indexes.

Index Name	Internal Use Only (yes/no)
Flow_Group_Statistics1	yes
Flow_Instance_Idx_*nnn*	yes
Idx*n*	yes
Metric_Dim_Dates_Idx_*nnn*	yes
Metric_Dim_Groups_Idx_*nnn*	yes
Metric_Events_Idx_*nnn*	yes
Metric_Fact_Groups_Idx_*nnn*	yes
Metric_Fact_Alerts_Idx_*nnn*	yes
Metric_Fact_Statistics_Idx_*nnn*	yes
Metric_Fact_Values_Idx_*nnn*	yes
Metric_Staging_Statistics_Idx_*nnn*	yes

Index Name	Internal Use Only (yes/no)
Node_Group_Statistics1	yes
Node_Instance_Comple_Idx_*nnn*	yes
Node_Instance_Idx_*nnn*	yes
Node_Instance_Starte_Idx_*nnn*	yes
Sys_*string*	yes

There are also stored procedures, database triggers and other database scripts defined by BPI and potentially defined by you. These also need to be taken into account for backup purposes or if you need to move the BPI data.

It is strongly recommended that you create an Oracle User or Microsoft SQL Server Database specifically for BPI, as it makes the task of identifying this BPI data much easier.

B

Migrating the BPI Server Installation to a 64-Bit System

This appendix describes how to migrate a BPI Server installation to a 64-bit system, either from a 32-bit system or a 64-bit system.

This chapter includes:

➤ Procedure Overview on page 287

➤ Preparation on page 288

➤ Installation on page 289

➤ Migration on page 290

➤ Start the Products on page 292

Procedure Overview

It is assumed that the BPI database will be re-used at its current location. If the BPI Instance database is to be migrated, see step 4 in "Preparation" on page 288 to migrate the database with either the MS SQL Server or Oracle tools.

For the purposes of this procedure, the source and target systems are identified as follows:

➤ Source BPI Server system name: > *<SourceHostName>*

➤ Target BPI Server system name: > *<TargetHostName>*

➤ Source BPI Instance database hostname > *<DBsourceHostName >*

If you are not migrating the BPI Instance database, the source and target BPI Instance database hostnames will be the same.

➤ Target BPI Instance database hostname > *<DBtargetHostName>*

Preparation

Complete the following steps to prepare for the migration:

1 On *<SourceHostName>,* be sure the BPI Instance database is referenced by its actual name, *<DBsourceHostName >*, and not as **localhost**:

   a Open the Administration Console and navigate to the *<MS SQL or Oracle>* **Server Access to BPI Instance Database** component.

   b Check to see if the *<MS SQL or Oracle>* **Server hostname** in the right panel is *<DBsourceHostName >*.

2 If the hostname shown is **localhost**, do the following:

   a Back up then edit the following files:

    ➤ *<BPI_INSTALLATION_HOME>***HP Business Process Insight\data\conf\bia\masterconfig.properties**

    ➤ *<BPI_INSTALLATION_HOME>***HP Business Process Insight\data\conf\bia\admin.cfg**

   b Modify references of **localhost** in both files to the appropriate Instance database server name (*<DBsourceHostName>*). The exception being the **local_host=localhost** line near the top of the **admin.cfg** file. Do not modify this line.

3 In the *<MS SQL or Oracle>* **Server Access to BPI Instance Database** component, verify that the appropriate server name is displayed for the BPI Instance database **hostname**.

4 If you are migrating the BPI Instance database to a new host, perform that operation here using the appropriate SQL Server or Oracle tools.

Essentially, you are replicating the original BPI Instance database <***DBsourceHostName***> to the desired target host <***DBtargetHostName***>.

If you are not migrating the Instance database to a different location, ignore this step.

5 On <***TargetHostName***>, be sure the database service is running.

6 On <***SourceHostName***>, be sure:

➤ Database service is running.

➤ BSM is **not** running. Verify using **myStatus.html** by selecting **Start**->**All Programs**->**HP Business Service Management**->**Administration**->**HP Business Service Management Status**. If BSM is running, disable it.

➤ BPI Server is **not** running — if it is, stop all BPI components through the BPI Administration Console, see "BPI Administration Console" on page 133.

7 On <***DBtargetHostName***>, be sure the database service is running.

Installation

Complete the following steps on <***TargetHostName***>:

1 Install BSM 9.x (if not already installed).

2 Start/Enable BSM. Be sure it is running using **myStatus.html**.

3 Install BPI Server 9.x. During installation, note the following:

a Do **not** set up the database through the installation wizard - select **No** in the BPI Database Creation Dialog.

b Be sure the LWSSO init strings are compatible between BPI and BSM.

c Do not launch or start up any BPI components.

Migration

Complete the following steps to perform the migration:

1 From the *<SourceHostName>* to the *<TargetHostName>* system, copy the entire file/directory structure under *<BPI_INSTALLATION_HOME>***HP Business Process Insight\data**** from the source to the target. Overwrite directories and files as necessary.

2 On *<TargetHostName>*, edit each of the following files (2a, 2b, and 2c).

Be sure the line items identified or tokens listed in these files are correctly set as specified below. Some of these tokens may not be present in the files copied over from the *<SourceHostName>* box, so you might need to add and set them within the file.

a Edit *<BPI_INSTALLATION_HOME>***HP Business Process Insight\data\conf\bia\admin.cfg** as follows:

1. Search for **DatabaseMSSQLJdbcUrl** or **DatabaseOracleJdbcUrl** depending on the BPI Instance database type being used.

2. Be sure that the server identified is the *<DBtargetHostName>* server.

b Edit *<BPI_INSTALLATION_HOME>***HP Business Process Insight\data\conf\bia\masterconfig.properties** as follows:

1. Add or modify the following tokens as shown:

ESC_BIA_OVIS_MSSQL_HOST=*<DBtargetHostName>*

BIA_OVIS_MSSQL_HOST=*<DBtargetHostName>*

BIA_ENGINE_MSSQL_HOST=*<DBtargetHostName>*

HPBPI_DASHBOARD_HOSTNAME=*<TargetHostName>*

BAC_SAMPLES_RECALCULATE_DELAY=0

BIA_MAX_GROUP_STATISTICS=100

BIA_INSTALLED_COMPONENTS=eventhandleradaptor,notification, modelrepos,engine,views,metrics,webservices,adaptorframework,b acdatasamplesprovider,jmseventhandleradaptor,modeler

BIA_METRIC_MAX_GROUP_STATISTICS=100

BIA_USE_USERDEF_LOG_VIEWER=true

> **ESC_HPBPI_PROCESS_REPOSITORY_MSSQL_HOST=<*TargetHostNa me*>**
>
> **HPBPI_PROCESS_REPOSITORY_MSSQL_DATABASE=<*BPI database name in Process Repository Database*>**
>
> **BIA_OVIS_IMPORTER_HOST=<*TargetHostName*>**
>
> **BIA_SMTP_HOST=<*TargetHostName*>**
>
> **BIA_METRIC_THRESHOLD_ACTIVE_CHECK_DELAY=300**
>
> **HPBPI_PROCESS_REPOSITORY_MSSQL_HOST=<*TargetHostName*>**
>
> **BIA_OVIS_ORACLE_HOST=<*DBtargetHostName*>**
>
> **BPI_SERVER_HOSTNAME=<*TargetHostName*>**
>
> **ESC_BIA_ENGINE_MSSQL_HOST=<*DBtargetHostName*>**

2. Verify any listed directory paths in this file conform to the appropriate BPI installation locations for <*SourceHostName*>, <*TargetHostName*> and <*DBtargetHostName*>. Be particularly aware of references to "Program Files (x86)" versus "Program Files" for the <*TargetHostName*>.

c Edit <*BPI_INSTALLATION_HOME*>**HP Business Process Insight\data\conf\bia\bia-model-repository-db.properties** as follows:

1. You may need to modify the specified **DatabaseName** to match the BPI **DatabaseName** on the <*TargetHostName*>. Be sure that the **hostname** and **Databasename** are correctly listed in this file.

2. Search for the parameter **DatabaseName=<*BPIDataBaseName on TargetHostName*>** to help identify the items that may need editing in the file.

3 On <*TargetHostName*>, execute:

<*BPI_INSTALLATION_HOME*>**HP Business Process Insight\bin\biaadmin applyconfig**

4 From the utilities provided by the new 9.x installation on <*TargetHostName*>, upgrade the reused BPI Instance database located on <*DBtargetHostName*>. From <*TargetHostName*>, execute:

a For SQL:

> *<BPI_INSTALLATION_HOME>***bin****bpidbsetup.bat**
> **-upgradeDatabase -mssql** *<DBtargetHostName>* *<databasePort>*
> *<bpiUser>* *<bpiUserPassword>* *<bpiSchema>*

  b For Oracle:

> *<BPI_INSTALLATION_HOME>***bin****bpidbsetup.bat**
> **-upgradeDatabase -oracle** *<DBtargetHostName>* *<databasePort>*
> *<bpiUser>* *<bpiUserPassword>* *<db_sid>*

  c Execute **bpidbsetup.bat -help** to see more detailed usage information.

Start the Products

Complete the following steps to start the product:

  1 On *<TargetHostName>*, start/enable BPI components using the BPI Administration Console.

  2 Verify the configurations in the Administration Console *<MS SQL or Oracle>* **Server Access to BPI Instance Database** component as described in "Preparation" on page 288.

C

Reporting Problems to Hewlett-Packard

Before reporting problems to Hewlett-Packard, make sure that you have the following information available.

This chapter includes:

➤ BPI Server Problem on page 293

➤ BPI Modeler Problem on page 294

BPI Server Problem

If the problem is with the BPI Server, supply the following information to HP:

➤ Hostname and host ID of the machine where the BPI component is running.

➤ Operating system and its version, plus patches on all the systems where the BPI components are running.

➤ Copy of the log files for all the BPI Server components.

➤ Copy of any relevant trace files.

➤ Note any error messages reported with the problem.

➤ Copy of your installation log files, which are located at:

  ➤ *bpi-install-dir***HP_Business_Process_Insight_InstallLog.log**

  ➤ *bpi-install-dir***HP_Business_Process_Insight_DBSetupLog.log**

➤ Copy of the latest InstallAnywhere installation log files, which are located at the root directory of the drive where you installed BPI; for example **C:**. The InstallAnywhere installation log files for BPI components all start with the string **hpbpi**.

➤ All BPI install and uninstall files in the root directory where you have installed BPI.

➤ Copy of the database log file.

➤ A copy of the BPI database tables at the time of the problem; see Appendix A, "Database Tables Defined for BPI," for a complete list of the database tables used by BPI.

➤ Screen captures of any error dialog boxes.

➤ A copy of the output from running the Integrity Checker. Using the Integrity Checker is described in section "Verifying Your Installation Using the Integrity Checker" on page 110.

BPI Modeler Problem

If the problem is with the BPI Modeler, supply the following information to HP:

➤ Copy of **modelclasses.cfg**, which is located in the *bpi-install-dir***data****conf****bia** directory.

➤ Copy of any Process, Data and Event definitions from the Modeler. These can be exported from the Modeler as described in the "Modeler Management" chapter in the *Using Business Process Insight Guide*.

Index

C

N

Index