HP Business Availability Center

for the Windows and Solaris operating systems

Software Version: 8.00

Business Process Insight Reference Information

hp ®

invent

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

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Third-Party Web Sites

HP provides links to external third-party Web sites to help you find supplemental information. Site content and availability may change without notice. HP makes no representations or warranties whatsoever as to site content or availability.

Copyright Notices

© Copyright 2007-2009 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Intel®, Pentium®, and Intel® XeonTM are trademarks of Intel Corporation in the U.S. and other countries.

JavaTM is a US trademark of Sun Microsystems, Inc.

Microsoft®, Windows®, Windows NT®, and Windows® XP are U.S registered trademarks of Microsoft Corporation.

Oracle® is a registered US trademark of Oracle Corporation, Redwood City, California.

Unix® is a registered trademark of The Open Group.

Documentation Updates

This guide's title page 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

You can 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 Support Online 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 HP Software 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 find more information about access levels, go to:

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

To register for an HP Passport ID, go to:

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

Contents

1 Welcome to This Guide

This guide provides reference information relating to Business Process Insight (BPI) and its integration as an Application with Business Availability Center.

BPI comprises a number of different components, some of which can be distributed to other systems; however, there are some BPI components that always need to be located together and these are collectively referred to as the BPI Server components.

For information about using and managing BPI, refer to *Using Business Process Insight*.

This chapter includes:

- How This Guide Is Organized on page 9
- Who Should Read This Guide on page 10
- Getting More Information on page 10

How This Guide Is Organized

This guide covers the following topics:

- Chapter 2, Introduction to Business Process Insight

  This chapter provides an overview of BPI and the BPI terminology.

- Chapter 3, BPI Server Architecture

  This chapter describes the architecture of the BPI application within BAC. It also describes how the BPI components relate to each other.

- Chapter 4, BPI and TransactionVision

  This chapter provides details of how to integrate BPI and TransactionVision. It describes how to import transaction event definitions into the BPI Server in order that you can link them into your Business Processes using the BPI Modeler.

- Chapter 5, Log Files

  This chapter describes the log files for the BPI components including properties and structure.

- Appendix A, Expression Grammar in Process, Data, and Filter Definitions

  This appendix lists the rules for the grammar that can be used for Business Process progression rules and expressions, and expressions within Business Process monitor filter definitions.

- Appendix B, Coercion Rules

  This appendix describes the rules for how properties are coerced when evaluating a binding, filter and assignment expressions within the BPI Modeler.

In addition, you can configure your BPI system to receive SOA business events by configuring an event source for SOA Manager. This is described in the *HP Business Process Insight Training Guide - Business Events*.

Who Should Read This Guide

This guide is intended for anyone who wants more detail about how the BPI components interrelate, and information about BPI's integration with TransactionVision.

The guide also provides information about the syntax and grammar used for expressions within the BPI Modeler.

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

Getting More Information

For a complete list of all online documentation included with HP Business Availability Center, additional online resources, information on acquiring documentation updates, and typographical conventions used in this guide, see the *HP Business Availability Center Deployment Guide* PDF.

2 Introduction to Business Process Insight

This chapter provides a description of the key concepts relating to BPI.

The chapter comprises the following topics:

The BPI Application Within Business Availability Center

BPI is a Business Availability Center Application that enables you to make better business decisions by improving the quality of the business information available to you. BPI provides you with visibility into the health and performance of the Business Processes that are running over your IT infrastructure. Using this information, you can assess the financial and business impact of delays or blockages in a process due to an IT performance problem or other incident such as an IT outage.

BPI achieves this by linking and relating the information held for your business and its IT infrastructure. It takes information, in the form of events, from both business applications and from the IT infrastructure, and presents these events in business terms; for example, rate of loss of orders, backlogs accumulating, value of these backlogs and potential impact on key customers. In other words, BPI enables you to visualize and then track the state, or health, of Business Processes within your organization.

Specifically you can use BPI to:

- model Business Processes and then measure key business data associated with these processes; for example, value of order or type of customer.

  To this end, BPI enables you to define Business Process monitors and instance thresholds, which then make it easier for you to set business objectives and receive instance threshold violations, based on these objectives.

- correlate each Step in the Business Process with the IT infrastructure on which it depends and provide IT with the visibility to show its contribution to the business.

To achieve this, BPI uses business data events and IT infrastructure events to monitor the status of all the Business Processes that you have modeled. As a result, when a problem occurs within the IT infrastructure, BPI reports on the instances of the processes that are being monitored and that are impacted; for example, BPI might report on the number of orders that are affected for a specific customer, for all customers or for a group of customers (gold, silver, bronze-level customers).

The key business data associated with the process are then used to determine the business impact of the problem. Business impact typically takes the form of financial impact, but it can also be the number of key customers affected, or the number of urgent orders affected, and so on.

BPI can also record Business Process monitors such as the time it takes a particular business event to move between steps in a Business Process. This data can also be used to report on Business Process Service Level Objectives (SLO). In addition, key performance indicators (KPIs) can be set on the BPI-related configuration items in the UCMDB. These KPIs are then be compared against objectives that you define.

When all this information is available, you can use the Business Availability Center SLM Application to raise alerts when BPI monitors do not achieve the objectives that are set. You then use then use the Dashboard Application to:

- proactively monitor your service level agreements (SLAs).

  The Dashboard Application enables you to see which SLAs are at risk and enables you to navigate to see which specific CI is potentially causing you to breach a service contract.

- calculate whether the measurements of an SLA are within the limits set out in a service level agreement (as defined in Service Level Management).

- calculate whether an SLA is in breach of contract, or might be in breach in the near future.

Figure 1 shows the structure of the BPI Business Process models as CIs within the UCMDB.

Figure 1   Business Availability Center Business Process Model
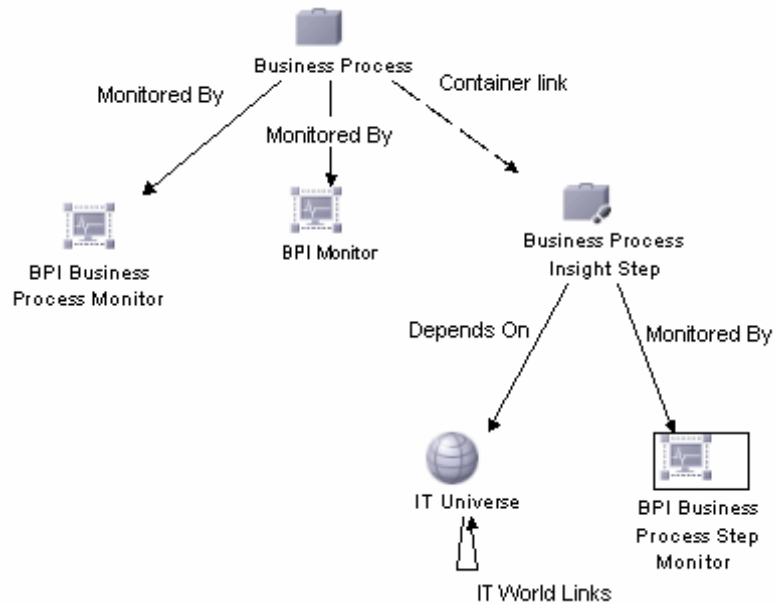
Each of the elements in Figure 1 is represented as a CI within the UCMDB. The BPI application sends data samples for the Business Process monitor CIs that you define and these data samples are then displayed on the process diagram and in the Business Process Application as KPIs.

What is a Business Process?

Business Processes are the business activities within your business; for example, for a telecom service provider a Business Process might include line provisioning, customer billing and equipment repairs. These Business Processes might not always be obvious, and they might not be formally recognized, but they always exist within a business. A Business Process (or a set of Business Processes) typically includes all the activities of the business.

For an insurance company, business activities might include setting up a new policy for a customer and processing an insurance claim. In addition to these line-of-Business Processes, a Business Process can also be an internal administration process, such as HR processes, payroll, and so on.

Business Processes are not necessarily exclusive to one company; many Business Processes involve customers and suppliers. Business Processes can be automated, using Business Process management software; however, only a fraction of Business Processes are fully automated. Most Business Processes are manual and some exist only in the behavior of the organization, that is, they are not formally documented.

What is an BPI Process?

A process (for BPI tracking purposes) does not need to include all the activities in a Business Process; it need contain only those activities that you need to understand in order to determine that your business is functioning optimally. As an example, a vehicle dashboard presents a summary of the key activities that you need to be aware of for the vehicle; it does not present all the information about all the vehicle operations.

The business activities that you need to model within BPI for your business are those required to obtain meaningful impact data about the health of your business. The fewer business activities that you track, the easier the BPI processes are to create and manage and visualize.

Structure of an BPI Process

Within BPI, a process is presented as a series of Steps and connections, which are called Arcs. The Steps represent the business activities and the arcs connect the Steps to show the expected direction of flow within the process; this is shown in Figure 2.

Figure 2   Basic BPI Process
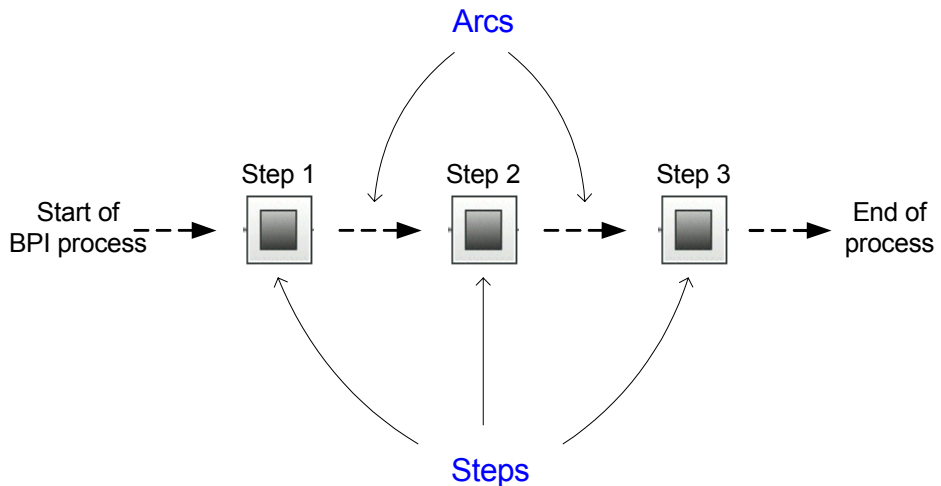
A Step in a Business Process might represent one or more business activities, according to how much detail of the Business Process you are modeling in order to track it.

BPI enables you to represent a business service in a graphical form. This allows you to show the logical structure of the elements of the service. The graphical representation also provides visual information relating to the progression of the process.

Using the graphical processes enables you to more easily connect business activities into a logical sequence using as few steps as possible; however, you do need to make sure that you model sufficient of the structure of the process to be able to report on the impact of problems in the process.

Example BPI Process

Figure 3 on page 18 is an example of a Business Process for an organization that provides building materials primarily for professionals, but it also has a store where the public can purchase building supplies. The organization is structured as a front office, middle office and back office:

- The front office includes:
    — a Web site for online access
    — a shop
    — a call center

    An ISP hosts the Web site and the shop has a dial-up connection through a point-of-sale checkout. The Call Center is out sourced.
- The middle office is where the orders are processed and validated.
- The back office is where the billing and the execution of the orders occur.

In this example, the middle and back office operational systems are managed using HP Business Availability Center.

Figure 3 shows the significant activities of the Business Process represented graphically in an BPI process. This is how they might appear within the BPI Modeler, which is the graphical tool for creating BPI processes.

As mentioned earlier, you do not need to include all the Steps in the Business Process. The Steps that you include are those that relate to the key business activities that need to be tracked.

Figure 3   Stock Control Process
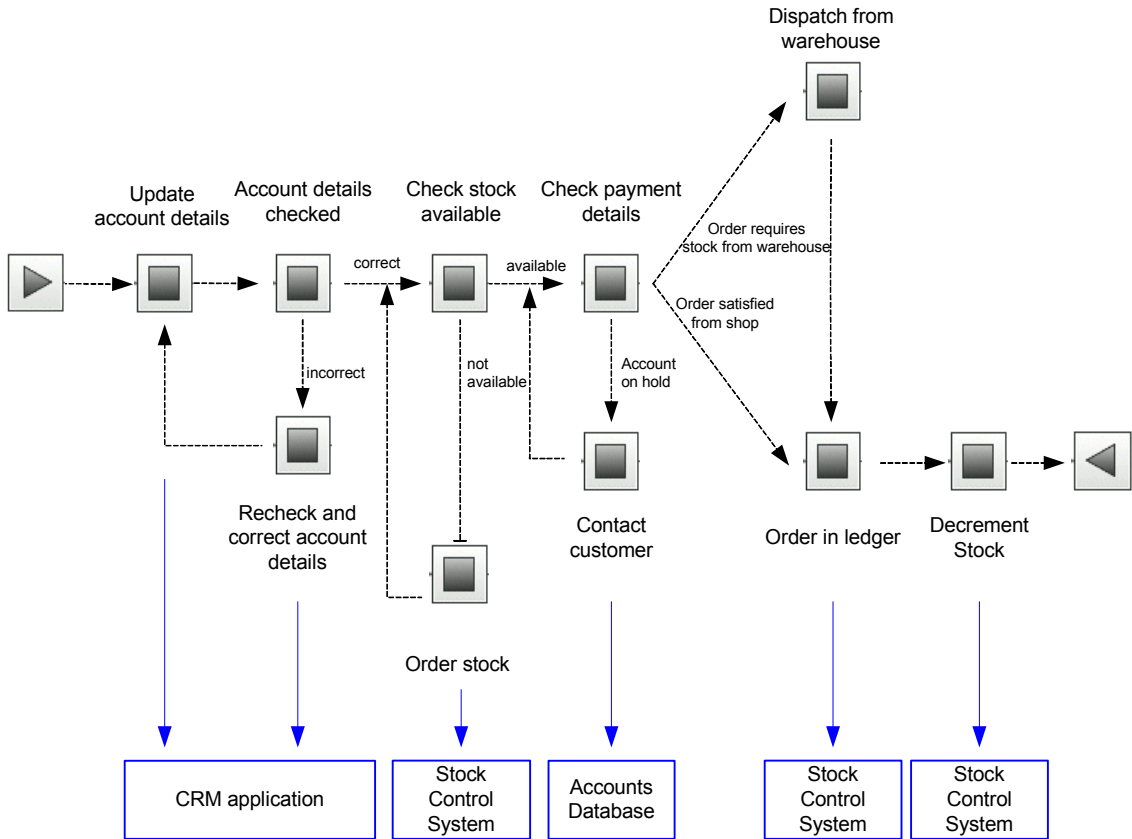
The process shown in Figure 3 utilizes underlying business applications and also includes IT operational resources. Some of the business applications might be operationally tracked and therefore represented as IT operational resource CIs within Business Availability Center. There might also be other business applications that are not being operationally tracked, but are still providing business application data for BPI. In this example, information

about the health of the Stock Control System is used at several points in the BPI Business Process and Figure 3 shows this dependency. Within the BPI Modeler, these resources are properties of the Steps in the process.

Figure 4 shows the hierarchy of the IT operational resources and how they relate to the applications that provide the resources.

Figure 4   Stock Control Resource Hierarchy
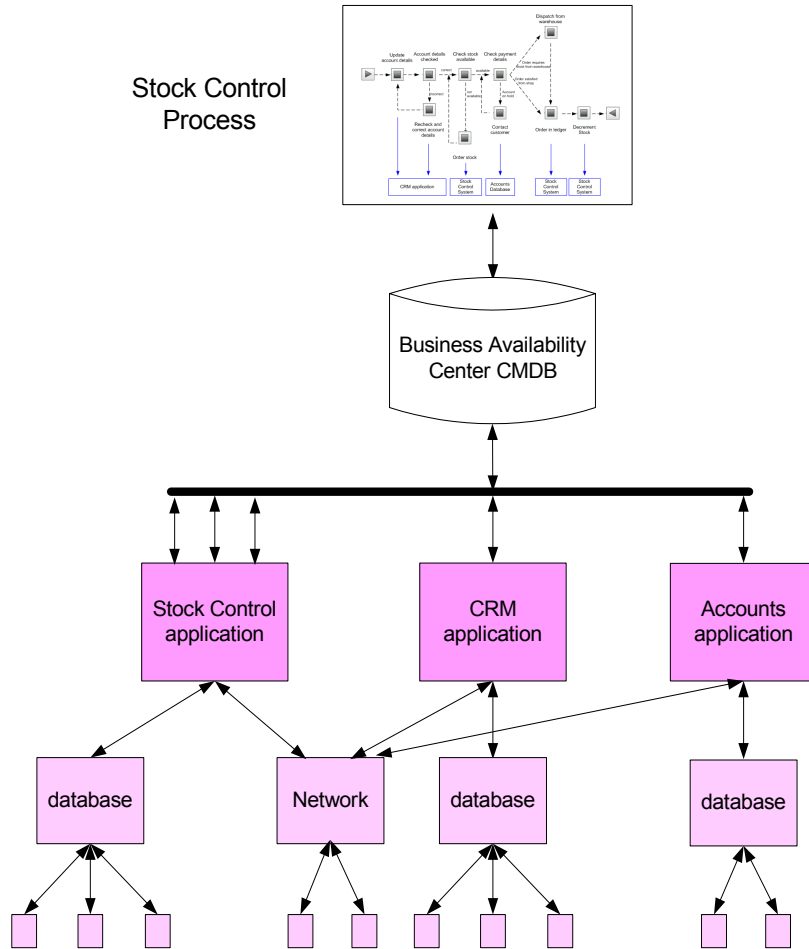
In this example, the Stock Control System is used at several points in the Business Process and the Stock Control process diagram shows this interaction.

Business Events

In order to integrate, or import, business event data from your business applications, you use the BPI adapter technology. This technology enables you to generate business events from sources such as files and databases, and business applications including SAP, PeopleSoft and JMS. These adapters enable business data to be accepted as the source data for BPI business events, which are received through the BPI Business Event Handler component.

In the context of BPI, an event is the notification of some change to the underlying system that is being monitored and where this state change causes a subsequent change in state of the modeled Process instances. This might be directly through the data within the event or through some transformation of the data.

The following sections describe business event components for BPI.

Business Events and openadaptor

The Business Event Handler is based on the openadaptor technology and openadaptor is the underlying mechanism for receiving business events. openadaptor is an open-source Java/XML-based software platform, which provides business system integration with little or no programming. openadaptor provides many ready-built interface components including adapters to capture business events from sources such as files, databases, JMS, LDAP, Tibco and MQ Series.

BPI also provides file and database adapters for use specifically with BPI; however, openadaptor is an open source product and easily configurable; therefore, you can extend the number of adapters for use with BPI as required.

For more information on openadaptor refer to:

http://www.openadaptor.org/

BPI uses a specific version of the openadaptor platform, and the version used is available as part of the BPI software.

BPI also includes specific adapters for SAP; see section BPI Accelerator for SAP Applications on page 21.

BPI Accelerator for SAP Applications

In addition to using the openadaptor framework to create solutions for integrating business applications, you can use the BPI Accelerator for SAP to enable you to track the status fields within an IDoc. Using the BPI Accelerator, you can create processes and business events that are based on the status changes within one or more IDocs.

The Accelerator is based on the openadaptor framework and uses the Business Event Handler to access the status information within an IDoc header. Refer to the *Business Process Insight Accelerator for SAP Guide* for full details of installing and using the Accelerator.

IT Processes

Within many organizations the IT processes and transactions are often hidden within the environment in which they have been automated. The processes might have started as fully documented process models, but as they are adapted to the needs of the organizations, the processes change and the related documentation becomes out of date. Using BPI helps you visualize IT processes that are hidden within the organizational structure and then set measurements on the IT operational resources that make up the processes in order to report on the status of the operation of these resources.

You model your IT operational resources within BAC (as CIs) and then use BPI to link the CIs to process Steps. BPI can then obtain status information relating to these operational resources for Business Processes that you have modeled within BPI.

TransactionVision

BPI can also utilize the transaction models produced by HP TransactionVision as input for a Business Process model. HP TransactionVision tracks and measures individual transactions as they progress through your organization. TransactionVision provides the link between low-level systems and network monitoring, and high-level Business Process monitoring in order to link your IT objectives with your business objectives.

HP TransactionVision does not assume that a model exists to describe all possible transaction flows. TransactionVision monitors transactions and discovers the transaction paths as it observes them being executed across the organization. You do not need to configure HP TransactionVision with information describing how you think your transactions flow; TransactionVision tells you how they actually flow based on observing real transactions as they are processed. You do have to configure TransactionVision to send the required transaction events to BPI.

The discovered transaction events are then available for use by BPI, as input for Business Process modeling.

Summary

By defining models for your Business Processes and mapping these models to the underlying IT operational resources modeled within Business Availability Center, you can be notified when Steps in the process are impacted by a failure in the underlying infrastructure.

There are significant benefits to being able to visualize your Business Processes and relate resources to these processes; for example, you can identify the impact a resource failure has on the process and how other Steps in the process might also be impacted.

However, understanding only the impact of resource status changes does not enable you to answer questions such as:

- What is the financial impact of the resource failure?
- How many transactions are waiting to be processed at a particular point in the process?

In order to answer these types of question, you also need business data from your business applications, plus you need the BPI system to maintain data on individual instances of the Business Processes.

Using rules, filters and database queries, BPI can provide monitor data based on each instance and groups of instances. This, plus the business data that BPI receives through business events, enables BPI to present business-related information related to how your business is impacted by a failure in an IT operational resource, or how your Business Processes are running with respect to the objectives set on KPIs.

Instances of Business Processes

In order to get more detail from your business systems, you need to be able to model the individual instances of the process and see how these are being impacted: a process instance is a specific occurrence of a particular process that you have defined, for example, a specific customer transaction, or a specific flight arrival.

To model instances of your processes, you need to define data objects. A data object usually represents a known object within your organization; for example, an order, an employee record or flight details, or indeed it might also be an aggregation of a number of properties. In the case of HP TransactionVision, these data objects are created automatically when the transaction events that you have identified are imported.

The data captured from the business events, which are being processed by your BPI system, can be used directly to change the status of a Business Process, or can be used to update the status of the data objects that you have defined. In the latter case, rules are used to evaluate how the Business Process is progressing, based on the changing status of the data object.

You can also add Business Process monitors to the instance-level information and use it to provide business measurements to indicate the health of the process instances that you are tracking.

Using the data collected for the defined Business Process monitors. BPI can present the state of the business back to the business manager. The data is presented according to the instances of the processes being tracked and the data collected. You can also use BPI to indicate to IT how they are meeting the needs of the business, for example, you can use the Business Process Insight Application Health page to show information about how many orders are waiting to be processed, or information about the value of orders that are waiting to have payment details checked.

You can also use the instance information that is recorded by BPI as follows:

- to view historical data recorded from BPI data using the Business Process Insight Application Reports page.

- monitor SLA violations using the reporting feature of the Service Level Management Application.

- create your own portlets using My BSM.

In conclusion, as a result of modeling these process instances and setting Business Process monitors, the business manager can see the impact of the critical business events on the Business Processes.

Modeling Business Processes

The most effective way to model a Business Process is to base it around the primary documents and Business Processes in your organization, for example, orders or insurance claim applications. As part of defining the process you need to identify the:

- Data that will be used to drive the process.

  These are the data that define the conditions under which a business task is deemed to have started and to have finished, and used to progress the process.

  You can also use data directly from your Business Events to start, stop and progress your Business Process.

- Data required to answer the business questions that you want to track.

  These data are obtained from your underlying business applications and can be thought of in terms of the data required for the business measurements that you want to report on for the process.

- Data required to characterize process instances.

  In addition to data that can be used to infer process Step state, or that is used in calculations against business criteria, you might also want to include Data properties that will be useful later when you want to distinguish between Business Process instances. As an example, you might want to group data by a specific location; in this case, the data relating to the location is not used to determine whether a process Step has started or completed. It is also not used to calculate any sort of business metric; however, it is useful to have the location as a category in analyzing business performance.

- Events, which are required to obtain the data to drive the process and answer the business questions.

  Events are defined in cooperation with the IT personnel who own the IT infrastructure and who you work with to obtain the data that you need.

- Key Business Process monitors that you want to measure for your Business Processes.

Operational and Business Data

This section discusses the types of data used by BPI in more detail.

Figure 5 shows how the Steps in a process can be related to the IT operational resources of your business, and that these resources, when combined with the business data and events received through the Business Event Handler, become the basis for tracking the health of your Business Processes.

Figure 5   Business Process and Events
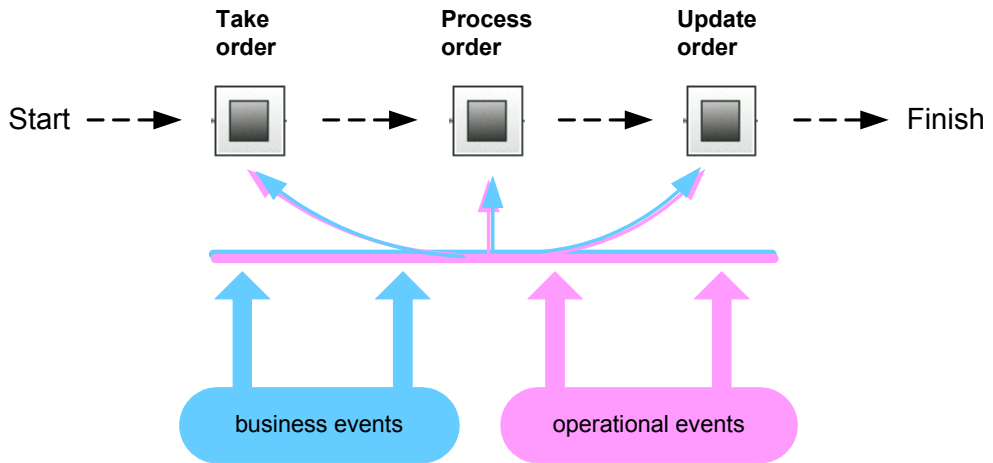
The BPI system also manages the data required for the process (the data models), as Data definitions. In the Stock Control example, it maintains information about the `account` business data and within that, information about `value`, `customer name`, and `customer type` data.

BPI does not need to know about all the data related to a process, just enough to give a business manager information on the business impact of the affected Business Processes. For example, in the Stock Control process (see Figure 3 on page 18), the `account` business data might also contain data relating to `customer address` and `customer preferences`, but this data is not required for process tracking purposes, so it does not need to be defined or tracked within BPI.

The following is a summary of the two types of events received into the BPI system:

- Events that carry the data to move the process through its Steps. These are the business events.

  You configure the business events that you need for the business applications that have access to (or hold) the data required for the Business Process.

- Events that report business impact information related to the IT infrastructure; these are the operational events, which are linked to the Business Process CIs.

You need to model business events for your solution. The events from underlying business applications can provide the business data focus.

The data from these underlying systems is defined as part of the Business Process and the resultant information used for reports and analysis.

A methodology, described later in this guide, takes you through the Steps that help you define and then refine your Business Process and its data.

Business Scenario for a Stock Control System and BPI

Continuing with the example of the Stock Control system, the following describes the scenario in more detail and shows how BPI can be used, by the business manager, to indicate the business impact of failures reported through Business Availability Center.

The scenario describes what happens when the stock control system fails - the business can accept orders, but it is unable to confirm that it can fulfill them. As can be seen from the Business Process diagram (Figure 6), the Stock Control System is accessed at the start of the Business Process to check the stock when the order is placed, and at the end of the Business Process, when the stock is dispatched.

Figure 6   Business Process
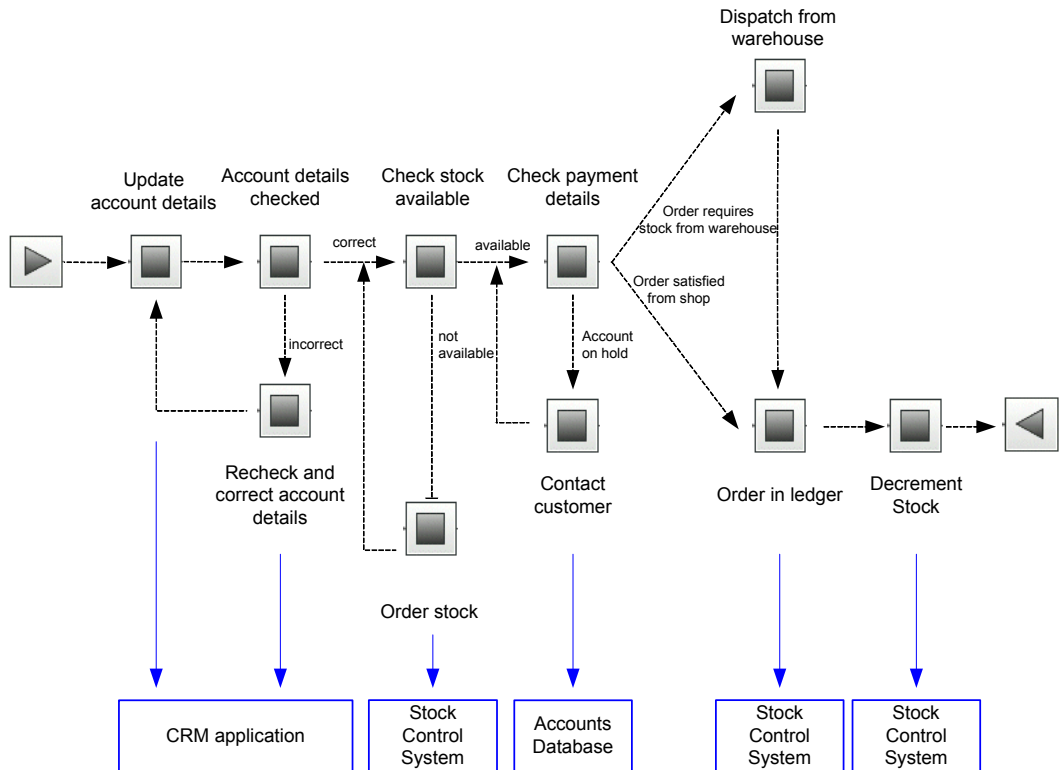
The business also has a Gold Customer scheme whereby key customers are guaranteed their deliveries within a specified time, or significant penalties in the way of discounts are incurred.

Using the BPI Application, the business manager receives information about:

- the impact of the IT operational resource failures
- the impact of the business service failures

Using both these sets of information, the business manager now has the Business Process monitor data that can be used to better understand the business impact of a failure. As an example, the business manager could define a monitor that records how long it takes a specific instance of a process to move between two activities in the process. The Business Process Insight Application shows this monitor information in terms of statistical information relating to the process instances and their progress. The business manager can view the information in tabular and graphical form according to their requirements.

In addition to, or in place of, monitors the business manager can request for annotations to be added to the processes as they are displayed on the Process Diagram on the Application Health page. These annotations can show information such as the current value of orders waiting to be processed.

The business manager also has the option to define instance thresholds for the Business Process monitors and have BPI generate instance threshold violations when the instance thresholds agreements are not achieved. The instance threshold violations can inform them of critical business events, due to event such as the failure of the stock control system. On receipt of a violation, the business manager uses the BPI Application Health page or My BSM portal (whatever their usual reporting tool might be) to query the BPI system and obtain more information.

As an example, they could determine the following:

- There are 200 customer orders waiting to be processed and 10 of these customers are Gold Account Customers.
- The impact of the current rate of orders, for example, given historical information about the rate of orders, the current number of 200 is likely to be 1,000 within a couple of hours, and 20,000 by the end of the day.
- Using historical information, the stock float should be good for today and tomorrow and most orders can be dispatched without incurring any delay.

- If required, temporary manual stock control can be put in place using the business data provided by the BPI system.

- Information about the stock levels

  The stock counters configured in the BPI system can be used temporarily whilst the stock control system is offline, so the business manager has an ongoing count of the stock and can monitor the situation. The business manager also knows the levels of stock in the shop, so they have the option to fulfill orders from the stock held at the shop if necessary.

The business manager is now in a position to make a business decision to prioritize the orders such that the Gold Customers and the high value orders are processed first. If any of the stock levels start to get low, the business manager can suspend the transactions, or inform customers that there might be some delay. They can also order additional stock manually based on the BPI data presented to them.

The introduction of BPI means that the business manager now understands the business impact of the stock control failure and as a result can:

- Continue to process from stock in shop and redirect that stock to Web orders if necessary.

- Use the BPI data to manage the stock float manually.

- Aggregate shop and warehouse stock as required.

- Monitor the Gold account customers and minimize the penalty payments.

The business events that provide this information come through the BPI Business Event Handler components and the operational events are defined in the UCMDB as IT operational resource CIs.

What BPI Provides

BPI provides the modeling tools that enable you to build a Business Process and associate key business services, data and events with the process. As part of defining the process, you need to consider what business measurements you are likely to want to take. Later, you can add Business Process monitors to these processes to record these business measurements.

Using Business Process monitor information, BPI can notify the business manager when significant events occur and when Business Process monitor instance thresholds are violated; this notification can be done through email, through the Notification Server, through the BPI Application or through a My BSM portlet. In addition, you can create reports using the Business Process Insight Application Reports pages.

Business managers can also receive the BPI impact reports through their own reporting tools, if they choose to. BPI writes all its Business Process monitor data to database tables, which are designed and structured to maximize the effectiveness of user and application queries.

BPI enables you to:

- Track instances of Business Processes defined and deployed within the BPI Modeler.

- Monitor the Business Processes and Business Process Steps using the instance-based data samples, which are sent by the BPI Server to BAC.

- Create monitors for your Business Processes and Business Process Steps using the Monitor Definer.

- View the health of your Business Processes, Business Process Steps and Business Process monitors according to objectives that you set using the Business Process Insight or Dashboard Application.

- View reports based on the history of the Business Process using the Business Process Insight Application.

- View details of the Business Process instances, Step instances and instance threshold instances using the Business Process Insight Application.

- Configure notification emails, HPOM messages and scripts for process instance violations using the Notification Administration.

- Use the BPI Administration Console to manage the BPI Server components.

The individual components of the BPI Server system, and their purpose, are more fully described in Chapter 3, BPI Server Architecture.

Summary

BPI is intended to help you raise the visibility of your Business Processes and better understand integration points between your Business Processes and the IT infrastructure. This provides you with improved business data and also enables you to make business decisions, based on this data, in a more timely manner; for example, understanding the business impact if a particular resource is made unavailable. This might not be direct financial loss, but other losses, such as reputation, regulation breaches, poor delivery times and so on.

Business disruptions such as these, and hundreds more that happen on a daily basis, unexpectedly and negatively impact the normal business operations of many companies. For example, if you are providing a resource and there is a sudden increase in the volume of traffic, perhaps due to a marketing promotion, it could have an impact on the network. The response time for customers might be seriously compromised, which in turn can cause anger, ill-will, and some customers taking their business elsewhere. This is not a failure in any of the underlying systems and applications, but a failure in the capacity of the system that has an impact on your customers.

3 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 BAC implementation. If you want an overall introduction to BPI as a BAC application, refer to Chapter 2, Introduction to Business Process Insight.

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

Specifically, the chapter covers the following topics:

- A high-level representation of the BPI component architecture; see section BPI as a Component of Business Availability Center on page 36.

- The BPI Application, which you use to track and monitor the impact information recorded on the BPI Server system; see section Architecture of the Business Process Insight Application on page 42.

- The BPI databases and how they are used by the BPI components; see section BPI Instance Database Schemas on page 50

- The BPI Server components; see section BPI Server on page 54.

- The BPI administration clients; see section BPI Administration on page 47.

- Integration points for BPI; see section BPI Integration Points on page 66

BPI as a Component of Business Availability Center

When you first install Business Availability Center, some BPI components are also installed as shown in Figure 7.

Figure 7   BPI Components Installed with Business Availability Center
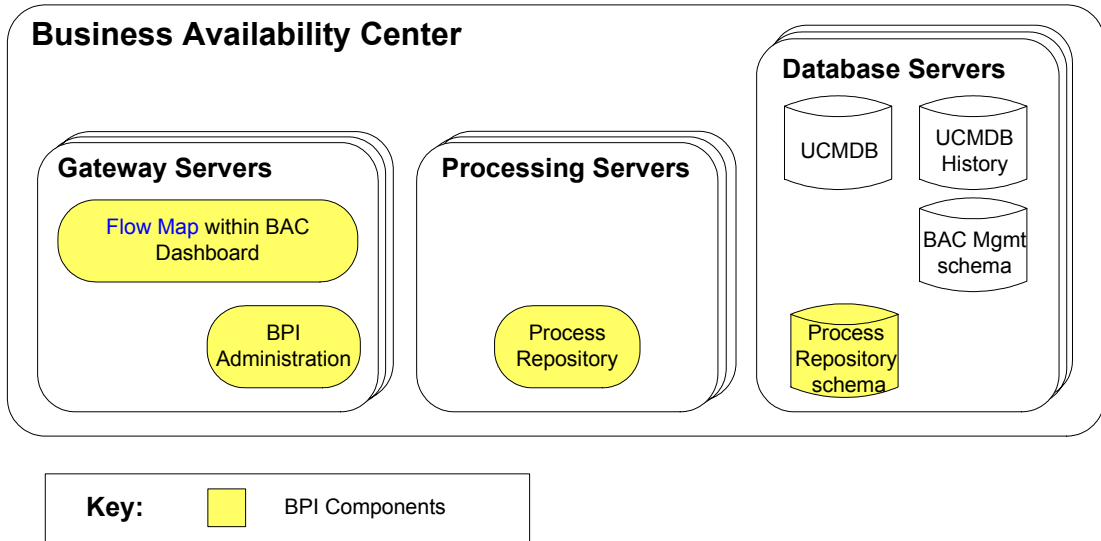
In this configuration, you can model a Business Process and map individual process Steps to the CIs representing the underlying IT infrastructure (IT operational resources).

The BPI Process Repository is installed as part of the Business Availability Center Processing Server. The BPI Process Repository manages Process Definitions, which you create using the BPI Modeler, to monitor IT operational resources defined within the CMDB. You create a database schema for the Process Repository as part of your BAC 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 Dashboard Application to view the status of the health of the IT operational resources from the Flow Map Tab.

The BPI components that are installed with Business Availability Center are:

- BPI Modeler, which you access from:

  `Admin > Business Process Insight`

- Process Repository and associated database schema.

- The BPI Flow Map, which is accessed from:

  `Applications > Dashboard > Flow Map`

When you subsequently install the BPI Server, the BPI Application is installed. Figure 8 shows the layout of the BAC system following the BPI Server installation, including the location of the BPI components.

The following list describes the relevant components of your Business Availability Center implementation:

- Gateway Server

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

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

- Database Server

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

Figure 8 shows your BAC system following an BPI Server installation.

Figure 8   BPI Components Following an BPI Server Installation
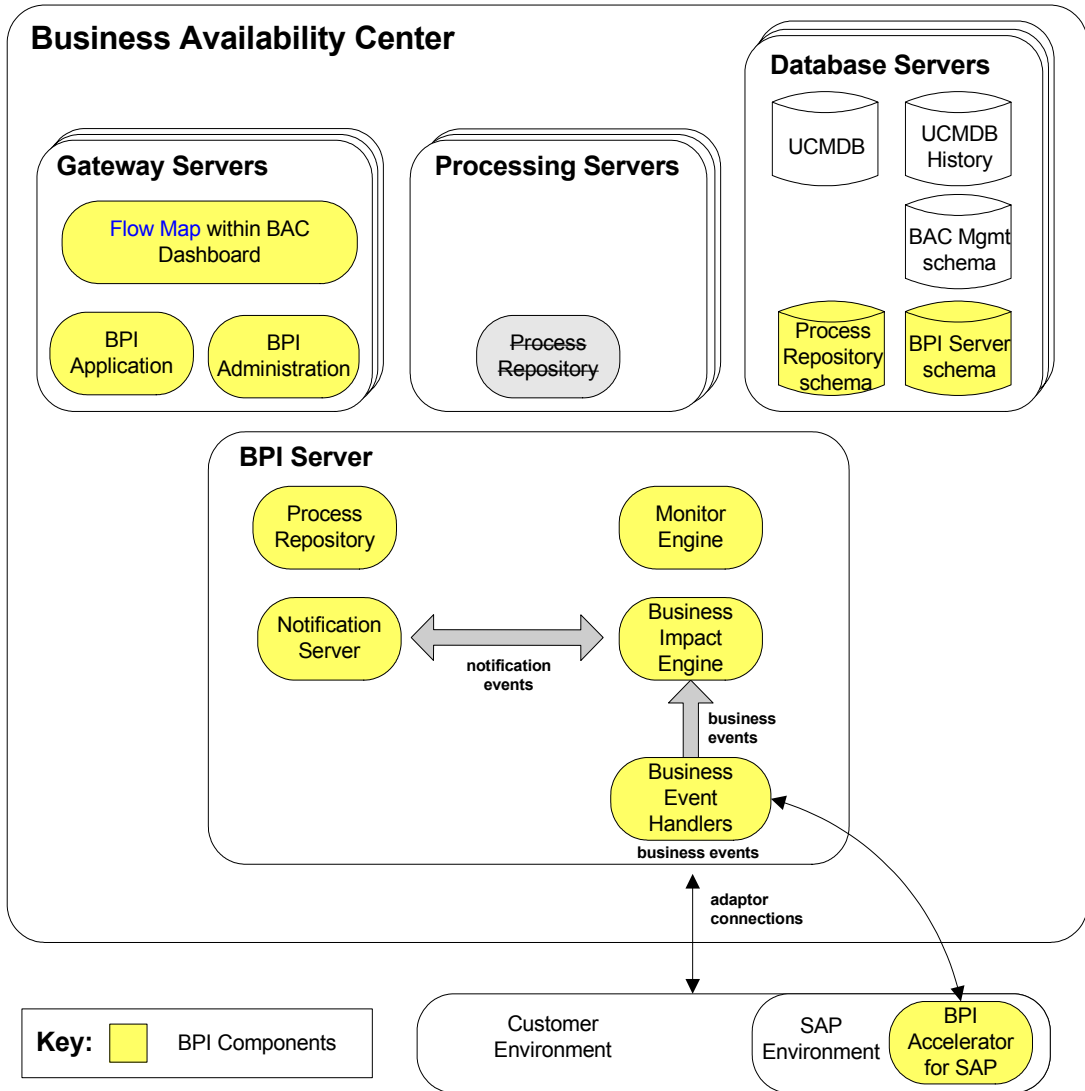
When you install the Server component of BPI, the existing Process Repository is replaced by a revised version, which is installed on the same system as the BPI Server. The revised Process Repository manages all Process Definitions from this point on, for processes that you create with instance tracking 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 Steps 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 Steps 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:

- The overall state of the process.

- The number of active process instances.

- Specific process instance statistics such as:

  — The start and stop times of each step.

  — Step 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:

- The time taken to process an order.

- The current backlog of flights waiting to get airborne.

- The 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, Business Process Insight appears as an option from the `Application` menu.

The Business Process Insight Application enables you to monitor and track Business Process instances, in real time, using a number of different views and graphical displays, for example the 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 Steps.

- monitor the business health of your Business Processes and Business Process Steps 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 section Business Process Insight Application Features on page 44.

The following section (section Architecture of the Business Process Insight Application on page 42) describes the architecture of the BPI Application.

Architecture of the Business Process Insight Application

The Business Process Insight Application is a BAC 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.

- UCMDB

  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.

Figure 9 shows how the Process Instance and Violation Pages are managed within BAC.

Figure 9   The BPI Server Pages within BAC
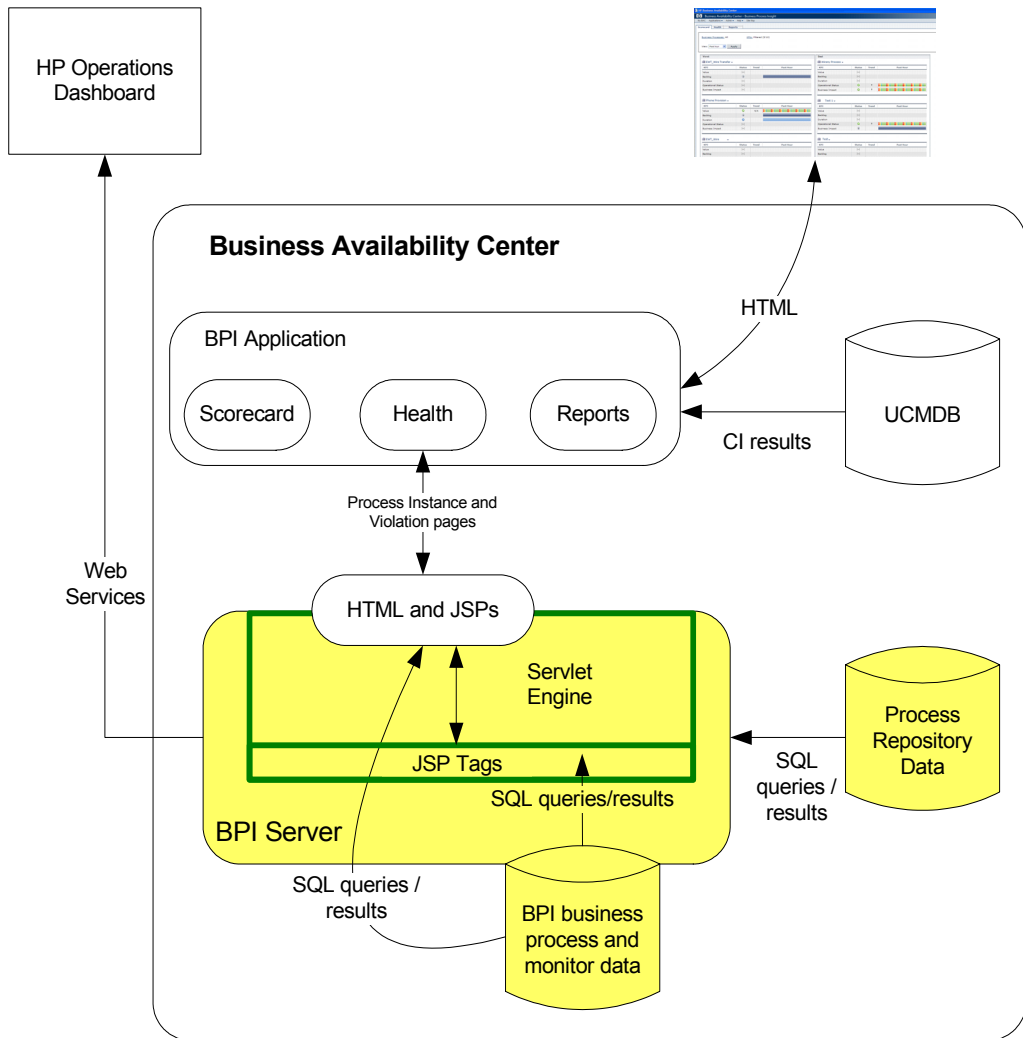
Business Process Insight Application Features

The BPI Application 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.

▶ If you do not have a BPI Application license, these options are not available to you from BAC.

The BPI Application comprises:

• 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 Flow Map 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 Step in the process; the diagram also shows IT operation resource status for each Step 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 Step 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 Step are shown.

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

– Instances

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

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

– Instance Violations

Displays lists of Monitor or instance Threshold violations for the selected Process or Step, 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 Step.

— 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 Step, 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 Step Summary

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

Email Notifications

In addition to proactively monitoring your Business Process health through the BPI Application, you can 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.

▶ The Notification Server sends alerts only for instance violations. Alerts that relate to your overall business health, and that are based on statistical calculations are configured and generated from business objective KPIs. You need to separately configure these KPIs for your Business Processes.

Email notifications are configured using the following option:

```
Admin > Business Process Insight > Notification
```

More information about managing your email notifications can be found in *Using Business Process Insight*.

BPI Administration

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

- Administration Console, which is accessed on the BPI Server system; see section BPI Server Administration Console on page 47.

- The administration option from BAC:

  Admin > Business Process Insight

  This is described in section BPI Application Administration on page 48.

BPI Server Administration Console

The BPI Server includes has an Administration Console that you use for managing the parameters of the BPI Server; for example, configuring the parameters required for the BPI connection to the Gateway Server and starting and stopping the BPI Server Components.

Using the BPI Administration Console to manage the Server components is described in *Using Business Process Insight*.

BPI Application Administration

To access the BPI Application pages, select:

`Admin > Business Process Insight`

The following are options from this page:

- Modeler

  You use the Modeler to create your Business Process definitions, Data definitions and Event definitions. The Modeler is a graphical editing tool and provides a unified interface for all the definition types that you are creating.

  The Modeler has two modes of operation, according to whether or not you have a valid BPI license and have installed the BPI Server. Refer to the *Business Process Insight Integration Training Guide - Modeling Processes* for more information.

- Monitor Definer

  You use the Monitor Definer to define Business Process Monitors and Monitor instance Threshold definitions for deployed processes.

  These Monitors and instance Thresholds are then used to gather statistical information, and instance thresholds relating to the health of your Business Processes. The results are shown using the BPI Application pages.

- Process Repository Explorer

  Use the Repository Explorer to view, print and manipulate the contents of the Process Repository schema. The Process Repository schema is a set of database tables that hold the data for the Business Processes that you have defined and modeled using the BPI Modeler.

- Process Intervention

  In specific circumstances, you can use the Process Intervention Client to access processes that you have deployed in order to modify or delete Process instances and their associated Data instances.

- Notification

  Use the Notification option to configure the notifications that you want to be generated for the Business Processes that you have deployed and the Business Process Monitors and Monitor instance Thresholds that you have defined.

  Once configured, the Notification Server can then notify you of the process-impact, out-of-sequence and monitor-threshold violations that are generated by the Business Impact Engine.

Refer to *Using Business Process Insight System* for more details of using these options.

BPI Instance 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.

Figure 10 shows the relationship between the database schemas and the BPI components. *Deploying Business Process Insight* provides details of BPI's database requirements.

Figure 10 BPI Database Schemas
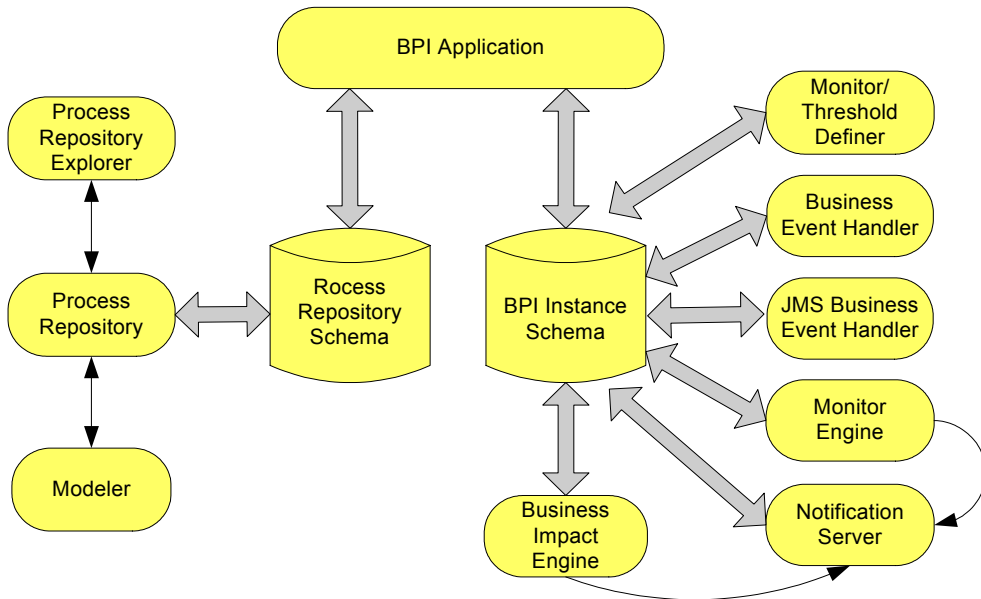
The BPI components' use of these schemas is described in the following sections:

- BPI Server Schema on page 51
- Process Repository Schema on page 53

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 51
- Monitor Instance Threshold Definitions on page 52
- Event Hospital Data on page 52
- Notification Data on page 52

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 Steps 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 Step 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.

The BPI Server is a separately licensed component of Business Availability Center. When installed, it 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 Availability Center.

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.

Figure 11 BPI Server Architecture
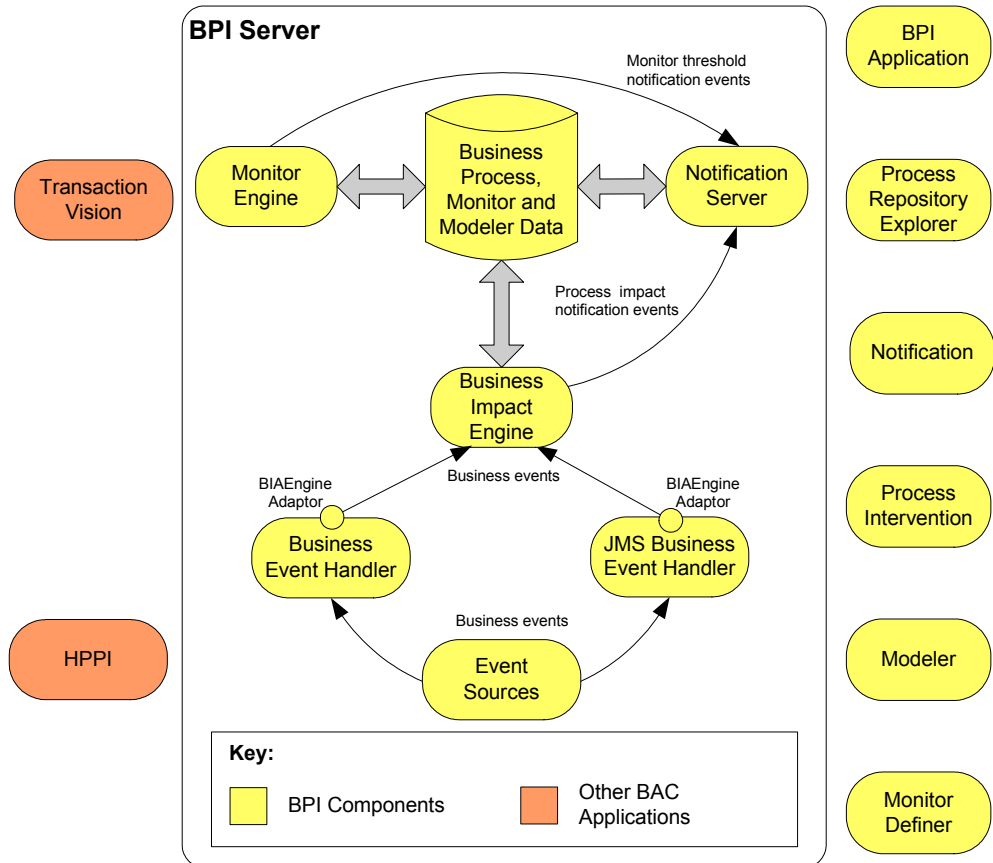
Business Impact Engine

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.

Figure 12 shows the Business Impact Engine architecture.

► Business Monitor Instance Threshold violations are managed by the Monitor Engine, which is described in *Using Business Process Insight*.

Figure 12 Business Impact Engine Architecture
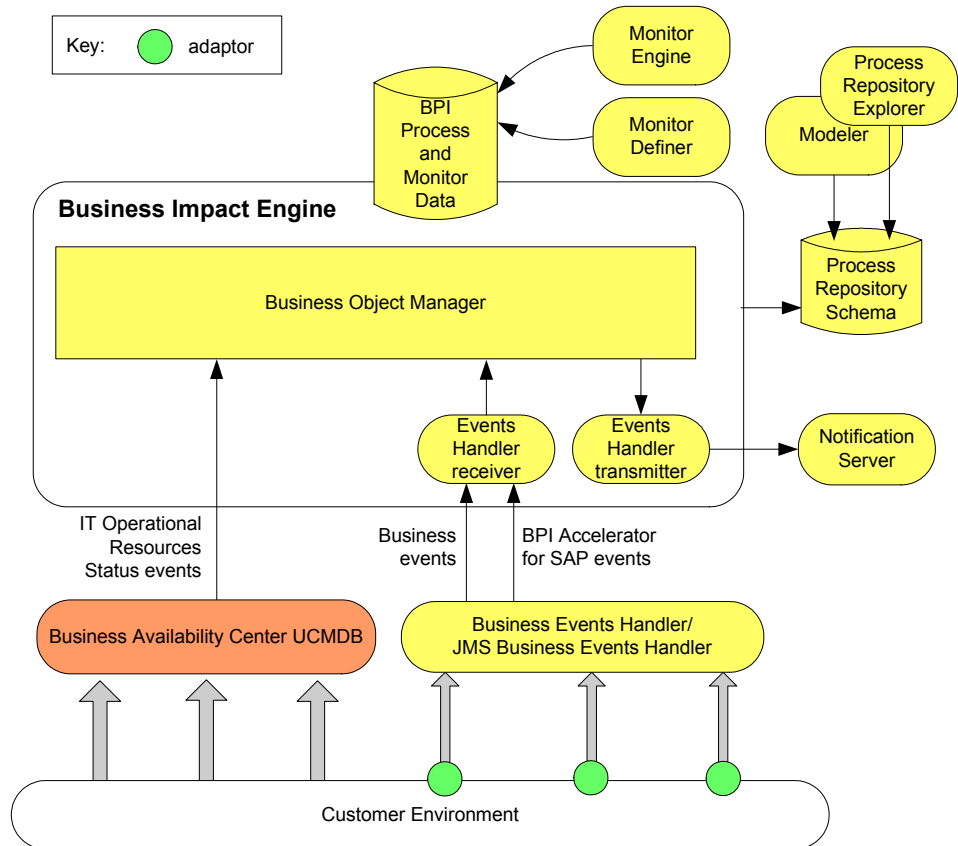
When an instance threshold violation is identified, the Business Impact Engine sends an impact violation to the Notification Server and writes data to the database, which can then be propagated to the BPI Application Health pages. This enables you either to be notified of the status of your Business 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 *Using Business Process Insight*.

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.

Figure 13 Business Event Handler
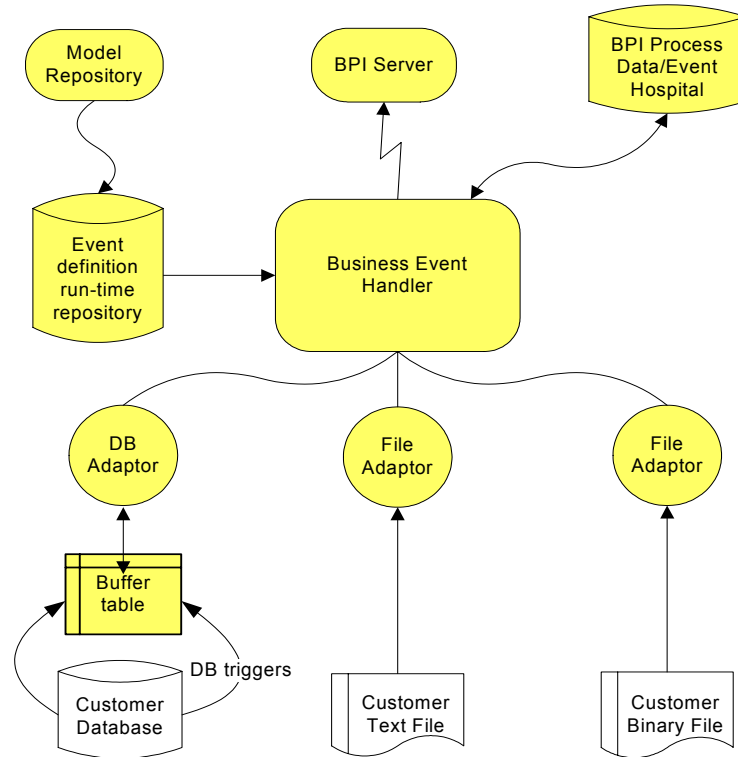
In the case of BPI, openadaptor is the framework that provides the following BPI event adaptors:

- 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 Steps.

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 58.

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.

► 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.

Configuring the JMS Business Event Handler is described in *Using Business Process Insight*.

Figure 14 JMS Business Event Handler
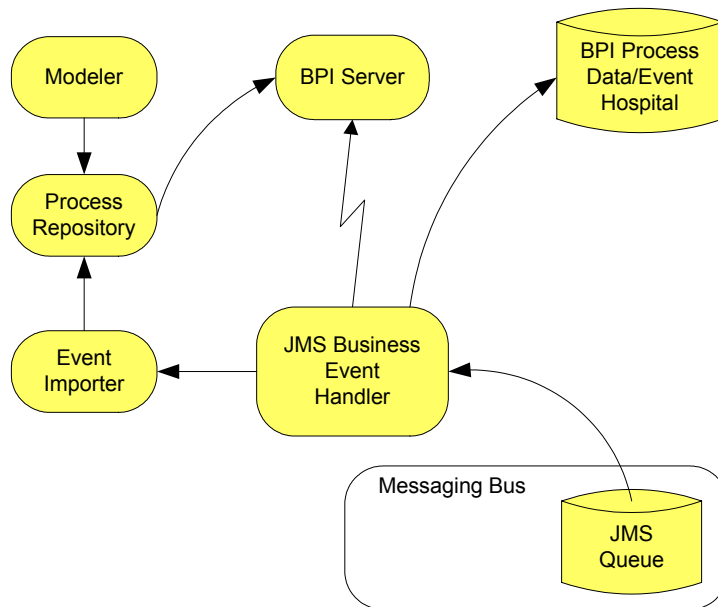
Figure 14 shows how Event definitions are imported from the JMS-enabled system through an Event Importer into the Process Repository. You can then access these Event Definitions from the Modeler, and add Data Definitions to reference these events. When the process, including the Event and Data definitions 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 *Using Business Process Insight* for more information about the features of the Repository Explorer.

The BPI Modeler and its features are described in more detail in the *Business Process Insight Integration Training Guide - Modeling Processes*.

BPI Integration Points

BPI can integrate with the following applications:

- TransactionVision; see section Integrating with HP TransactionVision on page 67

- SAP Applications; see section Integrating with SAP Applications on page 68

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

1. Import service definitions from other HP BTO Software components.

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

2. 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. The Business Event Handler is described in Chapter 3, BPI Server Architecture.

3. Configure JMS data sources through the JMS Business Event Handler in order to import JMS-based Business Events. The integration with TransactionVision is achieved using the JMS Business Event Handler.

   The JMS Business Event Handler is described in Chapter 3, BPI Server Architecture.

4. Configuring notifications through the Notification Server:

— You can configure email notifications that can be sent when a notification event is received by the Notification Server.

— You can create 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 is described in *Using Business Process Insight*.

5. Import 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. Refer to the *Business Process Insight Integration Training Guide - Importing BPEL*, for more information on the options for importing BPEL definitions.

Integrating with HP TransactionVision

You can configure BPI to integrate with TransactionVision. This 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 the displays reports based on the results of the analysis.

Once integrated, you can link your Business Process Steps to event definitions that have been created, or discovered, within TransactionVision. To do this you need to import the required TransactionVision events into the Process Repository, in order that they are available through the Modeler.

When you have created your Business Processes, BPI can receive events from TransactionVision and use the data within them to progress BPI Business Processes. You can then track and monitor the status of your Business Processes instances, which include the transaction details, through the Business Process Insight Application.

For details of configuring BPI to integrate with TransactionVision, refer to the Chapter 4, BPI and 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.

Figure 15 shows the integration between the BPI Accelerator and SAP.

Figure 15 BPI and SAP Integration Using BPI Accelerator
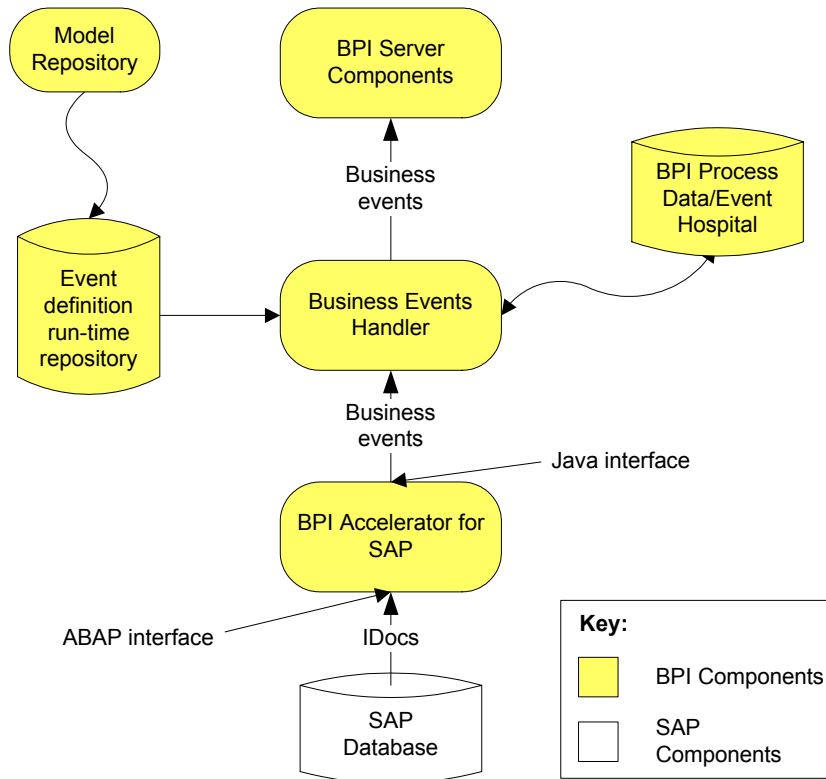
The *Business Process Insight Accelerator for SAP Guide* contains details of installing and setting up the Accelerator.

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 BAC 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 *Using Business Process Insight*.

4 BPI and TransactionVision

This chapter describes how you can integrate BPI and TransactionVision to have BPI show the status of instances of Business Processes, which are linked to specific business transactions within Transaction Vision.

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 integrating BPI and TransactionVision, you can link the business transactions into your Business Process Steps.

This chapter covers the following topics:

- Overview of TransactionVision Integration on page 72
- Configuring BPI to Integrate with TransactionVision on page 78

Overview of TransactionVision Integration

TransactionVision comprises a set of sensors. The sensors collect transaction events by monitoring the behavior of your organization's business applications. The sensors report these events to the TransactionVision analyzer, which correlates the events and further processes them. Specifically, the analyzer is responsible for assembling the individual transaction events into business transactions. It can then process and action these business transactions according to rules that you define. You can configure more than one analyzer for TransactionVision. Multiple analyzers are used for load balancing and all communicate with the same database. It therefore does not matter which analyzer you use to manage your data, as the data source is common to them all.

You can define rules within TransactionVision that are specific to BPI. These rules define the transaction events that you want to propagate to BPI. TransactionVision then collates these events and aggregates the data. Business, including the aggregated data are then placed on a JMS queue to be read by the JMS Business Event Handler.

In addition, you can use the TransactionVision Web User Interface to provide you with reports, based on the business transactions that TransactionVision has recorded.

The components of TransactionVision communicate using an application server, specifically, the message-oriented middleware component of the application server; this is referred to as a messaging bus in this chapter. You need a component of the same messaging bus on the machine where BPI is installed in order that BPI can communicate with TransactionVision. Typically, the messaging bus components are referred to as client and server components.

TransactionVision uses the SONIC MQ messaging bus and the examples and configuration described in this chapter are all based around SONIC MQ.

Figure 16 shows the high level architecture of the BPI and TransactionVision integration.

Figure 16 BPI and TransactionVision Integration Architecture
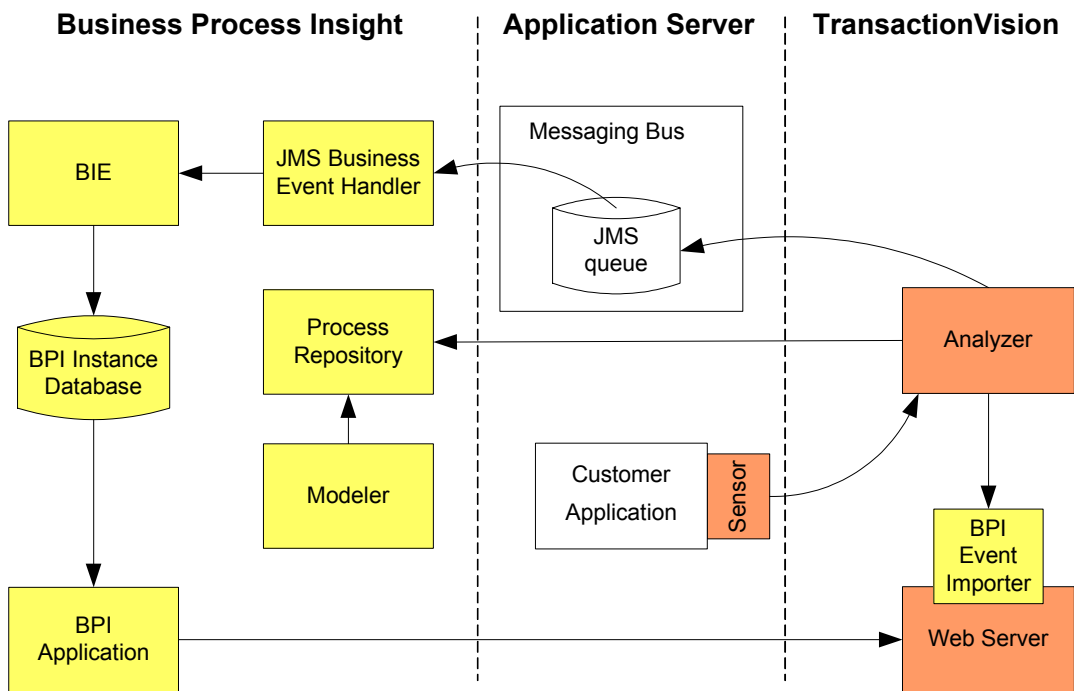
The integration between BPI and TransactionVision enables you to:

1. Import event definitions that have been defined within TransactionVision into the Process Repository,.

   You can then use the Modeler to create processes, which have Data definitions that subscribe to these imported events.

2. Receive events from TransactionVision and use the data within them to progress BPI Business Processes.

3. Link from the process instance details page within the BPI Application Health pages to the details of the instances within TransactionVision.

   This provides BPI users with more information relating to problems that might be occurring in a process at the transaction level.

Design-Time Integration

This section describes how data is moved between BPI and TransactionVision when you are designing your processes and configuring the BPI and TransactionVision systems.

Figure 17 shows the flow of information from TransactionVision to BPI. It assumes that:

- TransactionVision Sensors are already configured to collect the required transaction information.

- You have created a project within TransactionVision for the transactions that you want to collect.

- You have specified how you want to store the data that is extracted from the transactions that TransactionVision is collecting in its database; this is the structure of a database table.

- You have classified rules for how the transaction data should be saved in the database; that is, which data is stored in which database column.

Figure 17 TransactionVision Design-Time Integration
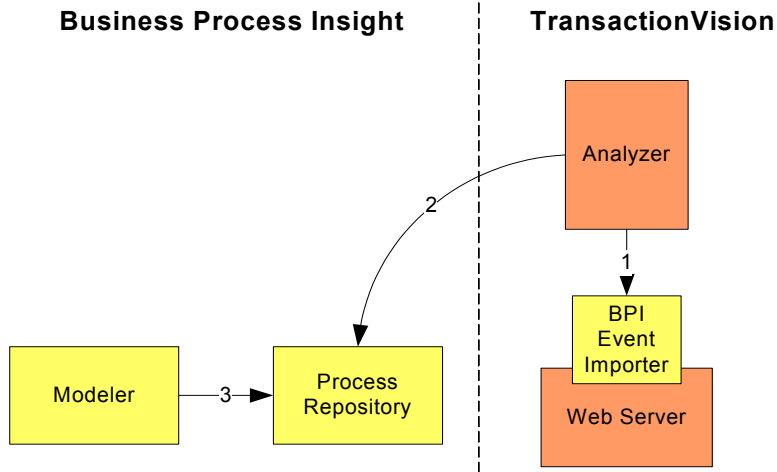
The flow of data, as shown in Figure 17, is described in the following steps:

Step 1: You classify rules for how the transaction data should be saved in the database; that is, which data is stored in which database column. The BPI Event Importer imports the required Business Event definitions into the Process Repository when you create or edit the TranscationVision Classification Rules.The events are then available to the BPI Modeler.

Step 2: Event and transaction rules have been defined to filter the incoming transactions, within TransactionVision and according to your requirements. A `BPI Action` has been defined that creates a Business Event from the transaction event data, The BPI Event Importer imports the required Business Event definitions into the Process Repository, where they are available to the Modeler.

Step 3: You define your processes as usual and create Data definitions, which can then subscribe to the Business Event definitions imported from TransactionVision.

Run-Time Integration

This section describes how data is processed between BPI and
TransactionVision when your system is operational.

Figure 18 TransactionVision Run-Time Integration
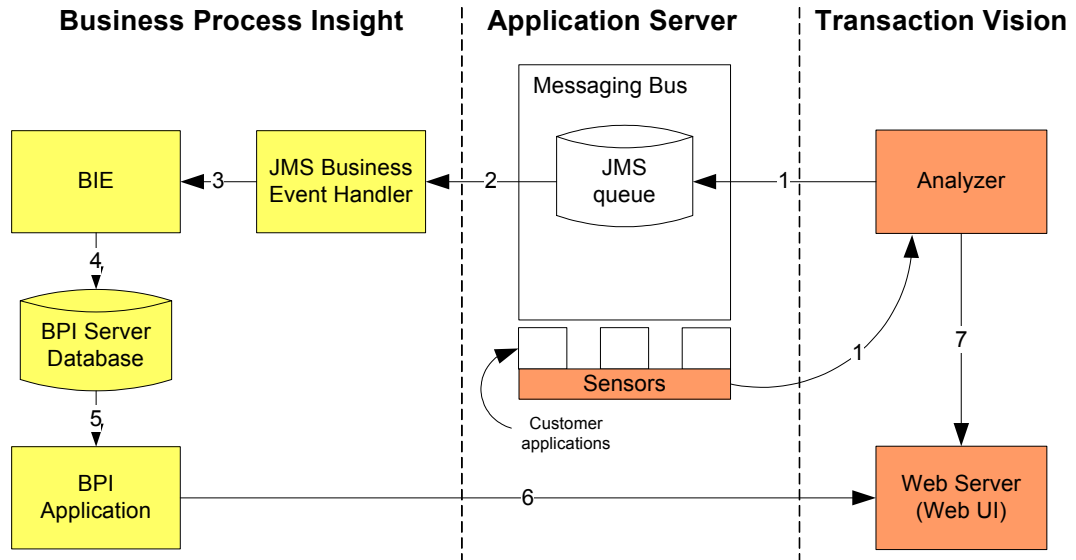
The flow of data, as shown in Figure 18, is described in the following steps:

Step 1: Using the rules that you have defined for the transaction events,
TransactionVision filters the appropriate events and completes the `BPI`
`Action`, which in turn places Business Events on the JMS queue that you
have defined.

Step 2: The JMS Business Event Handler is listening for new events that are
arriving on the queue.

Step 3: The JMS Business Event Handler accepts the business events and
sends them to the Business Impact Engine for processing.

Step 4: Details of the process and its progression are held in the BPI database
as usual.

Step 5: The BPI Application Health pages present the results of the process progression.

Step 6: The BPI Application Health pages provides links to the appropriate transaction events within the TransactionVision Web user interface.

Step 7: Details of the TransactionVision events are displayed within the TransactionVision Web user interface.

Configuring BPI to Integrate with TransactionVision

This section describes the tasks that you must complete in order that the BPI and TransactionVision integration can be enabled.

The configuration tasks described in these sections assume that the necessary BPI components have been installed on the TransactionVision and the BPI systems.

TransactionVision Configuration Tasks

The following tasks need to be completed on the TransactionVision server where a TransactionVision analyzer is installed. All analyzers can communicate with each other and all communicate with the same database. It therefore does not matter which analyzer you use to manage your data, as the data source is common to them all:

- Make sure that you have added the required TransactionVision Sensors to the applications you want to monitor; see the *TransactionVision Help*.

- Create a TransactionVision Project; see the *TransactionVision Help*.

- Define how to store the data extracted from the application within TransactionVision database; see the *TransactionVision Help*.

- Classify the rules that determine how and where to save the data extracted from the transaction events for BPI data in the TransactionVision database table, or tables; see section Classify Events on page 79.

- Create a JMS queue for BPI that is accessible from both TransactionVision and BPI; see the documentation for your messaging bus.

- Configure the JMS queue and JNDI repository for TransactionVision. The settings for the JMS queue need to be consistent across BPI and TransactionVision; see section Configure a JMS Queue and JNDI Repository On the TransactionVision System on page 79.

This document does not cover the full details of completing these tasks, it provides an overview of the key tasks; you need to refer to the TransactionVision documentation for details of using the TransactionVision

GUI to complete some of the required tasks. The TransactionVision integration with BPI is described in the *TransactionVision Administration Help*.

▶ There is a preconfigured Service Provider available for you to use, which means you do not have to configure the JMS queue and JNDI registry entries within TransactionVision. The preconfigured Service Provider, assumes you are using the SONIC MQ bus.

Configure a JMS Queue and JNDI Repository On the TransactionVision System

When you have created a JMS queue, you use the TransactionVision GUI to configure the details of this queue for the TransactionVision and BPI integration.

You need to configure a JNDI repository, on the TransactionVision system, and add the JNDI entries for accessing the newly configured JMS queue, which has been created on the messaging bus for the integration. Refer to the documentation for your JMS messaging bus for the details that you need for the JNDI repository configuration.

Using the TransactionVision interface, navigate to the page for the Business Process Insight integration option. You can access this page from:

`Admin > TransactionVision > Administration > BPI JMS Settings`

From the `Business Process Insight` page, enter the JNDI details that reference the JMS queue for the integration.

Classify Events

You need to classify the TransactionVision events that you want to have delivered to BPI. You classify TransactionVision events to label the classes and define the rules for extracting the required data from the events.

▶ You need to make sure that you specify appropriate filters to make sure that you are not monitoring events that are generated by TransactionVision.

TransactionVision provides a structured editor to enable you to manage transaction rules and associated attributes. You need to define rules which identify the transaction events that are sent to BPI. You use TransactionVision's `Transaction Definition Editor` to classify the events within TransactionVision.

Using the Transaction Definition Editor, you see an overview of all the transaction rules that are defined on your TransactionVision Server system. Using these rules, you can define:

- which events belong to which transaction class, or classes.
- how data is mapped between transaction events and business transactions.
- actions, based on criteria that you define.

You access the Transaction Definition Editor from:

```
Admin > TransactionVision > Administration > Transaction
Definition Editor
```

Refer to the *TransactionVision Help* in the section covering `Classifying Transactions` for full details of using the editor.

The following is an overview of the tasks that you need to complete:

1. From within the Transaction Definition Editor, select the `Common Section` for the event that you want to send to BPI.

2. Define the data that you want included in the event; you do this by defining `Attribute Rules` for each item of data.

3. Configure TransactionVision to generate a Business Event for BPI in a format and that can be imported into BPI. To do this, you need to specify:

   — the `Action` Java Bean provided by TransactionVision.

     This is a TransactionVision action, which is specific to BPI, and defines a standard action for communicating events from TransactionVision to BPI.

     To specify the Java Bean, you need to create an BPI `Action`. Use the option `New BPI Action` to create the `Action` for BPI.

     The Java Bean specification is as follows:

```
com.bristol.tvision.services.analysis.actions.OVBPIEventPublisher
```

   — the BPI business event name and event group that you have defined within the BPI Modeler for the TransactionVision process. This is the event name that you want to reference from within TransactionVision.

TransactionVision events are not sent to BPI until the transaction has been classified, at which point, the configured events are placed on the JMS event queue ready to be accessed by the BPI JMS Business Event Handler.

⚠ If you want to link different transaction events from TransactionVision into an BPI Business Process instance, there needs to be a common identifier in all TransactionVision events for BPI.

You must determine this unique identifier from a data item in the event payload, which is mapped to a transaction attribute. It is important to note that the business transaction identifier generated by TransactionVision cannot be used as a unique identifier. This is because TransactionVision business transactions identifiers can change as the transaction is processed.

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:

- If required, install the remote components of the application server's messaging bus, as used by TransactionVision, on the BPI server. You need to install the components that enable BPI to access the configured JMS queue. Refer to your messaging bus manuals for details of what is required.

  You need only do this if you are not using Sonic MQ.

- Import the TransactionVision events on the TransactionVision Server to create the required event definitions in the Model Repository; see section Import TransactionVision Events into Process Repository on page 82.

- 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 a Step in your Business Process to a TransactionVision event:

  ```
  Start and complete on TransactionVision events
  ```

  Refer to 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 integration from within the BPI Modeler.

- Configure BPI to receive Business Events from the JMS queue that you have created for BPI; see section Configure a JMS Business Event Handler on page 83.

- 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.

Import TransactionVision Events into Process Repository

You need to import the Business Events, defined in TransactionVision as BPI Actions, into BPI.This is in order that they are available within your process definitions. When you have imported them, you can then subscribe to the events from the Data definitions that you have created for the Business Process.

You import the definitions as follows:

```
Admin > TransactionVision > Administration > HP Business Process
Insight
```

There is an Import BPI Definitions button option on this page.

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 section Service Provider on page 83 and section Event Source on page 86.

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 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 TransactionVision integration. You will 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.

You configure the `Service Providers` settings using the BPI Administration Console on the BPI Server system as follows:

1. Start the Administration Console as follows:

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

2. Click the `JMS Business Event Handler` option in the left-hand navigation pane.

3. Click the `Service Provider` option, which is a sub option 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.

4. If this is the first integration with TransactionVision, select 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 *Using Business Process Insight*:

   — `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 TransactionVision is installed. In the value presented, this is `analyzerhost`. Substitute the correct fully-qualified host name for `analyzerhost`; for example:

 tcp://*your-host-name*: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.

5. Click the `OK` button to close the dialog.

6. Click the `Apply` button to apply your configuration changes.

If you want to create subsequent entries, select the `Add` button 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.

You configure the Event Sources settings using the Administration Console on the BPI Server machine as follows:

1. Start the Administration Console:

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

2. Click the JMS Business Event Handler option in the left-hand navigation pane.

3. Click the Event Source option, which is a sub 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 *Using Business Process Insight*:

   — 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 TransactionVision 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.

4. Click the `OK` button to close the dialog.

5. Click the `Apply` button 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.

You can now receive events from the TransactionVision system. 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.

Configure TransactionVision to Send Events to the JMS Queue

To configure TransactionVision to send events to the JMS queue, complete the steps described in sections:

- Configure a JMS Queue and JNDI Repository On the TransactionVision System on page 79.

  Use the information that you configured in the sections above.

- Classify Events on page 79.

JMS Service Provider Settings for Sonic MQ

If you are configuring the service provider details for a Sonic MQ JMS implementation within TransactionVision, the following is an example of the configuration settings you might set for BPI.

Service provider name	HP TransactionVision
Description	TV connection using Sonic MQ
Service Provider Classpath	
JNDI Provider Location	tcp://*hostname*:21110
JNDI Factory Class	com.sonicsw.jndi.mfcontext.MFContextFactory
JMS Factory Class	BPIQueueFactory

Event Source Settings for Sonic MQ

If you are configuring the service provider details for a Sonic MQ JMS implementation, the following is an example of the configuration settings you might set for BPI:

JMS Queue Subject	TV2BPI.EVENT.QUEUE
User Name	not required for Sonic MQ when running on the local system.
Password	not required for Sonic MQ when running on the local system

The example shows an event source for a `JMS Queue Subject` called `TV2BPI.EVENT.QUEUE`. There are no user name and password details required for this queue.

5 Log Files

This chapter describes the log files for the BPI components.

The chapter covers the following topics:

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 94.

The log files for the Administration Console are called `bia_adminserver`*n_n*`.log` and `bia_adminclient`*n_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 95.

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 95.

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.

A Expression Grammar in Process, Data, and Filter Definitions

This appendix lists the rules for the grammar that can be used for Business Process progression rules and expressions, and expressions within Business Process monitor filter definitions.

The progression rules and expressions are within the BPI Modeler for the start and complete conditions for Steps, and for filter expressions for Data definitions when subscribing to events.

Business Process monitor filter expressions are used within the Monitor Definer to optionally specify the conditions for collecting statistical data for the Business Process monitors.

The grammar used for expressions is similar to Java expressions. The following sections describe the grammar and its construct. Further examples of using this grammar are provided in the *Business Process Insight Integration Training Guide - Modeling Processes*.

Grammar

This following is an informal description of the grammar, it shows how expressions can be constructed. Note that this construct is subject to specific limitations according to how it is being used. Examples of the use of the grammar are provided in the following sections.

```
expression =>
 property-value = expression
|expression + expression
|expression - expression
|expression * expression
|expression / expression
|expression % expression
|- expression
|expression && expression
|expression || expression
|! expression
|expression == expression
|expression > expression
|expression >= expression
|expression < expression
|expression <= expression
|expression != expression
|(expression)
|expression ? expression : expression
| value

value =>
        constant_value
       |property_value
       |function_return_value
```

See section Function Return Values on page 102 for possible values for function_return_value.

```
constant_value =>
        string_contant
       |integer_constant
       |real_constant
       |true
       |false
       |null

property_value =>
        root_object.relationship_name.property_name
```

where:

- `root_object` is `this` or `event`.

  If no `root_object` is specified, the expression assumes `this`.

- `relationship_name` can be included zero or more times, depending on how the property relates to the root object.

Possible `property_value` expressions are:

- `this.property`

- `this.data.property`

- `event.property`

- `property`

The following are examples that you might use in your expressions. Further examples are given in the *Business Process Insight Integration Training Guide - Modeling Processes*:

- `this.OrderNumber`

- `this.Customer.Customer_ID`

- `event.Status`

Function Return Values

The following sections describe the data and string values that can be returned from the `value` expression.

These functions can be used in filter expressions, binding expressions, progression rules and assignment actions.

Functions for Dates and Times

The following functions convert expressions into milliseconds:

- `hours` (*number*)
- `minutes` (*number*)
- `seconds` (*number*)
- `days` (*number*)

These functions convert *number* to the equivalent number of milliseconds (which are the Business Impact Engine time units).

Functions for Identifying String Values

The following functions take a string parameter and return a boolean result.

- `string_property.contains(`*value*`)`

  Returns `true` if *value* is found in the property value.

- `string_property.starts(`*value*`)`

  Returns `true` if the property value starts with *value*.

- `string_property.ends(`*value*`)`

  Returns `true` if the property value ends with *value*.

Functions for All Property Types

The following function takes one or more parameters and returns a boolean result:

```
property.in(value,value,...)
```

This expression returns `True` if the value of `property` is one of the listed set of values (`value, value,...`), and `False` if it is not.

Process Progression Rules

The grammar for process progression rules is slightly different from the simple expression grammar, as it is used to describe deltas or changes to Data definition properties.

There are four styles of progression rule provided in BPI as follows:

- Complete on first assignment
- Complete on transition
- Start and complete on transitions
- Advanced conditions

The Advanced conditions style of progression rule requires you to enter the methods for the progression rules directly as described in section Methods for Progression Rules on page 106.

For the remaining progression styles, you do not need to enter the methods, you just select the style of the progression rule that you want for the Step. Ultimately, each style uses a subset, or selection of the methods described in section Methods for Progression Rules on page 106. You can see these methods if you define a progression rule using one of the styles and then change the style to Advanced conditions to view the methods created for the style.

More detailed information about using these progression rules in your processes is provided in the *Business Process Insight Integration Training Guide - Modeling Processes*.

The methods used for the progression rule styles are listed in Table 1 on page 104.

Table 1   Methods Used for Progression Rule Styles

Style	Methods Used	Comments
Complete on first assignment	`before()==null`	This style of progression rule uses the method listed to determine when the property value changes from Null to any other value.

Table 1   Methods Used for Progression Rule Styles

Style	Methods Used	Comments
Complete on transition	Complete Condition: `before().in()` and `after().in()`	This style of progression rule uses the methods listed to determine when the Step complete condition is met. It does this by testing the original and modified value of a property against one or more values in the rule. Use of the `in()` method enables you to enter a selection of values that can be tested. If you do not enter a value, any value for the property can satisfy the condition.
Start and complete on transitions	Start Condition: `before().in()` and `after().in()` Complete Condition: `before().in()` and `after().in()`	This style of progression rule uses the methods listed to determine when the Step start and Step complete conditions are met. It does this by testing the original and modified value of a property against one or more values in the rule. Use of the `in()` method enables you to enter a selection of values that can be tested. If you do not enter a value, any value for the property can satisfy the condition.

Methods for Progression Rules

In order to express progression rules using the `Advanced Conditions` option in the BPI Modeler, there are methods for Data definitions and their properties. A method represents the value relating to the data object or property at a point in time.

The methods are described in the following table and apply only to a property in a process progression rules.

Table 2   Methods for Progression Rules

Method	Description
`this.data.property.changed()`	This function is used to test when the property value changes. When the property value changes it returns a result of `True`.
`this.data.property.before()==` "*prop-value*"	This evaluates to the value of the property before a change and is executed only when the property values changes. It is implicit in this expression that the property value was as stated by *prop-value* and is now no longer that value. The type returned is the same as the type of the property.
`this.data.property.after()==` "*prop-value*"	This evaluates to the value of the property after a change and is executed only when the property value changes. It is implicit in this expression that the property value was a value other than *prop-value* and is now *prop-value*. The type returned is the same as the type of the property.

Table 2   Methods for Progression Rules

Method	Description
`this.data.created()`	This evaluates to true when the data definition is created. This method is executed only when a new Data definition is created. The method returns a Boolean result.

`this.data.property.in(`
`"prop-value1", "prop-value2",`
`...)`
	This evaluates to true when the value of the propery is set to any of the listed values.
	
	The method returns a Boolean result.

`this.data.terminated()`	This evaluates to true when the data definition is terminated, specifically when the Data definition received an event that is flagged as terminating it. This is specified in the Event Subscription dialog for the Data Definition; `Terminate this instance of the Data Definition after handling the event`. The method returns a Boolean result.
`onEvent("event-name")`	This evaluates to `true` when the specified event is processed by the Data definition. It evaluates to `false` if the event does not match.

Example of Complete Condition for Start Step When Data Definition Created

The following shows an example of the expression that you use when you want a process instance to be started when a Data definition for the process is created. In this case, a new process instance is started only when the start and complete conditions for the start Step in the process are met, for example:

```
this.order.created();
```

where:

- `this` is the reference to the process, in this case the `Order` process.
- `order` is the Data definition.
- `created()` is a method that evaluates to true when the `order` Data definition is created.

Example Start Condition for Start Step When Data Definition Modified

The following shows an example of the expression that you use when you want a process instance to be started when the Data definition for the process is modified. In this case a process instance is started not when the Data definition is created but when there is a change to the Data definition property.

```
this.order.priority.before() == null &&
this.order.priority.after() <= 3;
```

where:

- `this` is the reference to the process, in this case the `Order` process.
- `order` is the primary Data definition.
- `priority` is an integer property.
- `before()` is a method, which returns the value of the property before the change.
- `null` indicates that the property was not initialized before the change.
- `after()` is a method, which returns the value of the property after the change.

Example Start Condition using a Method on a Method

The following example is similar to the example above; however, shows an example of the expression that you use when you want a process instance to be started when the Data definition for the process is modified and the Data definitions contains a specific string.

```
this.order.customer_id.after().contains("X");
```

where:

- `this` is the reference to the process, in this case the `Order` process.
- `order` is the primary Data definition.
- `customer_id` is a string property that is a unique identifier for the customer.
- `after()` is a method, which returns the value of the property after the change.
- `contains()` is a string function, which returns a boolean result depending on whether the string "X" is present in the `customer_id`.

Example Start Condition for a Step Based on Event Properties

The following example shows an expression where a you want a process Step to be started when the business event `Order Updated` is received and the business events contains a `Status` attribute of `Credit Check Done`.

```
onEvent("Order Updated") && event.Status == "Credit Check Done";
```

Case Sensitivity for Expressions

This section describes case sensitivity for expression properties and for expressions containing string constants.

Expression Properties

Properties and methods in expressions are case sensitive, which means the following expression are not equivalent within your BPI system:

- ```
  this.order.priority.before() == null &&
  this.order.priority.after() <= 3;
  ```
- ```
  this.Order.Priority.before() == null &&
  this.Order.Priority.after() <= 3;
  ```

Expressions with String Constants

Expressions containing string constants are case sensitive, which means the following expressions are not equivalent within your BPI system:

- ```
  this.order.status.after() == "a";
  ```
- ```
  this.order.status.after() == "A";
  ```

Subtracting Dates in Filter or Binding Expressions

When you create a filter or binding expression where you want to subtract dates, you need to make sure you create expressions using the following format:

```
(this.dateProp + days (2) + hours (12) > event.dataProp
```

In this example, the duration is added to one date and tested against the second date to achieve the required result.

You cannot use expressions such as the following:

```
(event.dateProp - this.dateProp) > days (2) + hours (12)
```

This expression tests whether the difference between two dates is greater than two days and 12 hours.

B Coercion Rules

This appendix describes the rules for how properties are coerced when evaluating binding, filter and assignment expressions. This is an enforced evaluation by the BPI Modeler, based on the rules described in this Appendix.

Valid data types for use within BPI are:

- String
- Numeric, which can be one of:
  - Integer
  - Long
  - Double
  - Currency
- Boolean
- Date

Assignments

The following table shows how assignments are coerced:

Table 3   Assignment Coercion

Data Type...	Assigned from...
String	Any type
Numeric	Any other numeric value (with possible truncation if assignment is to an `integer`), or from `Strings` provided that the string is numeric. If the assignment does not produce a valid number, you receive an error similar to the following in the Business Impact Engine log file: `Cannot assign value prop-name to property prop-name as it is not a compatible type.`
Boolean	Other Boolean values, or from `Strings` provided that the string is either "true" or "false". If the assignment does not produce a valid boolean expression, you receive an error similar to the following in the Business Impact Engine log file: `Cannot assign value prop-name to property prop-name as it is not a compatible type.`
Date	Other `Dates` or a `Numeric` value (treated as number of milliseconds since 1970).

Expressions

The following table shows how values within expressions are coerced.

Table 4

Operator	Behavior
Arithmetic: • + • – • * • / • %	Supported between Numeric properties or values that can be coerced to numeric using the assignment rules described in Table 3.
Comparison: • < • <= • = • != • > • >=	The comparison operators can be used to provide a: • numeric comparison between Numeric properties or values that can be coerced to numeric using the assignment rules described in Table 3. • date comparison between Date properties. • alphabetic comparison (locale-specific) between String properties or values that can be coerced to Strings using the using the assignment rules described in Table 3 (including Boolean). The result of a comparison is always a boolean value.
String tests: • starts() • ends() • contains() • in()	Performs a test on a String property. Note that these tests cannot be applied to non-string properties.

Table 4

Operator	Behavior
Logical: • ! (Not) • && (And) • || (Or)	Supported between Boolean properties or values that can be coerced to boolean using the assignment rules described in Table 3.
Conditional value if, then, else: • ?:	First operand must be Boolean or be coerced to boolean, other operands can be any type; however they must be the same each other.

Using NULL Within an Expression

The BPI Modeler allows only the `equals` and `not-equals` tests with the `null` constant. However, there might be expressions using other operators involving values that are `null` when these expressions are executed, for example how is the following expression evaluated when *index* is `null`?

```
index + 1
```

The following table shows the rules that apply when comparing data items within an expression where one data item contains a `null` value.

Table 5

Operator	Comparison
• `==` • `!=`	Returns TRUE if the other operand is (or is not) `null`.
All other comparison operators	Always return `FALSE`. For example the following expressions both return `FALSE` if `index` is `null`: • `index > 10` • `index <= 10`
Arithmetic operators	Arithmetic operators for expressions containing a null value always return null as the result, for example: `index + 1` returns `null` if `index` is equal to `null`

Note that the `null` constant is case sensitive and when used in Java expressions, should always be lower case.

Index

T