IBM WebSphere Adapters
Version 7 Release 5 Fix Pack 2 (7.5.0.2)

*IBM WebSphere Adapter for Oracle
E-Business Suite User Guide
Version 7 Release 5 Fix Pack 2
(7.5.0.2)*

IBM

IBM WebSphere Adapters
Version 7 Release 5 Fix Pack 2 (7.5.0.2)

IBM WebSphere Adapter for Oracle E-Business Suite User Guide Version 7 Release 5 Fix Pack 2 (7.5.0.2)

IBM

May 2012

Contents

Chapter 1. Overview of IBM WebSphere Adapter for Oracle E-Business Suite

With IBM® WebSphere® Adapter for Oracle E-Business Suite, you can create service-oriented integrated applications, which can interact and exchange information with Oracle E-Business Suite. By using the adapter, you can send requests as well as receive events from the underlying Oracle database, without special coding.

WebSphere Adapter for Oracle E-Business Suite enables a two-way communication between the application running on WebSphere Application Server and the underlying Oracle database of Oracle E-Business Suite. Using the adapter, an application can send requests to read, create, modify, or delete data in the Oracle database, in many cases without writing any SQL code. To process requests received from an application, the adapter updates the Oracle database tables using SQL queries or stored procedures. An application can also receive events from the Oracle E-Business Suite. For example, it can be notified that specific objects in Oracle E-Business Suite are updated. To process events that result from changes to the Oracle E-Business Suite, the adapter delivers events to the application. Using the Oracle Business Event System and event notification, changes to the Oracle E-Business Suite can be automatically notified to other applications. By combining event processing by WebSphere Adapter for Oracle E-Business Suite and another adapter, updates can be automatically propagated to other enterprise applications such as Siebel, PeopleSoft, and SAP.

WebSphere Adapter for Oracle E-Business Suite provides a standard interface that integrates with the underlying Oracle database of the Oracle E-Business Suite applications; it supports the Oracle database server by using the Oracle Java Database Connectivity (JDBC) drivers. The Java Runtime Environment (JRE) version required by the Oracle JDBC driver must not be higher than the JRE version in the run time environment. The adapter uses business objects to exchange data between the application and the Oracle database, so the application does not need to use the JDBC application programming interface (API). *Business objects* are containers for application data that represent business functions or elements, such as an Oracle table or the result of an SQL query. The adapter understands the data format provided by the application, and can process the data, perform the operation, and send the results back in that format.

Hardware and software requirements

The hardware and software requirements for WebSphere Adapters are provided on the IBM Support website.

To view hardware and software requirements for WebSphere Adapters, see http://www.ibm.com/support/docview.wss?uid=swg27006249.

Additional information

The following links provide additional information you might need to configure and deploy your adapter:
* The compatibility matrix for WebSphere Business Integration Adapters and WebSphere Adapters identifies the supported versions of required software for

your adapter. To view this document, go to the WebSphere Adapters support page: http://www-947.ibm.com/support/entry/portal/Overview/Software/WebSphere/WebSphere_Adapters_Family.

- Technotes for WebSphere Adapters provide workaround and additional information that are not included in the product documentation. To view the technotes for your adapter, go to the following Web page, select the name of your adapter from the **Product category** list, and click the search icon: http://www.ibm.com/support/search.wss?tc=SSMKUK&rs=695&rank=8 &dc=DB520+D800+D900+DA900+DA800+DB560&dtm.

Technical overview of IBM WebSphere Adapter for Oracle E-Business Suite

The adapter supports integration of databases that are accessible through the JDBC application programming interface (API) with applications running on WebSphere Application Server. The adapter provides outbound and inbound processing under the Java Platform, Enterprise Edition (JEE) Connector Architecture (JCA) and integrates with other applications running on the WebSphere Application Server.

Outbound processing enables an application to access or modify data in a database. The adapter converts a request from the application to an outbound operation, which it runs to create, retrieve, update, or delete data in the database or to run a database program stored in the database. Processing these requests results in the creation, retrieval, update, or deletion of rows in the corresponding database tables. The adapter also enables you to run stored procedures or stored functions that are defined in the database, and to run user-defined SELECT, INSERT, UPDATE, and DELETE statements. You can use the adapter to integrate multiple applications with the same database.

An application running on WebSphere Application Server invokes a service in an outbound module, which sends a request to the adapter to process one or more business objects. The adapter uses the JDBC API to connect to the database server, which accesses the tables and other objects in the database.

Inbound processing enables an application to receive notification when objects in the database are changed. For example, an application can be notified when rows are created, updated, or deleted in selected database tables.

A database application changes tables in the database. The change causes a trigger, or another automated mechanism, to update the event store with information about the change. Periodically, the adapter polls the event store, retrieves, and processes events, and then delivers them to the export of a module that is part of an application that runs in WebSphere Application Server.

The adapter can process events in one of the following ways:
- Standard event processing, using an event store that is populated by the database application
- Custom event processing, using a user-defined database query

During *standard event processing*, when data is changed in the tables in the database, appropriate events are inserted into a database table called an event store, along with relevant information, such as key values. To capture the changed data in the database, you can place triggers on the respective tables, or use other methods such as Oracle Change Data Capture, which is provided for Oracle databases. To capture the changed data in Oracle E-Business Suite, configure the

Oracle Business Event system and Concurrent Program. The adapter polls the event store and retrieves a batch of events. The events can be filtered by business object type and time stamp and connector ID. The adapter uses each event to construct a business object that contains the business objects changed by that event. The business object is then dispatched to the exports that are configured to receive the specific business object.

During *custom event processing*, the adapter runs a query that was specified by the user as a standard SQL statement, a stored procedure, or a stored function. Any of these actions returns a result set for data returned by the query. Each row of the result set corresponds to a row in the event store. The adapter constructs a business object for each event and delivers it to the exports (also called endpoints) that are configured for (or have subscribed to) the specific business object.

For both standard and custom event processing, you can specify how often the adapter polls for events and how many events it retrieves each polling period.

Outbound processing

When an application component needs to query the existence of a record in the database or retrieve or modify data in the underlying Oracle database, the adapter acts as the connector between the application component and the database. The adapter provides a set of standard outbound operations, which process either the after-image style business objects. The adapter also supports both local and XA (distributed) transactions for outbound processing.

The adapter business object model uses the after-image style of business object for making updates. An *after-image* business object is one that contains the complete state of the business object after the required changes have been made to it.

Supported operations

Table 1 lists the outbound operations that are supported for each type of business object and indicates whether each supports after-image style processing.

Table 1. Outbound operations supported by type of business objects

Business objects supported	Operation	After-image support
Tables Views Synonyms - Nick Names	Create	Yes
	Update	Yes
	Delete	Yes
	Retrieve	Not applicable
	RetrieveAll	Not applicable
	Exists	Not applicable
Stored procedures	Execute	Not applicable
Queries	RetrieveAll	Not applicable

Transaction management

The adapter supports both local and XA (distributed) transactions for outbound processing. In the adapter, a transaction is an isolated interaction with the database. A transaction can consists of multiple operations on the database that are performed as an atomic unit. These operations are not affected by simultaneously occurring operations from other users of the database.

The adapter supports transactions only if the database server supports transactions. The types of transactions that are supported are local and XA transactions:

- A *local transaction* is one in which a component defines the start and end of the transaction with a single database. It uses a one-phase commit protocol. The transaction is managed and performed by the database.
- An *XA transaction* is one in which the transaction can span multiple heterogeneous databases. It uses a global or two-phase commit protocol. The transaction manager coordinates the transaction.

XA transactions

The adapter supports XA transactions for outbound processing. Choose one of these methods to configure the adapter for XA transactions:

- Specify a JNDI data source that supports XA transactions, using the XADataSourceJNDIName and PoolDataSourceJNDIName properties.
- Specify an XA data source, using the XADataSourceName property.

The XADataSourceJNDIName property represents a data source created within WebSphere Application Server. If you define a JNDI data source that supports XA transactions on the server, and specify that data source when you configure the adapter, the adapter participates in XA transactions. Optionally, if you use an XA data source, the adapter participates in XA transactions.

Outbound operations

Application components use operations to perform actions such as retrieving from a database. The adapter provides certain outbound operations. Details are provided on how the adapter processes business objects for each of the supported operations.

An operation can be performed using a standard SQL statement provided by the adapter or by a stored procedure that you define. You can run a stored procedure to perform the operation or to do custom processing before or after the operation. In each business object, you can configure how each operation is performed.

Create operation:

The Create operation creates rows in database tables corresponding to the business object in the request. When given a hierarchical business object, the Create operation recursively traverses the business object, creating rows corresponding to each business object in the hierarchy.

To process the Create operation, the adapter performs the following actions:

1. Recursively inserts each single-cardinality child business object contained with ownership into the database. In other words, the adapter creates the child and all child business objects that the child and its children contain.

   If the business object definition specifies that an attribute represents a child business object with single-cardinality and that attribute is empty, the adapter ignores the attribute. However, if the business object definition requires that the attribute represent a child, and it does not, the adapter returns an error and stops processing.

2. Retrieves and checks for the existence of each single-cardinality child business object contained without ownership. If the retrieval is unsuccessful, indicating that the child does not exist in the database, the adapter returns an error, and stops processing. If the Retrieve operation is successful, the adapter recursively

updates the child business object. If the retrieve operation is successful, the adapter continues the process of creating the parent business object; the adapter does not update the child business object without ownership.

Note: For this approach to work correctly when the child business object exists in the database, primary-key attributes in child business objects must be cross-referenced correctly on Create operations. If the child business object does not exist in the application database, the primary-key attributes must not be set.

3. Inserts the top-level business object in the database by performing the following actions:

   a. Sets each of the foreign-key values of the top-level business object to the primary key values of the corresponding child business object represented with single-cardinality. Because values in child business objects can be set by database sequences or counters or by the database itself during the creation of the child, this step ensures that the foreign-key values in the parent are correct before the adapter inserts the parent in the database.

   b. Generates a new, unique ID value for each attribute that is set automatically by the database. The name of the database sequence or counter is stored in the attribute application-specific information. If an attribute has an associated database sequence or counter, the value generated by the adapter overwrites any value passed in by the application server.

   c. Inserts the top-level business object into the database.

   Note: The adapter treats an empty complex column as null value irrespective of setting it to null or unset.

4. Processes each of its multiple-cardinality child business objects as follows:

   a. Sets the foreign-key values in each child to reference the value in the corresponding primary key attributes in the parent. Because the parent primary key values might have been generated during the creation of the parent, this ensures that the foreign-key values in each child are correct before the adapter inserts the child into the database.

   b. Inserts each of the multiple-cardinality child business objects into the database.

Retrieve operation:

The Retrieve operation extracts data from a database for a hierarchy of business objects.

To process the Retrieve operation, the adapter performs the following actions:

1. Removes all child business objects from the top-level business object it received. In other words, it makes a copy of the top-level business object without any children.

2. Retrieves the top-level business object from the database.

   - If the retrieval operation returns one row, the adapter continues processing.

   - If the retrieval operation returns no rows, indicating that the top-level business object does not exist in the database, the adapter returns the `RecordNotFoundException` error.

   - If the retrieval operation returns more than one row, the adapter returns the `MultipleMatchingRecordsException` error.

   The Retrieve operation uses only the primary key. Other columns are ignored.

3. Recursively retrieves all multiple-cardinality child business objects.

Note: The adapter does not enforce uniqueness when populating an array of business objects. The database is responsible for ensuring uniqueness of the business objects. If the database returns duplicate child business objects, the adapter returns duplicate children.

4. Recursively retrieves each of the single-cardinality children, regardless of whether the child business object is contained with or without ownership.

   Note: All single-cardinality child business objects are processed based on their occurrence in the business object and before the parent business object is processed.

Retrieving NULL data

The adapter can retrieve a record from a database table when the column value is NULL. For example, a Customer business object might have these columns: custid, ccode, fname, and lname, where custid and ccode form composite keys. Composite keys are primary keys that refer to more than one attribute and are used to define the uniqueness of the business object. You can retrieve a Customer record for which the value of ccode is NULL. The adapter generates a SELECT statement for the Retrieve operation as:

```
select custid, ccode, fname, lname from customer where custid=? and ccode is null
```

RetrieveAll operation:

The adapter uses the RetrieveAll operation to retrieve an array of business objects from the database. The adapter uses different processes depending on whether the RetrieveAll operation is for database table business objects or for user-specified SQL business objects.

For database table business objects

All of the key and non-key attributes populated in the incoming business object determine the selection criteria for the retrieval. The adapter may retrieve multiple rows for the top-level business object from the database, depending on the attributes selected. All values specified in the top-level business object are used. The settings in the child business object are ignored. If no attributes are populated in the incoming business object, all the rows are retrieved from the respective table in the database.

The name of a generated business object matches the name of the table in the database. For example, the Customer table in the database is represented as a business object named "Customer".

To retrieve an array of business objects, the adapter performs the following actions:

1. Constructs a container business object for all of the retrieved rows. The name of the container business object is the name of the business object with the string "Container" appended to it.

For query business objects

Business objects that are created for user-specified SELECT statements (query business objects) also support the RetrieveAll operation. The J2C Bean wizard generates the query business object by running the user-specified SQL SELECT statement and creating a hierarchy of query business objects.

To process the query business object generated by the J2C Bean wizard for the user-specified SELECT statement, the adapter performs the following actions:

1. Obtains the SELECT SQL statement from the query business object.
2. Determines whether a dynamic WHERE clause is specified in the query business object.
   - If there is a dynamic WHERE clause, the adapter replaces the default WHERE clause in the SELECT statement with the dynamic one.
   - If there is no dynamic WHERE clause, the adapter replaces parameters in the SELECT statement with the corresponding values specified in the query business object.
3. Runs the SELECT statement.
4. Obtains the result set that is returned and populates the query business object values with the data returned from the database, creating a container business object with the structure.
5. Retrieves the entire hierarchy (a *deep retrieve*) of each top-level query business object in the container, if any child business objects are defined for the query business objects.

   Note: A query business object can be a top-level business object only. A query business object cannot have child query business objects.

Retrieving NULL objects

The adapter can retrieve records from a database table when the column value is NULL. For example, a Customer business object might have these columns: custid, ccode, fname, and lname, where ccode need not be a primary key. You can query all the Customer records for which the ccode column is NULL. The adapter generates a select query for the RetrieveAll operation as:

```
select custid, ccode, fname, lname from customer where custid=? and ccode is NULL
```

Update operation:

In an Update operation, the source business object is compared to a business object that is retrieved from the database using the primary keys specified in the top-level, source business object.

When updating a hierarchical business object, the adapter performs the following actions:

1. Uses the primary key values of the source business object to retrieve the corresponding entity from the database. The retrieved business object is an accurate representation of the current state of the data in the database.

   If the retrieval fails, indicating that the top-level business object does not exist in the database, the adapter returns the RecordNotFoundException exception, and the update fails.

   If the retrieval succeeds, the adapter compares the retrieved business object to the source business object to determine which child business objects require changes in the database. The adapter does not, however, compare values in the source business object's simple attributes to those in the retrieved business object. The adapter updates the values of all non-key simple attributes.

   If all the simple attributes in the top-level business object represent keys, the adapter cannot generate an update query for the top-level business object. In this case, the adapter logs a warning and continues.

2. Recursively updates all single-cardinality children of the top-level business object.

   If ownership is true and the child is present in the source business object but not in the retrieved business object, the adapter recursively creates the child in the database.

   The adapter handles single-cardinality children contained with ownership in one of the following ways:

   - If the child is present in both the source and the retrieved business objects, instead of updating the existing child in the database, the adapter deletes the existing child and creates the child.
   - If the child is present in the source business object but not in the retrieved business object, the adapter recursively creates the child in the database.
   - If the child is present in the retrieved business object but not in the source business object, the adapter recursively deletes the child from the database.

   For single-cardinality children contained without ownership, the adapter attempts to retrieve every child that is present in the source business object from the database. If it successfully retrieves the child, the adapter populates the child business object but does not update it, because the adapter never modifies single-cardinality children contained without ownership. If the retrieval fails, the adapter returns an ObjectNotFound exception.

3. Updates all simple attributes of the retrieved business object, except those whose corresponding attribute in the source business object is not specified.

   Because the business object being updated must be unique, the adapter verifies that only one row is processed as a result. If more than one row is returned, the adapter returns an error.

4. Processes each multiple-cardinality child of the retrieved business object in one of the following ways:

   - If the child exists in both the source and the retrieved business object arrays, the adapter recursively updates it in the database.
   - If the child exists in the source array but not in the array of the retrieved business object, the adapter recursively creates it in the database.
   - If the child exists in the array of the retrieved business object but not in the source array, the adapter recursively deletes it from the database unless the application-specific information for the attribute that represents the child in the parent has the KeepRelationship property set to True. In this case, the adapter does not delete the child from the database.

NULL data and the Update operation

The adapter can update a record from a database table when the column value is NULL. For example, a Customer business object might have these columns: custid, ccode, fname, and lname, where custid and ccode form composite keys. Composite keys are primary keys that refer to more than one attribute and are used to define the uniqueness of the business object. You can update a Customer record for which the value of ccode is NULL. The adapter would generate an update query for the Update operation as:

```
update customer set fname=?, lname=? where custid=? and ccode is null
```

Note: The adapter treats an empty complex column as null value irrespective of setting it to null or unset.

Delete operation:

The Delete operation is performed by pruning the incoming business object and then retrieving the complete business object from the database. The Delete operation is then applied recursively on each business object in the hierarchy.

The Delete operation supports physical and logical deletes, depending on the StatusColumnName value in the application-specific information of the business object. If the StatusColumnName value is defined, the adapter performs a logical delete operation. If the StatusColumnName value is not defined, the adapter performs a physical delete operation.

Physical deletes

For physical deletes the adapter takes the following actions:
- It recursively deletes all multiple-cardinality child business objects.
- It deletes the top-level business object.
- It recursively deletes all single-cardinality child business objects contained with ownership.

Logical deletes

For logical deletes the adapter takes the following actions:
- It issues an update that sets the status attribute of the business object to the value specified by the business object-level application-specific information. The adapter ensures that only one database row is updated as a result, and it returns an error if this is not the case.
- It recursively logically deletes all single-cardinality children contained with ownership and all multiple-cardinality children. The adapter does not delete single-cardinality children contained without ownership.

NULL data and the Delete operation

The adapter can delete a record from a database table when the column value is null. For example, a Customer business object might have these columns: custid, ccode, fname, and lname, where custid and ccode form composite keys. Composite keys are primary keys that refer to more than one attribute and are used to define the uniqueness of the business object. You can delete a Customer record for which the value of ccode is null. The adapter generates a delete query for the Delete operation as:

```
delete from customer where custid=? and ccode is null
```

Execute operation:

The Execute operation is used to run stored procedures and stored functions. The J2C Bean wizard generates the required stored procedure business object that corresponds to the stored procedure or stored function definition in the database. The adapter uses the Execute operation to process the stored procedure business object.

The following information provides a simple example of a stored procedure, the business object that is constructed from it, and the steps the adapter uses to process the stored procedure business object with an Execute operation.

A simple example of a stored procedure:

```
PROCEDURE testSP(x IN int, msgSTR INOUT VARCHAR(10), status OUT int,
                 outrec OUT $structname,  retArr OUT $arrayname)
```

The procedure returns two result sets.

For this stored procedure, following is an example of the business object that is constructed:

```
BOLevel ASI
      SPName=testSP
      ResultSet=true
      MaxNumberOfResultSets=2
      ReturnValue = propName
                    Returned if the stored procedure is a function. function).
                    Will be property name corresponding to the child business
                    object if returned value is complex type(array/struct/resultset)
           Defined only if it is a Function

Properties
      x Type=IP
      msgStr Type=IO
      status Type=OP
      outrec Type OP - Child BO for outrec, ASI ChildBOType = struct
      retarr Type OP - n cardinality child BO for retArr, ASI ChildBOType = array
      childBOName1 - Child BO for 1st result set, ASI ChildBOType = resultset
      childBOName2 - Child BO for 2nd result set, ASI ChildBOType = resultset
```

To process this stored procedure business object with an Execute operation, the adapter:

1. Constructs the following stored procedure call: CALL testSP(x, msgStr, status, outrec, retArr).
2. Sets the input parameters x and msgStr on the callable statement.
3. Runs the callable statement.
4. Obtains the return value (if Function) and sets the value in the appropriate attribute if it is a scalar value, or in a child business object if it is a complex value (such as struct, array).
5. Obtains the first result set and creates the container for ResultSet1.
6. Obtains the second result set and creates the container for ResultSet2.
7. Obtains the output parameters msgStr and status, and sets the corresponding attributes on the business object.
8. Obtains the output parameter outrec and creates the child business object from the data returned in outrec. If outrec is a nested struct type, then the adapter recursively creates and stores data in the hierarchical child business object.
9. Obtains the output parameter retArr and creates a multiple cardinality child business object from the data returned in retArr. If retArr is a nested array type, then the adapter recursively creates and stores data in the hierarchical child business object.

Exists operation:

The Exists operation determines whether the database contains records that match the attributes set in a business object.

You can use both key and non-key attributes in the selection criteria.

Note: When using the J2C Bean wizard to discover table objects in a database, you can select multiple tables and add those tables to the selected objects portion of the Object Discovery and Selection screen. However, you cannot use the J2C Bean wizard to link or join the tables that you selected. If the goal of your business application requires the table business object to perform an Exists operation on

joined tables, you need to join the tables in the database to create a view of the joined tables. After you have created a view of the joined tables, you can then perform discovery on the view. The Exists operation would be supported on this view.

To process the Exists operation and send the results based on the specified business object attributes, the adapter performs the following actions:

1. The adapter receives a table business object from the import. This business object can be flat (simple with no child business objects) or hierarchical (complex, containing one, or more child business objects).

   If the business object is hierarchical, it is only for the top-level business object (the individual business object at the top of a hierarchical business object) for which the adapter builds the query.

   Note: The input business object that supports the Exists operation varies depending on the business object type. In addition to being supported by the table business object, the Exists operation is also supported by the views business object and the synonyms and nicknames business object.

2. The adapter uses the table business object to generate an SQL SELECT statement that it sends to the server.

   The SQL SELECT statement used is as follows:

   `select count(*) from TABLENAME where column1=? AND column2=?`

   Here is a sample SQL statement for our example:

   `select count(*) from CUSTOMER where fname='John' AND lname='Smith'`

   In this case, the SQL statement specifies *non-primary key* attributes `fname` and `lname`, with the assigned values of `John` and `Smith`.

   The adapter includes the attribute information from the table business object in the *where* clause of the SQL query.

3. The database server runs the SQL query and sends the results back to the adapter.

4. The adapter obtains the results of the SQL query from the database server and sets the `recordcount` and `status` attributes on the **ExistsResults** business object.

   For example, if the Exists operation determined there are two records that match the attribute and value settings in the business object, the adapter sets `status=true` and `recordcount=2`.

   If a record with the specified attributes is not found, the `status` output parameter is **false** and the `recordcount` output parameter is **0**.

5. The adapter returns the ExistsResult business object to the caller.

The following illustration shows how the adapter processes a table business object with an Exists operation.

Inbound processing

IBM WebSphere Adapter for Oracle E-Business Suite supports inbound event management with event delivery. Events are processed from an event store that is populated either by the database application or from the result of custom queries that you provide. You control how often the adapter polls for events and how many records are delivered to the export at one time.

The adapter polls for changes using one of these methods:

- Standard event processing, in which the adapter examines the event store for events that are stored there by the database application

- Custom event processing, in which the adapter runs user-defined queries, stored procedures, or stored functions

You can customize standard or custom event processing when you use the J2C Bean wizard to configure the adapter initially or at a later time by using the administrative console of the server to change the activation specification properties.

The database object that is the subject of the event is not retrieved until after a notification is delivered to the export. As a result, the detection and notification of any retrieval errors that occur is deferred until after notification of the export. This differs from the event processing in version 6.2.x of the adapter, where retrieval errors can be detected before the adapter notifies the export.

Standard event processing

In standard event processing, the adapter provides the SQL queries that poll for events and ensure that the event is delivered exactly one time.

When records are created, updated, or deleted in the tables in the database, an event program is run immediately. For Oracle database, the database triggers or tools such as Oracle Change Data Capture are run, and, for Oracle E-Business Suite, Oracle Business Event System and the concurrent programs are run. A trigger or other tool writes an event record into the *event store*, which is a persistent cache where event records are saved until a polling adapter can process them. The event store is implemented as a table in the same database as the user tables, which are the tables that contain the database objects accessed by the adapter.

You must define the triggers or set up other tools to report changes to the database tables about which you want to receive events.

The adapter offers assured once delivery, which guarantees that each event is delivered once and only once to the export. If you enable assured once delivery for the module, a transaction ID (XID) is set for each event in the event store. After an event is obtained for processing, the XID value for that event is updated in the event store. The event is then delivered to its corresponding export, and then deleted from the event store. If the database connection is broken or the application is stopped before the event can be delivered, the event cannot be processed completely. In this case, the XID column indicates that the event must be reprocessed and sent to the export again. After the database connection is reestablished or the adapter starts again, the adapter checks for events in the event store that have a value in the XID column. The adapter processes these events first, and then polls the other events during the poll cycles.

The adapter can process all events or filter events by business object type, time stamp, or connector ID. The filter is set by use of the activation specification property EventTypeFilter, FilterFutureEvents, or AdapterInstanceEventFilter. The EventFilterType property has a comma-delimited list of business object types. Only the types specified in the property are processed. If no value is specified for the property, no filter is applied, and all the events are processed. If the activation specification property FilterFutureEvents is set to true, the adapter filters events by timestamp and connector ID. The adapter compares the system time in each poll cycle to the time stamp on each event. If an event is set to occur in the future, it will not be processed until that time. If the AdapterInstanceEventFilter activation specification property is set, only the connector ID specified in the AdapterInstanceEventFilter property are processed.

Custom event processing

In custom event processing, you provide the SQL queries or stored procedures that poll for events.

With custom event processing, you control which events are delivered to the export by providing a database query (the *custom event query*) for the adapter to run in place of the SQL query it uses to poll the event store in standard event processing. The custom event query must perform any necessary filtering. You specify that you want custom event processing by selecting an option in the wizard or by setting the EventQueryType activation specification property in the administrative console.

Custom event processing supports assured once delivery if you create the standard event store for storing XID values. The adapter stores the events returned by the custom event query in the event store and it updates the events with XID values. The adapter processes the events in the same way as for standard event processing. Do not create a custom query that queries the standard event store, because that table temporarily holds the events when the adapter is configured for assured once delivery. In addition, in this situation the event store must not have an automatic generation of event ID values, because the adapter populates the event ID value it retrieves from the custom query in the event store.

Note: When you are using custom event processing, set the property "Ensure once-only event delivery (AssuredOnceDelivery)" on page 164 to `True`.

Turn custom event processing on by selecting an advanced option in the wizard when you configure your module to use the adapter or by setting the EventQueryType activation specification property.

Custom event query

You specify the custom event query to run by providing a user-defined event query in an advanced option in the wizard or by setting the CustomEventQuery activation specification property. Specify one of the following types of programs:

- Standard SQL statements
- A stored procedure
- A stored function

Any of these programs takes an input parameter containing the poll quantity, an activation specification property that the adapter provides at run time. The program can accept other input parameters as well. These programs must return a result set that has the poll quantity number of records and contains the following columns in order: event_id, object_key, object_name, and object_function. The adapter generates the event object from the result set and processes the events.

Standard SQL statements

You can provide an SQL SELECT statement that selects the events to process. The query can have input parameters in addition to the input parameter for poll quantity.

Stored procedure

The custom query can be a stored procedure that accepts the poll quantity as input and returns an output parameter of type result set. Use the following syntax to specify a stored procedure:

```
call procedure_name (?, ?)
```

Where *procedure_name* is the name of the stored procedure to run. The first parameter represents the poll quantity and the second parameter represents the result set.

The stored procedure can accept other input parameters as well, which you provide in the call statement itself, for example:

```
call procedure_name (25, ?, ?)
```

Stored function

The custom query can be a stored function that accepts the poll quantity as input and returns a result set. Use the following syntax to specify a stored function:

```
? = call function_name (?)
```

Where *function_name* is the name of the stored function. The first parameter represents the result set, and the second parameter represents the poll quantity.

The stored function can accept other input parameters, which you provide in the call statement itself, for example:

```
? = call function_name (?, 'abc')
```

Custom update and delete queries

Custom event processing also allows you to provide custom update and delete queries, which are run after each event is processed. You typically use an *update query* to ensure that a database record does not get picked up for processing during subsequent poll cycles. Use a *delete query* when database records need to be deleted after each event is processed. Both the update and delete queries are optional.

Update and delete queries are specified by the CustomUpdateQuery and CustomDeleteQuery activation specification properties. You can enter these queries as a standard SQL statement, a stored procedure, or a stored function. The syntax for the custom update or delete query is the same as the syntax for the custom query. Update and delete queries take an input parameter for the event ID. The adapter provides the value of event ID at run time. The queries can also have additional input parameters, which you provide in the query syntax itself, in the same way as described for the custom event query.

Event store

The event store is a persistent cache where event records are saved until the polling adapter can process them. The adapter uses the event store to track the inbound requests as they make their way through the system. Each time a database record is created, updated, or deleted, the adapter updates the status of the event in the event store. The status of each event is continually updated by the adapter for recovery purposes until the events are delivered to a configured export on the server.

The adapter polls the event records from the event store at regular intervals. In each poll call, a number of events are processed by the adapter. Events are

processed in ascending order of priority and ascending order of the event time stamp. In each poll cycle, the adapter picks up all new events. For each new event, the adapter retrieves the value set in the object key field for the event and then loads the business object that corresponds to the value specified in the object name field. After the object is loaded, the adapter sets the primary key values of the business object based on the value specified in the object key field. After setting the keys, adapter performs a retrieval of the object based on the keys. The business object is created from the retrieved information and is published to the export.

If you have associated a stored procedure with the Retrieve operation of the business object, you can define the mapping between the input parameters of the stored procedure and the business object attributes (generally, primary keys). If such a mapping is defined, then the adapter sets the input parameters for the stored procedure, invokes the stored procedure, and populates the object based on the results obtained from the stored procedure.

For stored procedures and functions, if you defined a mapping between the input parameters of the stored procedure or function and the business object attributes (generally using primary keys) using the RetrieveSP application-specific information, then the adapter sets the input parameters on the stored procedure, invokes the stored procedure, and populates the business object based on the results obtained from the stored procedure.

When the object_function column has the value `Delete`, which indicates that the object was deleted, the object is not retrieved from the database. The keys are set on the data object and the business object is created and delivered to the export.

If an event is successfully posted, the entry is deleted from the event store. For failed events, the entries remain in the event store and the event_status column is set to -1.

The table format and content of the event store are described Table 2.

Table 2. Definition of the event store database table

Column name	Type	Description
XID	String	The unique transaction ID (XID) value for assured once delivery.
event_id	Number	The unique event ID, which is a primary key for the table. This can have the same value as the object_key.
object_key	String	A string that contains keys of the record in the event store that is retrieved. This column cannot be `null`. Specify the value as one or more *key=value* pairs, separated by the semicolon character (;). Alternatively, you can specify only the values for the primary keys separated by the semicolon (;) character. In this case, the values must be specified in the same order as primary keys are defined in the business object.

Table 2. Definition of the event store database table (continued)

Column name	Type	Description
object_name	String	For runtime environments other than WebSphere Application Server: The name of the top-level Java bean that defines the business object. Each business object refers to a table or view. This column cannot be `null`.
object_function	String	The operation corresponding to the event (Delete, Create, Update, and so on). This column cannot be `null`.
event_priority	Number	Identifies the event priority. This value must be a positive integer. This column cannot be `null`.
event_time	Timestamp	Date and time when event was generated. The format is `mm/dd/yyyy hh:mm:ss`.
event_status	Number	The event status is initially set to a value for a new event and updated by the adapter as it processes the event. The status can have one of the following values: • 0: Identifies a new event. • 1: Identifies an event that has been delivered to an export. • -1: An error occurred while processing the event. This column cannot be `null`.
event_comment	String	Any comment associated with the event.
connector_ID	String	The unique identifier for the adapter instance that receives a specific event.

Monitoring inbound events

The adapter supports monitoring inbound events from the Oracle database, in addition to the other events you are monitoring using Business Monitor or WebSphere Business Events.

Monitoring inbound events using IBM Business Monitor:

You can use Rational® Application Developer for WebSphere Software and the adapter to send inbound events to WebSphere Application Server Common Event Infrastructure (CEI), where they are accessible to Business Monitor.

When you select the option to monitor inbound events in the Rational Application Developer for WebSphere Software J2C Bean wizard, the required artifacts are generated to monitor inbound events. These artifacts include the message-driven J2C bean, as well as the interface, Service Bean, interceptor class, helper class, and the event schemas that are required to create a monitor model. You can then deploy the resulting adapter inbound event monitoring application containing the message-driven bean (the adapter application) to a server, either a Business Monitor server or a remote server. The message-driven bean invokes the stateless session bean that makes the events accessible to the client. More importantly, it also listens for events coming in from the Oracle database (inbound events) and uses the interceptor to set up the intercepted inbound events as Common Base Events (CBE). Then, it posts these Common Base Events into a designated Java

Message Service (JMS) destination - Common Event Infrastructure queue, where they are accessible to Business Monitor for further processing.

Important: Inbound event monitoring is available to your application only if you have Business Monitor installed in your environment. For more information about installing Business Monitor, see http://publib.boulder.ibm.com/infocenter/dmndhelp/v8r0mx/topic/com.ibm.wbpm.mon.imuc.doc/inst/intro.html. For more information about the software requirements and configuration, see http://www-304.ibm.com/support/docview.wss?uid=swg27008414.

For enabling inbound event monitoring function, perform the following tasks:

1. Create the event store (see "Creating the event store" on page 43)
2. Launch the J2C Bean wizard (see "Launching the J2C Bean wizard" on page 43)
3. Configure the connector dependencies (see "Configuring the connector dependencies" on page 44)
4. Set the connection properties in J2C Bean wizard (see "Setting connection properties for the J2C Bean wizard" on page 45)
5. Configuring the module for inbound processing (see "Configuring the module for inbound processing" on page 65)

Related References

For a sample on enabling inbound event monitoring for Business Monitor, see http://publib.boulder.ibm.com/infocenter/radhelp/v8r5/topic/com.ibm.j2c.doc/topics/tcreatinginboundapps.html.

For information about how to disable the event monitor function, see http://publib.boulder.ibm.com/infocenter/radhelp/v8r5/topic/com.ibm.j2c.doc/topics/tdisablingwbe.html.

For a sample end-to-end scenario on publishing events to the Business Monitor, see http://publib.boulder.ibm.com/infocenter/dmndhelp/v8r0mx/topic/com.ibm.wbpm.mon.doc/scen/eis.html.

Monitoring inbound events using WebSphere Business Events:

You can use Rational Application Developer for WebSphere Software and the adapter to send inbound events to WebSphere Application Server JMS topic, where they are accessible to WebSphere Business Events.

Note: You cannot create a JMS connection to the remote server when the same connection factory name is duplicated. For more information, see http://publib.boulder.ibm.com/infocenter/wasinfo/v6r0/index.jsp?topic=/com.ibm.websphere.pmc.doc/tasks/tjn0033_.html

When you select the option to monitor inbound events in the Rational Application Developer for WebSphere Software J2C Bean wizard, the required artifacts are generated to monitor inbound events. These artifacts include the message-driven J2C bean, as well as the interface, Service Bean, interceptor class, helper class, and the event schemas that are required to create a monitor model. You can then deploy the resulting adapter inbound event monitoring application containing the message-driven bean (the adapter application) to a server. The message-driven bean invokes the stateless session bean that makes the events accessible to the client. More importantly, it also listens for events coming in from the Oracle database (inbound events) and uses the interceptor to set up the intercepted

inbound events as Common Base Events (CBE). Then, it posts these Common Base Events into a designated Java Message Service (JMS) destination - JMS topic, where they are accessible to WebSphere Business Events for further processing.

Important: Inbound event monitoring is available to your application only if you have WebSphere Business Events installed in your environment. For information about installing WebSphere Business Events, see http://publib.boulder.ibm.com/infocenter/wbevents/v6r2m1/index.jsp?topic=/com.ibm.wbe.install.doc/doc/install.html. The WebSphere Business Events works with WebSphere Application Server version 6.1; it is not supported in WebSphere Application Server version 7.0. For more information about the software requirements and configuration, see http://www.ibm.com/software/integration/wbe/requirements/.

For enabling inbound event monitoring function, perform the following tasks:

1. Create the event store (see "Creating the event store" on page 43)
2. Launch the J2C Bean wizard (see "Launching the J2C Bean wizard" on page 43)
3. Configure the connector dependencies (see "Configuring the connector dependencies" on page 44)
4. Set the connection properties in J2C Bean wizard (see "Setting connection properties for the J2C Bean wizard" on page 45)
5. Configuring the module for inbound processing (see "Configuring the module for inbound processing" on page 65)
6. Generate the eventBOTypeMapping.xml and eventBOTypeMapping.xsd files from the generated inbound session bean. The eventMapping file provides the mapping between the event and the business object schema that the WebSphere Business Event requires for monitoring the event. To generate the eventBOTypeMapping.xml and eventBOTypeMapping.xsd files:

   a. Right-click your session bean.
   b. Select **Source** > **Generate Event Mapping**.

   The EventMapping files get generated in the same folder as your business object schema files.

Related Reference

For integrating WebSphere Business Events with WebSphere Application Adapters, see http://publib.boulder.ibm.com/infocenter/wbevents/v6r2m1/topic/com.ibm.wbe.integrating.doc/doc/integratingusingwebsphereadapters.html.

Business objects

A business object is a structure that consists of data, the action to be performed on the data, and additional instructions, if any, for processing the data. WebSphere Adapter for Oracle E-Business Suite uses business objects to represent tables and views in the database as well as the results of database queries, stored procedures, and stored functions. Business objects can also create a hierarchy of objects from your database and group unrelated tables. Your component communicates with the adapter using business objects.

How the adapter uses business objects

An integrated application uses business objects to access a database. The adapter converts the business objects in outbound requests into JDBC API calls to access the database. For inbound events, the adapter converts the data in the events into business objects, which are returned to the application.

The adapter uses business objects to represent the following types of objects in a database:

- Tables and views
- Synonyms and nicknames
- Stored procedures and stored functions

Query business objects do not represent database objects. Query business objects represent a user-defined SQL query to run against the database.

Note: Before using the business objects to represent the objects types mentioned earlier, ensure that the Java keywords are not used to define the names of tables, views, stored procedures, and stored functions parameters.

Adapters use some business objects for output. These business objects include:

- Container business object, which contains the output from a RetrieveAll operation.
- ExistsResult business object, which contains the output from an Exists operation.

How data is represented in business objects

For table or view business objects

Each column in the table or view is represented by a simple attribute of the table or view business object. A *simple attribute* is an attribute that represents a single value, such as a String, Integer, or Date. Other attributes represent a child business object or an array of child business objects.

Simple attributes within the same business object cannot be stored in different database tables; however, the following situations are possible:

- The database table can have more columns than the corresponding business object has simple attributes; that is, some columns in the database are not represented in the business object. Only those columns needed for the processing of your business object must be included in your design.
- The business object can have more simple attributes than the corresponding database table has columns; that is, some attributes in the business object are not represented in the database. The attributes that do not have a representation in the database either have no application-specific information, are set with default values, or are parameters for stored procedures or stored functions.
- The business object can represent a view that spans multiple database tables. The adapter can use such a business object when processing events triggered by changes to the database, such as Create, Update, and Delete operations. When processing business object requests, however, the adapter can use such a business object only for Retrieve and RetrieveAll requests.

A table business object always has a primary key, even if the corresponding database table does not have a primary key. The adapter uses the column specified in the primary key attribute when it retrieves table business objects. The adapter supports tables that have composite, or multiple, primary keys. If a database table has one or more primary keys, the wizard sets the primary key property for those columns in the table business object. If the database table does not have a primary key, the J2C Bean wizard prompts you for primary key information when you configure that business object. Specify a column that contains unique data, such as a sequence or identity column.

Table and view business objects support the Create, Update, Delete, Retrieve, RetrieveAll, Exists, outbound operations. When running an Exists operation on a hierarchical table business object, only the top-level business object is queried.

For Oracle databases, the adapter supports complex data types such as array, table, structure, or nested structure in table business objects. The type name and the sub attribute details are automatically discovered and displayed for these types. The adapter processes these data types as child business objects of the table business object.

For stored procedure and stored function business objects

In a business object for a stored procedure or stored function, all the input and output parameters for the stored procedure or stored function have corresponding attributes in the business object. If any of the input or output parameters is of a complex type, such as an array or structure, then the corresponding business object attribute is a child business object type with the child business object containing the attributes of the array or structure. If the stored procedure returns a result set, a child business object is created that contains the attributes of the returned result set.

The business object for stored procedures and stored functions supports the Execute outbound operation.

The following Properties view shows business objects generated from a stored procedure that has one input type and two output types. One of the output parameters is of the Struct data type. The J2C Bean wizard generates a business object for the Struct type and adds it as a child object to the parent business object. For the attribute of type Struct in the parent business object, the ChildBOType application-specific information is set to Struct to indicate it is of type Struct.

In the Properties view, the ChildBOTypeName application-specific information is set to the value of the user-defined Struct type in the database.

For query business objects

A business object for a database query defines the SQL statement that performs the query and the parameters that the query requires. The query business object supports the RetrieveAll outbound operation.

As an example, assume a query business object to run the following SELECT statement:

```
select C.pkey, C.fname, A.city from customer C, address A
    WHERE (C.pkey = A.custid) AND (C.fname like ?)
```

The question mark (?) indicates an input parameter for the query. A query can have multiple parameters, each indicated in the SELECT statement by a question mark. Table 3 shows the attributes of the sample query business object. The query business object has simple attributes for each column to be extracted, a simple attribute for each parameter, and a "placeholder object" for the WHERE clause of the query, which holds the WHERE clause after parameter substitution.

Table 3. Attributes of a query business object

Business object attribute	Description
pkey	Corresponds to database column PKEY in the Customer table

Table 3. Attributes of a query business object (continued)

Business object attribute	Description
fname	Corresponds to database column FNAME in the Customer table
city	Corresponds to database column CITY in the Address table
parameter1	The parameter. There is one parameter for each ? (question mark) in the SELECT statement. In a SELECT statement with multiple parameters, subsequent parameters are named parameter2, parameter3, and so on.
jdbcwhereclause	A placeholder object for the WHERE clause

For Oracle databases, the adapter supports complex data types such as array, table, structure, or nested structure in the query result of the business object. The adapter does not support these complex types as parameters in query business objects.

How business objects are created

You create business objects by using the J2C Bean wizard, launched from Rational Application Developer for WebSphere Software. The wizard connects to the database, discovers database objects, and displays them to you. Select the database objects for which you want to create business objects. For example, specify which schemas you want to examine. In those schemas, select tables, views, stored procedures and functions, and synonyms and nicknames. In addition, you can create additional business objects. For example, you can create a business object to represent the results of user-defined SELECT, INSERT, UPDATE, or DELETE statements that are run against the database. The wizard helps you build a hierarchy of business objects, using parent-child relationships.

After you specify the business objects that you want and define the hierarchy of those objects, the wizard generates business objects to represent the objects that you have selected. It also generates other artifacts needed by the adapter.

In some instances, the wizard cannot completely configure a parent-child relationship. For these relationships, you use the business objects editor, launched from Rational Application Developer for WebSphere Software, to modify or complete the definition of a business object hierarchy that was created by the wizard.

Business object hierarchies
Define the relationships between database tables using parent-child relationships and data ownership in hierarchical business objects.

Business objects can either be flat or hierarchical. In a flat business object, all attributes are simple and represent one row in the database table. Hierarchies can contain related or unrelated business objects. Related business objects have parent-child relationships, with or without ownership.

The term *hierarchical* business object refers to a complete business object, including all the child business objects that it contains at any level. The term *individual* business object refers to one business object, independent of child business objects that it might contain or parent business objects that contain it. The individual business object can represent a view that spans multiple database tables. The term *top-level* business object refers to the individual business object at the top of the hierarchy, which does not itself have a parent business object.

A hierarchical business object has attributes that represent a child business object, an array of child business objects, or a combination of the two. In turn, each child business object can contain a child business object or an array of child business objects, and so on.

A *single-cardinality relationship* occurs when an attribute in a parent business object represents one child business object. In this case, the attribute is of the same type as the child business object. The adapter supports single-cardinality relationships, and single-cardinality relationships and data without ownership.

A *multiple-cardinality relationship* occurs when an attribute in the parent business object represents an array of child business objects. In this case, the attribute is of the same type as the child business objects.

Use the following types of relationships between business objects to define a hierarchy that represents your database tables:
- Single-cardinality relationships
- Single-cardinality relationships and data without ownership
- Multiple-cardinality relationships
- Child business objects with multiple parents

In each type of cardinality, the relationship between the parent and child business objects is described by the application-specific information of the key attributes in the business object storing the relationship.

Single-cardinality relationships in business objects:

In a single-cardinality relationship, an attribute in a parent business object represents one child business object. In this case, the attribute is of the same type as the child business object. The adapter supports single-cardinality relationships, and single-cardinality relationships and data without ownership.

Single-cardinality relationships

Typically, a business object that contains a single-cardinality child business object has at least two attributes that represent the relationship. The type of one attribute is the same as the child's type. The other attribute is a simple attribute that contains the child's primary key as a foreign key in the parent. The parent has as many foreign key attributes as the child has primary key attributes.

Because the foreign keys that establish the relationship are stored in the parent, each parent can contain only one child business object of a given type.

A parent business object can have a single-cardinality child with ownership and a single-cardinality child without ownership. Lookup tables are used for relationships without ownership. Ownership is indicated by the value of the Ownership application-specific information.

Single-cardinality relationships and data without ownership

Typically, each parent business object owns the data within the child business object that it contains. For example, if each Customer business object contains one Address business object, when a new customer is created, a new row is inserted into both the Customer and Address tables. The new address is unique to the new

customer. Likewise, when deleting a customer from the Customer table, the customer's address is also deleted from the Address table.

However, situations can occur in which multiple hierarchical business objects contain the same data, which none of them owns. For example, assume that the Address database table contains a reference to the StateProvince lookup table. Because the lookup table is rarely updated and is maintained independently of the address data, creating or modifying address data does not affect the state and province data in the lookup table. However, to be able to retrieve the StateProvince business object along with the Address business object, StateProvince must be a single-cardinality child of Address and the relationship must be defined without data ownership.

If your database design includes lookup tables, your business object design differs slightly from the database design. This is because the adapter retrieves data only for a table business object and its child table business objects. To use a lookup table, you need to create a single-cardinality parent-child relationship between the tables, without ownership. Although the StateProvince lookup table is not a child of the Address table in the database, the corresponding StateProvince business object is a single-cardinality child of the Address table business object because each address contains a single state or province. However, the Address business object does not "own" the StateProvince business object. Changes to an address do not result in a change to the list of states and provinces.

When the adapter receives a hierarchical business object with a Create, Delete, or Update request, the adapter does not create, delete, or update single-cardinality child business objects contained without ownership. The adapter performs only Retrieve operations on these business objects. If the adapter fails to retrieve such a single-cardinality business object, it returns an error and stops processing; it does not add or change values in the lookup table's business object.

Denormalized data and data without ownership

In addition to facilitating the use of static lookup tables, containment without ownership provides another capability: synchronizing normalized and denormalized data.

Synchronization of normalized to denormalized data: When the relationship is without ownership, you can create or change data when you synchronize from a normalized application to a denormalized one. For example, assume that a normalized source application stores data in two tables, A and B. Assume further that the denormalized destination application stores all the data in one table such that each entity A redundantly stores B data.

In this example, to synchronize a change in table B data from the source application to the destination application, you must trigger a table A event whenever table B data changes. In addition, because table B data is stored redundantly in table A, you must send a business object for each row in table A that contains the changed data from table B.

Note: When making updates to denormalized tables, ensure that each record has a unique key so that multiple rows are not modified as a result of one update. If such a key does not exist, the adapter provides an error stating that multiple records have been updated.

Synchronization of denormalized to normalized data: When synchronizing data from a denormalized source application to a normalized destination application, the adapter does not create, delete, or update data contained without ownership in the normalized application.

When synchronizing data to a normalized application, the adapter ignores all single-cardinality children contained without ownership. To create, remove, or modify such child data, you must process the data manually.

Multiple-cardinality relationships:

In a multiple-cardinality relationship, an attribute in the parent business object represents an array of child business objects. The attribute is of the same type as the child business object. The foreign key that describes the relationship is stored in the child, except when an application stores a single-child entity. Then the parent-child relationship is stored in the parent.

Typically, a business object that contains an array of child business objects has only one attribute that represents the relationship, and this attribute is normally the primary key. The type of the attribute is an array of the same type as the child business objects. For a parent to contain more than one child, the foreign keys that establish the relationship are stored in the child.

Therefore, each child has at least one simple attribute that contains the parent's primary key as a foreign key. The child has as many foreign key attributes as the parent has primary key attributes.

Because the foreign keys that establish the relationship are stored in the child, each parent can have zero or more children.

A multiple-cardinality relationship can be an N=1 relationship. Some applications store one child entity so that the parent-child relationship is stored in the child rather than in the parent. In other words, the child contains a foreign key whose value is identical to the value stored in the parent's primary key.

Applications use this type of relationship when child data does not exist independently of its parent and can be accessed only through its parent. Such child data requires that the parent and its primary key value exist before the child and its foreign key value can be created.

Database tables with multiple parent tables:

If a child table in the database has more than one parent table, you must manually configure the additional parent business objects using the editor. The J2C Bean wizard configures only one parent.

The business object schema

The business object schema is built out of database objects that you select when you run the J2C Bean wizard. Each database object translates into a top-level business object.

The schema defines business objects names and application-specific information. The business objects and their attributes and application-specific information are represented in the schema as follows:

- The business object maps to a complex type definition.

- The application-specific information for the business object is contained in annotations at the complex type.
- The attributes of the business object map to element type definitions.
- The application-specific information for each property in the business object is contained in annotations for the element types.

The template for the application-specific properties for the business object and for the attributes is defined in the metadata schema for the adapter. The name of the schema file is `OracleEBSASI.xsd`. The schema file generated for the adapter has a reference to this template in its annotations.

Tables, views, and synonyms overview

The Oracle database provides the common database objects such as tables, views, synonyms, and so on.

Table is a common database schema object in a database. The table objects are used to store data in the database. Oracle provides a number of table types for entering data into the database such as object, normal, nested, clustered, index-organized, and so on. Data can be entered into the database table either directly or through the Oracle open interface tables. The Oracle open interface tables interface between two systems. These tables import data from one system, validates data, identifies errors if any, and transfers the valid data to the Oracle base tables appropriately.

View is a database object that is considered as a virtual table or a stored query. Views do not store the data physically in them; instead they derive the data from the base table on which they are created. All the operations that are performed on the tables can be performed on views too.

Synonym is an alias or an alternative name given for the database objects such as the tables, views, and other database objects.

Stored procedure overview

A stored procedure can be a business object that your module runs with the Execute operation. It can run in place of the standard SQL for an operation on any business object, or it can perform additional actions before or after performing an operation.

A stored procedure is a group of SQL statements that form a logical unit and perform a particular task. A stored procedure encapsulates a set of operations or queries for the adapter to run on an object in a database server. The adapter uses stored procedures in the following ways:

- By creating a stored procedure business object to run against your database
- By enhancing a business object's operations by replacing the SQL statement provided for a business object's operation or by performing actions before or after the operation runs

Stored procedure business object overview

You can create a stored procedure business object that corresponds to a stored procedure or stored function in the database. You can then use the Execute operation to run the stored procedure against the data in the database.

The J2C Bean wizard helps you to build stored procedure business objects that run a stored procedure or stored function. The wizard examines the stored procedure

or stored function in the database to create the business object. A stored procedure business object has an attribute for each parameter.

For validating the stored procedure attributes, a sample value parameter is provided with each attribute. The sample value parameter is provided for both simple and complex data type attributes. The wizard uses the sample values that you provide to validate the stored procedure before saving it. The adapter uses the result that the stored procedure returns to validate the parameters, to obtain the maximum number of result sets returned, and to use the metadata of these result sets to generate child business objects. The wizard generates the hierarchy for stored procedure business objects automatically if you validate the stored procedure business object.

For both simple and complex data type attributes, the adapter automatically discovers and displays the data type and type name for each attribute of the stored procedure. If the stored procedure has input/output parameters or returns value parameters that are of complex data types such as the Struct or Array, the data type and the corresponding user-defined type name are saved in the SPComplexParameterTypeName property.

If the stored procedure contains the Oracle PL/SQL data type such as Record, the adapter creates a wrapper package with a wrapper stored procedure which converts the Record data type to Object data type so that the Oracle E-Business Suite adapter can support the Oracle PL/SQL Record data type. The names of the wrapper packages and stored procedures created for this purpose comply with the Oracle database object naming conventions along with the appropriate suffixes to differentiate them from the Oracle database object names. The names of the wrapper package and wrapper stored procedure consists of both the original package and stored procedure names along with the appropriate suffixes such as "_WPKG" and "_W" (for example, PKGA_PROC7_REC_TAB_WPKG.PKGA_PROC7_REC_TAB_W, where "PKGA" is the original package name, "PROC7_REC_TAB" is the original stored procedure name, "_WPKG" is the suffix for the package, and "_W" is the suffix for stored procedure). For each selected overloaded stored procedure and function which has PL/SQL Record type parameters, the adapter will create wrapper stored procedure in specific wrapper package with two different number tag (for example, "XXXX_WPKG01 and XXXX_WPKG02").

In the generated Wrapper for SP/SF with Record type parameter, if you select the boolean type parameter as TRUE or False during runtime, the adapter processes this datatype automatically.

WebSphere Adapter for Oracle E-Business Suite distinguishes the original SP/SF from the overloaded ones by a number tag that corresponds to an overload sequence in the Oracle database. These corresponding parameters for the selected SP/SF will be added based on the overload sequence.

If the stored procedure returns a result set, you need to set the number of result sets returned from this stored procedure in the MaxNumberOfResultSets property. This value represents the maximum number of result sets that are handled by the adapter run time.

During discovery and at run time, the WebSphere Adapter for Oracle E-Business Suite expects the returned result set from the stored procedure execution to contain columns with names. Some stored procedures return result set with unnamed

columns. For example, a stored procedure with the SQL statements like the examples that follow return result set with unnamed columns:

```
SELECT COUNT(*) FROM EMPLOYEE;
SELECT 111,222,333 FROM CUSTOMER;
```

Oracle processes such SQL SELECT statements by assigning "dummy" names to the table columns in the returned result set- like count(*) or d1, d2, d3 for the respective select statement examples shown above.

If the returned result set contains table columns with no names (because the database did not assign dummy names), the adapter creates dummy names for such columns.

Dummy column names, generated by either the database or by the adapter, are assigned to the attributes of the stored procedure business object.

The behavior (by the adapter or by the database) of assigning dummy names to unnamed table columns ensures that the stored procedure runs successfully during discovery and at run time.

For stored procedure business objects, the wizard supports nested Struct and Array objects, and can support any number of layers of nested hierarchy. The wizard can generate corresponding child business objects for all these nested Struct and Array objects.

Table 4. Complex data type properties for stored procedure business objects

Property name	Type	Description
SPComplexParameterType	String	Value can be one of: Array ResultSet Struct
SPComplexParameterTypeName	String	The name of the user-defined type. This property is required when the value of SPComplexParameterType is Struct or Array.
MaxNumberOfResultSets	Integer	The maximum number of returned result sets to be handled by the Adapter for Oracle E-Business Suite run time.

Stored procedures used in place of or in addition to operations

You can specify that the adapter use a stored procedure in the database in place of, before, or after the SQL statements that the adapter uses to perform an operation. Each business object can have a different set of stored procedures used with each operation.

The adapter can use simple SQL statements for Create, Update, Delete, Retrieve, or RetrieveAll operations. The column names used in the SQL statements are derived from an attribute application-specific information. The WHERE clause is constructed using key values specified in the business object. Each query spans one table only, unless posted to a view. However, you can replace or enhance the SQL statement provided by the adapter using stored procedures and stored functions.

WebSphere Adapter for Oracle E-Business Suite distinguishes the original SP/SF from the overloaded ones by a number tag that corresponds to an overload sequence in the Oracle database. These corresponding parameters for the selected SP/SF will be added based on the overload sequence.

The adapter can call a stored procedure or stored function in the following circumstances:

- Before processing a business object, to perform preparatory operational processes
- After processing a business object, to perform actions after the operation
- To perform a set of operations on a business object, instead of using a simple Create, Update, Delete, Retrieve, or RetrieveAll statement.

In a hierarchical business object, if you want the stored procedure to be performed for each business object in the hierarchy, you must separately associate a stored procedure with the top-level business object and each child business object or an array of business objects. If you associate a stored procedure with the top-level business object but do not associate it with each child business object, then the top-level business object is processed with the stored procedure, but the child business objects are processed using the standard SQL query.

Table 5 lists the application-specific information elements for a stored procedure and describes their purpose and use. A complete description of each element is provided in the sections that follow the table. A screen showing the stored procedure definition for a business object is shown in "View of business object with stored procedure definition" on page 31.

Table 5. Application-specific information for stored procedures in table and view business objects

Descriptive name	Element name	Purpose
Stored procedure type	StoredProcedureType	The stored procedure type defines the type of stored procedure to be used, and this determines when the stored procedure is called, for example, before processing a business object.
Stored procedure name	StoredProcedureName	The name of the stored procedure that is associated with the appropriate StoredProcedureType.
Result set	ResultSet	This value specifies whether the stored procedure returns a result set. If the result set is returned, a multiple-cardinality child for the current business object is created using the values returned in the result set rows.
Parameters	Parameters	Each Parameters element describes one parameter for a stored procedure or stored function.
Return value	ReturnValue	A value that indicates it is a function call, not a procedure call, because the value is returned by the stored procedure.

Stored procedure type

The stored procedure type defines the type of stored procedure to be used, and this determines when the stored procedure is called, for example, before processing a business object.

Table 6. Stored procedure type element characteristics

Required	Yes
Default	None
Possible values	Can be one of: • Before*Operation*SP • After*Operation*SP • *Operation*SP *Operation* specifies one of the operation names: Create, Update, Delete, Retrieve, or RetrieveAll.
Bidirectional transformation supported	No
Property type	String
Usage notes	Stored procedure types associated with RetrieveAll apply to top-level business objects only. You can remove any selected application-specific information from the StoredProcedureType property. All the corresponding operation application-specific information property groups are also removed.
Examples	• CreateSP: Performs the create operation • UpdateSP: Performs the update operation • BeforeCreateSP: Runs before creating a business object • AfterCreateSP: Runs after creating a business object • AfterDeleteSP: Runs after deleting a business object

Stored procedure name

The name of the stored procedure that is associated with the appropriate StoredProcedureType.

Table 7. Stored procedure name element characteristics

Required	Yes
Default	None
Bidirectional transformation supported	Yes
Property type	String

Result set

This value determines whether the stored procedure returns a result or not. If the result set is returned, a multiple-cardinality child for the current business object is created using the values returned in the result set rows.

Table 8. Result set element characteristics

Required	Yes
Default	None
Possible values	True False
Bidirectional transformation supported	No
Property type	Boolean
Usage notes	If your stored procedure returns a result set, use the business object editor after finishing the J2C Bean wizard to verify that this attribute is set to true. The Oracle JDBC driver does not always return this value correctly.

Parameters

There is one Parameters element for each parameter for a stored procedure or stored function. Each Parameters element defines the name and type of one parameter.

Table 9. Parameters element characteristics

Required	Yes
Default	None

Contents	Each Parameters element specifies the following information:
	• PropertyName: Specifies the name of the business object attribute to pass as the parameter.
	• Type: Specifies the type of the parameter, one of the following values:
	  – IP for input only
	  – OP for output only
	  – IO for input and output
	  – RS for result set

Bidirectional transformation supported	No
Property type	String
Usage notes	A result set can be returned only as an output parameter. In that case, one of the parameters must have the type RS, to indicate a result set.

Return value

A value that indicates it is a function call, not a procedure call, because a value is returned.

Table 10. Return value element characteristics

Required	No
Default	None
Possible values	Can be RS or the name of a business object attribute or child business object.
Bidirectional transformation supported	No

Table 10. Return value element characteristics (continued)

Property type	String
Usage notes	If the returned value is RS, the returned value is a result set and is used to create the multiple-cardinality container corresponding to this business object. If the returned value is the name of an attribute, the value is assigned to that particular attribute in the business object. If the attribute is another child business object, the adapter returns an error. When you associate a stored procedure with a business object that is generated from a table or view, and if the stored procedure is a function, a value is returned from this stored procedure. One ReturnValue application-specific information value is added to the operation application-specific information. The existence of this application-specific information implies that it is a function call and not a procedure call, because a value is being retuned by the function. If the value of this application-specific information is a business object attribute name, the returned value is assigned to that particular attribute in the business object. If the value of this application-specific information is another child business object, the adapter run time returns an error. In summary, if the returned value is of a simple data type, the wizard enables you to bind one business object attribute to it, and the value of this application-specific information is set to the name of that business object attribute. But if the returned value is a result set, the wizard sets the value of this application-specific information to RS. **Note:** A result set can be returned as an output parameter or as a returned value if it is a stored function. The type of the output parameter is set to RS to indicate that this parameter is used to return a result set.

View of business object with stored procedure definition

The following Properties view screen shows the customer business object that has the associated stored procedure information for RetrieveSP and AfterRetrieveSP for the Retrieve operation. The adapter runs the RTASSER.RETR_CUSTNAME stored procedure in place of the standard SQL to retrieve a table business object. After the business object is retrieved, the adapter runs the RTASSER.RETR_CUSTINFO stored procedure.

Stored functions overview

The Oracle database supports stored functions in addition to stored procedures. Stored functions are like stored procedures except that they always return a value. The adapter supports them in a similar manner.

The adapter supports stored functions that a user creates with the CREATE FUNCTION statement. Although this type of function is sometimes called a *user-defined function* (UDF), that term more typically refers to a Java stored function or procedure, which the adapter does not support.

A function call has the following syntax:

```
? = call FunctionName parameter_list
```

Contrast this to a stored procedure call, which has the following syntax:

```
call SPName parameter_list
```

You specify the attribute that contains the returned value by using the ReturnValue business object application-specific information.

Query business object overview

Query business objects run a user-defined SELECT statement against the database and return the matching records in business objects.

The J2C Bean wizard helps you to build query business objects that run user-defined SELECT statements against the database. Specify the SELECT statement, using ? (the question mark) in place of any parameter that can be substituted in the SELECT statement. The wizard then provides an area where you can specify the data type of each parameter and provide a sample value. The sample value must match data in the database because wizard uses the SELECT statement results to create the query business object.

You must validate the configuration of the query in the wizard before saving it. When you validate, the wizard runs the SELECT statement using the sample values. After obtaining the result set, the wizard analyzes the metadata to obtain the column name and column type of all columns. For each column of the returned result set, the wizard generates one corresponding attribute in the query business object.

In the query business object, the **jdbcwhereclause** takes precedence over the values specified for the **parameter <n>**, where 'n' is the parameter index. This means that, if the query business object contains the **parameter <n>** and the value is specified for the **jdbcwhereclause**, then the dynamic query is formed with the WHERE clause using the **jdbcwhereclause** value. This will ignore the values in the **parameter <n>**. However, if you want to use the default WHERE clause then you can specify the WHERE clause parameters by specifying values to **parameter <n>** and un-setting or setting **jdbcwhereclause** to null.

For example, assume that you have specified the following SELECT statement:

```
select * from customer where fname=? and age=?
```

The WHERE clause has two parameters, which are indicated by question marks (?). The first parameter has the data type **string**, to match the data type of the fname column. The second parameter has the data type **int**, matching the age column. If your database has a customer record where the fname column contains the string Mike and the age column contains the integer 27, you can specify those values as sample values when configuring the query business object. The wizard configures the business object to correspond to the returned result set.

You can also create query business objects to retrieve the interface status of the Oracle open interface tables. The interface status column in the interface table contains the error codes. These codes indicate if the data processed from the Oracle open interface table to the Oracle base table is successful or not. The error code can be used to troubleshoot the issue in case of failure during the data transfer.

Service Bean

The business data exchanged between the client application and the resource adapter is represented as Service Bean. The metadata describing the business data is defined as business objects and represented as the XSD schemas. The Service Bean is generated from these XSDs and is the realization of the business objects.

A Service Bean is a structure that consists of data and, in some cases, metadata with additional instructions, for processing the data. It is a generated, hierarchical, Java objects implementing Record interface. The data can represent a business entity, such as an invoice or an employee record.

You create Service Bean by using the J2C Bean wizard, launched from Rational Application Developer for WebSphere Software connector tools. The wizard connects to the Oracle E-Business Suite, discovers data structures in the EIS, and generates Service Bean to represent them. The adapter supports records that are hierarchically structured. Information about the processed object is stored in the application-specific information for the object and each of its attributes.

The J2C Bean wizard

The J2C Bean wizard is a tool you use to build application interacting with the Oracle E-Business Suite system before deploying it to WebSphere Application Server. The wizard establishes a connection to the Oracle E-Business Suite server, discovers business objects and services (based on search criteria you provide), and generates Java data bindings based on the services or functions discovered.

Using Rational Application Developer for WebSphere Software you establish a connection to the Oracle E-Business Suite server to browse the metadata repository on the server. In the wizard, you select the business objects you need. The wizard automatically generates the XSD schemas that represent the business objects and the corresponding Java data bindings.

The result of running the J2C Bean wizard is a library that contains the Oracle Java data bindings representing business objects and the J2C Bean providing interface and implementation for the functionality of the EIS system. The Java data bindings are used as input and output arguments of the methods of J2C Bean.

Log and Trace Analyzer

The adapter creates log and trace files that can be viewed with the Log and Trace Analyzer.

The Log and Trace Analyzer can filter log and trace files to isolate the messages and trace information for the adapter. It can also highlight the adapter's messages and trace information in the log viewer.

The adapter's component ID for filtering and highlighting is a string composed of the characters OEBSRA plus the value of the adapter ID property. For example, if the adapter ID property is set to 001, the component ID is OEBSRA001.

If you run multiple instances of the same adapter, ensure that the first seven characters of the adapter ID property are unique for each instance so that you can correlate the log and trace information to a particular adapter instance. By making the first seven characters of an adapter ID property unique, the component ID for multiple instances of that adapter is also unique, allowing you to correlate the log and trace information to a particular instance of an adapter. For example, when you set the adapter ID property of two instances of WebSphere Adapter for Oracle E-Business Suite to 001 and 002. The component IDs for those instances, OEBSRA001 and OEBSRA002, are short enough to remain unique, enabling you to distinguish them as separate adapter instances. However, instances with longer adapter ID properties cannot be distinguished from each other. If you set the adapter ID properties of two instances to Instance01 and Instance02, you will not be able to

examine the log and trace information for each adapter instance because the component ID for both instances is truncated to OEBSRAInstanc.

For outbound processing, the adapter ID property is located in both the resource adapter and managed connection factory property groups. If you update the adapter ID property after using the J2C Bean wizard to configure the adapter for outbound processing, be sure to set the resource adapter and managed connection factory properties consistently. It prevents inconsistent marking of the log and trace entries. For inbound processing, the adapter ID property is located only in the resource adapter properties, so this consideration does not apply.

For more information, see the "Adapter ID (AdapterID)" on page 131 property.

Standards Compliance

This product is compliant with several government and industry standards, including accessibility standards and Internet Protocol standards.

Accessibility

Administration
> The run time administrative console is the primary interface for deployment and administration of enterprise applications. These consoles are displayed within a standard web browser. By using an accessible Web browser, such as Microsoft Internet Explorer or Netscape Browser, you are able to:
> - Use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen
> - Use voice recognition software, such as IBM via Voice, to enter data and to navigate the user interface
> - Operate features by using the keyboard instead of the mouse
>
> You can configure and use product features by using standard text editors and scripts or command line interfaces instead of the graphical interfaces that are provided. When appropriate, the documentation for specific product features contains additional information about the accessibility of the features.

J2C Bean wizard
> The J2C Bean wizard is the primary component used to create application accessing EIS systems. This wizard is implemented as an Eclipse plug-in that is available through Rational Application Developer for WebSphere Software is fully accessible.

Keyboard navigation
> This product uses standard Microsoft Windows® navigation keys.

IBM and accessibility
> See the IBM Accessibility Center website http://www.ibm.com/able/ for more information about the commitment that IBM has to accessibility.

Internet Protocol, Version 6 (IPv6)

WebSphere Application Server, version 6.1.0 and later and its JavaMail component support dual-stack Internet Protocol Version 6.0 (IPv6). For more information about this compatibility in WebSphere Application Server, see IPv6 support in the WebSphere Application Server information center. For more information about IPv6, see http://www.ipv6.org.

Chapter 2. Planning for adapter implementation

Before using WebSphere Adapter for Oracle E-Business Suite, make sure that you understand the experience you need and the server environment in which it runs. Learn the considerations for deploying the adapter in your server environment, and find out how to improve the performance and availability of the adapter by using a clustered server environment.

Before you begin

Before you begin to configure and deploy the module, you must possess a thorough understanding of the business integration concepts, Java Database Connectivity (JDBC), the Oracle database knowledge, the Oracle E-Business Suite products in your environment, and the features and capabilities of Rational Application Developer for WebSphere Software and WebSphere Application Server.

To configure and deploy WebSphere Adapter for Oracle E-Business Suite you must understand and have experience with the following concepts, tools, and tasks:

- The business requirements of the solution you are building.
- Oracle E-Business Suite, JDBC, and the database products in your environment. This includes Oracle E-Business Suite integration interfaces, Oracle Business Event System, Oracle Workflow, Oracle data access issues, transactional models, and connections across heterogeneous relational databases, queues, and Web services.
- The capabilities and requirements of the server you plan to use for the integration solution such as the WebSphere Application Server. You must know how to configure and administer the host server and how to use the administrative console to set and modify property definitions, configure connection factories, and manage events.
- The tools and capabilities provided by Rational Application Developer for WebSphere Software. You must know how to use these tools to create modules, wire and test components, and complete other integration tasks.

Security

The adapter uses the J2C authentication data entry, or the authentication alias feature of Java 2 security to provide secure user name and password authentication. For more information about security features, see the documentation for WebSphere Application Server.

Support for protecting sensitive user data in log and trace files

You can configure the adapter to prevent sensitive or confidential data, in the log and trace files, from being viewed by users without authorization.

Log and trace files for the adapter can contain data from your Oracle database, which might contain sensitive or confidential information. Sometimes these files might be seen by individuals without authorization to view sensitive data. For example, a support specialist must use the log and trace files to troubleshoot a problem.

To protect the data in such situations, the adapter lets you specify whether you want to prevent confidential user data from displaying in the adapter log and trace files. You can select this option in the J2C Bean wizard or change the HideConfidentialTrace property. When this property is enabled, the adapter replaces the sensitive data with XXX's.

See "Managed connection factory properties" on page 134 for information about this optional property.

The following types of information are considered potentially sensitive data and are disguised:
* The contents of a business object
* The contents of the object key of the event record
* User name, Password, Environment, and Role
* The URL used to connect to the Oracle database

The following types of information are not considered user data and are not hidden:
* The contents of the event record that are not part of the event record object key, for example, the XID, event ID, business object name, and event status
* Business object schemas
* Transaction IDs
* Call sequences

User authentication

The adapter supports several methods for supplying the user name and password that are needed to connect to the Oracle database. By understanding the features and limitations of each method, you can pick a method that provides the appropriate level of security and convenience for your application.

To integrate an adapter into your application, a user name and password are needed at the following times:
* When the J2C Bean wizard connects to Oracle database to extract, or *discover*, information about the objects and services that you can access with the adapter.
* At run time on WebSphere Application Server, when the adapter connects to Oracle database to process outbound requests and inbound events.

Authentication in the wizard

The J2C Bean wizard prompts for the connection information for the discovery process, and then reuses it as the default values of the adapter properties that specify the connection information used at run time. You can use a different user name and password while running the wizard than you use when the application is deployed to the server. You can even connect to a different Oracle database, although the schema name must be the same in both databases. For example, while developing and integrating an application that uses WebSphere Adapter for Oracle E-Business Suite, you might not use the production database; using a test database with the same data format but fewer, simulated records lets you develop and integrate the application without impacting the performance of a production database and without encountering restrictions caused by the privacy requirements for customer data.

The wizard uses the user name and password that you specify for the discovery process only during the discovery process; they are not accessible after the wizard is completed.

Authentication at run time

At run time, the adapter needs to provide the user name and password to connect to the Oracle database. To connect without user intervention, the adapter must access a saved copy of the user information. In a server environment, there are several methods for saving user information. You can configure the adapter to get your user information, through any of the following methods:
- Adapter properties
- J2C authentication alias

Saving the user name and password in adapter properties is a direct way to provide this information at run time. You provide this user name and password when you use the J2C Bean wizard to configure your module. Although directly specifying the user name and password seems the most straightforward method, it has important limitations. Adapter properties are not encrypted; the password is stored as clear text in fields that are accessible to others on the server. Also, when the password changes, you must update the password in all instances of the adapter that access that Oracle database. This includes the adapters embedded in application EAR files as well as adapters that are separately installed on the server.

Using a J2C authentication data entry, or authentication alias, created with the Java Authentication and Authorization Service (JAAS) feature of Java 2 security is a robust, secure way to deploy applications. An administrator creates the authentication alias that is used by one or more applications that need to access a system. The user name and password must be known only to that administrator, who can change the password in a single place, when a change is required.

Deployment options

There are two ways to deploy the adapter. You can either embed it as part of the deployed application, or you can deploy it as a stand-alone RAR file. The requirements of your environment affect the type of deployment option you choose.

The following are the deployment options:
- When you deploy the adapter as an embedded component, the adapter is included within an enterprise application archive (EAR) file and is available only to the application in the EAR file.
- When you deploy the adapter as a stand-alone component, the adapter is represented by a stand-alone resource adapter archive (RAR) file. When it is deployed, it is available to all applications deployed in the server instance.
- **With module for use by single application**: With the adapter files embedded in the module, you can deploy the module to any application server. Use an embedded adapter when you have a single module using the adapter or if multiple modules need to run different versions of the adapter. Using an embedded adapter enables you to upgrade the adapter in a single module without the risk of destabilizing other modules by changing their adapter version.
- **On server for use by multiple applications**: If you do not include the adapter files in a module, you must install them as a stand-alone adapter on each

application server where you want to run the module. Use a stand-alone adapter when multiple modules can use the same version of the adapter and you want to administer the adapter in a central location. A stand-alone adapter can also reduce the resources required by running a single adapter instance for multiple modules.

While creating the project for your application using Rational Application Developer for WebSphere Software, you can choose how to package the adapter [either bundled with the (EAR) file or as a stand-alone (RAR) file]. Your choice affects how the adapter is used in the run time environment, as well as how the properties for the adapter are displayed on the administrative console.

Choosing either to embed an adapter with your application or to deploy the adapter as a stand-alone module depends on how you want to administer the adapter. If you want a single copy of the adapter and do not care about disruption to multiple applications when you upgrade the adapter, then you would be more likely to deploy the adapter as a stand-alone module.

If you plan to run multiple versions, and if you care more about potential disruption when you upgrade the adapter, you would be more likely to embed the adapter with the application. Embedding the adapter with the application allows you to associate an adapter version with an application version and administer it as a single module.

To deploy the RAR file to the application server, you must obtain and install Adapter for Oracle E-Business Suite. This provides the RAR file that you install following instructions supplied with WebSphere Application Server.

Considerations for embedding an adapter in the application

Consider the following items if you plan to embed the adapter with your application:
- An embedded adapter has class loader isolation.

  A class loader affects the packaging of applications and the behavior of packaged applications deployed on run time environments. *Class loader isolation* means that the adapter cannot load classes from another application or module. Class loader isolation prevents two similarly named classes in different applications from interfering with each other.
- Each application in which the adapter is embedded must be administered separately.

Considerations for using a stand-alone adapter

Consider the following items if you plan to use a stand-alone adapter:
- Stand-alone adapters have no class loader isolation.

  Because stand-alone adapters have no class loader isolation, only one version of any given Java artifact is run and the version and sequence of that artifact is undetermined. For example, when you use a stand-alone adapter there is only *one* resource adapter version, *one* adapter foundation class (AFC) version, or *one* third-party JAR version. All adapters deployed as stand-alone adapters share a single AFC version, and all instances of a given adapter share the same code version. All adapter instances using a given third-party library must share that library.

- If you update any of these shared artifacts, all applications using the artifacts are affected.

  For instance, if you have an adapter that is working with server version X, and you update the version of the client application to version Y, your original application might stop working.
- Adapter Foundation Classes (AFC) is compatible with previous versions, but the latest AFC version must be in every RAR file that is deployed in a stand-alone manner.

  If more than one copy of any JAR file is in the class path in a stand-alone adapter, the one that is used is random; therefore, they all must be the latest version.

Considerations while deploying adapters with different versions

When you install multiple adapters with different versions of CWYBS_AdapterFoundation.jar, and if a lower version of the CWYBS_AdapterFoundation.jar is loaded during runtime, the adapter will return the ResourceAdapterInternalException error message, due to a version conflict. For example, when you install Oracle E-Business Suite adapter version 7.0.0.3 and WebSphere Adapter for Oracle E-Business Suite version 7.5.0.2, the following error message is displayed "The version of CWYBS_AdapterFoundation.jar is not compatible with IBM WebSphere Adapter for Oracle E-Business Suite" as IBM WebSphere Adapter for Oracle E-Business Suite loads file:/C:/IBM/WebSphere/ ProcServer7/profiles/ProcSrv01/installedConnectors/CWYOE_OracleEBS.rar/ CWYBS_AdapterFoundation.jar with version 7.0.0.3. However, the base level of this jar required is version 7.5.0.2. To overcome this conflict, you must ensure that all adapters are at same version level. For further assistance, contact WebSphere Adapters Support for help.

There are occasions when you have to work with embedded adapters that do not need a client-server communication, standalone adapters that need a server connection, or a hybrid mix of adapter connections.

The following scenarios cover the different behaviors of AFC version conflict detection, when you are deploying two or more adapters and at least one of the adapter version is 7.5 or higher.

Deploying a standalone Adapter
1. Install WebSphere Adapter for Siebel Business Applications version 7.0.1.0 through the WebSphere Application Server administrative console.
2. Install WebSphere Adapter for SAP Software version 7.5.0.0 through the administrative console.
3. Create ActivationSpec for an ALE passthrough inbound operation.
4. Create an application in Rational Application Developer for WebSphere Software for a standalone ALE passthrough inbound operation.
5. Install and start the application through the administrative console.
6. Verify the error.

Note: An error message will be generated in the log/trace area of WebSphere Application Server, to indicate an AFC version conflict.

Deploying an embedded Adapter

1. Import a build of WebSphere Adapter for PeopleSoft Enterprise version 7.0.1.0, using a RAR file.
2. Create a Peoplesoft Inbound EMD operation.
3. Import a build of WebSphere Adapter for Oracle E-Business Suite version 7.5.0.0, using a RAR file.
4. Create an Oracle inbound EMD operation, in the same module where you have created the Peoplesoft Inbound EMD operation.
5. Deploy the module to WebSphere Application Server.
6. Check the trace.

At step 5, the deployment will fail. At step 6, you will get an internal error message due to the AFC version conflict.

Note: To avoid a name conflict between the business object generated by the two adapters, you may need to generate the artifacts into different folders.

Deploying a combination of standalone and embedded Adapters
1. Install WebSphere Adapter for Oracle E-Business Suite version 7.0.1.0 through theWebSphere Application Server administrative console.
2. Create an ActivationSpec for a Oracle inbound operation.
3. Create an application in Rational Application Developer for WebSphere Software for a Oracle inbound operation, for the standalone Adapter deployment.
4. Deploy the Oracle inbound application and trigger your inbound events.
5. Create an application in Rational Application Developer for WebSphere Software for a WebSphere Adapter for SAP Software version 7.5.0.0 inbound embedded Adapter deployment.
6. Deploy an SAP inbound application, and trigger your inbound events.

Note: You can resolve the AFC version conflict by using different class loaders for the standalone and embedded deployments. You can start both Oracle and SAP inbound applications successfully, and process Inbound events without exception.

For further assistance, visit http://www-947.ibm.com/support/entry/portal/ Overview/Software/WebSphere/WebSphere_Adapters_Family.

WebSphere Adapters in clustered environments

You can improve adapter performance and availability by deploying a module on a clustered server environment. Clusters are groups of servers that are managed together to balance workloads and to provide high availability and scalability.

The module you deployed is replicated across all servers in a cluster, regardless of whether you deploy the module using a stand-alone or an embedded adapter. The following IBM products support WebSphere Adapters in a clustered environment:

- WebSphere Application Server
- WebSphere Application Server Network Deployment
- WebSphere Extended Deployment

To deploy and configure WebSphere Adapter for Oracle E-Business Suite in a clustered environment, see: "Deploying the module in a clustered environment" on page 85. When you set up a server cluster, you create a Deployment Manager

profile. The HAManager, a subcomponent of the Deployment Manager, notifies the Java Platform, Enterprise Edition (JEE) Connector Architecture (JCA) container to activate an adapter instance.

Using WebSphere Extended Deployment, you can optionally enhance the performance of adapter instances in your clustered environment. WebSphere Extended Deployment extends the WebSphere Application Server Network Deployment capabilities by using a dynamic Workload Manager instance instead of a static Workload Manager. The dynamic Workload Manager instance can optimize the performance of adapter instances in the cluster by dynamically balancing the load of the requests. This means that application server instances can be automatically stopped and started based on the load variations, allowing systems with different capacities and configurations to handle load variations evenly.

In clustered environments, adapter instances can handle both inbound and outbound processes.

High availability for inbound processes

Inbound processes are based on events triggered as a result of updates to data in the Oracle database. WebSphere Adapter for Oracle E-Business Suite is configured to detect updates by polling an event table. The adapter then publishes the event to its endpoint.

When you deploy a module to a cluster, the Java Platform, Enterprise Edition (JEE) Connector Architecture (JCA) container checks the enableHASupport resource adapter property. If the value for the enableHASupport property is true, which is the default setting, all of the adapter instances are registered with the HAManager with a policy 1 of N. This policy means that only one of the adapter instances starts polling for events. Although other adapter instances in the cluster are started, they remain dormant with respect to the active event until the active adapter instance finishes processing the event. If the server on which the polling thread was started shuts down for some reason, an adapter instance that is running on one of the backup servers is activated.

Note: In the active-passive configuration mode of the adapters, the endpoint application of the passive adapter instance also listens to the events/messages even if the `enableHASupport` property is set to `True`. This is because the `alwaysactivateAllMDBs` property in the JMS activation specification is set to `True`. To stop the endpoint application of the passive adapter instance from listening to the events, you must set the `alwaysactivateAllMDBs` property value to `False`. For more information, see "Disabling end point applications of the passive adapter" on page 106.

Note: In clustered environments, when the adapter works in a Active-Active configuration, it provides both high availability and load balancing support. This functionality is useful in production environments where high performance is needed.

Important: Do not change the setting of the enableHASupport property.

High availability for outbound processes

In clustered environments, multiple adapter instances are available to perform outbound process requests. Accordingly, if your environment has multiple applications that interact with WebSphere Adapter for Oracle E-Business Suite for

outbound requests, then you might improve performance by deploying the module to a clustered environment. In a clustered environment, multiple outbound requests can be processed simultaneously, as long as they are not attempting to process the same record.

If multiple outbound requests are attempting to process the same record, such as a Customer address, the workload management capability in WebSphere Application Server Network Deployment distributes the requests among the available adapter instances in the sequence they were received. As a result, these types of outbound requests in a clustered environment are processed in the same manner as those in a single server environment: one adapter instance processes only one outbound request at a time.

Chapter 3. Configuring the module for deployment

To configure the adapter so that it can be deployed on WebSphere Application Server, use Rational Application Developer for WebSphere Software to create a module, which is exported as an EAR file when you deploy the adapter. You then specify the business objects you want to discover and the system on which you want to discover them.

Creating the event store

You need to create the event store in the Oracle database before the adapter can process inbound events. You can set triggers on user tables as needed to populate the event store. For Oracle E-Business Suite, you need to configure the Oracle Business Event System and Concurrent Program to capture the event.

About this task

Perform this task only if you need inbound processing of events. Create the event store in the database that contains the tables for which events are reported.

Procedure

1. Create the event store. Sample scripts are provided to create the event store for the Oracle database, as follows:
   - `ibm_websphere_event_table_create.sql`

   The samples are located in the *RAD_installation_dir*/ResourceAdapters/ OracleEBS_*version*/Scripts directory, where *RAD_installation_dir* is the installation directory for Rational Application Developer for WebSphere Software, and *version* identifies the version of the adapter, for example, 7.5.5.
2. For Oracle database, set up triggers on user tables so that the changes to the user tables can automatically generate events that is stored in the event store. If using triggers is not an option, then you can populate the event store by using custom SQL code or stored procedures that are not invoked through triggers. Instead, the triggers can be invoked through a batch program that runs periodically or that you run manually. For Oracle E-Business Suite, it is recommended that you configure the Oracle Business Event System and Concurrent Programs to define the events and automatically capture the events. The Concurrent Programs can be scheduled to run periodically or run manually when required.

Results

The event store is available for event processing.

Launching the J2C Bean wizard

To begin the process of creating and deploying a module, you start the J2C Bean wizard in Rational Application Developer for WebSphere Software. The wizard creates a connector project, which is used to organize the files associated with the module.

Before you begin

Ensure that you have gathered the information you need to establish a connection to the Oracle database. For example, you need the name or IP address of the Oracle database and the user ID and password to access it.

About this task

If you have an existing project, you can use it instead of creating a new one. Select it before you start the wizard.

Procedure

1. To start the J2C Bean wizard, go to the Enterprise Explorer of Rational Application Developer for WebSphere Software and click **File** > **New** > **Other** > **J2C** > **J2C Bean**. Click **Next**.
2. In the Resource Adapter Selection window, expand the Oracle node. Select **IBM WebSphere Adapter for Oracle E-Business Suite (IBM :** *version***)**, where *version* is the version of the adapter you want to use, and then click **Next**.
3. In the **Connector Import** window, accept the default project name in the **Connector project** field or type a different name.
4. In the **Target server** field, select the type of server where you want to deploy the module. The wizard creates the artifacts that are appropriate to that server.
5. Click **Next**. The Connector Settings window is displayed.

What to do next

Continue working in the J2C Bean wizard. The next step is to add dependent JAR files to the project.

Configuring the connector dependencies

The J2C Bean wizard needs a copy of certain files from the Oracle E-Business Suite to be able to communicate with it. Specify the location of the JAR files that contain the JDBC driver for Oracle and any native system library files that are needed.

Before you begin

You must run the J2C Bean wizard in Rational Application Developer for WebSphere Software to perform this task.

About this task

In addition to performing this task when you configure the module, you also need to deploy files on WebSphere Application Server.

Procedure

1. Obtain the Oracle JDBC driver-specific files or native libraries for your database software and operating system from your database administrator or from the database software website. The following table lists the Oracle JDBC driver files needed for the Oracle database software.

Table 11. Oracle JDBC driver files for Oracle database software

Database software	Driver	JDBC driver files	Native System Libraries
Oracle	Thin driver	ojdbc6.jar or ojdbc14.jar depending on the JRE version	None

> **Note:** The Java Runtime Environment (JRE) version required by the Oracle JDBC driver must not be higher than the JRE version in the runtime environment. For example, if the JRE version is 1.5 in the runtime environment, then the Oracle JDBC Driver must be "ojdbc14.jar" which supports both Java Development Kit (JDK) 1.4 and JDK 1.5; if the JRE version is 1.6 in the runtime environment, then the Oracle JDBC Driver must be "ojdbc6.jar" which supports JDK version 1.6.

2. Obtain the Oracle JMS driver-specific files or native libraries for your database software and operating system from your database administrator or from the database software website. The following table lists the Oracle JMS files needed for Oracle database software.

Table 12. Oracle JMS driver files for Oracle database software

Database software	Driver	JMS driver files	Native System Libraries
Oracle	JDBC Thin	ojdbc6.jar or ojdbc14.jar depending on the JRE version	None
Oracle	JMS	apaqi*.jar	None
Oracle	JMS	jmscommon*.jar	None

> **Note:** The jdbc driver and jms driver should be compatible with Java Runtime Environment (JRE) version and Oracle Database version.

3. In the Connector Settings window, specify the location of the Oracle JDBC driver-specific files that are required by the adapter.

   a. In **Oracle JDBC driver JAR files**, click **Add** and select the Oracle JDBC driver files or Oracle JMS driver files.

   b. If you use the Oracle JDBC type 2 driver, click **Add** in **System libraries** to add the native system libraries that are required to access the database server. If you use Oracle JDBC type 4 driver, leave this field empty.

4. Click **Next**. The wizard displays the Adapter Style window.

Results

The wizard has the files it needs to communicate with the database server.

What to do next

Continue working in the J2C Bean wizard. The next step is to provide the information that the wizard needs to communicate with the Oracle database.

Setting connection properties for the J2C Bean wizard

For the J2C Bean wizard to connect to the database instance to discover database objects, connection properties must be specified.

Before you begin

Before you can configure the connection properties, you must have started the J2C Bean wizard.

About this task

The J2C Bean wizard requires these properties to connect to the database for discovery and for creating the service description. For more information about the properties, see "Connection properties for the wizard" on page 130.

Procedure

1. In the Adapter Style window, select **Outbound** to pass data from your service import to the adapter or **Inbound** to pass data from the adapter to your service export, and then click **Next**. The Discovery Configuration window is displayed.

2. Optional: Event monitoring is available to your application only if you have Business Monitor or WebSphere Business Events installed in your environment.

   • Use the following procedure, to generate and monitor common base events and manage these events with the Common Event Infrastructure (CEI) services using Business Monitor:

      a. Select the Enable Inbound Event Monitor check box and then click **Next**. The **Event and JMS configuration** window is displayed.

      b. In the **Event type** field, select WebSphere Business Monitor.

      c. In the **Queue connectionFactory JNDI name** field, accept the default value, jms/cei/EventQueueConnectionFactory.

      d. In the **Queue JNDI name** field, accept the default value, jms/cei/EventQueue.

      e. Click **Advanced** to add advanced properties.

      f. In the Remote JNDI provider configuration section, specify the naming provider URL host name and port name values in the **Naming provider URL Host** and **Naming provider URL port** fields to connect to the remote WebSphere Application Server from the wizard.

      g. In the Connection authentication configuration section, type the user name in the **User name** field to connect to the server from the wizard.

      h. In the Connection authentication configuration section, type the password in the **Password** field to connect to the server from the wizard.

      i. Click **Next**. The Discovery Configuration window is displayed.

   • Use the following procedure, to generate and monitor common base events and manage these events using WebSphere Business Events:

      a. Select the Enable Inbound Event Monitor check box and then click **Next**. The **Event and JMS configuration** window is displayed.

      b. In the **Event type** field, select WebSphere Business Events.

      c. In the **Topic connectionFactory JNDI name** field, accept the default value, jms/WbeTopicConnectionfactory .

      d. In the **Topic JNDI name** field, accept the default value, jms/WBE/CbeListener.

      e. Click **Advanced** to add advanced properties.

f. The **Remote JNDI provider configuration** field is used to configure the remote topics.

If the bus in the local cell has the same name as the bus in remote cell, the application always connect to the local cell. It does not use any of the provider endpoints specified on the connection factory, so the Remote Topic Configuration information that you enter is ignored. For more information about remote Topic Configuration, refer to http://publib.boulder.ibm.com/infocenter/wasinfo/v6r0/index.jsp?topic=/com.ibm.websphere.pmc.doc/tasks/tjn0033_.html .

g. In the **Naming provider URL Host** field, specify the naming provider URL host name to connect to the remote WebSphere Application Server from the wizard.

h. In the **Naming provider URL port** field, specify the naming provider URL port name to connect to the remote WebSphere Application Server from the wizard.

i. In the Connection authentication configuration section, type the user name in the **User name** field to connect to the server from the wizard.

j. In the Connection authentication configuration section, type the password in the **Password** field to connect to the server from the wizard.

k. Click **Next**. The Discovery Configuration window is displayed.

3. In the Discovery Configuration window, specify the connection properties for the wizard to connect to the Oracle database.

a. In the list of database software, select your product and version. The **Properties** area displays fields where you specify database-specific connection properties.

b. In the **JDBC driver type** field, select the type of JDBC driver you want to use.

c. In the **System ID** field, specify the database name.

d. In the **Host name** field, specify the host name or IP address of the database server. If you specify the IP address in IPv6 format, enclose the address within square brackets ([]).

e. In the **Port number** field, specify the port number for connecting to the database. If you select a named driver in the **JDBC driver type** field, the wizard provides a default value for the **Port number** field. If you select the driver Other, the port number is not enabled.

f. If you select a named driver in the **JDBC driver type** field, the wizard provides a default value for the **JDBC driver class name** field and builds the value for the **Database URL** field from other connection fields. If you select the driver Other, you must specify the driver class name and database URL (although part of the database URL might be filled in for you).

g. In the **Additional JDBC driver connection properties** field, specify additional properties to be set when connecting to the database. Specify one or more *name*:*value* pairs, separated by the semicolon character (**;**). For example:

```
loginTimeout:20;readOnly:true;
securityMechanism:USER_ONLY_SECURITY
```

The connection information is used only for the discovery process. Later in the wizard, you can specify different connection information to be used at run time.

4. In the **User name** and **Password** fields, type the user name and password to use to connect to the database from the wizard. This user name is used only

during the discovery process and is not saved. Later in the wizard, you can specify a different user name and password or a different authentication method to be used at run time.

5. In the **Prefix for business object names** field, type a string to be placed at the beginning of business object names. This prefix field does not require a value, and can be left blank.

6. Optional: To set additional advanced properties, click **Advanced**.

7. Optional: To enable bidirectional support for the adapter at run time:

   a. In the **Bidi properties** area, select **Bidi transformation**.

   b. Set the ordering schema, text direction, symmetric swapping, character shaping, and numeric shaping properties to control how bidirectional transformation is performed.

8. Optional: If you want to turn on AUTOCOMMIT for the database, select the **Set auto commit on database connection** check box.

9. To change the location of the log files or the amount of information included in the logs, select the **Specify the level of logging desired** check box, and then provide the following information:

   • In the **Log file output location** field, specify the location of the log file for the wizard.

   • In the **Logging level** field, specify the severity of errors that you want logged.

   This log information is for the wizard only; at run time, the adapter writes the error messages and the trace information into the standard log and trace files for the server.

10. Click **Next**.

    If the wizard generates the com.ibm.adapter.framework.BaseException exception, the adapter cannot connect to the database server. The message contains additional information about a possible cause for the problem. In addition, you can check the logs, which are located in the directory specified in the **Log file output location** field. Make sure that the connection information is correct.

Results

The J2C Bean wizard connects to the database and displays the Object Discovery and Selection window.

What to do next

Continue working in the wizard. The next step is to examine the database to locate the objects for which you want the wizard to create business objects.

Configuring the module for outbound processing

To configure a module to use the adapter for outbound processing, use the J2C Bean wizard in Rational Application Developer for WebSphere Software to find and select business objects and services from the Oracle database, and to generate the business object definitions and related artifacts.

Discovering database objects

After connecting to the database, run a query to search for database objects. Browse through the discovered objects navigation tree to understand the structure

of objects in the Oracle database and use filters to display only the database objects you want to see. Define how many business objects you want to create for user-defined database queries.

Before you begin

You must understand the data requirements of the program that needs to access the database. For example, you need the following information about the database:

- Which schemas your module needs to access
- What type of database objects you need to access in those schemas
- Which tables, views, synonyms or nicknames, or stored procedures or stored functions you need to access
- How many query business objects you need to create, including parameter values and sample database values for the parameters

About this task

This task starts in the Object Discovery and Selection window of the J2C Bean wizard.

Procedure

1. In the Object Discovery and Selection window, click **Edit Query**. The Query Filter Parameters window is displayed.

   Use the Query Filter Parameters window to perform the following tasks:

   - Reduce the search time by searching a subset of database schemas
   - Omit one or more types of database objects from the search
   - Make the wizard prompt you for application-specific information that cannot be automatically determined based on information in the database
   - Specify the number of query business objects you want to create
   - Map the Oracle data types Date and Timestamp to date and dateTime

2. To limit the number of database schemas that are retrieved, type the name of the schema or a name pattern in **Schema name or pattern**. Use the question mark or underscore (? or _) character to match a single character and the asterisk or percentage sign (* or %) to match multiple characters. Only schemas that start with that string or match that pattern are displayed when you run the query. If you do not specify a schema name pattern, all schemas in the database are displayed. Using a filter can speed up the discovery process if your database contains many schemas.

3. To omit one or more types of objects from the search, select the types of objects that you want to omit (tables, views, stored procedures and stored functions, and synonyms or nicknames) in **Supported database object types**, and then click **Remove**. If you change your mind, click **Add** to add the object type back. If you need to access only specific types of database objects and speed the discovery process, omit the ones you do not need.

4. The table, stored procedure, and stored function objects with the date and timestamp data types are mapped to the string data type by default. If you want to map these objects to the actual data types that are supported by the JDBC driver such as the date and datetime data types, select the **Map Oracle JDBC Date/Timestamp types to date/dateTime** check box.

5. Select the **Prompt for additional configuration settings when adding business object** check box. Then, when you add a database object to the list of business objects to create, the wizard automatically prompts you for all

user-configurable application-specific information for the object. For example, if you select this option, the wizard guides you through the process of building a simple parent-child hierarchy of business objects.

6. To create business objects to run the user-defined database queries, select **Create a query business object to build user-defined Select statements** and then type the number of query business objects you want to create. In the **Query Properties** window, specify only the number of query business objects you want to create; the wizard prompts for the name and other details about the business objects at a later time.

   You can create user-defined queries for various business requirements such as to track the interface status of the Oracle open interface table. The interface status ensures if the data transferred from the interface tables to the base tables is successful. The error code stored in the interface status column helps in troubleshooting any failure in data transfer.

7. Click **OK** to save your changes to the database query.

8. In the Object Discovery and Selection window, click **Execute Query** to use the query to discover database objects and to create a template for query business objects.

   The **Objects discovered by query** pane lists the database objects that were discovered.

9. In the **Objects discovered by query** list, click **+** (the plus sign) to expand a schema node and then expand **Tables**, **Views**, **Stored Procedures**, and **Synonyms - Nicknames** nodes beneath it, to see the database objects discovered by the wizard, or search a subset of each object type by using filtering.

   To limit the number of objects displayed for a particular object type, highlight a node but do not expand it. Click the **Filtering** icon. Use the question mark or underscore (? or _) character to match a single character, and the asterisk or percentage sign (* or %) to match multiple characters. Only object types, such as table or view, which start with that string or match that pattern are displayed under the node.

10. Click **+** (the plus sign) to expand the node for **Query Statements** to display the template for query business objects.

Results

The wizard displays the database objects you can access using the adapter and business object template for query business objects.

What to do next

Continue working in the J2C Bean wizard. The next step is to select the objects you want to use in your module, configure each business object, and create hierarchies of business objects.

Selecting and configuring business objects

Using the list of database objects discovered by the J2C Bean wizard, and the query object templates you specified, continue using the wizard to select the database objects that you need to access in your module. Then provide the configuration information for your new business objects.

About this task

The Object Discovery and Selection window allows you to select and configure objects in any order, with the exception that you must select and configure a parent table before you can select and configure its child tables. Aside from this restriction, you have the flexibility to add objects individually or to add them several at a time. You can mix objects from the various nodes of the **Objects discovered by query** list. For example, you can select several table objects, a stored procedure object, and a Query statement, and add and configure them at the same time.

The high-level flow of selecting and configuring business objects is as follows:

1. Select one or more objects in the **Objects discovered by query** list of the Object Discovery and Selection window.
2. Click the > (Add) button.
3. The wizard opens the Configuration Parameters window.
    - If you select a single object, a single Configuration Parameters window is displayed.

       Complete that window, specifying the user-configurable attributes and other information that the wizard cannot discover by examining the database, and click **OK** to save the configuration.
    - If you select multiple objects, the Configuration Parameters window displays with one page for each object selected.

       Click the name of each object in turn. The window displays the same information you would see had you selected this object individually.

       Important: Do not click **OK** until you complete the configuration pages for all the objects. The wizard does not close the notebook until you provide all the required fields, but you can close the window before providing optional fields. If you do not configure the optional fields in the wizard, you must use the business object editor to configure them after exiting the wizard.
4. The wizard adds the configured object to the **Objects to be imported** list.

As long as you do not exit from the wizard, you can work iteratively to select and configure the business objects you need in your module. However, before you start the wizard to add objects to an existing module, understand the requirements of the program that uses the business objects carefully. The wizard overwrites the existing business objects in the same path.

Selecting and configuring tables, views, and synonyms or nicknames for outbound processing

Select and configure business objects for tables, views, and synonyms or nicknames to be used in your module. For inbound processing, these are the business objects that are delivered in events.

Before you begin

To perform this task, you need to understand the structure of the data in the database and know what database objects the module needs to access. Specifically, you need to know the following information:

- The structure of the tables, views, and synonyms or nicknames, including the columns you need, and column attributes such as data type
- The relationships between the tables, including the cardinality and ownership of parent-child relationships

About this task

This task is performed through the J2C Bean wizard. Start in the Object Discovery and Selection window and then work in a Configuration Parameters window that is specific to the business object you are configuring.

Procedure

1. In the **Objects discovered by query** list of the Object Discovery and Selection window, select one or more tables, views, or synonyms from the **Objects discovered by query** list, and click the **>** (Add) button to add the object or objects to the **Objects to be imported** list. Alternatively, you can also filter the tables, views, or synonyms by specifying a valid name or pattern for at least one of the filter fields in the Filter Properties window.

   a. Click the object you want to filter, and then click the ![icon] (Create or edit filter.) button, located at the top of the **Objects discovered by query** pane.

   b. In the Filter Properties window, type a name or pattern in the **Object name or pattern** field. Use the question mark or underscore (? or _) to match a single character, and the asterisk or percentage (* or %) to match multiple characters. The name is not case sensitive.

   c. Click **OK**. The object that matches the given filter condition is displayed.

   d. Select one or more objects from the discovered list, and click the **>** (Add) button to add the object to the **Selected objects** list.

2. If the table has a column that is used to indicate logical deletes:

   a. Select the column name in the **Name of the column used to perform logical deletes** field.

   b. In the **Value used to indicate a deleted object** field, type the value that indicates that a row is logically deleted. You can get this value from your database administrator.

3. If the **Select primary key for table** *table_name* area is displayed, click **Add**, select the column to be used as the primary key for the table business object, and then click **OK**. If the table has a composite key, you can select multiple columns. The **Select primary key for table** *table_name* area is displayed only when the database table does not have a column designated as the primary key. Each table business object must have a primary key, even if the associated database table does not have a key. If the primary key is defined in the database, this section of the window is not displayed.

4. Optional: Define a parent-child relationship between business objects.

   To build a parent-child hierarchy, configure the parent table first, and return to the Object Discovery and Selection window to select and configure the child tables.

   Configure the parent-child relationship using the area of the Configuration Parameters window. These fields are not displayed for the first table you configure.

   a. In the **Choose parent table** field, select the name of the parent table you are configuring. If you do not see the parent table in the list, the parent table has not yet been configured. Go back and configure the parent object before configuring the child objects.

   b. Specify the cardinality of the relationship:

      • If the table has a single-cardinality relationship with the parent table, select the **Single cardinality** check box. In a single cardinality relationship, a parent can have only one child business object of this

type. A single-cardinality relationship can be used with ownership to represent a true child or without ownership to represent lookup tables or other peer objects in a database.

- If the table has a multiple-cardinality relationship, do not select the **Single cardinality** check box. In a multiple-cardinality relationship, a parent can have an array of child business objects of this type.

c. Build the foreign key relationship between the parent and child by specifying for each child column whether it is a foreign key in the parent table.

- If the child column is not a foreign key, select NONE.
- If a child column is a foreign key, select the column in the parent table that corresponds to the child column.

   Note: The wizard can configure only a single parent table. If the child table has multiple parent tables, you must use the business object editor to configure the remaining parent tables after exiting the wizard.

d. If the parent object owns the child object, then the child objects in the database are deleted when the parent is deleted. To indicate that this child is owned by its parent, select the **Parent object owns child object (cascade delete)** check box. Otherwise, clear this option to prevent child objects, such as lookup tables, from being deleted when their parent is deleted.

e. If you do not want the child objects to be deleted as part of an Update operation, select the **Preserves *child_table_name* when the parent is updated** check box.

   When a parent table is updated, the adapter compares the child business objects present in the input with the child business objects returned from the database. By default, the adapter deletes any child objects returned from the database that are not present in the input business object.

f. By default, you can perform operations on parent business objects without specifying the child business objects. If you want to ensure that a parent business object specifies its child business objects when the parent is submitted for a change, select the *Child_table_name* **required for operations on parent** check box.

5. An operation can be performed using either a standard SQL statement generated by the adapter or using stored procedures or stored functions from the database. If you want to use stored procedures or stored functions:

a. Click **Add**.

b. In the Add window, select the type of the stored procedure you want to run. For each operation, you can select a stored procedure that performs the operation, as well as stored procedures that run before or after the operation. For example, for the Create operation, you can specify any of these stored procedures: CreateSP, BeforeCreateSP, and AfterCreateSP.

   Note: If you configure the table with RetrieveAllSP, ensure that at least one parameter of the stored procedure is a Cursor and the ResultSet ASI for the stored procedure is set to `true` to avoid the "No resultset found associated with the stored procedure" exception being generated at run time.

c. Click **OK**. The Configuration Parameters window now shows the stored procedure types you selected and expands to display an area where you configure each one. It might be necessary to scroll down to see the new areas.

Note: In a hierarchical business object, if you want the stored procedure to be performed for each business object in the hierarchy, you must separately associate a stored procedure with the top-level business object and each child business object or an array of business objects. If you associate a stored procedure with the top-level business object but do not associate it with each child business object, then the top-level business object is processed with the stored procedure, but the child business objects are processed using the standard SQL query.

6. For each stored procedure type that you selected, specify the name of the stored procedure in the database and then configure the business object.

   a. In the **Schema name** field, select the name of the schema that contains the stored procedure.

      1) Click **Select**.
      2) In the Select Value window, select the name of the schema you want to work with.
      3) Click **OK**.

   b. Specify the name of the stored procedure or stored function.

      1) In the **Stored procedure name or pattern** field, either type the name of the stored procedure or stored function, or type a name pattern. Use the question mark or underscore (? or _) to match a single character, and the asterisk or percentage sign (* or %) to match multiple characters.
      2) In the **Stored procedure name** field, select the name of the procedure you want. If the stored procedure list contains many items, the **Select** button is displayed next to the **Stored procedure name** field. Click **Select** to open the Select window and select the name of the stored procedure or stored function.

      The Configuration Parameters window expands to provide an area where you configure the stored procedure. The wizard automatically generates the list of parameters by examining the stored procedure in the database.

   c. For each parameter in the stored procedure (on the left), select the table column (on the right) to pass to the stored procedure in that parameter.

7. To specify the data type mapping for each column in the table:

   a. Click **Advanced**.
   b. Expand **Table columns**. For each column in the table, the default data type mapping is displayed. For Oracle databases, if the table contains any complex data type such as an array, structure, nested structure or table, the type name and the child attribute details are also automatically discovered and displayed. The following figure displays the type name and child attribute details of an Oracle table containing complex data types.
   c. Review the mapping and make changes if required.

   Note: If the primary key in a table is of the date or timestamp type, then the object_key in the event_table must be in the 'yyyy-mm-dd hh-mm-ss' format.

8. When all the fields in the window are completed, click **OK** to save the configuration of the business object. The table, view, synonym, and nickname business objects you defined are now listed in the Object Discovery and Selection window.

9. To change the configuration of an object from the **Selected objects** list, select the object name and then click the ✎ (Edit) icon.

10. When you have selected and configured all business objects that you need, click **Next** to set global properties and configure wrapper business objects.

What to do next

Continue working in the Object Discovery and Selection window to select and configure other types of business objects.

Selecting and configuring stored procedures and stored functions

To select and configure business objects that correspond to stored procedures and stored functions in the database, you filter the database objects, and specify the configuration properties for the database object.

Before you begin

To select and configure business objects for stored procedures or stored functions, you need to understand the structure of the data in the database and know what objects the module needs to access. In particular, you need to know the parameters passed to the stored procedures or stored functions that your module needs to access.

About this task

This task is performed through the J2C Bean wizard. Start in the Object Discovery and Selection window and then work in a Configuration Parameters window that is specific to the business object you are configuring.

Procedure

1. In the **Objects discovered by query** list of the Object Discovery and Selection window, expand the node for the schema that contains the stored procedure or stored function you want to work with.

2. Filter the stored procedures by specifying a valid name or pattern for at least one of the filter fields in the Filter Properties window.

   a. Click **Stored Procedures** and then click the ![icon] (Create or edit filter) button, located at the top of the **Objects discovered by query** pane.

   b. In the Filter Properties window, type a name or pattern in the **Object name or pattern** field. Use the question mark or underscore (? or _) to match a single character and the asterisk or percentage (* or %) to match multiple characters. The name is case-sensitive.

   c. In the **Catalog name or pattern** field, type the name or a pattern. Use the question mark or underscore (? or _) to match a single character and the asterisk or percentage (* or %) to match multiple characters.

   d. Click **OK**. The Stored Procedures node displays all the stored procedures that match the given filter condition.

      WebSphere Adapter for Oracle E-Business Suite distinguishes the original SP/SF from the overloaded ones by a number tag that corresponds to an overload sequence in the Oracle database. These corresponding parameters for the selected SP/SF will be added based on the overload sequence.

3. Select one or more objects from the **Stored Procedures** list, and click the **>** (Add) button to add the object to the **Objects to be imported** list.

   Stored procedures that are defined in the PL/SQL packages are displayed in the format *SPName*(*PackageName*). For example, if the EMP_MGMT package

contains the CREATE_DEPT stored procedure, the stored procedure is displayed in the list as CREATE_DEPT(EMP_MGMT). The Configuration Parameters window lists the attributes of the stored procedure business object, which include the names and data types of the parameters of the stored procedure, and information about the result sets that are returned.

4. If the stored procedure returns any result set, make sure that the value for the **The maximum number of ResultSets returned from the stored procedure** field reflects the maximum number expected. The wizard creates the required number of result set business objects to hold the results.

5. Configure each parameter:

   a. The **Data type** field displays the data type of the parameter.

   b. In the **Sample Value** field, type a valid value.

   Note: If the stored procedure or function contains a VArray or an Object attribute with a Date data type, the sample value must be given in a specific format to successfully validate the stored procedure or function. The Date data type can also be mapped to the String data type. Therefore, the format for the sample value depends on the data type to which the Date is mapped. The following table provides the format to be followed for each Date data type mapping.

Table 13. Date data type formats

Data type mapping	SP/SF with Date data type	VArray/Object in SP/SF with Date data type
Date to String	Format: Day-Month-Year (Example: '01-JAN-0001')	Format: Year-Month-Day Hour:Minute:Second.[Millisecond] (Example: '0001-01-01 01:00:00.000000000').
Date to Date	Format:Year-Month-Day (Example: '0001-01-01')	Format:Year-Month-Day (Example '0001-01-01')

   • The Year parameter can be specified from a minimum of one digit to a maximum of four digits in all the formats.

   • The Month parameter can be specified from a minimum of one digit to a maximum of two digits in all the formats, except in Date to String mapping of the SP/SF with the Date data type where the Month parameter is specified in three alphabetic characters as specified in the example in the previous table.

   • The Day, Hour, Minute, and Second parameters can be specified from a minimum of one digit to a maximum of two digits in all the formats.

   • The Milliseconds, which is an optional parameter, can be specified from a minimum of one digit to a maximum of nine digits in all the formats.

6. To validate the syntax of the stored procedure using the sample values, click **Validate**. The result of the validation is displayed in the **Result** area.

   Note: Ensure that the number of result sets is correct after you validate the syntax of the stored procedure because the Oracle driver does not always return the expected result set information. If the number is not correct after validation, set the correct number, and then click **OK** to save and close the window. After you close the wizard, you might verify the setting in the MaxNumOfRetRS application-specific parameter for the business object.

   If the **Result** area displays the `Validation failed` message, use the error message that is displayed before the `Validation failed` message, to correct the definition. Ensure that the data type of the parameters and the sample values are correct.

The OracleMetadataDiscovery.log file in the .metadata folder of your workspace contains the additional information about the problem if the logging level includes a WARNING.

When you see the message `Validation was successful`, click **OK** to save the definition of the stored procedure business object.

Important: If the stored procedure or stored function returns a result set, do not click **OK** until the validation succeeds. The wizard uses the results returned during validation to create business objects to hold the result. If the stored procedure validation is not successful, the adapter does not return the result set at run time.

7. To change the configuration of an object from the **Selected objects** list, select the object name and then click the ✎ (Edit) icon.

Results

The business objects you configured for stored procedures and stored functions are listed in the Object Discovery and Selection window.

What to do next

In the Object Discovery and Selection window, continue to select and configure other types of business objects. When you are finished, click **Next** to set the global properties.

Selecting and configuring query business objects

Select and configure query business objects for user-defined SELECT statements for use in your module.

Before you begin

To configure query business objects, you must know the structure of the data in your database, including the tables and views. You need to know the name and data type of the columns that your module needs to access. You must also be able to write the SQL SELECT statements.

About this task

This task is performed through the J2C Bean wizard. Start in the Object Discovery and Selection window and then work in a Configuration Parameters window that is specific to the business object you are configuring.

Procedure

1. In the **Objects discovered by query** list of the Object Discovery and Selection window, expand the **Query Statements** node. This node contains an object template, named **Select Statement** *n*, for each query business object you requested in the window. For example, if you specified a count of two query business objects in that window, the **Objects discovered by query** list contains two object templates.

2. Select one or more of the object templates and click the **>** (Add) button to add the objects to the **Objects to be imported** list.

3. In **Name of the business object**, type a name for the business object. The name can contain spaces and national language characters.

4. In **Select statement**, type the SELECT statement you want to run. Indicate each parameter with a question mark (**?**). The following sample SELECT statements illustrate the flexibility of the query business object:

- `select * from customer where ccode=?`
- `select * from customer where id=? and age=?`
- `select * from customer where lname like ?`
- `select C.pkey, C.fname, A.city from customer C, address A WHERE (C.pkey = A.custid) AND (C.fname like ?)`

   As you type each **?**, the window expands to display an area where you define the WHERE clause for that parameter.

5. In **Where clause parameter** *n*, provide information about each parameter in the SELECT statement.

   a. In **Parameter type**, select the data type of the parameter. For Oracle databases, the adapter does not support the complex types such as array, table, structure, or nested structure as parameters in the query business objects.

   b. In **Sample value**, type a sample value for the parameter.

   For example, for a parameter corresponding to a column containing the family name of the customer, you might select `string` as the data type and provide a sample value of `Smith`.

6. Click the **Validate** button to validate the syntax of the select statement using the sample values. **Result** displays the result of the validation.

   If the **Result** area displays the `Validation failed` message, use the error message that is displayed before the `Validation failed` message, to correct the definition. Ensure that the data type of the parameters and the sample values are correct.

   The OracleMetadataDiscovery.log file in the .metadata folder of your workspace contains the additional information about the problem if the logging level includes a WARNING.

7. To specify the data type mapping for each column in the result set returned by the select statement:

   a. Click **Advanced**.

   b. Expand **Result set returned by the Select statement**. For each column in the result set, the default data type mapping is displayed. For Oracle databases, if the query result contains any complex data type, such as an array, structure, nested structure or table, the type name and the child attribute details are also automatically discovered and displayed.

   c. Review the mapping and make changes if required.

8. Click **OK** to save the definition of the query business object.

Results

The query business objects you defined are now listed in the Object Discovery and Selection window.

What to do next

In the Object Discovery and Selection window, continue to select and configure other types of business objects. When you are finished, click **Next** to set the global properties.

Setting global properties for operations

After you select database objects in the J2C Bean wizard, you need to specify properties that apply to all business objects.

Procedure

1. When the **Objects to be imported** list in the Object Discovery and Selection window contains all the business objects you want to use in your application, click **Next**.

2. In the Configure Composite Properties window, review the list of operations.

   This window lists all the operations that the adapter supports for the outbound services for all business objects that you selected on the previous window. Not all operations are supported by each business object. For example, query business objects support only the RetrieveAll operation. Stored procedure business objects support only the Execute operation.

3. To remove an operation that you do not need, select the operation name and click **Remove**. If you change your mind, click **Add** and restore a removed operation.

4. Specify how you want to retrieve the records.
   - If you want the RetrieveAll operation to return all records matching the query, select the **Return all records for RetrieveAll operation** check box or enter -1 in the **Maximum records for RetrieveAll operation** field.
   - If you want to specify the maximum number of records, the RetrieveAll operation must return, enter a value in the **Maximum records for RetrieveAll operation** field. The default value is 100. For more information about this property, see "Maximum records for RetrieveAll operation" on page 149.

     Note: The "Maximum records for RetrieveAll operation" on page 149 property applies only if you are using the RetrieveAll operation. This property field is disabled if you remove the RetrieveAll operation in step 3 or if you select the **Return all records for RetrieveAll operation** check box.

5. In **Business object namespace**, accept the default namespace or type the full name of another namespace.

   The namespace is prefixed to the business object name to keep the business object schemas logically separated.

6. Optionally, in **Folder**, type the relative path to the folder where the generated business objects are to be stored.

   Note: If you are creating multiple adapter artifacts within a module, ensure that you specify different business object folders for each adapter within the module. For example, if you are creating artifacts for Oracle, JDBC, SAP, and JDE within a module, you need to create different relative folders for each of these adapters. If you do not specify different relative folders, the existing artifacts are overwritten when you generate new artifacts.

7. Click **Next**.

Results

You have provided information that applies to all business objects in the module.

What to do next

Continue working in the wizard. The next step is to specify deployment information to use at run time and information for saving the service as a module.

Setting deployment properties and generating artifacts

After you select and configure business objects for your module, use the J2C Bean wizard to configure properties that the adapter uses to connect to a specific database. The wizard creates a project where all the artifacts and property values are saved.

About this task

This task is performed through the J2C Bean Creation and Deployment Configuration window of the J2C Bean wizard.

Procedure

1. In the **Project name** field, select or create a project into which the J2C Bean is generated.
   - To select an existing project, click **Browse**. If the project name you want appears in the **Project Selection** list, select its name.
   - Otherwise, click **New** to create a project. In the New Source Project Creation window, select the type of project you want to create.
     a. Java project: For information about creating a Java project, see Creating a Java project.
     b. EJB project: For information about creating a EJB project, see Creating an EJB project.
     c. Web project: For information about creating a web project, see Creating a web project.
2. In the **Package name** field, select or create a package into which the J2C Bean is generated.
   - To select an existing package, click **Browse**. If the package name you want appears in the **Project Selection** list, select its name.
   - Otherwise, create a package:
     a. Click **New**.
     b. In the Java Package window, type a name for the package. For example, myadapteroutboundpkg.
     c. Click **Finish**.
3. In the **Interface name** field, specify the interface name you want to use for your business objects. For example, MyAdapterOutboundInterface. The implementation name is automatically generated by suffixing "Impl" to the interface name, and the name is displayed in the **Implementation name** field. For example, MyAdapterOutboundInterfaceImpl.
4. Optional: In the **Generate Command Bean** section, select the operation for which you want to generate a command bean. If you create command bean, you need to specify the command bean name as well as the input and output names.
5. In the Connection properties area, specify how you want the adapter to connect to the database.
   - In the managed connection mode, you can specify connection properties in the Java Naming and Directory Interface (JNDI) name of the managed

connection factory defined on the application server. The managed connection factory is EIS adapter-specific and contains all of the required connection information. Managed connection implies that the resource adapter is installed directly on the server and, therefore, the JNDI name of the managed connection factory is visible to all EAR (enterprise application archive) files installed on the application server. To obtain the connection through JNDI, select the **Managed Connection (recommended)** check box. This type of connection is managed by the application server.

- In the non-managed connection mode, the application communicates directly to the EIS and manages all the connections. Therefore, all of the connection information in this setup is encapsulated within the application. For a J2C bean, this means that all of the connection information is specified in the generated J2C bean. To obtain the connection directly from the resource adapter, select the **Non-managed Connection** check box.

Note: If you check both choices, the Managed Connection is executed first. If the JNDI name cannot be found, the non-managed connection parameters is used.

6. If you select **Managed Connection (recommended)** check box, specify how you want the adapter to specify the connection properties.
   - To select an existing **JNDI Lookup Name**, click **Browse**.
   - To create a new one, click **New**.
     a. In the Server selection window, specify the server on which to deploy the adapter and click **Next**.
     b. In the New J2C Connection Factory window, specify the name in the **JNDI Name** field. For example, `com/eis/AdapterCF`.
     c. Specify the advanced properties by clicking **Advanced**. The RAR file is installed on IBM WebSphere Application Server and the connection factory is created using the JNDI name specified.
     d. Expand each of the following group sections to review the advanced properties:
        - **Additional connection configuration**
          1) If you want to turn on AUTOCOMMIT for the database, select **Set auto commit on database connection**. For more information about the property, see "Auto commit (AutoCommit)" on page 137.
          2) Set **Additional JDBC driver connection properties**. For more information about the property, see "Additional JDBC driver connection properties [name:value;name:value] (JDBCDriverConnectionProperties)" on page 137.
          3) Set **SQL query to verify the connection**, See more information about the property in "SQL query to verify the connection (PingQuery)" on page 142.
          4) In **Query timeout**, type the length of time, in seconds, that the adapter must wait for a response to a database query. See more information about the property in "Query timeout (seconds) (QueryTimeOut)" on page 141.
          5) Set **Return business object even when the stored procedure result set is empty**. See more information about the property in "Return business object even when the stored procedure result set is empty (ReturnDummyBOForSP)" on page 141.
        - **Logging and tracing**

- If you have multiple instances of the adapter, set Adapter ID to a value that is unique for this instance.
- If you want to disguise potentially sensitive user information in log and trace files, select the **Disguise user data as "XXX" in log and trace files** check box.
  – **Bidi properties**
    - To enable bidirectional support for the adapter at run time, select the **Bidi transformation** check box.
    - Set the properties, ordering schema, text direction, symmetric swapping, character shaping, and numeric shaping to control how bidirectional transformation is performed.
  – **Connection retry settings**
    - To specify the number of times the adapter can attempt to reconnect to the EIS in case of connection failure, set **Maximum retries on connection failure** to a value greater than or equal to zero. For more information, see "Maximum retries on connection failure (connectionRetryLimit)" on page 143.
    - To specify the time interval between retries if connection fails, set **Connection Retry interval (in milliseconds)** to a value in milliseconds. This property is enabled only when the connectionRetryLimit property has a value greater than zero. For more information, see "Connection Retry interval (in milliseconds) (connectionRetryInterval)" on page 144.

  e. Click **Finish**.

  Note: When you try to look up an existing JNDI in the JNDI Lookup Wizard screen, the wizard looks up only the JNDI that is created using the first RAR file. If you have more than one RAR file for the same adapter in the server, you can view only the JNDI generated using the first RAR file. The JNDI generated from the other RAR files are not looked up by the wizard.

7. If you select the **Non-managed Connection** check box, the Database system connection information area is expanded to show the connection information. Review the connection information and change the values if required.
   - To use a predefined XA data source on the server (for XA connection):
     a. Select the **Join the global transaction** check box.
     b. From the **Database connection information** list, select **Specify predefined XA DataSource**.
     c. In the **Database system connection information** area, enter the value in the **XA DataSource JNDI name** field. This value must be set to a JNDI data source that supports XA transactions that is created on the WebSphere Application Server. For more information about this property, see "XA DataSource JNDI name (XADataSourceJNDIName)" on page 148.
   - To specify the connection information to be saved in the adapter properties (for XA connection):
     a. Select the **Join the global transaction** check box.
     b. From the **Database connection information** list, select **Specify XA database connection information**.
     c. In the **Database system connection information** area, enter values in the **XA DataSource name** and **Database URL** fields. For more information

about these properties, see "Database URL (DatabaseURL)" on page 137 and "XA DataSource name (XADataSourceName)" on page 143.

- To use a predefined connection pool data source (for local connection):

  a. Clear the **Join the global transaction** check box.

  b. From the **Database connection information** list, select **Specify predefined connection pool DataSource**.

  c. In the **Database system connection information** area, enter the name of the existing JNDI data source in the **Connection pool DataSource JNDI name** field. For more information about this property, see "Connection pool DataSource JNDI name (PoolDataSourceJNDIName)" on page 146.

- To specify connection information to be saved in the adapter properties (for local connection):

  a. Clear the **Join the global transaction** check box.

  b. From the **Database connection information** list, select **Specify local database connection information**.

  c. In the **Database system connection information** area, enter values in the **Database URL** and **JDBC driver class name** fields. For more information about these properties, see "Database URL (DatabaseURL)" on page 137 and "JDBC driver class (JDBCDriverClass)" on page 139.

8. Specify the advanced properties by clicking **Advanced**.

9. When you are finished setting properties, select the **Click to launch the J2C deployment** checkbox to immediately start the wizard.

10. When you are finished setting properties, click **Finish**.

Results

The new project is added to the Enterprise Explorer perspective. The module is created in the project and artifacts are generated.

The generated artifacts allow you to build an enterprise application that accesses the EIS. You can use the J2C Bean and Java data bindings directly, in the non-managed mode or generate JSP or EJB that uses the J2C Bean. Rational Application Developer provides tools to automate this generation. You can access these tools from the **New>Others** menu, for their description refer to the Rational Application Developer documentation.

Completing the configuration

In some situations, manual configuration steps are needed to complete the configuration of your business objects.

About this task

Perform this task when you need to customize the artifacts generated by the wizard. You might do this in the following situations:

- To set the CopyAttribute parameter for a column so that its value is set to the same value as another column.
- To add or remove attributes from a business object. For example, you can simplify your business object design by removing the simple attribute corresponding to any database column that you do not need to reference.
- To configure additional parents for a table business object that has multiple parents. The wizard configures only one parent for a table business object.

This topic provides detailed instructions for setting the CopyAttribute parameter to a table business object. Other changes to business object structures can be accomplished using similar techniques.

The CopyAttribute parameter is contained in the properties of the attribute for the column that you want to populate with values and application-specific information from another column. For example, if you want the `contact` column of a new row in the table to contain the same value as the `e-mail` column, set the CopyAttribute parameter of the `contact` attribute to `e-mail`.

Procedure

1. In the Enterprise Explorer of Rational Application Developer for WebSphere Software, expand the Java project, and then locate the table business object. The business object name is the name of the database schema plus the name of the database table. An optional namespace might be included at the beginning of the name.
2. Right-click the business object name, select **Open**. In the editor, go to the Types area and double-click the name of the business object. The editor displays the business object, which has a field for each column.
3. In the editor, select the column you want to set to match another column.
4. In the Properties view, select **Extensions**. If the Properties view is not visible, right-click the column name and click **Show Properties**.
5. Right-click **OracleBusinessObjectTypeMetadata** and then select **New > oracleasi:CopyAttribute**.
6. Select the **CopyAttribute** property.
7. In the Extension Details area, set the text value to the name of the column that contains the information to copy. The column can be in the current business object or its parent business object. To copy from a column in the current business object, set the value to the column name, for example, `phoneid`. To copy from a column in the parent business object, prefix the column name with two periods (..), for example, `..phone`.
8. In the Enterprise Explorer in Rational Application Developer for WebSphere Software, expand the Java project module name, expand `src`, and then locate the table business object Java file. The business object Java file name is the name of the database schema plus the name of the database table Java file name. For example, J2caCustomer.java.
9. Right-click the business object Java file name and select **Open**.
10. Find the column attribute code line. For example, for `phoneid`, the corresponding line is: `annotation = new LinkedHashMap(); annotation.put("ColumnName", "PHONEID");`
11. Below the code mentioned in the previous step, add the code for new attributeType. For example, if you want to copy from a column in the parent business object, prefix the column name with two periods(..). For example, `annotation.put("CopyAttribute", "..phone");`
12. Save the file.
13. Regenerate the EJB or WEB, and EAR project.

Results

The business object is configured to use the CopyAttribute property to set the business object attribute and properties for a database column based on information in another column.

What to do next

You can now test and deploy the module.

Generating the EJB or JSP project

After you create the J2C bean, use the Web Page, Web Service, or EJB from J2C Java wizard to create the JSP, EJB or web service.

About this task

You can create a JSP, an EJB, or a web service to deploy your J2C bean.

Procedure

1. To start the J2C Bean wizard, go to the Enterprise Explorer of Rational Application Developer for WebSphere Software and click **New>Other>J2C> Web Page, Web Service, or EJB from J2C Java Bean**.
2. Click **Next**.
3. In the J2C Java bean selection window, click **Browse**.
4. In the Find J2C bean window, type the first letter of the implementation name generated earlier or type the full name and press the Enter key.
5. Select the implementation name from the list and click **OK**.
6. Click **Next**.
7. In the Deployment Information window, select the Java EE Resource Type as **EJB**, **Simple JSP** or **Web Service** and click **Next**.

   Note: In the Deployment Information window, the **Configure Resource Adapter Deployment** check box is available for selection only if you have selected the **Non-managed Connection** check box when specifying the deployment settings.

8. If you select **EJB**, the Enterprise bean creation wizard is displayed. This wizard creates the Java project as an EJB. For information about creating an EJB, see Deploying to an EJB.

   Note: You can deploy an EJB only when the J2C bean is in the EJB project. Otherwise, this option is not available for selection.

9. If you select **Simple JSP**, the Simple JSP Creation wizard is displayed. For details about creating a new web project, see Creating a web project. For information about creating a Simple JSP, see Deploying to a Simple JSP.
10. If you select **Web Service**, the Web Service Creation wizard is displayed. For information about creating a web service, Deploying a J2C application as a web service.
11. Export the project as an EAR file for deployment.

Configuring the module for inbound processing

To configure a module to use the adapter for inbound processing, use the J2C Bean wizard in Rational Application Developer for WebSphere Software to find and select business objects and services from the Oracle database, and to generate business object definitions and related artifacts.

Discovering database objects

After configuring the connection properties, run a query to search for database objects. Browse through the discovered objects navigation tree to understand the structure of objects in the Oracle database and use filters to display only the database objects you want to see.

Before you begin

You must know the data requirement of the program that needs to access the database. For example, you need the following information about your database:
- Which schemas your module needs to access
- What type of database objects you need to access in those schemas

About this task

This task starts in the Object Discovery and Selection window of the J2C Bean wizard.

Procedure

1. In the Object Discovery and Selection window, click **Edit Query**. The Query Filter Parameters window is displayed.

   Note: The **Create a query business object to build user-defined select statements** is available only for outbound processing.

   Use the Query Filter Parameters window to perform the following tasks:
   - Reduce the search time by searching a subset of database schemas
   - Omit one or more types of database objects from the search
   - Make the wizard prompt you for application-specific information that cannot be automatically determined based on information in the database
   - Map the Oracle data types Date and Timestamp to date and dateTime

2. To limit the number of database schemas that are retrieved, type the name of the schema or a name pattern in **Schema name or pattern**. Use the question mark or underscore (? or _) character to match a single character and the asterisk or percentage sign (* or %) to match multiple characters. Only schemas that start with that string or match that pattern are displayed when you run the query. If you do not specify a schema name pattern, all schemas in the database are displayed. Using a filter can speed up the discovery process if your database contains many schemas.

3. To omit one or more types of objects from the search, select the types of objects that you want to omit (tables, views, and synonyms or nicknames) in **Supported database object types**, and then click **Remove**. If you change your mind, click **Add** to add the object type back. If your database contains object types that you do not need to access, omitting them can speed up the discovery process.

4. Select the **Prompt for additional configuration settings when adding business object** check box. Then, when you add a database object to the list of business objects to create, the wizard automatically prompts you for all user-configurable application-specific information for the object. For example, if you select this option, the wizard guides you through the process of building a simple parent-child hierarchy of business objects.

5. Click **OK** to save your changes to the query.

6. In the Object Discovery and Selection window, click Execute Query to use the query to discover database objects.

   The **Objects discovered by query** pane lists the objects that were discovered. The tables, views, and synonyms/nicknames are sorted by schema name.

7. In the **Objects discovered by query** list, click **+** (the plus sign) to expand a schema node and the **Tables**, **Views**, and **Synonyms - Nicknames** nodes to see the database objects discovered by the wizard.

Results

The wizard has discovered the database objects you can access using the adapter.

What to do next

Continue working in the J2C Bean wizard. The next step is to select the objects you want to use in your module, configure each business object, and create hierarchies of business objects.

Selecting and configuring business objects

Using the list of database objects discovered by the J2C Bean wizard, and the query object templates you specified, continue using the wizard to select the database objects that you need to access in your module. Then provide the configuration information for your new business objects.

About this task

The Object Discovery and Selection window allows you to select and configure objects in any order, with the single exception that you must select and configure a parent table before you can select and configure its child tables. Apart from this restriction, you have the flexibility to add objects individually or to add them several at a time. You can mix objects from the various nodes of the **Objects discovered by query** list. For example, you can select several tables and view objects and a stored procedure object, and add them at the same time.

The high-level flow of selecting and configuring business objects is as follows:

1. Select one or more objects in the **Objects discovered by query** list of the Object Discovery and Selection window.
2. Click the **>** (Add) button.
3. The wizard opens the Configuration Parameters window.
   - If you select a single object, a single Configuration Parameters window is displayed.

     Complete that window, specifying the user-configurable attributes and other information that the wizard cannot discover by examining the database, and click **OK** to save the configuration.
   - If you select multiple objects, the Configuration Parameters window displays with one page for each object selected.

     Click the name of each object in turn. The window displays the same information you would see had you selected this object individually.

     Important: Do not click **OK** until you complete the configuration pages for all the objects. The wizard does not close the notebook until you provide all the required fields, but you can close the window before providing optional

fields. If you do not configure the optional fields in the wizard, you must use the business object editor to configure them after exiting the wizard.

4. The wizard adds the configured object to the **Objects to be imported** list.

As long as you do not exit from the wizard, you can work iteratively to select and configure the business objects you need in your module. However, before you start the wizard to add objects to an existing module, understand the requirements of the program that uses the business objects carefully. The wizard overwrites the existing business objects in the same path.

Selecting and configuring tables, views, and synonyms or nicknames for inbound processing

Select and configure business objects for tables, views, and synonyms or nicknames for use in your module. For inbound processing, the tables, views, and synonyms are the business objects that are delivered in events.

Before you begin

To perform this task, you need to understand the structure of the data in the database and know what database objects the module needs to access. Specifically, you need to know the following information:

- The structure of the tables, views, and synonyms or nicknames, including the columns you need, and column attributes such as data type
- The relationships between the tables, including the cardinality and ownership of parent-child relationships

About this task

This task is performed through the J2C Bean wizard. Start in the Object Discovery and Selection window and then work in a Configuration Parameters window that is specific to the business object you are configuring.

Procedure

1. In the **Objects discovered by query** list of the Object Discovery and Selection window, select one or more tables, views, or synonyms. Click the **>** (Add) button to add the object or objects to the **Objects to be imported** list. Alternatively, you can also filter the tables, views, or synonyms by specifying a valid name or pattern for at least one of the filter fields in the **Filter Properties** window.

   a. Click the object you want to filter, and then click the 𝌆 (Create or edit filter.) button, located at the top of the **Objects discovered by query** pane.
   b. In the Filter Properties window, type a name or pattern in the **Object name or pattern** field. Use the question mark or underscore (? or _) to match a single character and the asterisk or percentage (* or %) to match multiple characters. The name is not case sensitive.
   c. Click **OK**. The object that matches the given filter condition is displayed.
   d. Select one or more objects from the discovered list, and click the **>** (Add) button to add the object to the **Objects to be imported** list.

2. If the table has a column that is used to indicate logical deletes:

   a. Select the column name in the **Name of the column used to perform logical deletes** field.

    b. In the **Value used to indicate a deleted object** field, type the value that indicates that a row is logically deleted. You can get this value from your database administrator.

3. If the **Select primary key for table** *table_name* area is displayed, click **Add**, select the column to be used as the primary key for the table business object, and then click **OK**. If the table has a composite key, you can select multiple columns. The **Select primary key for table** *table_name* area is displayed only when the database table does not have a column designated as the primary key. Each table business object must have a primary key, even if the associated database table does not have a key. If the primary key is defined in the database, this section of the window is not displayed.

4. Optional: Define a parent-child relationship between business objects.

   To build a parent-child hierarchy, configure the parent table first, and return to the Object Discovery and Selection window to select and configure the child tables.

   Configure the parent-child relationship using the area of the Configuration Parameters window. These fields are not displayed for the first table you configure.

    a. In the **Choose parent table** field, select the name of the parent table you are configuring. If you do not see the parent table in the list, the parent table has not yet been configured. Go back and configure the parent object before configuring the child objects.

    b. Specify the cardinality of the relationship:

   - If the table has a single-cardinality relationship with the parent table, select the **Single cardinality** check box. In a single cardinality relationship, a parent can have only one child business object of this type. A single-cardinality relationship can be used with ownership to represent a true child or without ownership to represent lookup tables or other peer objects in a database.

   - If the table has a multiple-cardinality relationship, do not select the **Single cardinality** check box. In a multiple-cardinality relationship, a parent can have an array of child business objects of this type.

    c. Build the foreign key relationship between the parent and child by specifying for each child column whether it is a foreign key in the parent table.

   - If the child column is not a foreign key, select NONE.

   - If a child column is a foreign key, select the column in the parent table that corresponds to the child column.

     Note: The wizard can configure only a single parent table. If the child table has multiple parent tables, you must use the business object editor to configure the remaining parent tables after exiting the wizard.

    d. If the parent object owns the child object, then the child objects in the database are deleted when the parent is deleted. To indicate that this child is owned by its parent, select the **Parent object owns child object (cascade delete)** check box. Otherwise, clear this option to prevent child objects, such as lookup tables, from being deleted when their parent is deleted.

    e. If you do not want the child objects to be deleted as part of an Update operation, select the **Preserves** *child_table_name* **when the parent is updated** check box.

   When a parent table is updated, the adapter compares the child business objects present in the input with the child business objects returned from

the database. By default, the adapter deletes any child objects returned from the database that are not present in the input business object.

  f. By default, you can perform operations on parent business objects without specifying the child business objects. If you want to ensure that a parent business object specifies its child business objects when the parent is submitted for a change, select the *Child_table_name* **required for operations on parent** check box.

5. An operation can be performed using either a standard SQL statement generated by the adapter or using stored procedures or stored functions from the database. If you want to use stored procedures or stored functions:

  a. Click **Add**.

  b. In the Add window, select the type of the stored procedure you want to run. For each operation, you can select a stored procedure that performs the operation, as well as stored procedures that run before or after the operation. For example, for the Create operation, you can specify any of these stored procedures: CreateSP, BeforeCreateSP, and AfterCreateSP.

     Note: If you configure the table with RetrieveAllSP, ensure that at least one parameter of the stored procedure is a Cursor and the ResultSet ASI for the stored procedure is set to true to avoid the "No resultset found associated with the stored procedure" exception being generated at run time.

  c. Click **OK**. The Configuration Parameters window now shows the stored procedure types you selected and expands to display an area where you configure each one. It might be necessary to scroll down to see the new areas.

     Note: In a hierarchical business object, if you want the stored procedure to be performed for each business object in the hierarchy, you must separately associate a stored procedure with the top-level business object and each child business object or an array of business objects. If you associate a stored procedure with the top-level business object but do not associate it with each child business object, then the top-level business object is processed with the stored procedure, but the child business objects are processed using the standard SQL query.

6. For each stored procedure type that you selected, specify the name of the stored procedure in the database and then configure the business object.

  a. In the **Schema name** field, select the name of the schema that contains the stored procedure.

     1) Click **Select**.

     2) In the Select Value window, select the name of the schema you want to work with.

     3) Click **OK**.

  b. Specify the name of the stored procedure or stored function.

     1) In the **Stored procedure name or pattern** field, either type the name of the stored procedure or stored function, or type a name pattern. Use the question mark or underscore (? or _) to match a single character, and the asterisk or percentage sign (* or %) to match multiple characters.

     2) In the **Stored procedure name** field, select the name of the procedure you want. If the stored procedure list contains many items, the **Select**

button is displayed next to the **Stored procedure name** field. Click **Select** to open the Select window and select the name of the stored procedure or stored function.

The Configuration Parameters window expands to provide an area where you configure the stored procedure. The wizard automatically generates the list of parameters by examining the stored procedure in the database.

   c. For each parameter in the stored procedure (on the left), select the table column (on the right) to pass to the stored procedure in that parameter.

7. To specify the data type mapping for each column in the table:

   a. Click **Advanced**.

   b. Expand **Table columns**. For each column in the table, the default data type mapping is displayed. For Oracle databases, if the table contains any complex data type, such as an array, structure, nested structure or table, the type name and the sub attribute details are also automatically discovered and displayed. The following figure displays the type name and sub attribute details of an Oracle table containing complex data types.

   c. Review the mapping and change them if required.

Note: If the primary key in a table is of the date or timestamp type, then the object_key in the event_table must be in the 'yyyy-mm-dd hh-mm-ss' format.

8. When all the fields in the window are completed, click **OK** to save the configuration of the business object. The table, view, synonym, and nickname business objects you defined are now listed in the Object Discovery and Selection window.

9. To change the configuration of an object from the **Selected objects** list, select the object name and then click the ✎ (Edit) icon.

10. When you have selected and configured all business objects that you need, click **Next** to set global properties.

What to do next

Continue working in the Object Discovery and Selection window to select and configure other types of business objects.

Setting global properties for operations

After you have selected database objects in the J2C Bean wizard, you need to specify properties that apply to all business objects.

Procedure

1. When the **Objects to be imported** list in the Object Discovery and Selection window contains all the business objects you want to use in your application, click **Next**.

2. In the Configure Composite Properties window, review the list of operations. This list contains the operations that the adapter supports for the inbound services. To add to the list of operation includes operations for all business objects you selected on the previous window.

   The specified operations are set for all business objects that are generated.

3. To remove an operation that you do not need, select the operation name and click **Remove**. If you change your mind, click **Add** and restore a removed operation.

4. In **Business object namespace**, accept the default namespace or type the full name of another namespace.

The namespace is prefixed to the business object name to keep the business object schemas logically separated. For more information about this property, see "Business object namespace (BusinessObjectNameSpace)" on page 160.

5. Optionally, in **Folder**, type the relative path to the folder where the generated business objects are to be stored.

    Note: If you are creating multiple adapter artifacts within a module, ensure that you specify different business object folders for each adapter within the module. For example, if you are creating artifacts for Oracle, JDBC, SAP, and JDE within a module, you need to create different relative folders for each of these adapters. If you do not specify different relative folders, the existing artifacts are overwritten when you generate new artifacts.

6. When you are finished, click **Next**.

Results

You have provided information that applies to all business objects in the module.

What to do next

Continue working in the wizard. The next step is to specify deployment information to use at run time and information for saving the service as a module.

Setting deployment properties and generating artifacts

After you select and configure business objects for your module, use the J2C Bean wizard to configure properties that the adapter uses to connect to a specific database. The wizard creates a project where all the artifacts and property values are saved.

About this task

This task is performed through the J2C Bean Creation and Deployment Configuration window of the J2C Bean wizard.

Procedure

1. In the **EJB Project name** field, select or create a new EJB project.
    - To select an existing project, click **Browse**. If the required project name appears in the **Project Selection** list, select its name.
    - Otherwise, create a new project:
        a. Click **New**.
        b. In the EJB Project window, type a project name. For example, MyAdapterInboundEJB.
        c. In the EAR Membership area, click **New** to create a new ear project.
        d. In EAR Application Project window, type an EAR project name. For example, MyAdapterInboundEJBEAR
        e. Click **Finish** to return to the EJB Project window.
        f. Click **Finish**.
2. In the **Package name** field, select or create a new package into which the J2C Bean is generated.
    - To select an existing package, click **Browse**. If the required package name appears in the **Project Selection** list, select its name.
    - Otherwise, create a new package:

a. Click **New**.

b. In the New Java Package window, type a name for the package. For example, myadapterinboundejbpkg.

c. Click **Finish**.

3. In the **Stateless Session EJB's local business interface name** field, specify the interface name you want to use for your business object. For example, MyAdapterInboundInterface. The interface name is suffixed with "MDB" and it is displayed automatically in the **Message Driven EJB name** field. For example, MyAdapterInboundInterfaceMDB. Similarly, the interface name is suffixed with "SB" and it is displayed automatically in the **Stateless Session EJB name** field. For example, MyAdapterInboundInterfaceSB.

4. In the Inbound Connection configuration section, specify the following:

   - Type the JNDI name for an existing activation specification in WebSphere Application Server or click **New** to create a JNDI name. For example, `com/eis/AdapterAS`.

     Note: When you try to look up an existing JNDI in the JNDI Lookup Wizard screen, the wizard looks up only the JNDI that is created using the first RAR file. If you have more than one RAR file for the same adapter in the server, you can view only the JNDI generated using the first RAR file. The JNDI generated from the other RAR files are not looked up by the wizard.

   - To create a new activation specification, click **New**.

     a. In the Server selection window, select the server on which the resource adapter will be deployed and click **Next**.

     b. In the New J2C Activation Specification window, type a JNDI name for the activation specification.

     c. Specify the advanced properties by clicking **Advanced**. The RAR file is installed on IBM WebSphere Application Server and the activation specification is created using the JNDI name specified.

     d. Expand each of the following group sections to review the advanced properties:

        – **Event polling configuration**

          - In **Interval between polling periods**, type the number of milliseconds that the adapter waits between polling periods. For more information, see "Interval between polling periods (PollPeriod)" on page 170.

          - In **Maximum events in polling period**, type the number of events to deliver in each polling period. For more information, see "Maximum events in polling period (PollQuantity)" on page 171.

          - In **Time between retries in case of system connection failure (in milliseconds)**, type the number of milliseconds to wait before trying to connect after a connection failure during polling. For more information, see "Time between retries in case of system connection failure (RetryInterval)" on page 171.

          - In **Maximum number of retries in case of system connection failure**, type the number of times to retry the connection before reporting a polling error. For more information, see "Maximum number of retries in case of system connection failure (RetryLimit)" on page 172.

          - If you want the adapter to stop if polling errors occur, select **Stop the adapter when an error is encountered while polling**. If you do not select this option, the adapter logs an exception but continues to

run. For more information, see "Stop the adapter when an error is encountered while polling (StopPollingOnError)" on page 173.

- You can select **Retry EIS connection on startup** if you want the adapter to retry an inbound connection that was not made to the Oracle database when starting. Only communication failures to the Oracle database are considered. For more information, see "Retry EIS connection on startup (RetryConnectionOnStartup)" on page 172.

– **Event delivery configuration**
  - In **Type of delivery**, select the delivery method. The methods are described in "Delivery type (DeliveryType)" on page 164.
  - If you want to ensure that events are delivered only once and to only one export, select **Ensure assured-once event delivery**. This option might reduce performance but does not result in duplicate or missing event delivery. For more information, see "Ensure once-only event delivery (AssuredOnceDelivery)" on page 164.
  - In the **Event types to process** field, type a comma-separated list of the business objects for which you want events delivered. Leave this field blank to receive events for all business object types. For example, if you want to receive events only when the Customer and Order tables, but not other tables, are changed in the enterprise information system, set this field to Customer,Order. For more information, see "Event types to process (EventTypeFilter)" on page 166.
  - In the **Adapter Instance for event filtering** field, specify the connector ID where the events will be delivered. For more information, see "Adapter Instance for event filtering (AdapterInstanceEventFilter)" on page 158.
  - In the **Retry limit for failed events** field, specify the number of times to try to deliver an event after a delivery failure. For more information, see "Retry limit for failed events (FailedEventRetryLimit)" on page 167.
  - Under **Number of connections for event delivery**, specify the minimum and maximum number of connections to use to deliver events. For more information, see "Minimum connections (MinimumConnections)" on page 169 and "Maximum connections (MaximumConnections)" on page 169.

– **Additional JDBC driver connection properties**
  - Set **Additional JDBC driver connection properties**. See more information about the property in "Additional JDBC driver connection properties [name:value;name:value] (JDBCDriverConnectionProperties)" on page 159.
  - Set **SQL query to verify the connection**. See more information about the property in "SQL query to verify the connection (PingQuery)" on page 170.
  - In **Query timeout**, type the length of time, in seconds, that the adapter must wait for a response to a database query. See more information about the property in "Query timeout (seconds) (QueryTimeOut)" on page 171.
  - Set **Return business object even when the stored procedure result set is empty**. See more information about the property in "Return business object even when the stored procedure result set is empty (ReturnDummyBOForSP)" on page 173.

– **Event configuration**
1) In **Event order by**, indicate the order in which events are retrieved and processed. This is a comma-separated list of column names from the event table, plus the keywords that control the sort order for each column. Use `asc` for ascending order and `desc` for descending order. For more information, see "Event order by (EventOrderBy)" on page 165.
2) In **Event table name**, accept the default name of the table that contains the event store, or type a different table name. For more information, see "Event table name (EventTableName)" on page 166.
3) In **Stored procedure to run before poll**, the name of a stored procedure or stored function to run before the actual poll query is called. For more information, see "Stored procedure to run before polling (SPBeforePoll)" on page 174.
4) In **Stored procedure to run after poll**, specify the name of a stored procedure or stored function to run after each polling cycle. For more information, see "Stored procedure to run after polling (SPAfterPoll)" on page 174.
5) In **Event query type for processing events**, select the type of event processing you want to use:
   - To use the standard event processing provided by the adapter, select **Standard**.
   - To provide your own queries to customize event processing, select **User-Defined (Dynamic)**. If you select this option, complete the additional fields described in the following table.

Field	What to specify	For more information
User-defined delete query	The name of the query, stored procedure, or stored function that is run after each event is processed to delete records that can be deleted after the event is delivered.	"User-defined delete query (CustomDeleteQuery)" on page 160
User-defined event query	The name of the query, stored procedure, or stored function that performs the polling for events.	"User-defined event query (CustomEventQuery)" on page 161

Field	What to specify	For more information
User-defined update query for failed event delivery	The name of the query, stored procedure, or stored function that is run when an event is not delivered successfully.	"User-defined update query for failed event delivery (CustomUpdateQueryForFailedEvent)" on page 161
User-defined update query	The name of the query, stored procedure, or stored function that is run after each event is processed to prevent the event from being picked up for processing in a subsequent event cycle.	"User-defined update query (CustomUpdateQuery)" on page 162

    e. Click **Finish** to return to the J2C Bean Creation and Deployment Configuration window.

5. Type the existing Java Authentication and Authorization Services alias. The alias is used to retrieve the user name and password set on the configured J2C activation specification. The name is case sensitive and includes the node name.

6. In the Service Operations section, click **Edit Operations** to review the names of operations or add a description about the operations to be generated in the interface file.

7. Click **Finish**.

What to do next

Deploy the module and test

Completing the configuration

In some situations, manual configuration steps are needed to complete the configuration of your business objects.

About this task

Perform this task when you need to customize the artifacts generated by the wizard. You might do this in the following situations:

- To set the CopyAttribute parameter for a column so that its value is set to the same value as another column.

- To add or remove attributes from a business object. For example, you can simplify your business object design by removing the simple attribute corresponding to any database column that you do not need to reference.
- To configure additional parents for a table business object that has multiple parents. The wizard configures only one parent for a table business object.

This topic provides detailed instructions for setting the CopyAttribute parameter to a table business object. Other changes to business object structures can be accomplished using similar techniques.

The CopyAttribute parameter is contained in the properties of the attribute for the column that you want to populate with values and application-specific information from another column. For example, if you want the `contact` column of a new row in the table to contain the same value as the `e-mail` column, set the CopyAttribute parameter of the `contact` attribute to `e-mail`.

Procedure

1. In the Enterprise Explorer of Rational Application Developer for WebSphere Software, expand the Java project, and then locate the table business object. The business object name is the name of the database schema plus the name of the database table. An optional namespace might be included at the beginning of the name.
2. Right-click the business object name, select **Open**. In the editor, go to the Types area and double-click the name of the business object. The editor displays the business object, which has a field for each column.
3. In the editor, select the column you want to set to match another column.
4. In the Properties view, select **Extensions**. If the Properties view is not visible, right-click the column name and click **Show Properties**.
5. Right-click **OracleBusinessObjectTypeMetadata** and then select **New** > **oracleasi:CopyAttribute**.
6. Select the **CopyAttribute** property.
7. In the Extension Details area, set the text value to the name of the column that contains the information to copy. The column can be in the current business object or its parent business object. To copy from a column in the current business object, set the value to the column name, for example, `phoneid`. To copy from a column in the parent business object, prefix the column name with two periods (..), for example, `..phone`.
8. In the Enterprise Explorer in Rational Application Developer for WebSphere Software, expand the project module name, expand `ejbModule`, and then locate the table business object java file. The business object java file name is the name of the database schema plus the name of the database table java file name. For example, J2caCustomer.java.
9. Right-click the business object Java file name and select **Open**.
10. Find the column attribute code line. For example, for `phoneid`, the corresponding line is: `annotation = new LinkedHashMap(); annotation.put("ColumnName", "PHONEID");`
11. Below the code mentioned in the previous step, add the code for new attributeType. For example, if you want to copy from a column in the parent business object, prefix the column name with two periods(..). For example, `annotation.put("CopyAttribute", "..phone");`
12. Save the file.

Results

The business object is configured to use the CopyAttribute property to set the business object attribute and properties for a database column based on information in another column.

What to do next

You can now test and deploy the module.

Chapter 4. Deploying the module

Deploy a module to place the files that make up your module and adapter into an operational environment for testing. In Rational Application Developer for WebSphere Software, the integrated test environment features runtime support for WebSphere Application Server, depending on the test environment profiles that you selected during installation.

Deployment environments

There are test and production environments into which you can deploy modules and adapters.

In Rational Application Developer for WebSphere Software, you can deploy your modules to one or more servers in the test environment. This is typically the most common practice for running and testing projects. However, you can also export modules for server deployment on WebSphere Application Server as EAR files using the administrative console or command-line tools.

Deploying the module for testing

In Rational Application Developer for WebSphere Software, you can deploy a module that includes an embedded adapter to the test environment and work with server tools that enable you to perform such tasks as editing server configurations, starting, and stopping servers and testing the module code for errors. The testing is generally performed on the interface operations of your components, which enables you to determine whether the components are correctly implemented and the references are correctly wired.

Configuring the connector dependencies

Dependent JARs have to be added to the libraries directory or packaged in the EAR.

About this task

The JARs are set in the class path and these dependent libraries have to be made available for run time when the module is deployed. There are two ways to make the dependent libraries available, one for either stand-alone deployment or embedded deployment and the other for embedded deployment only.

Configuring the connector dependencies on the server

The adapter needs certain files installed on WebSphere Application Server server to be able to communicate with the Oracle database.

Before you begin

To obtain the required files and copy them to WebSphere Application Server, use the following procedure.

About this task

An adapter needs the Oracle JDBC driver-specific files to be able to communicate with it.

Procedure

1. Obtain the Oracle JDBC driver-specific files or native libraries for your database software and operating system from your database administrator or from the database software Website

   The following table lists the Oracle JDBC driver files needed for Oracle database software.

Table 14. Oracle JDBC driver files for Oracle database software

Database software	Driver	JDBC driver files	Native System Libraries
Oracle	Thin driver	ojdbc6.jar or ojdbc14.jar depending on the JRE version	None

   Note: The Java Runtime Environment (JRE) version required by the Oracle JDBC driver must not be higher than the JRE version in the runtime environment. For example, if the JRE version is 1.5 in the runtime environment, then the Oracle JDBC Driver should be "ojdbc14.jar" which supports both Java Development Kit (JDK) 1.4 and JDK 1.5; if the JRE version is 1.6 in the runtime environment, then the Oracle JDBC Driver should be "ojdbc6.jar" which supports JDK version 1.6.

2. Copy the files to the server.

   • In a testing environment in Rational Application Developer for WebSphere Software, copy the files to the ${WAS_INSTALL_ROOT}/runtimes/bi_v7/lib/ext directory.

   • In a production environment, copy the files to the ${WAS_INSTALL_ROOT}/lib/ext directory of WebSphere Application Server.

Configuring the connector dependencies when the adapter is bundled

You must copy the dependent JAR files to the EAR application before you can run your adapter applications. You must use this method only for embedded deployment.

About this task

To obtain the required files and copy them to the EAR application, use the following procedure:

Procedure

1. From the appropriate module, go to the workspace and copy the JAR files to the directory. For example if the name of the module is ModuleName, then go to the workspace and copy the JAR files to the `ModuleNameApp/EarContent` directory.

2. Modify the adapter RAR's manifest file, `manifest.mf`, with the list of JAR files required by the adapter. Add the JAR files in the following format: Class-Path: dependantjar1.jar, dependantjar2.jar.

3. Copy the native libraries to the run time bin directory and deploy the application.

Results

The vendor software libraries are now a part of your run time environment.

Preparing to test outbound operations

Before you can test your module's outbound processing with the Rational Application Developer for WebSphere Software test client, you might need to modify some of your business objects.

About this task

This step is performed in the Rational Application Developer for WebSphere Software test client. Open Rational Application Developer for WebSphere Software it from the Enterprise Explorer view by right-clicking the name of your project and then clicking **Test** > **Test Module**.

Procedure

- **Query business objects**

  If your query business object was created without a WHERE clause (for example, it was defined with a SELECT statement like `Select * from Customer`), unset the jdbcwhereclause attribute of the query business object before testing in the test client.

- **Table, view, and synonyms or nicknames business objects**

  Before testing the RetrieveAll operation, you need to unset any attribute whose value you are not setting as part of your test.

- **Query business objects**

  Before testing the RetrieveAll operation, you need to unset any attribute whose value you are not setting as part of your test.

Adding the module to the server

In Rational Application Developer for WebSphere Software, you can add modules to one or more servers in the test environment.

Before you begin

If the module you are testing uses an adapter to perform inbound processing, generate and wire a *target component* to which the adapter sends the events.

About this task

In order to test your module and its use of the adapter, you need to add the module to the server.

Procedure

1. *Conditional:* If there are no servers in the **Servers** view, add and define a new server by performing the following steps:
   a. Place your cursor in the **Servers** view, right-click, and select **New** > **Server**.
   b. From the Define a New Server window, select the server type.
   c. Configure servers settings.
   d. Click **Finish** to publish the server.
2. Add the module to the server.
   a. Switch to the servers view. In Rational Application Developer for WebSphere Software, select **Windows** > **Show View** > **Servers**.

a. Start the server. In the **Servers** tab in the lower-right pane of the Rational Application Developer for WebSphere Software screen, right-click the server, and then select **Start**.

3. When the server status is *Started*, right-click the server, and select **Add and Remove** .

4. In the Add and Remove Projects screen, select your project and click **Add**. The project moves from the **Available projects** list to the **Configured projects** list.

5. Click **Finish**. This deploys the module on the server.

   The Console tab in the lower-right pane displays a log while the module is being added to the server.

What to do next

Test the functionality of your module and the adapter. For information about testing a module using the test client, see the *Testing applications on a server* topic in the Rational Application Developer for WebSphere Software information center.

Deploying the module for production

Deploying a module created with the J2C Bean wizard to WebSphere Application Server in a production environment is a two-step process. First, you export the module in Rational Application Developer for WebSphere Software as an enterprise archive (EAR) file. Second, you deploy the EAR file using the WebSphere Application Server administrative console.

Configuring the connector dependencies on the server

The adapter needs certain files installed on WebSphere Application Server server to be able to communicate with the Oracle database.

Before you begin

To obtain the required files and copy them to WebSphere Application Server, use the following procedure.

About this task

An adapter needs the Oracle JDBC driver-specific files to be able to communicate with it.

Procedure

1. Obtain the Oracle JDBC driver-specific files or native libraries for your database software and operating system from your database administrator or from the database software Website

   The following table lists the Oracle JDBC driver files needed for Oracle database software.

Table 15. Oracle JDBC driver files for Oracle database software

Database software	Driver	JDBC driver files	Native System Libraries
Oracle	Thin driver	ojdbc6.jar or ojdbc14.jar depending on the JRE version	None

Note: The Java Runtime Environment (JRE) version required by the Oracle JDBC driver must not be higher than the JRE version in the runtime environment. For example, if the JRE version is 1.5 in the runtime environment, then the Oracle JDBC Driver should be "ojdbc14.jar" which supports both Java Development Kit (JDK) 1.4 and JDK 1.5; if the JRE version is 1.6 in the runtime environment, then the Oracle JDBC Driver should be "ojdbc6.jar" which supports JDK version 1.6.

2. Copy the files to the server.
   - In a testing environment in Rational Application Developer for WebSphere Software, copy the files to the ${WAS_INSTALL_ROOT}/runtimes/bi_v7/lib/ext directory.
   - In a production environment, copy the files to the ${WAS_INSTALL_ROOT}/lib/ext directory of WebSphere Application Server.

Installing the RAR file (for modules using stand-alone adapters only)

If you chose not to embed the adapter with your module, but instead choose to make the adapter available to all deployed applications in the server instance, you need to install the adapter in the form of a RAR file to the application server. A RAR file is a Java archive (JAR) file that is used to package a resource adapter for the Java 2 Connector (J2C) architecture.

Before you begin

You must set **Deploy connector project** to **On server for use by multiple adapters** in the J2C Bean Creation and Deployment Configuration window of the J2C Bean wizard.

About this task

Installing the adapter in the form of a RAR file results in the adapter being available to all Java EE application components running in the server run time.

Procedure

1. If the server is not running, right-click your server in the **Servers** view and select **Start**.
2. When the server status changes to **Started**, right-click the server and select **Administration** > **Run administrative console**.
3. Log on to the administrative console.
4. Click **Resources** > **Resource Adapters** > **Resource adapters**.
5. In the Resource adapters page, click **Install RAR**.
6. In the Install RAR file page, click **Browse** and navigate to the RAR file for your adapter.

   The RAR files are typically installed in the following path:
   `RAD_installation_directory/ResourceAdapters/adapter_name/adapter.rar`
7. Click **Next**.
8. Optional: In the Resource adapters page, change the name of the adapter and add a description.
9. Click **OK**.
10. Click **Save** in the **Messages** box at the top of the page.

What to do next

The next step is to export the module as an EAR file that you can deploy on the server.

Exporting the module as an EAR file

Using Rational Application Developer for WebSphere Software, export your module as an EAR file. By creating an EAR file, you capture all of the contents of your module in a format that can be easily deployed to WebSphere Application Server.

Before you begin

Before you can export a module as an EAR file, you must have created a module to communicate with your service. The module should be displayed in the Rational Application Developer for WebSphere Software Enterprise Explorer view.

About this task

To export the module as an EAR file, perform the following procedure.

Procedure

1. Right-click the module and select **Export**.
2. In the Select window, expand **Java EE**.
3. Select **EAR file** and click **Next**.
4. Optional: Select the correct EAR application. The EAR application is named after your module, but with "App" added to the end of the name.
5. Browse for the folder on the local file system where the EAR file will be placed.
6. Optional: To export the source files, select the **Export source files** check box. This option is provided in case you want to export the source files in addition to the EAR file. Source files include files associated with Java components, data maps, and so on.
7. Optional: To overwrite an existing file, click **Overwrite existing file**.
8. Click **Finish**.

Results

The contents of the module are exported as an EAR file.

What to do next

Install the module in the administrative console. This deploys the module to WebSphere Application Server.

Installing the EAR file

Installing the EAR file is the last step of the deployment process. When you install the EAR file on the server and run it, the adapter, which is embedded as part of the EAR file, runs as part of the installed application.

Before you begin

You must have exported your module as an EAR file before you can install it on WebSphere Application Server.

About this task

To install the EAR file, perform the following procedure. For more information about clustering adapter module applications, see the http://www.ibm.com/ software/webservers/appserv/was/library/.

Procedure

1. If the server is not running, right-click your server in the **Servers** view and select **Start**.
2. When the server status changes to **Started**, right-click the server and select **Administration** > **Run administrative console**.
3. Log on to the administrative console.
4. Click **Applications** > **New Application** > **New Enterprise Application**.
5. Click **Browse** to locate your EAR file and click **Next**. The EAR file name is the name of the module followed by "App."
6. Optional: If you are deploying to a clustered environment, complete the following steps.
   a. On the **Step 2: Map modules to servers** window, select the module and click **Next**.
   b. Select the name of the server cluster.
   c. Click **Apply**.
7. Click **Next**. In the Summary page, verify the settings and click **Finish**.
8. Optional: If you are using an authentication alias, complete the following steps:
   a. Expand **Security** and select **Business Integration Security**.
   b. Select the authentication alias that you want to configure. You must have administrator or operator rights to change the authentication alias configurations.
   c. Optional: If it is not already specified, type the **User name**.
   d. If it is not already specified, type the **Password**.
   e. If it is not already specified, type the password again in the **Confirm Password** field.
   f. Click **OK**.

Results

The project is now deployed and the Enterprise Applications window is displayed.

What to do next

If you want to set or reset any properties or you would like to cluster adapter project applications, make those changes using the administrative console before configuring troubleshooting tools.

Deploying the module in a clustered environment

In Rational Application Developer for WebSphere Software, you can deploy the IBM WebSphere Adapter for Oracle E-Business Suite in a clustered environment.

To deploy the module in a clustered environment, use any of the following approaches.

- **Embedded module:** The adapter is embedded in the application and specific to it. The adapter cannot be shared between multiple applications.
- **Node level module with embedded activation specification:** The adapter is deployed at the node level, with the activation specification created during module creation. The adapter can be shared across multiple applications.
- **Node level module with JNDI activation specification reference:** The adapter is deployed at the node level, and the application provides a JNDI reference to the activation specification. You must create the reference at the cluster scope from the administrative console, with the same JNDI name. The adapter can be shared across multiple applications.

Deploying module embedded in the application

The adapter is deployed embedded in the application and specific to it. The adapter cannot be shared between multiple applications.

Before you begin

The following steps are a necessary prerequisite to configure and deploy the module.
- Rational Application Developer for WebSphere Software version 7.5.0.0 or above.
- A clustered topology deployment environment on the WebSphere Application Server available from Rational Application Developer for WebSphere Software.
- Create a clustered topology deployment environment, as shown in the following **Gold Topology** configuration figure.
- Deploy the adapter and the adapter applications (EAR files) in the AppTarget (the target that hosts the SCA container).

About this task

To create an application with the embedded adapter, use the J2C Bean wizard.

Procedure

1. In the Service Configuration Properties window, from the **Deploy connector project** property list, select **With module for use by single application**.
2. Create the module as described in the Business process management samples for WebSphere Adapters.
3. In the **Dependencies** option for the module, after the module is created, ensure that the **Deploy with module** option is selected for the adapter.
4. If the server is not running, right-click your server in the **Servers** view and select **Start**.
5. When the server status changes to **Started**, right-click the server and select **Administration** > **Run administrative console**.
6. From the **Deployment Manager Admin Console**, click **Install applications** to deploy the application.
7. On the **Step 2: Map modules to servers** window, select the module and click **Next**. For the embedded adapter option, the adapter is deployed as part of the application.
8. In the Enterprise Applications view, select the new application *<adapter_name>***EmbeddedModuleApp**. The new application is displayed after the application is deployed at the deployment manger level.
9. Select the node and click **Installed applications** to view the deployed application on each individual node.

Results

The resource adapter is embedded and deployed as part of the application.

Deploying module at node level with embedded activation specification

The adapter is deployed at the node level, with the activation specification created during module creation. The adapter can be shared across multiple applications.

Before you begin

The following steps are a necessary prerequisite to configure and deploy the module.

- Rational Application Developer for WebSphere Software version 7.5.0.0 or above.
- A clustered topology deployment environment on the WebSphere Application Server available from Rational Application Developer for WebSphere Software.
- Create a clustered topology deployment environment, as shown in the following **Gold Topology** configuration figure.
- Deploy the adapter and the adapter applications (EAR files) in the AppTarget (the target that hosts the SCA container).

About this task

To create an application with the node level adapter and activation specification properties specified in the module itself, use the J2C Bean wizard.

Procedure

1. In the Service Configuration Properties window, from the **Deploy connector project** property list, select **On server for use by multiple applications**.
2. From the **Connection properties** list, select **Use properties below**.
3. Create the module as described in the Business process management samples for WebSphere Adapters.
4. In the **Dependencies** option for the module, ensure that the **Deploy with module** option is not selected for the adapter. Here, the adapter is not part of the module, therefore you must deploy the adapter before deploying the application.
5. If the server is not running, right-click your server in the **Servers** view and select **Start**.
6. When the server status changes to **Started**, right-click the server and select **Administration** > **Run administrative console**. Log on to the administrative console.
7. To deploy the adapter at individual nodes, click **Resources** > **Resource Adapters** > **Resource adapters**. In the clustered environment, you must install the adapter in each node separately.
8. In the Resource adapters page, click **Install RAR**.
9. In the Install RAR file page, click **Browse** and navigate to the RAR file for your adapter. Deploy the RAR on each node.

   The RAR files are typically installed in the following path:
   `RAD_installation_directory/ResourceAdapters/adapter_name/adapter.rar`
10. For deployment at node level, do not select any **Scope** because the scope is always **Node**. Click **Next**.

11. Optional: In the Resource adapters page, change the name of the adapter and add a description. Click **OK**.

12. Click **Save** in the **Messages** box at the top of the page.

13. For node level deployment, check if the adapter RAR is deployed at the node level.

14. To deploy the adapter at the cluster level, click **Resources** > **Resource Adapters** > **Resource adapters**.

15. In the Resource adapters window, set the **Scope** to **Cluster**, and then click **New**.

16. Select the RAR deployed at the node level.

17. Check if the adapter RAR is now deployed at the cluster level. Deploy the application after the adapter is deployed at the node level on the individual nodes, and then at the cluster level.

18. From the **Deployment Manager Admin Console**, click **Install applications** to deploy the application.

19. On the **Step 2: Map modules to servers** window, select the module and click **Next**. The adapter is not part of the deployed application.

20. In the **Admin Console**, click **Resources** > **Resource Adapters** > **IBM WebSphere Adapter for Oracle** > **J2C activation specifications** to view the activation specification from the adapter deployed at the cluster level.

Results

The resource adapter is deployed at the node level, with the activation specification.

Deploying module at node level with JNDI activation specification

The adapter is deployed at the node level, and the application provides a JNDI reference to the activation specification. You must create the activation specification with the same JNDI name at the cluster scope from the administrative console. The adapter can be shared across multiple applications

Before you begin

The following steps are a necessary prerequisite to configure and deploy the module.

- Rational Application Developer for WebSphere Software version 7.5.0.0 or above.
- A clustered topology deployment environment on the WebSphere Application Server available from Rational Application Developer for WebSphere Software.
- Create a clustered topology deployment environment, as shown in the following **Gold Topology** configuration figure.
- Deploy the adapter and the adapter applications (EAR files) in the AppTarget (the target that hosts the SCA container).

About this task

To create an application with the node level adapter and activation specification properties specified in the module itself, use the J2C Bean wizard.

Procedure

1. In the Service Configuration Properties window, from the **Deploy connector project** property list, select **On server for use by multiple applications**.

2. From the **Connection properties** list, select **Use JNDI lookup name configured on server**.

3. In the **JNDI lookup name** property field, specify the JNDI name. Use this same JNDI name when you create the activation specification from the Admin Console.

4. Create the module as described in the Business process management samples for WebSphere Adapters.

5. In the **Dependencies** option for the module, ensure that the **Deploy with module** option is not selected for the adapter.

6. If the server is not running, right-click your server in the **Servers** view and select **Start**.

7. When the server status changes to **Started**, right-click the server and select **Administration** > **Run administrative console**. Log on to the administrative console.

8. To install the adapter at the node level, click **Resources** > **Resource Adapters** > **Resource adapters**. In the clustered environment, you must install the adapter in each node separately.

9. In the Resource adapters page, click **Install RAR**.

10. In the Install RAR file page, click **Browse** and navigate to the RAR file for your adapter. Deploy the RAR on each node.

    The RAR files are typically installed in the following path: *IID_installation_directory*/ResourceAdapters/*adapter_name*/*adapter*.rar

11. For deployment at node level, do not select any **Scope** because the scope is always **Node**. Click **Next**.

12. Optional: In the Resource adapters page, change the name of the adapter and add a description. Click **OK**.

13. Click **Save** in the **Messages** box at the top of the page.

14. To install the RAR at the cluster level, click **Resources** > **Resource Adapters** > **Resource adapters**

15. In the Resource adapters page, set the **Scope** to **Cluster**, and then click **New**.

16. Select the RAR deployed at the node level, and then check if the adapter RAR is now deployed at the cluster level. Deploy the application after the adapter is deployed at the node level on the individual nodes, and then at the cluster level.

17. From the **Deployment Manager Admin Console**, click **Install applications** to deploy the application.

18. In the **Admin Console**, click **Resources** > **Resource Adapters** > **IBM WebSphere Adapter for Oracle** > **J2C activation specifications** > **New** to create the activation specification from the adapter deployed at the cluster level.

19. When installing the adapter, in the **Name** field, you must enter the same name as defined in the RAR.

20. In the **JNDI name** field, you must enter the same name as given during the module creation.

21. Click **Resources** > **Resource Adapters** > **IBM WebSphere Adapter for Oracle** > **J2C activation specifications** to check if the JNDI reference on the adapter is same as the one specified for the module.

22. Click **Resources** > **Resource Adapters** > **IBM WebSphere Adapter for Oracle** > **J2C activation specifications** > **Custom properties** to set values for the activation specification in the Admin Console.

23. From the **Deployment Manager Admin console**, click **Install applications** to deploy the application after you deploy the RAR and create the activation specification.

24. On the **Step 2: Map modules to servers** page, select the module and click **Next**. The adapter is not part of the deployed application.

Results

The resource adapter is deployed at the node level, with the JNDI activation specification reference.

Chapter 5. Configuring the application on WebSphere Application Server

When you are running the adapter in a stand-alone deployment, use the administrative console of the server to start, stop, monitor, and troubleshoot the adapter module. In an application that uses an embedded adapter, the adapter module starts or stops when the application is started or stopped.

Configuring logging and tracing

Configure logging and tracing to suit your requirements. Enable logging for the adapter to control the status of event processing. Change the adapter log and trace file names to separate them from other log and trace files.

Configuring logging properties

Use the administrative console to enable logging and to set the output properties for a log, including the location, level of detail, and output format of the log.

About this task

Before the adapters can log monitored events, you must specify the service component event points that you want to monitor, what level of detail you require for each event, and format of the output used to publish the events to the logs. Use the administrative console to perform the following tasks:

* Enable or disable a particular event log
* Specify the level of detail in a log
* Specify where log files are stored and how many log files are kept
* Specify the format for log output

  If you set the output for log analyzer format, you can open trace output using the Log Analyzer tool, which is an application included with your application server. This is useful if you are trying to correlate traces from two different server processes, because it allows you to use the merge capability of the Log Analyzer.

Note: For more information about monitoring on a application server, including service components and event points, see http://publib.boulder.ibm.com/infocenter/dmndhelp/v8r0mx/topic/com.ibm.wbpm.admin.doc/topics/welcome_wps_mon.html.

You can change the log configuration statically or dynamically. Static configuration takes effect when you start or restart the application server. Dynamic or run time configuration changes apply immediately.

When a log is created, the detail level for that log is set from the configuration data. If no configuration data is available for a particular log name, the level for that log is obtained from the parent of the log. If no configuration data exists for the parent log, the parent of that log is checked, and so on, up the tree, until a log with a non-null level value is found. When you change the level of a log, the change is propagated to the child logs, which recursively propagate the change to their child log, as necessary.

To enable logging and set the output properties for a log, use the following procedure.

Procedure

1. In the navigation pane of the administrative console, select **Servers** > **WebSphere application servers**.
2. Click the name of the server that you want to work with.
3. Under **Troubleshooting**, click **Logging and tracing**.
4. Click **Change log detail levels**.
5. Specify when you want the change to take effect:
   - For a static change to the configuration, click the **Configuration** tab.
   - For a dynamic change to the configuration, click the **Runtime** tab.
6. Click the names of the packages whose logging level you want to modify. The package names for WebSphere Adapters start with **com.ibm.j2ca.***:
   - For the adapter base component, select **com.ibm.j2ca.base.***.
   - For the adapter base component and all deployed adapters, select **com.ibm.j2ca.***.
   - For the core component that is common between WebSphere Adapter for JDBC and WebSphere Adapter for Oracle E-Business Suite, select the **com.ibm.j2ca.dbadapter.core.***
   - For the WebSphere Adapter for Oracle E-Business Suite only, select the **com.ibm.j2ca.oracleebs.*** package.
7. Select the logging level.

Logging Level	Description
Fatal	The task cannot continue or the component cannot function.
Severe	The task cannot continue, but the component can still function. This logging level also includes conditions that indicate an impending fatal error, that is, situations that strongly suggest that resources are on the verge of being depleted.
Warning	A potential error has occurred or a severe error is impending. This logging level also includes conditions that indicate a progressive failure, for example, the potential leaking of resources.
Audit	A significant event has occurred that affects the server state or resources.
Info	The task is running. This logging level includes general information outlining the overall progress of a task.
Config	The status of a configuration is reported or a configuration change has occurred.
Detail	The subtask is running. This logging level includes general information detailing the progress of a subtask.

8. Click **Apply**.
9. Click **OK**.
10. Optional: To have static configuration changes take effect, stop and then restart the application server.

Results

Log entries from this point forward contain the specified level of information for the selected adapter components.

Changing the log and trace file names

To keep the adapter log and trace information separate from other processes, use the administrative console to change the file names. By default, log and trace information for all processes and applications on a application server is written to the SystemOut.log and trace.log files.

Before you begin

You can change the log and trace file names at any time after the adapter module has been deployed to an application server.

About this task

You can change the log and trace file names statically or dynamically. Static changes take effect when you start or restart the application server. Dynamic or run time changes apply immediately.

Log and trace files are in the *install_root*/profiles/*profile_name*/logs/ *server_name* folder.

To set or change the log and trace file names, use the following procedure.

Procedure

1. In the navigation pane of the administrative console, select **Applications > Application Types > WebSphere application servers**.
2. In the Enterprise Applications list, click the name of the adapter application. This is the name of the EAR file for the adapter, but without the ear file extension. For example, if the EAR file is named Accounting_OutboundApp.ear, then click **Accounting_OutboundApp**.
3. In the Configuration tab, select **Modules>Manage Modules**.
4. In the list of modules, click IBM WebSphere Adapter for Oracle.
5. In the Configuration tab, under Additional Properties, click **Resource Adapter**.
6. In the Configuration tab, under Additional Properties, click **Custom properties**.
7. In the Custom Properties table, change the file names.
   a. Click either **logFilename** to change the name of the log file or **traceFilename** to change the name of the trace file.
   b. In the Configuration tab, type the new name in the **Value** field. By default, the log file is called SystemOut.log and the trace file is called trace.log.
   c. Click **Apply** or **OK**. Your changes are saved on your local machine.
   d. To save your changes to the master configuration on the server, use one of the following procedures:
      • **Static change**: Stop and restart the server. This method allows you to make changes, but those changes do not take effect until you stop and start the server.
      • **Dynamic change**: Click the **Save** link in the Messages box above the Custom properties table. Click **Save** again when prompted.

Changing configuration properties for embedded adapters

To change the configuration properties after you deploy the adapter as part of a module, you use the administrative console of the runtime environment. You can update resource adapter properties (used for general adapter operation), managed connection factory properties (used for outbound processing), and activation specification properties (used for inbound processing). For information about configuring logging properties and changing the log and trace file names, see "Configuring logging and tracing" on page 91.

Setting resource adapter properties for embedded adapters

To set resource adapter properties for your adapter after it has been deployed as part of a module, use the administrative console. You select the name of the property you want to configure and then change or set the value.

Before you begin

Your adapter module must be deployed on WebSphere Application Server.

About this task

Custom properties are default configuration properties shared by all IBM WebSphere Adapters.

To configure properties using the administrative console, use the following procedure:

Procedure

1. If the server is not running, right-click your server in the **Servers** view and select **Start**.
2. When the server status changes to **Started**, right-click the server and select **Administration** > **Run administrative console**.
3. Log on to the administrative console.
4. Select **Applications** > **Application Types** > **WebSphere enterprise application**.
5. From the Enterprise Applications list, click the name of the adapter module whose properties you want to change. The **Configuration** page is displayed.
6. Under **Modules**, click **Manage Modules**.
7. Click **IBM WebSphere Adapter for Oracle**.
8. From the **Additional Properties** list, click **Resource Adapter**.
9. On the next page, from the **Additional Properties** list, click **Custom properties**.
10. For each property you want to change, perform the following steps.

    Note: See "Resource adapter properties" on page 130 for more information about these properties.
    a. Click the name of the property. The **Configuration** page for the selected property is displayed.
    b. Change the contents of the **Value** field or type a value, if the field is empty.
    c. Click **OK**.
11. In the Messages area, click **Save**.

Results

The resource adapter properties associated with your adapter module are changed.

Setting managed (J2C) connection factory properties for embedded adapters

To set managed connection factory properties for your adapter after it has been deployed as part of a module, use the administrative console. You select the name of the property you want to configure and then change or set the value.

Before you begin

Your adapter module must be deployed on WebSphere Application Server.

About this task

You use managed connection factory properties to configure the target Oracle database instance.

Note: In the administrative console, the properties are referred to as "J2C connection factory properties."

To configure properties using the administrative console, use the following procedure.

Procedure

1. If the server is not running, right-click your server in the **Servers** view and select **Start**.
2. When the server status changes to **Started**, right-click the server and select **Administration** > **Run administrative console**.
3. Log on to the administrative console.
4. Select **Applications** > **Application Types** > **WebSphere enterprise application**.
5. In the Enterprise Applications list, click the name of the adapter module whose properties you want to change.
6. Under **Modules**, click **Manage Modules**.
7. Click **IBM WebSphere Adapter for Oracle**.
8. In the **Additional Properties** list, click **Resource Adapter**.
9. On the next page, from the **Additional Properties** list, click **J2C connection factories**.
10. Click the name of the connection factory associated with your adapter module.
11. In the **Additional Properties** list, click **Custom properties**.

    Custom properties are those J2C connection factory properties that are unique to IBM WebSphere Adapter for Oracle E-Business Suite. Connection pool and advanced connection factory properties are properties you configure if you are developing your own adapter.
12. For each property you want to change, perform the following steps.

    Note: See "Managed connection factory properties" on page 134 for more information about these properties.
    a. Click the name of the property.

b. Change the contents of the **Value** field or type a value, if the field is empty.

c. Click **OK**.

13. In the Messages area, click **Save**.

Results

The managed connection factory properties associated with your adapter module are changed.

Setting activation specification properties for embedded adapters

To set activation specification properties for your adapter after it has been deployed as part of a module, use the administrative console. You select the name of the message endpoint property you want to configure, and then change or set the value.

Before you begin

Your adapter module must be deployed on WebSphere Application Server.

About this task

You use activation specification properties to configure the endpoint for inbound processing.

To configure properties using the administrative console, use the following procedure.

Procedure

1. If the server is not running, right-click your server in the **Servers** view and select **Start**.

2. When the server status changes to **Started**, right-click the server and select **Administration** > **Run administrative console**.

3. Log on to the administrative console.

4. Select **Applications** > **Application Types** > **WebSphere enterprise application**.

5. From the Enterprise Applications list, click the name of the adapter module whose properties you want to change.

6. Under **Modules**, click **Manage Modules**.

7. Click **IBM WebSphere Adapter for Oracle**.

8. From the **Additional Properties** list, click **Resource Adapter**.

9. On the next page, from the **Additional Properties** list, click **J2C activation specifications**.

10. Click the name of the activation specification associated with the adapter module.

11. From the **Additional Properties** list, click **J2C activation specification custom properties**.

12. For each property you want to change, perform the following steps.

    Note: See "Activation specification properties" on page 155 for more information about these properties.

a. Click the name of the property.
    b. Change the contents of the **Value** field or type a value, if the field is empty.
    c. Click **OK**.
13. In the Messages area, click **Save**.

Results

The activation specification properties associated with your adapter module are changed.

Changing configuration properties for stand-alone adapters

To set configuration properties after you install a stand-alone adapter, use the administrative console of the runtime environment. Provide the general information about the adapter and then set the resource adapter properties (which are used for general adapter operation). If the adapter is used for outbound operations, create a connection factory and then set the properties for it. If the adapter is used for inbound operations, create an activation specification and then set the properties for it. For information about configuring logging properties and changing the log and trace file names, see "Configuring logging and tracing" on page 91.

Setting resource adapter properties for stand-alone adapters

To set resource adapter properties for your stand-alone adapter after it has been installed on WebSphere Application Server, use the administrative console. You select the name of the property you want to configure and then change or set the value.

Before you begin

Your adapter must be installed on WebSphere Application Server.

About this task

Custom properties are default configuration properties shared by all IBM WebSphere Adapters.

To configure properties using the administrative console, use the following procedure:

Procedure

1. If the server is not running, right-click your server in the **Servers** view and select **Start**.
2. When the server status changes to **Started**, right-click the server and select **Administration** > **Run administrative console**.
3. Log on to the administrative console.
4. Click **Resources** > **Resource Adapters** > **Resource adapters**.
5. In the Resource adapters page, click **IBM WebSphere Adapter for Oracle**.
6. In the **Additional Properties** list, click **Custom properties**.
7. For each property you want to change, perform the following steps.
    a. Click the name of the property.

b. Change the contents of the **Value** field or type a value, if the field is empty.

c. Click **OK**.

8. In the Messages area, click **Save**.

Results

The resource adapter properties associated with your adapter are changed.

Setting managed (J2C) connection factory properties for stand-alone adapters

To set managed connection factory properties for your stand-alone adapter after it has been installed on WebSphere Application Server, use the administrative console. You select the name of the property you want to configure and then change or set the value.

Before you begin

Your adapter must be installed on WebSphere Application Server.

About this task

You use managed connection factory properties to configure the target Oracle database instance.

Note: In the administrative console, the properties are referred to as "J2C connection factory properties."

To configure properties using the administrative console, use the following procedure:

Procedure

1. If the server is not running, right-click your server in the **Servers** view and select **Start**.

2. When the server status changes to **Started**, right-click the server and select **Administration** > **Run administrative console**.

3. Log on to the administrative console.

4. Click **Resources** > **Resource Adapters** > **Resource adapters**.

5. In the Resource adapters page, click **IBM WebSphere Adapter for Oracle**.

6. In the **Additional Properties** list, click **J2C connection factories**.

7. If you are going to use an existing connection factory, skip ahead to select from the list of existing connection factories.

   Note: If you have selected **Specify connection properties** when you use the J2C Bean wizard to configure the adapter module, you do not need to create a connection factory.

   If you are creating a connection factory, perform the following steps:

   a. Click **New**.

   b. In the **General Properties** section of the **Configuration** tab, type a name for the connection factory. For example, you can type `AdapterCF`.

   c. Type a value for **JNDI name**. For example, you can type `com/eis/AdapterCF`.

d. Optional: Select an authentication alias from the **Component-managed authentication alias** list.

e. Click **OK**.

f. In the Messages area, click **Save**.

   The newly created connection factory is displayed.

8. In the list of connection factories, click the one you want to use.

9. In the **Additional Properties** list, click **Custom properties**.

   Custom properties are those J2C connection factory properties that are unique to WebSphere Adapter for Oracle E-Business Suite. Connection pool and advanced connection factory properties are properties you configure if you are developing your own adapter.

10. For each property you want to change, perform the following steps.

    Note: See "Managed connection factory properties" on page 134 for more information about these properties.

    a. Click the name of the property.

    b. Change the contents of the **Value** field or type a value, if the field is empty.

    c. Click **OK**.

11. After you have finished setting properties, click **Apply**.

12. In the Messages area, click **Save**.

Results

The managed connection factory properties associated with your adapter are set.

Setting activation specification properties for stand-alone adapters

To set activation specification properties for your stand-alone adapter after it has been installed on WebSphere Application Server, use the administrative console. You select the name of the message endpoint property you want to configure, and then change or set the value.

Before you begin

Your adapter must be installed on WebSphere Application Server.

About this task

You use activation specification properties to configure the endpoint for inbound processing.

To configure properties using the administrative console, use the following procedure.

Procedure

1. If the server is not running, right-click your server in the **Servers** view and select **Start**.

2. When the server status changes to **Started**, right-click the server and select **Administration** > **Run administrative console**.

3. Log on to the administrative console.

4. Click **Resources** > **Resource Adapters** > **Resource adapters**.
5. In the Resource adapters page, click **IBM WebSphere Adapter for Oracle**.
6. In the **Additional Properties** list, click **J2C activation specifications**.
7. If you are going to use an existing activation specification, skip ahead to select from an existing list of activation specifications.

   Note: If you have selected **Use predefined connection properties** when you use the J2C Bean wizard to configure the adapter module, you must create an activation specification.

   If you are creating an activation specification, perform the following steps:
   a. Click **New**.
   b. In the **General Properties** section of the **Configuration** tab, type a name for the activation specification. For example, you can type `AdapterAS`.
   c. Type a value for **JNDI name**. For example, you can type `com/eis/AdapterAS`.
   d. Optional: Select an authentication alias from the **Authentication alias** list.
   e. Select a message listener type.
   f. Click **OK**.
   g. Click **Save** in the **Messages** box at the top of the page.
      The newly created activation specification is displayed.
8. In the list of activation specifications, click the one you want to use.
9. In the Additional Properties list, click **J2C activation specification custom properties**.
10. For each property you want to set, perform the following steps.
    a. Click the name of the property.
    b. Change the contents of the **Value** field or type a value, if the field is empty.
    c. Click **OK**.
11. After you have finished setting properties, click **Apply**.
12. In the Messages area, click **Save**.

Results

The activation specification properties associated with your adapter are set.

Adding dependency libraries to the deployed resource adapter

The deployed resource adapter running in the WebSphere Application Server requires the same dependency libraries as it does in Rational Application Developer for WebSphere Software to process requests. The method for adding these library files depends on the mode of the resource adapter deployment: standalone or embedded in the EAR file.

Stand-alone deployment

The dependency libraries may be added to the resource adapter deployed stand-alone either during initial deployment of the RAR file or by configuring the Resource Adapter properties after deployment. To set the values during initial deployment of the RAR file, specify Class path and Native path locations. Class path is used to point to JAR files, and Native path is used to point to native libraries, such as *.dll, *.so. To set the dependency library path files after the

adapter has been installed on WebSphere Application Server, use the administrative console to modify the values for the Resource Adapter.

EAR deployment

For the rare case when the connector needs to be embedded in the EAR file, the dependant libraries are added as shared libraries. Define the appropriate shared library containing software dependency files and associate them with the EAR file.

About this task

There are two methods to do this task:
- Using enhanced EAR editor in Rational Application Developer for WebSphere Software
- Using administrative console of the WebSphere Application Server

Using enhanced EAR editor

You can use the EAR editor in Rational Application Developer for WebSphere Software to add the dependency libraries.

About this task

To create shared libraries using the EAR editor, use the following procedure.

Procedure

1. Open Enhanced EAR editor.
2. Click the **Deployment** tab.
3. Expand **Shared Library** section.
4. Click **Add** to add new shared library.
5. Specify the shared library parameters and click **OK**.
6. Deploy the EAR to the server.

Results

The dependent libraries are added as shared libraries.

Using administrative console of the WebSphere Application Server

You can use the administrative console of the WebSphere Application Server to add the dependency libraries.

Before you begin

Ensure that the dependent files are available on the server machine in the separate folder. If needed, copy the dependent files on the server machine.

Procedure

1. Define WebSphere variables to point to appropriate folders.
2. Define the shared library through the server administrative console; you can specify it using WebSphere variables defined in step 1.
3. Deploy the EAR to the server.
4. Configure the EAR to reference defined shared library.

Results

The dependent libraries are added as shared libraries.

Chapter 6. Troubleshooting and support

Common troubleshooting techniques and self-help information help you identify and solve problems quickly. For information about configuring logging properties and changing the log and trace file names, see "Configuring logging and tracing" on page 91.

Techniques for troubleshooting problems

Troubleshooting is a systematic approach to solving a problem. The goal is to determine why something does not work as expected and how to resolve the problem. Certain common techniques can help with the task of troubleshooting.

The first step in the troubleshooting process is to describe the problem completely. Without a problem description, neither you or IBM® can know where to start to find the cause of the problem. This step includes asking yourself basic questions, such as:
- What are the symptoms of the problem?
- Where does the problem occur?
- When does the problem occur?
- Under which conditions does the problem occur?
- Can the problem be reproduced?

The answers to these questions typically lead to a good description of the problem, and that is the best way to start down the path of problem resolution.

What are the symptoms of the problem?

When starting to describe a problem, the most obvious question is "What is the problem?" Which might seem like a straightforward question; however, you can break it down into several more-focused questions that create a more descriptive picture of the problem. These questions can include:
- Who, or what, is reporting the problem?
- What are the error codes and messages?
- How does the system fail? For example, is it a loop, hang, lock up, performance degradation, or incorrect result?
- What is the business impact of the problem?

Where does the problem occur?

Determining where the problem originates is not always simple, but it is one of the most important steps in resolving a problem. Many layers of technology can exist between the reporting and failing components. Networks, disks, and drivers are only a few components to be considered when you are investigating problems.

The following questions can help you to focus on where the problem occurs in order to isolate the problem layer.
- Is the problem specific to one platform or operating system, or is it common for multiple platforms or operating systems?
- Is the current environment and configuration supported?

Remember that if one layer reports the problem, the problem does not necessarily originate in that layer. Part of identifying where a problem originates is

understanding the environment in which it exists. Take some time to completely describe the problem environment, including the operating system and version, all corresponding software and versions, and hardware information. Confirm that you are running within an environment that is a supported configuration; many problems can be traced back to incompatible levels of software that are not intended to run together or have not been fully tested together.

When does the problem occur?

Develop a detailed timeline of events leading up to a failure, especially for those cases that are one-time occurrences. You can most simply do this by working backward: Start at the time an error was reported (as precisely as possible, even down to the millisecond), and work backward through the available logs and information. Typically, you need to look only as far as the first suspicious event that you find in a diagnostic log; however, this is not always simple to do and takes practice. Knowing when to stop looking is especially difficult when multiple layers of technology are involved, and when each has its own diagnostic information.

To develop a detailed timeline of events, answer the following questions:
- Does the problem happen only at a certain time of day or night?
- How often does the problem happen?
- What sequence of events leads up to the time that the problem is reported?
- Does the problem happen after an environment change, such as upgrading or installing software or hardware?

Responding to these types of questions can provide you with a frame of reference in which to investigate the problem.

Under which conditions does the problem occur?

Knowing what other systems and applications are running at the time that a problem occurs is an important part of troubleshooting. These and other questions about your environment can help you to identify the root cause of the problem:
- Does the problem always occur when the same task is being performed?
- Does a certain sequence of events need to occur for the problem to surface?
- Do any other applications fail at the same time?

Answering these types of questions can help you explain the environment in which the problem occurs and correlate any dependencies. Remember that just because multiple problems might have occurred around the same time, the problems are not necessarily related.

Can the problem be reproduced?

From a troubleshooting standpoint, the "ideal" problem is one that can be reproduced. Typically with problems that can be reproduced, you have a larger set of tools or procedures at your disposal to help you investigate. Consequently, problems that you can reproduce are often simpler to debug and solve. However, problems that you can reproduce can have a disadvantage: If the problem is of significant business impact, you do not want it to recur! If possible, re-create the problem in a test or development environment, which typically offers you more flexibility and control during your investigation.

Tip: Simplify the scenario to isolate the problem to a suspected component.

The following questions can help you with reproducing the problem:
- Can the problem be re-created on a test machine?
- Are multiple users or applications encountering the same type of problem?
- Can the problem be re-created by running a single command, a set of commands, a particular application, or a stand-alone application?

First-failure data capture (FFDC) support

The adapter supports first-failure data capture (FFDC), which provides persistent records of failures and significant software incidents that occur during run time in WebSphere Application Server.

The FFDC feature runs in the background and collects events and errors that occur at run time. The feature provides a means for associating failures to one another, allowing software to link the effects of a failure to their causes, and thereby facilitate the quick location of the root cause of a failure. The data that is captured can be used to identify exception processing that occurred during the adapter run time.

When a problem occurs, the adapter writes exception messages and context data to a log file, which is located in the *install_root*/profiles/*profile*/logs/ffdc directory.

For more information about first-failure data capture (FFDC), see the WebSphere Application Server documentation.

Tracing the XML Gateway Web Service status

Use Oracle E-Business Suite administration console to find the XML Gateway Web Service status.

The messages in return Business Object indicate whether the XML Gateway Web Services are invoked successfully or not. If succeeded, it means the submitted document is received by XML Gateway Web Services server side and put into the queue for consequent asynchronous process. In order to check whether the requests are successfully processed by XML Gateway internal implementation, log on to Oracle E-Business Suite administration console and find out the details in the Transaction Monitor.

Resolving namespace conflict in Web Service/JMS interface

When using the Web Service/JMS interface to connect to two different Oracle E-Business Suite servers, if the selected root DTD and the root Element are similar, the generated artifacts will have the same namespace in the top BO, which results in a naming conflict.

To avoid this problem, set a different BO prefix.

Resolving connection information lost during EMD

When using different interfaces of WebSphere Adapter for Oracle E-Business Suite and accessing the EMD several times, the connection information may be lost at the J2C Bean Creation and Deployment Configuration window. For example, when you are configuring the module for the first time, and choose the XML Gateway JMS processing option, to reach the J2C Bean Creation and Deployment Configuration window, and then go back to the Discovery Configuration window,

to choose Advanced Queue for processing, and reach the J2C Bean Creation and Deployment Configuration window again, you may find that the "Database URL" is empty, when you choose "Specify XA database connection information" as "Database connection information".

To avoid this, choose only one type of interface during the configuration, and not switch between interface options.

Disabling end point applications of the passive adapter

Problem

In the active-passive configuration mode of the adapters, the endpoint application of the passive adapter instance also listens to the events or messages even if the enableHASupport property is set to `True`.

Cause

By default, in WebSphere Application Server, the `alwaysactivateAllMDBs` property in the JMS activation specification is set to `True`. This enables the endpoint application of all the adapter (active or passive) instances to listen to the events.

Solution

To stop the endpoint application of the passive adapter instance from listening to the events, you must set the `alwaysactivateAllMDBs` property value to `False`. The JMS activation specification is associated with one or more MDBs and provides the necessary configuration to receive events. If the `alwaysActivateAllMDBs` property is set to `False`, then the endpoint application of only the active adapter instance receives the events.

Perform the following procedure, to set the `alwaysActivateAllMDBs` property to `False`.
1. Log on to the administrative console.
2. Go to **Resources> JMS > Activation specifications**.
3. Click the activation specification corresponding to the application from the list.
4. Click **Custom properties** under **Additional properties**.
5. Click **alwaysActivateAllMDBs**.
6. Change the value to **False**.
7. Click **Apply** and **OK**.

Result

The endpoint application of only the active adapter instance listens to the events.

Solutions to common problems

Some of the problems you encounter while running WebSphere Adapter for Oracle E-Business Suite with your database are described, along with solutions and workaround. These problems and solutions are in addition to those documented as technotes on the software support Web site.

For a complete list of technotes about WebSphere Adapters, see
http://www.ibm.com/support/search.wss?tc=SSMKUK&rs=695&rank=8
&dc=DB520+D800+D900+DA900+DA800+DB560&dtm.

RecordNotFoundException for RetrieveAll operation in test client

Problem

When performing a RetrieveAll operation in the Rational Application Developer
for WebSphere Software test client, the RecordNotFoundException exception is
generated when data is expected from the query. The following message is
generated: `RecordNotFoundException: Record not found in EIS`.

Cause

This exception can occur if the WHERE clause for the SELECT statement does not
set all the attributes of the business object. Leaving the attribute blank, which is the
default value, is not the same as explicitly unsetting the value.

Solution

In the test client, set the values of the attributes that are required to `<unset>`.
Repeat the RetrieveAll operation. If the exception is generated again, it is likely
that no matching records exist in the database table.

CLOB datatypes of 4K or larger cannot be inserted into Oracle 9i or 10g databases

Problem

You get the following exception when inserting CLOB (character large object)
values of 4 K and larger into Oracle 9i or 10g databases:

- Oracle 9i: `ResourceAdapt E`
  `com.ibm.j2ca.dbadapter.core.runtime.DBOperationHandler.`
  `executePreparedCUDStatement CWYDB0301E: An operation on the database`
  `failed with a SQL exception for the following reason: No more data to`
  `read from socket.`
- Oracle 10g: `ResourceAdapt E`
  `com.ibm.j2ca.dbadapter.core.runtime.DBOperationHandler.`
  `executePreparedCUDStatement CWYDB0301E: An operation on the database`
  `failed with a SQL exception for the following reason: ORA-01460:`
  `unimplemented or unreasonable conversion requested`

Cause

You are using an older driver that does not correctly support CLOBs larger than
4K.

Solution

Use the Oracle Thin driver from Oracle 10.1.0.2 or a later release.

Some generated business objects have no attributes for Oracle database objects

Problem

For some business objects that are generated from an Oracle database object, the generated business object has no attributes for the table columns.

Cause

Under certain conditions, the Oracle JDBC driver does not return the column information for a database object. The following bugs are currently filed with Oracle for these issues:

- 2281705. DATABASEMETADA.GETCOLUMNS does not return underlying table if there is a synonym
- 2696213. JDBC GETPROCEDURECOLUMNS does not return columns for the synonym of a procedure

Additionally, column information is not returned if a private synonym that references an object in another schema is used.

Solution

For tables that have a synonym, generate the business object using the synonym for the table.

For synonyms of a procedure, generate the business object using the original procedure that the synonym is based on.

For private synonyms that reference an object in another schema, either use the original table or create a synonym in the current schema.

ResourceException exceptions during outbound processing

If you get a ResourceException exception, examine the root cause field to determine the cause. Common problems have the following root causes:

- SQLException exception

  If the SQLException exception includes the text `UserID or password is invalid`, then the user ID or password specified for the outbound connection is not correct.

  For example:

  ```
  javax.resource.ResourceException: ORA-01017; invalid username/password;
  logon denied.
  ```

- ConnectException exception

  If the text included with the ConnectException exception includes text like `is not reachable` or `could not establish the connection`, then the database server might not be operational or there might be a network problem that prevents a connection.

  For example:

  ```
  java.sql.SQLException: Io exception: The Network Adapter could not
  establish the connection
  ```

ResourceException exception during inbound processing

This exception indicates that there is a repeated problem connecting to the database. To polls for events, the adapter must connect top the database. If the connection fails, the adapter waits a configurable amount of time before trying to connect again. The adapter tries a configurable number of times before it stops polling. When the adapter stops polling, it generates the ResourceException exception.

Class loader violation occurs when starting the J2C Bean wizard

Problem

It is not possible to use the J2C Bean wizard after using a connection to the database in the Data perspective. At the end of the second panel of the wizard, the following exception is generated:

com.ibm.adapter.framework.api.ImportException
Reason: class loading constraint violated (class:
oracle/jdbc/driver/OracleConnection
method: getWrapper()Loracle/jdbc/OracleConnection;) at pc:0

The error occurs in both of the following situations:
- When you establish a connection to the database through the J2C Bean wizard, an error occurs when you attempt to connect to the database from the Data perspective.
- When you establish a connection to the database through the Data perspective, the error occurs when you attempt to connect to database through the J2C Bean wizard.

Cause

The error occurs because the Data perspective and the wizard use their own class loaders. Once the DLL, which is the native library used by the JDBC driver, is loaded in the Data perspective, it cannot be loaded again in the wizard. JVMs have an inherent restriction that only allows one class loader to load native libraries at any given time. So if class loader A loads DLL B, then no other class loader can load DLL B until class loader A is released and garbage is collected. Because you cannot really control garbage collection, it typically means that if you want to load DLL B with another class loader, you need to restart the JVM. This limitation is a known one and is documented for WebSphere Application Server.

Solution

The only solution is to restart Rational Application Developer for WebSphere Software when this error occurs.

Closed connection error occurs when using XA with Oracle 10g

Problem

When the Adapter for Oracle E-Business Suite is used to perform an XA transaction using Oracle 10g, the adapter generates a closed connection exception: javax.resource.ResourceException: Closed Connection.

Cause

This is a known issue with the Oracle 10g database driver. The following bug has been filed with Oracle for this issue: 3488761 Connection closed error from OracleConnection.getConnection() - 10G drivers.

Solution

The bug has been fixed in the Oracle 10g Release 2 driver. As a workaround, you can use the Oracle 9i JDBC Thin drivers to connect to the database for XA transactions.

Error occurs while starting a transaction on Oracle

Problem

When the Adapter for Oracle E-Business Suite is used to perform an XA transaction using Oracle database, the following error is generated: `WTRN0078E: An attempt by the transaction manager to call start on a transactional resource has resulted in an error. The error code was XAER_RMERR.`

Cause

For the Oracle database server to support XA transactions, some commands need to be run.

Solution

Two scripts that are included in the Oracle directory should be run. This activity mostly needs to be performed by your Oracle database administrator, because you must be logged on to Oracle as SYSOPER or SYSDBA in order to have the necessary permissions for these scripts to work. The scripts are:

```
<ORACLE_HOME>javavm\install
file: initxa.sql
file: initjvm.sql
```

The initxa.sql script configures the database for XA. Once it runs successfully, your database is configured for XA. The script might run successfully the first time you try. Unfortunately, it probably will not run successfully because some of the database memory spaces are too small.

To fix this, run the initjvm.sql script. It probably fails too, but in doing so, it indicates which parameters need to be adjusted. The parameters are stored in this file:

```
<ORACLE_HOME>\database
file: init<DATABASE_SID>.ora
```

Table 16 shows two parameters that typically need to be increased. Your particular database configuration might require adjustment of different parameters.

Table 16. Typical parameter sizes

Parameter Name	Minimum Value
java_pool_size	12000000
shared_pool_size	24000000

A closer look at the transaction (XID) column in the event table

If the adapter is configured for assured once delivery, use the status column with the XID column to determine whether the event has been processed:

- If the XID column contains 0, the event has not yet been picked up for processing.
- If the XID column contains a transaction ID (that is, it does not contain 0), then the adapter has started to process the event but has not completed processing. You might see this combination when the adapter or application server stops while the event is being processed. The transaction manager will either COMMIT or ROLLBACK these transactions during recovery.

Handling unexpected results from a query SQL statement

Should you receive unexpected results from a query, turn on tracing and look at the query SQL in the log. Turning on tracing is especially helpful when you are in the test client to see if you forgot to *unset* all the unnecessary attributes. It is also practical to turn on tracing to determine if you did not specify your business object correctly.

Capturing the RollBack exception during a global transaction deployment

If an error occurs during the deployment of an inbound operation for global transactions, the adapter does not display the RollBack exception and the EventStatus in the event table is not updated to the value "-1" as expected. To enable the inbound process to capture the RollBack exception, modify the Enterprise JavaBeans (EJB) version 3.0 code manually.

1. Modify the local business interface file (*ss.java) of the Stateless session EJB as follows:

   ```
   Import javax.ejb.SessionContext;
   ... ...
   public class OracleInboundssSB implements OracleInboundEJBpkg.OracleInboundss {
     @Resource
     private SessionContext ctx;
     SiebelOutboundInterface siebelbean;
     @InboundMethodBinding(nativeMethod="emitCreateAfterImageAdministratorCustomer")
   public void createAdministratorCustomer(com.ibm.xmlns.prod.websphere.j2ca.jdbc.administratorcu
     ... ...
     Catch (Exception e){
     ctx.setRollbackOnly();
     e.printStackTrace();
     }
   ```

2. Modify the Stateless session EJB file (*SB.java) as follows:

   ```
   public void createAdministratorCustomer(com.ibm.xmlns.prod.websphere.j2ca.jdbc.administratorcu
   ```

Adapter returns version conflict exception message

Problem

When you install multiple adapters with different versions of CWYBS_AdapterFoundation.jar, and if a lower version of the CWYBS_AdapterFoundation.jar is loaded during runtime, the adapter will return the ResourceAdapterInternalException error message, due to a version conflict. For example, when you install Oracle E-Business Suite adapter version 7.0.0.3 and WebSphere Adapter for Oracle E-Business Suite version 7.5.0.2, the following error message is displayed "The version of CWYBS_AdapterFoundation.jar is not

compatible with IBM WebSphere Adapter for Oracle E-Business Suite" as IBM WebSphere Adapter for Oracle E-Business Suite loads file:/C:/IBM/WebSphere/ProcServer7/profiles/ProcSrv01/installedConnectors/CWYOE_OracleEBS.rar/CWYBS_AdapterFoundation.jar with version 7.0.0.3. However, the base level of this jar required is version 7.5.0.2. To overcome this conflict, you must ensure that all adapters are at same version level. For further assistance, contact WebSphere Adapters Support for help.

Solution

Ensure that all adapters are at the same version level.

For further assistance, visit http://www-947.ibm.com/support/entry/portal/Overview/Software/WebSphere/WebSphere_Adapters_Family.

Problem

After configuring the adapter, you might note that:

1. the webservices client based on the WSDL is not getting generated properly in Rational Application Developer for WebSphere Software.
2. you may not be able to invoke the WSDL using certain webservices client, such as RESTUI firefox plugin and soapUI tool.

Solution

Perform the following steps to enable the adapter module to load the ASI file.

1. Create a library project.
2. Change to the Enterprise Explorer view in Rational Application Developer for WebSphere Software.
3. Locate the ASI file of the adapter in **connector project -> connectorModule**.
4. Copy the ASI file of the adapter and paste it in the library project.
5. Add the library project to the list of dependencies for the adapter module.
6. Clean project.

Support

This section provides information about how to troubleshoot a problem with your IBM® software, including instructions for searching knowledge bases, downloading fixes, and obtaining support.

Searching knowledge bases (Web search)

You can often find solutions to problems by searching IBM knowledge bases. You can optimize your results by using available resources, support tools, and search methods.

About this task

You can find useful information by searching the information center for Product X. However, sometimes you need to look beyond the information center to answer your questions or resolve problems.

To search knowledge bases for information that you need, use one or more of the following approaches:

- Search for content by using the IBM® Support Assistant (ISA).

  ISA is a no-charge software serviceability workbench that helps you answer questions and resolve problems with IBM software products. You can find instructions for downloading and installing ISA on the ISA website.
- Find the content that you need by using the IBM Support Portal.

  The IBM Support Portal is a unified, centralized view of all technical support tools and information for all IBM systems, software, and services. The IBM Support Portal lets you access the IBM electronic support portfolio from one place. You can tailor the pages to focus on the information and resources that you need for problem prevention and faster problem resolution. Familiarize yourself with the IBM Support Portal by viewing the demo videos (https://www.ibm.com/blogs/SPNA/entry/the_ibm_support_portal_videos) about this tool. These videos introduce you to the IBM Support Portal, explore troubleshooting and other resources, and demonstrate how you can tailor the page by moving, adding, and deleting portlets.
- Search for content by using the IBM masthead search. You can use the IBM masthead search by typing your search string into the Search field at the top of any ibm.com® page.
- Search for content by using any external search engine, such as Google, Yahoo, or Bing. If you use an external search engine, your results are more likely to include information that is outside the ibm.com domain. However, sometimes you can find useful problem-solving information about IBM products in newsgroups, forums, and blogs that are not on ibm.com.

Tip: Include "IBM" and the name of the product in your search if you are looking for information about an IBM product.

Getting Fixes

A product fix might be available to resolve your problem.

About this task

To get product fixes, perform the following steps.

Procedure

1. Determine which fix you need. Check the list of IBM WebSphere Adapter for Oracle E-Business Suite recommended fixes to confirm that your software is at the latest maintenance level. Check the list of problems fixed in the IBM WebSphere Adapter for Oracle E-Business Suite fix readme documentation that is available for each listed fix pack to see if IBM has already published an individual fix to resolve your problem. To determine what fixes are available using IBM Support Assistant, run a query on fix from the search page.

   Individual fixes are published as often as necessary to resolve defects in WebSphere Application Server IBM WebSphere Adapter for Oracle E-Business Suite. In addition, two kinds of cumulative collections of fixes, called fix packs and refresh packs, are published periodically for IBM WebSphere Adapter for Oracle E-Business Suite, in order to bring users up to the latest maintenance level. You should install these update packages as early as possible in order to prevent problems.

   Note: A list of recommended, generally available (GA) fixes for the WebSphere Java™ Connector Architecture (JCA) and WebSphere Business Integration adapters are available here. If a Fix Pack is not available for an adapter, it

implies that the GA version is the recommended version and details about that version of the adapter can be found in the Release notes.

2. Download the fix. Open the download document and follow the link in the Download package section. When downloading the file, ensure the name of the maintenance file is not changed. This includes both intentional changes and inadvertent changes caused by certain web browsers or download utilities.

3. Apply the fix. Follow the instructions in the Installation Instructions section of the download document.

4. Optional: To receive weekly notification of fixes and updates, subscribe to My Support e-mail updates.

Self-help resources

Use the resources of IBM software support to get the most current support information, obtain technical documentation, download support tools and fixes, and avoid problems with WebSphere Adapters. The self-help resources also help you diagnose problems with the adapter and provide information about how to contact IBM software support.

Support website

The WebSphere Adapters software support website at http://www-947.ibm.com/support/entry/portal/Overview/Software/WebSphere/WebSphere_Adapters_Family provides links to many resources to help you learn about, use, and troubleshoot WebSphere Adapters, including:
- Flashes (alerts about the product)
- Technical information including the product information center, manuals, IBM Redbooks®, and whitepapers
- Educational offerings
- Technotes

Recommended fixes

A list of recommended fixes you must apply is available at the following location: http://www.ibm.com/support/docview.wss?fdoc=aimadp&rs=695 &uid=swg27010397.

Technotes

Technotes provide the most current documentation about WebSphere Adapter for Oracle E-Business Suite, including the following topics:
- Problems and their currently available solutions
- Answers to frequently asked questions
- How to information about installing, configuring, using, and troubleshooting the adapter
- *IBM Software Support Handbook*

For a list of technotes for WebSphere Adapter for Oracle E-Business Suite, see http://www-01.ibm.com/support/docview.wss?uid=swg27020209.

For a list of technotes for all adapters, see http://www.ibm.com/support/search.wss?tc=SSMKUK&rs=695&rank=8 &dc=DB520+D800+D900+DA900+DA800+DB560&dtm.

Chapter 7. Reference

Detailed information about business objects, adapter properties (enterprise service discovery properties, resource adapter properties, managed (J2C) connection factory properties, activation specification properties, and interaction specification properties), messages, and related product information is provided for your reference.

Business object information

A business object is a structure that contains application-specific information (metadata) about how the adapter must process the business object as well as the operation to be performed on the business object. The name of the business object is generated by the J2C Bean wizard in accordance with the naming convention for the adapter.

Business object attributes

Business object attributes define the content of a business object and are built from the list of columns in the database object.

A business object is simply a container for the data specified in the attributes. Each attribute has a name, type, cardinality, and several other properties. The J2C Bean wizard sets the attribute name to the name of the column. The adapter adds the attribute cardinality, type, and application-specific information. The structure of the data in the database is defined by the business object, but data in the database is in the business object attributes.

Table 17 lists the properties of a business object attribute and describes their interpretation and settings.

Table 17. Attribute properties

Properties	Interpretation and settings
Cardinality	An integer specifying the cardinality of a business object. Each business object attribute that represents a child or an array of child business objects has the value of single or multiple (an unbounded integer) cardinality. In both single- and multiple-cardinality relationships, the relationship between the parent and child business objects is described by the application-specific information of the key attribute in the business object storing the relationship.
Foreign Key	When arrays of child business objects whose cardinality is n are retrieved, foreign keys are used in the WHERE clause of SELECT statements. The RetrieveAll operation overrides the use of keys and foreign keys. **Note:** The adapter does not support specifying an attribute that represents a child business object as a foreign key.
Name	This property represents the unique name of the attribute, if it is a simple attribute, or the name of the business object, if it is a child business object.

Table 17. Attribute properties (continued)

Properties	Interpretation and settings
MinOccurs MaxOccurs	If the column is not a primary key and is not null able, the MinOccurs and MaxOccurs attributes are required, and their values are set to at least 1.
Primary Key	Indicates whether this attribute is a primary key. At least one simple attribute in each business object must be specified as the primary key. If the primary key property is set to `true` for a simple attribute, the adapter adds that attribute to the WHERE clause of the SELECT statement and SQL UPDATE statements that it generates while processing the business object. The RetrieveAll operation overrides the use of primary and foreign keys. **Note:** The adapter does not support specifying an attribute that represents a child business object or an array of child business objects as a primary key attribute.
Required	Specifies whether an attribute must contain a value. If this property is set to `true` for a container whose cardinality is single (1), then the adapter requires that the parent business object contain a child business object for this attribute. Business objects that are passed to the adapter for Create, Update, and Delete operations must also contain a child business object. Cardinality is single (1) for simple attributes and multiple (n) for container attributes. The adapter causes a Create operation to fail if a business object does not have a valid value or a default value for a required attribute. It also fails if no data is available upon retrieval from the database for this object.
Type	For simple attributes, this property specifies the type of the attribute, such as Integer, String, Date, Timestamp, Boolean, Double, or Float. The supported types for simple attributes and their mapping to the Oracle type of a database object are described in Table 18 on page 117. For attributes that specify a child business object, this property specifies the name of the business object.

The type of each database object, returned as the Oracle metadata, maps to the business object attribute types as listed in Table 18 on page 117. Only the Oracle types listed are supported by the adapter. Any columns with types that are not listed are not added to the business object. An informational message is produced explaining the problem, for example, `The column named xxxx in the table named yyyy is not of a supported type and is not added to the business object.`

Note: The default data type mapping varies based on the different Oracle JDBC driver versions. If the Oracle metadata does not map to the same data type during the configuration of the Oracle database objects, select the appropriate data type manually in the Configuration Parameters window. After the generation of the business object, if you find the Oracle metadata not mapped to the same business object attribute type, update the attribute data type manually in the XSD file for the business object.

Table 18. Oracle metadata column type and business object attribute types

Oracle metadata column type	Business object attribute type
CHAR LONG VARCHAR2	String
NUMBER	Decimal
TIMESTAMP	DateTime (String data type is displayed by default)
DATE	Date (String data type is displayed by default)
FLOAT	Double
BLOB	hexBinary
CLOB	String
NCHAR NVARCHAR2	String
RAW LONG RAW	hexBinary
STRUCT or ARRAY	The adapter processes these data types as child business objects of the table or query business objects. **Note:** The adapter supports complex types for the Oracle table and query business objects only. If the table contains any complex data type, such as an array, structure, nested structure or table, the type name and the sub attribute details are also automatically discovered and displayed. **Note:** The adapter treats an empty complex column as null irrespective of setting it to null or unset.
BOOLEAN	The adapter supports the boolean datatype for SP/SF with Record type parameter.

Attribute application-specific information

Application-specific information (ASI) for business object attributes differs depending on whether the attribute is a simple attribute, or is an attribute that represents a child or an array of child business objects. The application-specific information for an attribute that represents a child differs depending on whether the parent-child relationship is stored in the child or in the parent.

Application-specific information for simple attributes

For simple attributes, the format for application-specific information consists of a number of parameters and their values. The only parameter that is required for a simple attribute is the column name. The application-specific information for simple attributes is described in Table 19 on page 118.

Table 19. Application-specific information for simple attributes

Parameter	Type	Description	Default value
BLOB	Boolean	Indicates whether the database column that corresponds to this attribute has the BLOB data type. While displaying BLOB data, the adapter displays the number of bytes as a hexadecimal value. The attribute type is hexBinary. If `True`, the column data type is BLOB.	None
ByteArray	Boolean	Specifies whether the column is a binary data type. If `True`, the adapter reads and writes binary data to the database and sends that data as a string to the application server. The adapter sets binary data on the business object. The attribute type is hexBinary.	`False`
ChildBOType	String	If the attribute is a complex data type, use this application-specific information to specify the actual type: • Struct • Array • ResultSet	None
ChildBOTypeName	String	When the value of the ChildBOType application-specific information is either Struct or Array, this parameter value represents the name of the user-defined type. This value is case sensitive.	
CLOB	Boolean	Indicates whether the database column that corresponds to this attribute has the CLOB data type. This value applies only to attributes of type String. If `True`, the column data type is CLOB. The CLOB attribute has a String Type whose length is used to define the length of the CLOB.	None
ColumnName	String	The name of the database column corresponding to this attribute. This is the only required parameter.	None

Table 19. Application-specific information for simple attributes (continued)

Parameter	Type	Description	Default value
CopyAttribute	String	A user-specified value that refers to another attribute name from within the same business object or parent business object. If the value set in the application-specific information refers to the name of another attribute within the same business object, then the adapter uses the value of the other attribute to set the value of this attribute (on which application-specific information is defined) before it adds the business object to the database during a Create operation. For example, if you want the `contact` column of a new row in the table to contain the same value as the `email` column, set the CopyAttribute parameter of the `contact` attribute to `email`. The value cannot reference an attribute in a child business object, but it can reference an attribute in the parent business object by preceding the name with two periods. For example, you can reference the `ccode` attribute in a parent business object as `..ccode`. If you do not include this parameter in the application-specific information, the adapter uses the value of the current attribute without copying the value from another attribute.	None
DateType	String	Specifies that the corresponding element is a date or time stamp type. Specify one of the following values: • Date • Timestamp When setting the value of an attribute of the DateType type, use the following formats: • For Date, use yyyy-MM-dd • For Timestamp, use yyyy-MM-dd hh:mm:ss.fffffffff **Note:** The adapter uses the Timestamp value present in the database. To learn more about the Timestamp method, go to the Sun Web site at http://java.sun.com/j2se/1.5.0/docs/api/ and search for *Timestamp*.	None

Parameter	Type	Description	Default value

FixedChar	Boolean	Specifies whether the attribute is of fixed length when the columns in the table are of type CHAR, not VARCHAR. For example, when set to `true`, if a particular attribute is linked to a column that is of type CHAR, the adapter pads the attribute value with blanks to the maximum length of the attribute when querying the database.
		
		This parameter must be updated manually in the XSD file for the business object. Open the business object by using an XML or text editor to edit the XSD file. Two changes must be made, as follows:
		
		1. Remove the `type="string"` added by default to the `<element>` tag for the object attribute.
		
		2. Add a new `<simpletype>` section before the `</element>` tag, as in this example:
		
		`<xsd:simpleType>`
		`<xsd:restriction base="xsd:string">`
		`<xsd:maxLength value="10"/>`
		`</xsd:restriction>`
		`</xsd:simpleType>`
		
		Save the object definition, and ensure that no validation errors occur in the XSD file after it has been updated.
		
		Following this table, see the Example of FixedChar parameter in the business object XSD file.
			false

ForeignKey	String	The value of this property depends on whether the parent/child relationship is stored in the parent business object or in the child. If the relationship is stored in the parent, the value includes both the type of the child business object and the name of the attribute in the child to be used as the foreign key *(Child_BO_name/Child_Property_Name)*. If the relationship is stored in the child, set the value to include only the name of the attribute in the parent to be used as the foreign key. If an attribute is not a foreign key, do not include this parameter in the application-specific information.	None
OrderBy	String	If a value is specified and the attribute is in a child business object, the adapter uses the value of the attribute in the ORDER BY clause of retrieval queries. The adapter can retrieve child business objects in either ascending order (ASC) or descending order (DESC). If you do not include this parameter in the application-specific information, the adapter does not specify the retrieval order.	None
PrimaryKey	Boolean	If the column associated with this attribute is a primary key in the corresponding table in the database, PrimaryKey is to `True`,	None

Table 19. Application-specific information for simple attributes (continued)

Parameter	Type	Description	Default value
SPParameterType	String	Specifies the type of stored procedure Possible values are: • IP (input only) • OP (output only) • IO (input and output) • RS (result set)	None
UniqueIdentifier (UID)	String	The adapter uses this parameter to generate the unique ID for the business object. It supports the generation of sequences and identity columns. The format of this parameter is as follows: UID=AUTO|*Sequence_Name* For a sequence, set the UID attribute to the name of the sequence. For an identity column, set the UID attribute to AUTO. If the attribute does not require a unique ID, do not include this parameter in the application-specific information.	None
DateFormat	String	Allows you to customize the format of the Date and Timestamp data types in the **Application Info** section of the **Properties view** in Rational Application Developer for WebSphere Software.	None

The format of attribute application-specific information is shown in the following example section of an XSD file:

Example section of an XSD file

```
<element name="pkey" nillable="true"
                    minOccurs="0" maxOccurs="1">
                    <annotation xml:space="preserve">
             <appinfo

     source="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/
metadata">

                    <oracleasi:OracleAttributeTypeMetadata
     xmlns:oracleasi="http://www.ibm.com/xmlns/prod/websphere/j2ca/
oracle/metadata">

     <oracleasi:ColumnName>PKEY</oracleasi:ColumnName>

     <oracleasi:PrimaryKey>true</oracleasi:PrimaryKey>

     <oracleasi:FixedChar>true</oracleasi:FixedChar>
                    </oracleasi:OracleAttributeTypeMetadata>
             </appinfo>
         </annotation>
         <simpleType>
         <restriction base="string">
         <maxLength value="10"/>
         </restriction>
      </simpleType>
             </element>
             <element name="ccode" type="string" nillable="true"
```

```
                        minOccurs="0" maxOccurs="1">
                        <annotation xml:space="preserve">
<appinfo
source="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/metadata">
<oracleasi:OracleAttributeTypeMetadata
xmlns:oracleasi="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/
metadata">
<oracleasi:ColumnName>CCODE</oracleasi:ColumnName>
<oracleasi:PrimaryKey>false</oracleasi:PrimaryKey>
<oracleasi:ForeignKey>custinfoobj/ccode</oracleasi:ForeignKey>
</oracleasi:OracleAttributeTypeMetadata>
</appinfo>
</annotation>
                    </element>
                    <element name="fname" type="string" nillable="true"
                        minOccurs="0" maxOccurs="1">
                        <annotation xml:space="preserve">
<appinfo
source="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/metadata">
<oracleasi:OracleAttributeTypeMetadata
xmlns:oracleasi="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/
metadata">
<oracleasi:ColumnName>FNAME</oracleasi:ColumnName>
<oracleasi:PrimaryKey>false</oracleasi:PrimaryKey>
</oracleasi:OracleAttributeTypeMetadata>
</appinfo>
</annotation>
                    </element>
                    <element name="lname" type="string" nillable="true"
                        minOccurs="0" maxOccurs="1">
                        <annotation xml:space="preserve">
<appinfo
source="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/metadata">
<oracleasi:OracleAttributeTypeMetadata
xmlns:oracleasi="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/
metadata">
<oracleasi:ColumnName>LNAME</oracleasi:ColumnName>
<oracleasi:PrimaryKey>false</oracleasi:PrimaryKey>
</oracleasi:OracleAttributeTypeMetadata>
</appinfo>
</annotation>
                    </element>
                    <element name="custinfoobj"
                        type="rtassercustinfo:RtasserCustinfo"
nillable="true" minOccurs="0"
                        maxOccurs="1">
                        <annotation xml:space="preserve">
<appinfo
source="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/metadata">
<oracleasi:OracleAttributeTypeMetadata
xmlns:oracleasi="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/
metadata">
<oracleasi:Ownership>true</oracleasi:Ownership>
<oracleasi:KeepRelationship>false</oracleasi:KeepRelationship>
</oracleasi:OracleAttributeTypeMetadata>
</appinfo>
</annotation>
                    </element>
```

Example of FixedChar parameter in the business object XSD file

```
                    <element name="pkey" nillable="true"
                                minOccurs="0" maxOccurs="1">
                                <annotation xml:space="preserve">
                        <appinfo

                source="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/
```

```
                    metadata">
                                    <oracleasi:OracleAttributeTypeMetadata

            xmlns:oracleasi="http://www.ibm.com/xmlns/prod/websphere/j2ca/
oracle/metadata">

            <oracleasi:ColumnName>PKEY</oracleasi:ColumnName>

            <oracleasi:PrimaryKey>true</oracleasi:PrimaryKey>

            <oracleasi:FixedChar>true</oracleasi:FixedChar>
                            </oracleasi:OracleAttributeTypeMetadata>
                    </appinfo>
                </annotation>
                <simpleType>
                <restriction base="string">
                <maxLength value="10"/>
                </restriction>
            </simpleType>
</element>
```

Application-specific information for attributes of type child business object

Two application-specific information parameters are used for attributes that refer to child business objects (complex, as opposed to simple, attributes). When you set this application-specific information, specify the parameters listed in Table 20.

Table 20. Application-specific information for attributes of type child business object

Parameter	Type	Description	Default value
KeepRelationship	Boolean	If True, this parameter prevents the deletion of a child business object during an Update operation.	None
Ownership	Boolean	This parameter specifies that a child business object is owned by the parent. If True, then Create, Update, and Delete operations on the child business object are allowed. If False, then no updates can be applied to the child business object. When its parent is created, the existence of the child is validated to ensure that relationship integrity is maintained in the database.	None

Example of ownership in the business object XSD file

```
<element name="addressobj"
                        type="rtasseraddress:RtasserAddress"
nillable="true"
                        minOccurs="0"
                        maxOccurs="unbounded">
                        <annotation xml:space="preserve">
<appinfo
source="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/metadata">
<oracleasi:OracleAttributeTypeMetadata
xmlns:oracleasi="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/
metadata">
<oracleasi:Ownership>true</oracleasi:Ownership>
</oracleasi:OracleAttributeTypeMetadata>
</appinfo>
</annotation>
</element>

<element name="custinfoobj"
                        type="rtassercustinfo:RtasserCustinfo"
```

```
nillable="true"                                        minOccurs="0"
                          maxOccurs="1">
                          <annotation xml:space="preserve">
<appinfo
source="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/metadata">
<oracleasi:OracleAttributeTypeMetadata
xmlns:oracleasi="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/
metadata">
<oracleasi:Ownership>false</oracleasi:Ownership>
</oracleasi:OracleAttributeTypeMetadata>
</appinfo>
</annotation>
</element>
```

An example of the XSD definition file for single- and multiple-cardinality child
business objects is provided here. The element custInfoObj is a single-cardinality
child business object, and addressObj is a multiple-cardinality child business object.

**Another example XSD file for single- and multiple-cardinality child business
objects**

```
<element name="addressobj"
                        type="rtasseraddress:RtasserAddress"
nillable="true"
                        minOccurs="0"
                        maxOccurs="unbounded">
                        <annotation xml:space="preserve">
<appinfo
source="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/metadata">
<oracleasi:OracleAttributeTypeMetadata
xmlns:oracleasi="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/
metadata">
<oracleasi:Ownership>true</oracleasi:Ownership>
</oracleasi:OracleAttributeTypeMetadata>
</appinfo>
</annotation>
</element>

<element name="custinfoobj"
                        type="rtassercustinfo:RtasserCustinfo"
nillable="true"                                        minOccurs="0"
                        maxOccurs="1">
                        <annotation xml:space="preserve">
<appinfo
source="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/metadata">
<oracleasi:OracleAttributeTypeMetadata
xmlns:oracleasi="http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/
metadata">
<oracleasi:Ownership>false</oracleasi:Ownership>
</oracleasi:OracleAttributeTypeMetadata>
</appinfo>
</annotation>
</element>
```

Application-specific information for operations

The adapter uses application-specific information at the operation level to perform
operations, such as to retrieve and update information in the database. The adapter
retrieves and updates database tables using SQL queries, stored procedures, or
stored functions, as specified in the business objects.

If you choose to add stored procedures or stored functions to the business objects,
set the operation application-specific information (ASI) as specified in Table 21 on
page 125.

Table 21. Operation application-specific information

Operation ASI for StoredProcedure parameters element	Set by wizard	Description
Parameters	Yes	Lists the stored procedure parameters.
PropertyName	Yes	Set to the name of the business object attribute that you select.
ResultSet	No	If the stored procedure returns a result set, set this parameter to True in the business object definition.
ReturnValue	Yes	If the stored procedure has a return value, this parameter contains one of these values: • The string RS. This value indicates that the procedure returns a result set, which is used to create the multiple-cardinality container corresponding to this business object. • The name of the business object attribute. This value indicates that procedure returns the value that is to be assigned to that particular attribute in the business object at run time. If the attribute is another child business object, the adapter returns an error.
StoredProcedure	Yes	Set to the stored procedure name.
StoredProcedureType	Yes	You choose from a list of types. For information about valid stored procedure types, see ."Stored procedure type" on page 29
Type	Yes	Set to the type of the stored procedure parameter. Possible values are: • IP (input only) • OP (output only) • IO (input and output) • RS (result set)

Business object-level application-specific information

Application-specific information in business object definitions provides the adapter with application-dependent instructions for how to process business objects. The adapter parses the application-specific information from the business object or from its attributes or operations to generate queries for Create, Update, Retrieve, and Delete operations.

Application-specific information for table and view business objects

Application-specific information at the business-object level is used to specify the name of the corresponding database table and to provide information necessary to perform a physical or logical Delete operation.

The J2C Bean wizard sets the TableName application-specific information attribute to a value in the form of *SchemaName.TableName*. It prompts you for the information necessary to perform a physical or logical Delete operation and then sets the business object-level application-specific information shown in Table 22 on page 126.

Table 22. Business object application-specific information (ASI) for a table business object

Application-specific information	Type	Description
TableName	String	The name of the database table corresponding to this business object.
StatusColumnName	String	Indicates whether the adapter is to logically or physically delete data in the table. If the StatusColumnName parameter is not set, data is physically deleted. If the parameter is set, it specifies the name of the column that indicates a logically deleted row. You specify this parameter when you select the table object in the J2C Bean wizard. This parameter applies to both Update and Delete operations.
StatusValue	String	The value that indicates that a column is logically deleted. You specify this value when you select the table object in the J2C Bean wizard.

To illustrate how the adapter determines whether to do a logical or physical delete in response to an Update or Delete operation, assume that a Customer business object has the business object application-specific information shown in Table 23.

Table 23. Sample parameters for business object application-specific information for a table business object

Application-specific information	Value
TableName	customer
StatusColumnName	status
StatusValue	deleted

Assume that the adapter receives a request to delete a customer. Because the business object includes the StatusColumnName parameter in its application-specific information, the adapter performs a logical delete operation. It does this by placing the string "deleted", which is specified in the StatusValue parameter, in the status column, which is the column specified in the StatusColumnName parameter.

Such a request causes the adapter to issue the following SQL statement:

```
UPDATE customer set status = 'deleted' where pkey = . . . .
```

However, if the StatusColumnName parameter is not set, the customer records are physically deleted. The adapter issues the following SQL statement:

```
DELETE from customer where pkey = . . . .
```

Application-specific information for stored procedure business objects

For business objects that are based on stored procedures, the J2C Bean wizard sets the business object-level application-specific information SPName to a value in the form of *SchemaName + SPName*. It sets the business object-level application-specific information, which is listed in Table 24 on page 127. The attributes of the business object are created based on the stored procedure input/output parameters. If the stored procedure has one returned value, a corresponding business object attribute

is created. If the returned value or any of the input/output parameters are complex data types, the wizard creates child business objects for them.

The discovery of database objects in the J2C Bean wizard can support nested structures and arrays. If these child business objects are generated from returned result sets, their names are in the form of *Prefix* + *SchemaName* + *SPName* + RetRS + *Number*. For example, if one stored procedure returns two result sets, the wizard creates two child business objects for them. Their names are *Prefix* + *SchemaName* + *SPName* + RetRS1 and *Prefix* + *SchemaName* + *SPName* + RetRS2.

When the child business objects are generated from input/output parameters with a complex data type of ResultSet, Struct, or Array, these child business object names are in the form of *Prefix+SchemaName+SPName+ParameterName*. For those child business objects that correspond to nested structures and arrays, their business object names are in the form of *Prefix+SchemaName+SPName+ParameterName+ColumnName*.

Table 24. Business object application-specific information (ASI) for business objects based on stored procedures

Application-specific information	Type	Description
SPName	String	The name of the stored procedure or stored function
ResultSet	Boolean	Indicates whether the stored procedure or stored function returns a result set. If true, the stored procedure returns one or more result sets. If false, the stored procedure or stored function does not return a result set.
MaxNumberOfRetRS	String	The maximum number of returned result sets that are handled by the adapter run time
ReturnValue	String	Set to the name of the corresponding business object attribute if the stored procedure has a return value. If the returned value is of simple data type, the attribute is also of simple data type. If the returned value is a result set, this attribute points to a child business object.

Application-specific information for query business objects

For query business objects, there is one business object-level application-specific information, as shown in Table 25.

Table 25. Business object application-specific information (ASI) for query business objects

Application-specific information	Type	Description
SelectStatement	String	The complete SELECT statement that performs the query. You specify the statement in the J2C Bean wizard.

When the wizard generates a stored procedure business object, it creates a child business object if necessary, such as for ResultSet, Struct, and Array. Creating parent-child relationships between table business objects is done manually using the Business Object Editor.

The wizard handles business objects based on synonym/nicknames like objects based on tables and views, even when a synonym is of a stored procedure.

Naming conventions

When the J2C Bean wizard generates a business object, it provides a name for the business object that reflects the naming convention for the adapter. Typically, the business object name indicates the structure of the business object.

When the J2C Bean wizard creates names for a business object, it replaces any special character except the underscore (_) in the business object name with U followed by its Unicode number. For example, the business object name for the Order_Item table in the database is Order_Item. The business object name for the Shipping-Address table is ShippingU45Address.

Business object names have no semantic value to the adapter or the database; that is, they derive no information nor meaning from the business object name. If one name is replaced by another, the adapter behavior remains the same.

Business object names can carry database-specific metadata. A name can use a string like `Oracle` or `%AppName%` as a prefix to help distinguish between two types of business objects: application-specific and generic. The remainder of the name can describe the table or stored procedure that the business object represents. For example, if the business object definition is generated for the Employee Table in a database application, such as Human Resources (HR), then the respective business object name is `HREmployee`.

For business objects that do not correspond to database objects, such as business objects for database queries, if you give the business object the same name as a table or stored procedure business object, a different number is appended at the end of each name to differentiate them and avoid overwriting.

Globalized characters are supported in any business object name.

You can rename business objects by using the refactoring functionality in Rational Application Developer for WebSphere Software. For more details, refer to the Rational Application Developer for WebSphere Software documentation.

The following table describes the naming conventions that the wizard uses for the business object.

Table 26. Business object naming conventions

Element	Naming convention
Business objects for: • Tables • Views • Stored procedures • Stored functions • Synonyms and nicknames	For those business objects that are based on tables, views, stored procedures, and synonyms and nicknames, the J2C Bean wizard generates the name of the business object in the form of *Prefix + SchemaName + ObjectName*, where: • *Prefix* is the value as specified in the J2C bean connection property named Prefix. A prefix is not required, and if not specified, no prefix is added to the business object name. • *SchemaName* is the name of the schema to which the object belongs. • *ObjectName* is the name of the table, view, stored procedure, stored function, or synonym/nickname. A number is appended if necessary to differentiate the business object from another business object with the same name. For example, using the prefix Campaign12 for the Customer table in the Sales schema, the business object name is Campaign12SalesCustomer.

Table 26. Business object naming conventions (continued)

Element	Naming convention
Query business objects	For query business objects, the J2C Bean wizard generates the name of the business object in the form of *Prefix + QueryBOName*, where: • *Prefix* is the prefix you specify in the wizard. A prefix is not required, and if not specified, no prefix is added to the business object name. • *QueryBOName* is the value you specified when you configured the business object in the wizard. A number is appended if necessary to differentiate the business object from another business object with the same name.

Configuration properties

IBM WebSphere Adapter for Oracle E-Business Suite has several categories of configuration properties, which you set with the J2C Bean wizard while generating or creating objects and services. You can change the resource adapter, managed connection factory, and activation specification properties after you deploy the application to WebSphere Application Server.

Outbound configuration properties

IBM WebSphere Adapter for Oracle E-Business Suite has several categories of outbound connection configuration properties, which you set with the J2C Bean wizard while generating or creating objects and services. You can change the resource adapter and managed connection factory properties after you deploy the module to WebSphere Application Server using Rational Application Developer for WebSphere Software or the administrative console, but connection properties for the J2C Bean wizard cannot be changed after deployment.

Guide to information about properties

The properties used to configure WebSphere Adapter for Oracle E-Business Suite are described in detail in tables included in each of the configuration properties topics, such as Resource adapter properties, Managed connection factory properties, and so on. To help you use these tables, information about each row you might see is explained here.

The following table explains the meaning of each row that might be displayed in the table for a configuration property.

Row	Explanation
Required	A required field (property) must have a value in order for the adapter to work. Sometimes the J2C Bean wizard provides a default value for required properties. Removing a default value from a required field on the J2C Bean wizard *will not change that default value*. When a required field contains no value at all, the J2C Bean wizard processes the field using its assigned default value, and that default value is displayed on the administrative console. Possible values are **Yes** and **No**. Sometimes a property is required only when another property has a specific value. When this is the case, the table notes this dependency. For example, • Yes, when the EventQueryType property is set to Dynamic • Yes, for Oracle databases
Possible values	Lists and describes the possible values that you can select for the property.

Row	Explanation
Default	The predefined value that is set by the J2C Bean wizard. When the property is required, you must either accept the default value or specify one yourself. If a property has no default value, the table states No default value. The word None is an acceptable default value, and does not mean that there is no default value.
Unit of measure	Specifies how the property is measured, for example in kilobytes or seconds.
Property type	Describes the property type. Valid property types include: • Boolean • String • Integer
Usage	Describes usage conditions or restrictions that might apply to the property. For instance, here is how a restriction would be documented: For Rational Application Developer for WebSphere Software version 6.40 or earlier, the password: • Must be uppercase • Must be 8 characters in length For versions of Rational Application Developer for WebSphere Software later than 6.40, the password: • Is not case-sensitive • Can be up to 40 characters in length. This section lists other properties that affect this property or the properties that are affected by this property and describes the nature of the conditional relationship.
Example	Provides sample property values, for example: "If Language is set to JA (Japanese), code page number is set to 8000".
Globalized	If a property is globalized, it has national language support, meaning that you can set the value in your national language. Valid values are **Yes** and **No**.
Bidi supported	Indicates whether the property is supported in bidirectional (bidi) processing. Bidirectional processing refers to the task of processing data that contains both right-to-left (Hebrew or Arabic, for example) and left-to-right (a URL or file path, for example) semantic content within the same file. Valid values are **Yes** and **No**.

Connection properties for the wizard

J2C Bean wizard connection properties are used to establish a connection between the J2C Bean wizard, a tool that is used to create business objects, and the database.

Resource adapter properties

The resource adapter properties control the general operation of the adapter, such as specifying the namespace for business objects. You set the resource adapter properties using the J2C Bean wizard when you configure the adapter. After deploying the adapter, use the administrative console to change these properties.

The following table lists the resource adapter properties and their purpose. A complete description of each property is provided in the sections that follow the

table. For information about how to read the property details table, see "Guide to information about properties" on page 129.

Table 27. Resource adapter properties for the Adapter for Oracle E-Business Suite

Name		Description
In the wizard	**In the administrative console**	
Adapter ID	AdapterID	Identifies the adapter instance for PMI events and for logging and tracing.
Database vendor	DatabaseVendor	The type of database that the adapter uses for special processing.
Disguise user data as "XXX" in log and trace files	HideConfidentialTrace	Specifies whether to disguise potentially sensitive information by writing X strings instead of user data in the log and trace files.
Query time out	QueryTimeOut	The maximum number of seconds a query can take for all SQL statements.
Return business object even when the stored procedure result set is empty	ReturnDummyBOForSP	Specifies whether to return output parameters when the result set is empty.
(Not available)	enableHASupport	Do not change this property.
(Not available)	LogFileSize	Deprecated
(Not available)	LogFilename	Deprecated
(Not available)	LogNumberOfFiles	Deprecated
SQL query to verify the connection	PingQuery	The SQL query used to test the reliability of the connection to the database.
(Not available)	TraceFileSize	Deprecated
(Not available)	TraceFileName	Deprecated
(Not available)	TraceNumberOfFiles	Deprecated

Adapter ID (AdapterID)

This property identifies a specific deployment or instance of the adapter.

Table 28. Adapter ID details

Required	Yes
Default	001
Property type	String

Table 28. Adapter ID details (continued)

Usage	This property identifies the adapter instance in the log and trace files, and also helps identify the adapter instance while monitoring adapters. The adapter ID is used with an adapter-specific identifier, OEBSRA, to form the component name used by the Log and Trace Analyzer tool. For example, if the adapter ID property is set to 001, the component ID is OEBSRA001. If you run multiple instances of the same adapter, ensure that the first seven characters of the adapter ID property are unique for each instance so that you can correlate the log and trace information to a particular adapter instance. By making the first seven characters of an adapter ID property unique, the component ID for multiple instances of that adapter is also unique, allowing you to correlate the log and trace information to a particular instance of an adapter. For example, when you set the adapter ID property of two instances of WebSphere Adapter for Oracle E-Business Suite to 001 and 002. The component IDs for those instances, OEBSRA001 and OEBSRA002, are short enough to remain unique, enabling you to distinguish them as separate adapter instances. However, instances with longer adapter ID properties cannot be distinguished from each other. If you set the adapter ID properties of two instances to Instance01 and Instance02, you will not be able to examine the log and trace information for each adapter instance because the component ID for both instances is truncated to OEBSRAInstanc. For inbound processing, the value of this property is set at the resource adapter level. For outbound processing, the value can be set both at the resource adapter level and the managed connection factory level. After you use the J2C Bean wizard to configure the adapter for outbound processing, you can set the resource adapter and managed connection factory properties independently.
Globalized	Yes
Bidi supported	No

Database vendor (DatabaseVendor)

This property specifies the type of database that is used. The type is determined by the database vendor name.

Table 29. Database vendor details

Required	Yes
Possible values	Oracle
Default	ORACLE
Property type	String
Usage	Some SQL statements require special processing, which varies by database type. For example, the Struct and Array data types in Oracle require special processing. This property specifies the RDBMS that is used, which determines the database type. For other databases, the adapter does not perform any special processing. Make sure that the correct driver is specified in the JDBCDriverClass property. Specify Oracle as the value that corresponds to your database vendor.
Globalized	No
Bidi supported	No

Disguise user data as "XXX" in log and trace files (HideConfidentialTrace)

This property specifies whether to replace user data in log and trace files with a string of X's to prevent unauthorized disclosure of potentially sensitive data.

Table 30. Disguise user data as "XXX" in log and trace files details

Required	No
Possible values	True False
Default	False
Property type	Boolean
Usage	If you set this property to True, the adapter replaces user data with a string of X's when writing to log and trace files. For inbound processing, the value of this property is set at the resource adapter level. For outbound processing, the value can be set both at the resource adapter level and the managed connection factory level. After you use the J2C Bean wizard to configure the adapter for outbound processing, you can set the resource adapter and managed connection factory properties independently.
Globalized	No
Bidi supported	No

Enable high availability support (enableHASupport)

Do not change this property. It must be set to true.

Query timeout (seconds) (QueryTimeOut)

This property specifies the maximum number of seconds a query can take to run all SQL statements.

Table 31. Query timeout details

Required	No
Default	No default value
Unit of measure	Seconds
Property type	Integer
Usage	If the query takes longer than the number of seconds specified, the database generates an SQL exception that is captured. The associated message is logged in the log file. If a value is not specified, no timeout is set on the query.
Globalized	No
Bidi supported	No

Return business object even when the stored procedure result set is empty (ReturnDummyBOForSP)

This property specifies whether to return output parameters when the result set is empty.

Table 32. Return business object even when the stored procedure result set is empty details

Required	No
Possible values	True False
Default	False
Property type	Boolean
Usage	The Retrieve Stored Procedure (RetrieveSP) operation returns a result set. If the result set is empty and the ReturnDummyBOForSP property is set to False, no business objects are created, and the output parameters returned by the procedure call cannot be retrieved. However, if the ReturnDummyBOForSP property is set to True, a dummy business object is created and the returned parameters (including the out and inout parameters) from the stored procedure are populated into the corresponding attributes
Globalized	No
Bidi supported	No

SQL query to verify the connection (PingQuery)

This property specifies the SQL query that is used to test the reliability of the connection to the database.

Table 33. Ping query details

Required	No
Property type	String
Default	No default value
Usage	This property contains the SQL query statement that you want to run to determine whether the adapter can connect to the database. The adapter runs the ping query every time it receives a SQLException exception while performing the outbound operation. The adapter does not try to recover the connection. If the ping query indicates that the connection to the database is no longer valid, the adapter notifies the container. It is the responsibility of the connection pool manager to remove the stale connection from the pool, which allows subsequent outbound requests to be processed.
Globalized	No
Bidi supported	No

Managed connection factory properties
Managed connection factory properties are used by the adapter at run time to create an outbound connection instance with the Oracle E-Business Suite.

You set managed connection factory properties using the J2C Bean wizard during adapter configuration. You can change them using the Rational Application

Developer for WebSphere Software editor or, after deployment, with the
WebSphere Application Server administrative console.

The following table lists and describes the managed connection factory properties.
A complete description of each property is provided in the sections that follow the
table. For information about how to read the property detail tables in the sections
that follow, see "Guide to information about properties" on page 129.

Note: The J2C Bean wizard refers to these properties as managed connection
factory properties, while the administrative console refers to them as J2C
connection factory properties.

Table 34. Managed connection factory properties for Adapter for Oracle E-Business Suite

Property name		
In the wizard	In the administrative console	Description
Adapter ID	AdapterID	Identifies the adapter instance for PMI events and for logging and tracing.
Additional JDBC driver connection properties [name:value;name:value]	JDBCDriverConnectionProperties	Additional properties for connecting to the database using the JDBC driver, which are used in addition to the UserName and Password properties.
Auto commit	AutoCommit	The AutoCommit value to use on the connection
DataSource JNDI name	DataSourceJNDIName	This property is deprecated. Use the XADataSourceJNDIName or PoolDataSourceJNDIName properties in place of this
Database URL	DatabaseURL	The database URL that is used to connect to the database.
Database vendor	DatabaseVendor	The type of database that the adapter uses for special processing.
Disguise user data as "XXX" in log and trace files	HideConfidentialTrace	Specifies whether to disguise potentially sensitive information by writing X strings instead of user data in the log and trace files.
JDBC driver class	JDBCDriverClass	The class name of the JDBC driver that is used to connect to the database.
Password	Password	Password for the corresponding user name
Query timeout	QueryTimeOut	The maximum number of seconds a query can take for all SQL statements.
Return business object even when the stored procedure result set is empty	ReturnDummyBOForSP	Specifies whether to return output parameters when the result set is empty.
SQL query to connection	PingQuery	The SQL query used to test the reliability of the connection to the database.
User name	UserName	The database user name.
"XA DataSource name (XADataSourceName)" on page 143	XADataSourceName	The name of the XA data source to use to establish a connection to the database for XA (distributed) transactions

Property name		
In the wizard	In the administrative console	Description
"Database connection information (ConnectionType)" on page 145	ConnectionType	Specifies how the adapter establishes connection to the database.
"Connection pool DataSource JNDI name (PoolDataSourceJNDIName)" on page 146	PoolDataSourceJNDIName	The JNDI name of the connection pool data source to be used to establish a connection to the database.
"XA DataSource JNDI name (XADataSourceJNDIName)" on page 148	XADataSourceJNDIName	The JNDI name of the XA data source used to establish a connection to the database.

Adapter ID (AdapterID)

This property identifies a specific deployment or instance of the adapter.

Table 35. Adapter ID details

Required	Yes
Default	001
Property type	String
Usage	This property identifies the adapter instance in the log and trace files, and also helps identify the adapter instance while monitoring adapters. The adapter ID is used with an adapter-specific identifier, OEBSRA, to form the component name used by the Log and Trace Analyzer tool. For example, if the adapter ID property is set to 001, the component ID is OEBSRA001.

If you run multiple instances of the same adapter, ensure that the first seven characters of the adapter ID property are unique for each instance so that you can correlate the log and trace information to a particular adapter instance. By making the first seven characters of an adapter ID property unique, the component ID for multiple instances of that adapter is also unique, allowing you to correlate the log and trace information to a particular instance of an adapter.

For example, when you set the adapter ID property of two instances of WebSphere Adapter for Oracle E-Business Suite to 001 and 002. The component IDs for those instances, OEBSRA001 and OEBSRA002, are short enough to remain unique, enabling you to distinguish them as separate adapter instances. However, instances with longer adapter ID properties cannot be distinguished from each other. If you set the adapter ID properties of two instances to Instance01 and Instance02, you will not be able to examine the log and trace information for each adapter instance because the component ID for both instances is truncated to OEBSRAInstanc.

For inbound processing, the value of this property is set at the resource adapter level. For outbound processing, the value can be set both at the resource adapter level and the managed connection factory level. After you use the J2C Bean wizard to configure the adapter for outbound processing, you can set the resource adapter and managed connection factory properties independently. |
| Globalized | Yes |
| Bidi supported | No |

Additional JDBC driver connection properties [name:value;name:value] (JDBCDriverConnectionProperties)

This property contains additional information for connecting to the database using the JDBC driver.

Table 36. Additional JDBC driver connection properties details

Row	Explanation
Required	No
Possible values	Database connection properties are database-specific.
Default	No default value
Property type	String
Usage	These connection properties are used in addition to the UserName and Password properties to customize the database connection used by the adapter. Specify the connection properties as one or more *name:value* pairs separated by the semicolon character (;).
Example	The following value of this property specifies a login timeout interval: `loginTimeout:20; ConnectionRetryCount:5; ConnectionRetryDelay:5`
Globalized	Yes
Bidi supported	No

Auto commit (AutoCommit)

This property specifies whether AutoCommit is set for the connection.

Table 37. Auto commit details

Required	No
Possible values	`True` `False`
Default	`False`
Property type	Boolean
Usage	This property is ignored for XA (distributed) transactions.
Globalized	No
Bidi supported	No

Database URL (DatabaseURL)

This property specifies the JDBC driver-specific URL for creating a connection to the database.

Table 38. Database URL details

Required	Yes, when the connection is of the type LocalConnectionProps or ConnectionProps.
Default	No default value
Property type	String

Table 38. Database URL details (continued)

Usage	In the J2C Bean wizard, compose the database URL by entering in database-specific fields. For example, the database URL for an Oracle database is composed of the system ID (SID), server host name, and database port number. In the administrative console, type the complete database URL value.
	If your database server supports IPv6, you can specify the host name portion of the database URL in IPv6 format.
	If you specify the host name as an IP address in IPv6 format, enclose the IP address in square brackets ([]).
Examples	The following example illustrates typical database URL values for a common database:`jdbc:oracle:thin:@9.26.248.148:1521:dev`
Globalized	Yes
Bidi supported	Yes

Database vendor (DatabaseVendor)

This property specifies the type of database that is used. The type is determined by the database vendor name.

Table 39. Database vendor details

Required	Yes
Possible values	`Oracle`
Default	`ORACLE`
Property type	String
Usage	Some SQL statements require special processing, which varies by database type. For example, the Struct and Array data types in Oracle require special processing. This property specifies the RDBMS that is used, which determines the database type.
	For other databases, the adapter does not perform any special processing. Make sure that the correct driver is specified in the JDBCDriverClass property.
	Specify `Oracle` as the value that corresponds to your database vendor.
Globalized	No
Bidi supported	No

Disguise user data as "XXX" in log and trace files (HideConfidentialTrace)

This property specifies whether to replace user data in log and trace files with a string of X's to prevent unauthorized disclosure of potentially sensitive data.

Table 40. Disguise user data as "XXX" in log and trace files details

Required	No
Possible values	True False

Default	False
Property type	Boolean
Usage	If you set this property to True, the adapter replaces user data with a string of X's when writing to log and trace files. For inbound processing, the value of this property is set at the resource adapter level. For outbound processing, the value can be set both at the resource adapter level and the managed connection factory level. After you use the J2C Bean wizard to configure the adapter for outbound processing, you can set the resource adapter and managed connection factory properties independently.
Globalized	No
Bidi supported	No

Throw exception when no records are found (ErrorOnEmptyResultSet)

This property specifies whether to return an empty business object when no records are found.

Table 41. Throw exception when no records are found details

Required	No
Possible values	True False
Default	True
Property type	Boolean
Usage	If you set this property to False, the RetrieveAll operation does not generate a RecordNotFoundException exception when no records are found.
Globalized	No
Bidi supported	No

JDBC driver class (JDBCDriverClass)

This property specifies the class name of the JDBC driver that is used to connect to the database.

Table 42. JDBC driver class details

Row	Explanation
Required	Yes, when the connection is of the type LocalConnectionProps or ConnectionProps.
Possible values	The value specified when you set the connection properties to connect to the database instance.
Default	No default value
Property type	String

Table 42. JDBC driver class details (continued)

Usage	In the J2C Bean wizard, the JDBC driver class is displayed if you select a common database software and driver combination, such as type 4 drivers for recent versions of Oracle. For type 2 drivers , you must type the database class name. If you select a type 2 driver or a generic driver, you must type the JDBC driver class name. In the administrative console, type the database-specific name of the driver.
Examples	Values for the JDBC driver class is displayed in both the J2C Bean wizard and the administrative console. The following examples depict the JDBC driver class properties in both the J2C Bean wizard and the administrative console.**In the J2C Bean wizard:** • To connect to an Oracle 10 database using the type 4 driver, select `Oracle Thin Driver`. **In the administrative console:** **Oracle Thin JDBC driver** `oracle.jdbc.driver.OracleDriver`
Globalized	No
Bidi supported	No

Password (Password)

This property specifies the password for the database user name.

Table 43. Password details

Required	No. For inbound processing, if you set the Authentication alias or the DataSourceJNDIName, the password is not mandatory. However, if you set the DataSourceJNDIName, and the Password field, the value specified for the Password takes precedence. For outbound processing, if you set the Authentication alias, XADataSourceJNDIName or PoolDataSourceJNDIName property, the password is not mandatory. However, if you set the XADataSourceJNDIName or PoolDataSourceJNDIName property, and the Password field, the value specified for the Password takes precedence.
Default	No default value
Property type	String
Usage	For inbound processing, setting this property overrides the password specified for the data source on the server using the Authentication alias or DataSourceJNDIName property. For outbound processing, setting this property overrides the password specified for the data source on the server using the Authentication alias, or XADataSourceJNDIName or PoolDataSourceJNDIName property. If you specify JAAS as the security credential, the authentication alias will override this property.

Table 43. Password details (continued)

Globalized	Yes
Bidi supported	Yes

Query timeout (seconds) (QueryTimeOut)

This property specifies the maximum number of seconds a query can take to run all SQL statements.

Table 44. Query timeout details

Required	No
Default	No default value
Unit of measure	Seconds
Property type	Integer
Usage	If the query takes longer than the number of seconds specified, the database generates an SQL exception that is captured. The associated message is logged in the log file. If a value is not specified, no timeout is set on the query.
Globalized	No
Bidi supported	No

Return business object even when the stored procedure result set is empty (ReturnDummyBOForSP)

This property specifies whether to return output parameters when the result set is empty.

Table 45. Return business object even when the stored procedure result set is empty details

Required	No
Possible values	True False
Default	False
Property type	Boolean
Usage	The Retrieve Stored Procedure (RetrieveSP) operation returns a result set. If the result set is empty and the ReturnDummyBOForSP property is set to False, no business objects are created, and the output parameters returned by the procedure call cannot be retrieved. However, if the ReturnDummyBOForSP property is set to True, a dummy business object is created and the returned parameters (including the out and inout parameters) from the stored procedure are populated into the corresponding attributes
Globalized	No
Bidi supported	No

SQL query to verify the connection (PingQuery)

This property specifies the SQL query that is used to test the reliability of the connection to the database.

Table 46. Ping query details

Required	No
Property type	String
Default	No default value
Usage	This property contains the SQL query statement that you want to run to determine whether the adapter can connect to the database.
	The adapter runs the ping query every time it receives a SQLException exception while performing the outbound operation.
	The adapter does not try to recover the connection. If the ping query indicates that the connection to the database is no longer valid, the adapter notifies the container. It is the responsibility of the connection pool manager to remove the stale connection from the pool, which allows subsequent outbound requests to be processed.
Globalized	No
Bidi supported	No

User name (UserName)

This property specifies the database user name that is used to access the database.

Table 47. User name details

Required	No. For inbound processing, if you set the Authentication alias or the DataSourceJNDIName, the User name property is not mandatory. However, if you set the DataSourceJNDIName, and User name field, the value specified for the User name takes precedence.
	For outbound processing, if you set the Authentication alias, XADataSourceJNDIName or PoolDataSourceJNDIName property, the User name is not mandatory. However, if you set the XADataSourceJNDIName or PoolDataSourceJNDIName property, and the User name field, the value specified for the User name takes precedence.
Default	No default value
Property type	String
Usage	For inbound processing setting this property overrides the user name specified for the data source on the server using the DataSourceJNDIName property or the Authentication alias.
	For outbound processing, setting this property overrides the user name specified for the data source on the server using the XADataSourceJNDIName property, PoolDataSourceJNDIName property or the Authentication alias.
	If you specify JAAS as the security credential, the authentication alias will override this property.

Table 47. User name details (continued)

Globalized	Yes
Bidi supported	Yes

XA DataSource name (XADataSourceName)

This property specifies the name of the XA data source used to establish a connection to the database for XA (distributed) transactions.

Table 48. XA data source name details

Required	No
Default	No default value
Property type	String
Usage	This property is used to make an XA connection to an Oracle database.
Examples	A typical value for an Oracle database: `oracle.jdbc.xa.client.OracleXADataSource`
Globalized	No
Bidi supported	No

Maximum retries on connection failure (connectionRetryLimit)

This property specifies the number of times the adapter attempts to reestablish a connection to the EIS, when the adapter encounters an error related to the outbound connection.

Table 49. Maximum retries in case of system connection failure details

Required	No
Possible values	Integer equal to and greater than zero
Default	0
Property type	Integer

Usage	When the adapter encounters an error related to the outbound connection, it retries to establish a physical connection (when physical connection is not established) for the number of times specified for this property with a time delay specified in the property "Connection Retry interval (in milliseconds) (connectionRetryInterval)." If the value is set to 0, the adapter does not validate the connection to the EIS and it executes the outbound operation. If the EIS connection is invalid, the outbound operation fails. The adapter does not attempt to reconnect to the EIS. If the value is set to greater than 0, then during each request the adapter validates if the connection to the EIS is active. • If the connection is valid, the operation is completed. • If connection is invalid, the adapter terminates the current managed connection and a new managed connection is created (new physical connection). If the adapter is successful in reestablishing the connection to the EIS, the outbound operation is completed otherwise the adapter generates the ResourceException after trying to reconnect for the specified number of times.
Globalized	No
Bidi supported	No

Connection Retry interval (in milliseconds) (connectionRetryInterval)

This property specifies the time interval between attempts to reconnect to the EIS if the connection fails.

Table 50. Retry interval if connection failure details

Required	No
Possible values	Integer equal to and greater than zero
Default	60000
Unit of measure	Milliseconds
Property type	Integer
Usage	When the adapter encounters an error while establishing a connection to the EIS, this property specifies the time interval the adapter waits between attempts to reestablish a connection. By default, this property disabled and it is enabled only when the value of the "Maximum retries on connection failure (connectionRetryLimit)" on page 143property is set to greater than zero.
Globalized	No
Bidi supported	No

Database connection information (ConnectionType)

This property specifies how the adapter establishes connection to the database.

Table 51. Database connection information

Required	Yes
Possible values	XADataSourceJNDI, XAConnectionProps, PoolDataSourceJNDI, or LocalConnectionProps
Default	No default value
Property type	String

Table 51. Database connection information (continued)

Usage	This property specifies how the adapter establishes database connection at run time. This property can have the following values:
	• XADataSourceJNDI- Indicates that the database connection is established using the XADataSourceJNDIName property that corresponds to the predefined XA data source.
	• XAConnectionProps - Indicates that the database connection is established using the XADataSourceName and DatabaseURL or XADatabaseName properties.
	• PoolDataSourceJNDI - Indicates that the database connection is established using the poolDataSourceJNDIName property that corresponds to the predefined data source.
	• LocalConnectionProps - Indicates that the database connection is established using the DatabaseURL and JDBCDriverClass properties.
	For a new application, this property is automatically set by J2C Bean wizard.
	• If the DataSourceJNDIName property is set, the value of this property is set to XADataSourceJNDI.
	• If the XADataSourceName is set, the value of this property is set to XAConnectionProps.
	• Otherwise, the value of this property is set to LocalConnectionProps.
	If this property is not set, adapter uses backward compatibility mode to establish database connection. In the backward compatibility mode, the properties for connecting to the database are used in the following order:
	1. If the DataSourceJNDIName property is set, the adapter uses it to establish the connection to the database.
	2. If the DataSourceJNDIName property is not set, and the XADataSourceName and XADatabaseName properties are set, the adapter uses them to establish the connection. The DataSourceJNDIName property represents an XA or connection pool data source. If you define a JNDI data source on the server that supports XA transactions and then specify that data source when you configure the adapter, you can connect to any type of database that supports XA transactions. If you use an XA data source and database, the adapter supports XA transaction only for DB2® and Oracle databases.
	3. If the DataSourceJNDIName, XADataSourceName, and XADatabaseName properties are not set, then the adapter uses the DatabaseURL, JDBCDriverClass, UserName, and Password properties to establish the connection.
Globalized	No
Bidi supported	No

Connection pool DataSource JNDI name (PoolDataSourceJNDIName)

The JNDI name of the connection pool data source to be used to establish a connection to the database.

Table 52. Connection pool data source

Required	No
Default	No default value
Property type	String
Usage	Use this property to specify the JNDI name of a connection pool data source in WebSphere Application Server that specifies connection information for the target database. If the "Database connection information (ConnectionType)" on page 145 property is set to "PoolDataSourceJNDI" the adapter use this property to establish the connection to the database. To improve the performance of outbound operations, specify the name of a data source that is enabled for prepared statement caching. If the other valid authentication properties are also set, they override authentication properties in the data source.
Globalized	No
Bidi supported	No

JDBC driver class (JDBCDriverClass)

This property specifies the class name of the JDBC driver that is used to connect to the database.

Table 53. JDBC driver class details

Row	Explanation
Required	Yes, when the connection is of the type LocalConnectionProps or ConnectionProps.
Possible values	The value specified when you set the connection properties to connect to the database instance.
Default	No default value
Property type	String
Usage	In the J2C Bean wizard, the JDBC driver class is displayed if you select a common database software and driver combination, such as type 4 drivers for recent versions of Oracle. For type 2 drivers , you must type the database class name. If you select a type 2 driver or a generic driver, you must type the JDBC driver class name. In the administrative console, type the database-specific name of the driver.
Examples	Values for the JDBC driver class is displayed in both the J2C Bean wizard and the administrative console. The following examples depict the JDBC driver class properties in both the J2C Bean wizard and the administrative console.**In the J2C Bean wizard:** • To connect to an Oracle 10 database using the type 4 driver, select `Oracle Thin Driver`. **In the administrative console:** **Oracle Thin JDBC driver** `oracle.jdbc.driver.OracleDriver`

Table 53. JDBC driver class details (continued)

Globalized	No
Bidi supported	No

XA DataSource JNDI name (XADataSourceJNDIName)

This property specifies the JNDI name of the XA data source used to establish a connection to the database.

Table 54. XA Data source JNDI name details

Required	No
Default	No default value
Property type	String
Usage	Use this property to specify the JNDI name of a XA data source in WebSphere Application Server that specifies connection information for the target database. If the "Database connection information (ConnectionType)" on page 145 is set to "XADataSourceJNDI" the adapter uses this property to establish the connection to the database. To improve the performance of outbound operations, specify the name of a data source that is enabled for prepared statement caching. If the other valid authentication properties are also set, they override authentication properties in the data source.
Globalized	Yes
Bidi supported	No

Password (Password)

This property specifies the password for the database user name.

Table 55. Password details

| Required | No. For inbound processing, if you set the Authentication alias or the DataSourceJNDIName, the password is not mandatory. However, if you set the DataSourceJNDIName, and the Password field, the value specified for the Password takes precedence.

For outbound processing, if you set the Authentication alias, XADataSourceJNDIName or PoolDataSourceJNDIName property, the password is not mandatory. However, if you set the XADataSourceJNDIName or PoolDataSourceJNDIName property, and the Password field, the value specified for the Password takes precedence. |
|---|---|
| Default | No default value |
| Property type | String |

Table 55. Password details (continued)

Usage	For inbound processing, setting this property overrides the password specified for the data source on the server using the Authentication alias or DataSourceJNDIName property.
	For outbound processing, setting this property overrides the password specified for the data source on the server using the Authentication alias, or XADataSourceJNDIName or PoolDataSourceJNDIName property.
	If you specify JAAS as the security credential, the authentication alias will override this property.
Globalized	Yes
Bidi supported	Yes

Interaction specification properties

Interaction specification, or InteractionSpec, properties control the interaction for an operation. The J2C Bean wizard sets the interaction specification properties when you configure the adapter. Typically, you do not need to change these properties. However, some properties for outbound operations can be changed by the user. For example, you might increase the value of the interaction specification property that specifies the maximum number of records to be returned by a RetrieveAll operation, if your RetrieveAll operations do not return complete information. To change these properties after the application is deployed, use the editor in Rational Application Developer for WebSphere Software.

Table 56 lists and describes the interaction specification property that you set. For information about how to read the property detail tables in the sections that follow, see "Guide to information about properties" on page 129.

Table 56. Interaction specification property for the Adapter for Oracle E-Business Suite

Property name	Description
"Maximum records for RetrieveAll operation"	Maximum number of result sets to return during a RetrieveAll operation

Maximum records for RetrieveAll operation

This property specifies the maximum number of records to return for a RetrieveAll operation.

Table 57. Maximum records for RetrieveAll operation details

Required	Yes
Default	100
Property type	Integer
Globalized	No
Bidi supported	No

Inbound configuration properties

WebSphere Adapter for Oracle E-Business Suite has several categories of inbound connection configuration properties, which you set with the J2C Bean wizard while generating or creating objects and services. You can change the resource adapter and activation specification properties after you deploy the module using Rational

Application Developer for WebSphere Software or the administrative console, but connection properties for the J2C Bean wizard cannot be changed after deployment.

Guide to information about properties

The properties used to configure WebSphere Adapter for Oracle E-Business Suite are described in detail in tables included in each of the configuration properties topics, such as Resource adapter properties, Managed connection factory properties, and so on. To help you use these tables, information about each row you might see is explained here.

The following table explains the meaning of each row that might be displayed in the table for a configuration property.

Row	Explanation
Required	A required field (property) must have a value in order for the adapter to work. Sometimes the J2C Bean wizard provides a default value for required properties.
	Removing a default value from a required field on the J2C Bean wizard *will not change that default value*. When a required field contains no value at all, the J2C Bean wizard processes the field using its assigned default value, and that default value is displayed on the administrative console.
	Possible values are **Yes** and **No**.
	Sometimes a property is required only when another property has a specific value. When this is the case, the table notes this dependency. For example,
	• Yes, when the EventQueryType property is set to Dynamic
	• Yes, for Oracle databases
Possible values	Lists and describes the possible values that you can select for the property.
Default	The predefined value that is set by the J2C Bean wizard. When the property is required, you must either accept the default value or specify one yourself. If a property has no default value, the table states No default value.
	The word None is an acceptable default value, and does not mean that there is no default value.
Unit of measure	Specifies how the property is measured, for example in kilobytes or seconds.
Property type	Describes the property type. Valid property types include:
	• Boolean
	• String
	• Integer

Row	Explanation
Usage	Describes usage conditions or restrictions that might apply to the property. For instance, here is how a restriction would be documented: For Rational Application Developer for WebSphere Software version 6.40 or earlier, the password: • Must be uppercase • Must be 8 characters in length For versions of Rational Application Developer for WebSphere Software later than 6.40, the password: • Is not case-sensitive • Can be up to 40 characters in length. This section lists other properties that affect this property or the properties that are affected by this property and describes the nature of the conditional relationship.
Example	Provides sample property values, for example: "If Language is set to JA (Japanese), code page number is set to 8000".
Globalized	If a property is globalized, it has national language support, meaning that you can set the value in your national language. Valid values are **Yes** and **No**.
Bidi supported	Indicates whether the property is supported in bidirectional (bidi) processing. Bidirectional processing refers to the task of processing data that contains both right-to-left (Hebrew or Arabic, for example) and left-to-right (a URL or file path, for example) semantic content within the same file. Valid values are **Yes** and **No**.

Connection properties for the wizard

J2C Bean wizard connection properties are used to establish a connection between the J2C Bean wizard, a tool that is used to create business objects, and the database.

Resource adapter properties

The resource adapter properties control the general operation of the adapter, such as specifying the namespace for business objects. You set the resource adapter properties using the J2C Bean wizard when you configure the adapter. After deploying the adapter, use the administrative console to change these properties.

The following table lists the resource adapter properties and their purpose. A complete description of each property is provided in the sections that follow the table. For information about how to read the property details table, see "Guide to information about properties" on page 129.

Table 58. Resource adapter properties for the Adapter for Oracle E-Business Suite

Name		Description
In the wizard	**In the administrative console**	
Adapter ID	AdapterID	Identifies the adapter instance for PMI events and for logging and tracing.
Database vendor	DatabaseVendor	The type of database that the adapter uses for special processing.

Name		Description
In the wizard	**In the administrative console**	
Disguise user data as "XXX" in log and trace files	HideConfidentialTrace	Specifies whether to disguise potentially sensitive information by writing X strings instead of user data in the log and trace files.
Query time out	QueryTimeOut	The maximum number of seconds a query can take for all SQL statements.
Return business object even when the stored procedure result set is empty	ReturnDummyBOForSP	Specifies whether to return output parameters when the result set is empty.
(Not available)	enableHASupport	Do not change this property.
(Not available)	LogFileSize	Deprecated
(Not available)	LogFilename	Deprecated
(Not available)	LogNumberOfFiles	Deprecated
SQL query to verify the connection	PingQuery	The SQL query used to test the reliability of the connection to the database.
(Not available)	TraceFileSize	Deprecated
(Not available)	TraceFileName	Deprecated
(Not available)	TraceNumberOfFiles	Deprecated

Adapter ID (AdapterID)

This property identifies a specific deployment or instance of the adapter.

Table 59. Adapter ID details

Required	Yes
Default	001
Property type	String

Table 59. Adapter ID details (continued)

Usage	This property identifies the adapter instance in the log and trace files, and also helps identify the adapter instance while monitoring adapters. The adapter ID is used with an adapter-specific identifier, `OEBSRA`, to form the component name used by the Log and Trace Analyzer tool. For example, if the adapter ID property is set to `001`, the component ID is `OEBSRA001`.
	If you run multiple instances of the same adapter, ensure that the first seven characters of the adapter ID property are unique for each instance so that you can correlate the log and trace information to a particular adapter instance. By making the first seven characters of an adapter ID property unique, the component ID for multiple instances of that adapter is also unique, allowing you to correlate the log and trace information to a particular instance of an adapter.
	For example, when you set the adapter ID property of two instances of WebSphere Adapter for Oracle E-Business Suite to `001` and `002`. The component IDs for those instances, `OEBSRA001` and `OEBSRA002`, are short enough to remain unique, enabling you to distinguish them as separate adapter instances. However, instances with longer adapter ID properties cannot be distinguished from each other. If you set the adapter ID properties of two instances to `Instance01` and `Instance02`, you will not be able to examine the log and trace information for each adapter instance because the component ID for both instances is truncated to `OEBSRAInstanc`.
	For inbound processing, the value of this property is set at the resource adapter level. For outbound processing, the value can be set both at the resource adapter level and the managed connection factory level. After you use the J2C Bean wizard to configure the adapter for outbound processing, you can set the resource adapter and managed connection factory properties independently.
Globalized	Yes
Bidi supported	No

Database vendor (DatabaseVendor)

This property specifies the type of database that is used. The type is determined by the database vendor name.

Table 60. Database vendor details

Required	Yes
Possible values	`Oracle`
Default	`ORACLE`
Property type	String
Usage	Some SQL statements require special processing, which varies by database type. For example, the Struct and Array data types in Oracle require special processing. This property specifies the RDBMS that is used, which determines the database type.
	For other databases, the adapter does not perform any special processing. Make sure that the correct driver is specified in the JDBCDriverClass property.
	Specify `Oracle` as the value that corresponds to your database vendor.
Globalized	No
Bidi supported	No

Disguise user data as "XXX" in log and trace files (HideConfidentialTrace)

This property specifies whether to replace user data in log and trace files with a string of X's to prevent unauthorized disclosure of potentially sensitive data.

Table 61. Disguise user data as "XXX" in log and trace files details

Required	No
Possible values	True False
Default	False
Property type	Boolean
Usage	If you set this property to True, the adapter replaces user data with a string of X's when writing to log and trace files. For inbound processing, the value of this property is set at the resource adapter level. For outbound processing, the value can be set both at the resource adapter level and the managed connection factory level. After you use the J2C Bean wizard to configure the adapter for outbound processing, you can set the resource adapter and managed connection factory properties independently.
Globalized	No
Bidi supported	No

Enable high availability support (enableHASupport)

Do not change this property. It must be set to true.

Query timeout (seconds) (QueryTimeOut)

This property specifies the maximum number of seconds a query can take to run all SQL statements.

Table 62. Query timeout details

Required	No
Default	No default value
Unit of measure	Seconds
Property type	Integer
Usage	If the query takes longer than the number of seconds specified, the database generates an SQL exception that is captured. The associated message is logged in the log file. If a value is not specified, no timeout is set on the query.
Globalized	No
Bidi supported	No

Return business object even when the stored procedure result set is empty (ReturnDummyBOForSP)

This property specifies whether to return output parameters when the result set is empty.

Table 63. Return business object even when the stored procedure result set is empty details

Required	No
Possible values	True False
Default	False
Property type	Boolean
Usage	The Retrieve Stored Procedure (RetrieveSP) operation returns a result set. If the result set is empty and the ReturnDummyBOForSP property is set to False, no business objects are created, and the output parameters returned by the procedure call cannot be retrieved. However, if the ReturnDummyBOForSP property is set to True, a dummy business object is created and the returned parameters (including the out and inout parameters) from the stored procedure are populated into the corresponding attributes
Globalized	No
Bidi supported	No

SQL query to verify the connection (PingQuery)

This property specifies the SQL query that is used to test the reliability of the connection to the database.

Table 64. Ping query details

Required	No
Property type	String
Default	No default value
Usage	This property contains the SQL query statement that you want to run to determine whether the adapter can connect to the database. The adapter runs the ping query every time it receives a SQLException exception while performing the outbound operation. The adapter does not try to recover the connection. If the ping query indicates that the connection to the database is no longer valid, the adapter notifies the container. It is the responsibility of the connection pool manager to remove the stale connection from the pool, which allows subsequent outbound requests to be processed.
Globalized	No
Bidi supported	No

Activation specification properties

Activation specification properties are properties that hold the inbound event processing configuration information for an export.

You set activation specification properties using the J2C Bean wizard during adapter configuration and can change them using the Rational Application

Developer for WebSphere Software editor or, after deployment, the WebSphere Application Server administrative console.

The following table lists and describes the activation specification properties. A complete description of each property is provided in the sections that follow the table. For information about how to read the property detail tables in the sections that follow, see "Guide to information about properties" on page 129.

Table 65. Activation specification properties for Adapter for Oracle E-Business Suite

Property name		
In the wizard	**In the administrative console**	**Description**
Adapter Instance for event filtering	AdapterInstanceEventFilter	Identifier that determines whether this adapter instance processes specific events in the event store.
Additional JDBC driver connection properties [name:value;name:value]	JDBCDriverConnectionProperties	Additional properties for connecting to the database using the JDBC driver, which are used in addition to the UserName and Password properties.
Business object namespace	BusinessObjectNameSpace	The namespace for the business object definitions
User-defined delete query	CustomDeleteQuery	The name of the query, stored procedure, or stored function that is run after each event is processed to delete records that can be deleted after the event is delivered.
User-defined event query	CustomEventQuery	The name of the query, stored procedure, or stored function that performs the polling for events.
User-defined update query	CustomUpdateQuery	The name of the query, stored procedure, or stored function that is run after each event is processed to prevent the event from being picked up for processing in a subsequent event cycle.
User-defined update query for failed event delivery	CustomUpdateQueryForFailedEvent	The name of the query, stored procedure, or stored function that is run when an event is not delivered successfully.
Data source JNDI name	DataSourceJNDIName	The name of the JNDI data source used to establish a connection to the database.
Database URL	DatabaseURL	The database URL that is used to connect to the database.
Database vendor	DatabaseVendor	The type of database that the adapter uses for special processing.

Table 65. Activation specification properties for Adapter for Oracle E-Business Suite (continued)

Property name		
In the wizard	**In the administrative console**	**Description**
Do not process events that have a timestamp in the future	FilterFutureEvents	Specifies whether the adapter filters out future events by comparing the timestamp on each event with the system time.
Ensure once-only event delivery	AssuredOnceDelivery	Specifies whether the adapter provides assured once delivery of events.
Event order by	EventOrderBy	The order in which events are retrieved and processed
Event query type	EventQueryType	Determines whether to use the standard event store or custom query
Event table name	EventTableName	Name of the database table that contains events generated by the database for inbound processing
Event types to process	EventTypeFilter	A delimited list of event types that indicates to the adapter which events it should deliver. The adapter supports event filtering based on operations, in addition to the BO Name.
Retry limit for failed events	FailedEventRetryLimit	The number of times the adapter attempts to redeliver an event before marking the event as failed.
Interval between polling periods	PollPeriod	The length of time that the adapter waits between polling periods.
JDBC driver class	JDBCDriverClass	The class name of the JDBC driver that is used to connect to the database.
Maximum connections	MaximumConnections	The maximum number of connections that the adapter can use for inbound event delivery.
Minimum connections	MinimumConnections	The minimum number of connections that the adapter can use for inbound event delivery.
Maximum number of retries in case of system connection failure	RetryLimit	The number of times the adapter tries to reestablish an inbound connection after an error.
Password	Password	Password for authorizing the user to retrieve events from the database.
Poll quantity	PollQuantity	The number of events the adapter delivers to the export during each poll period.
Query time out	QueryTimeOut	The maximum number of seconds a query can take for all SQL statements.

Table 65. Activation specification properties for Adapter for Oracle E-Business Suite (continued)

Property name		Description
In the wizard	**In the administrative console**	
Retry EIS connection on startup	RetryConnectionOnStartup	Controls whether the adapter retries the connection to the Oracle database if it cannot connect at startup.
Time between retries in case of system connection failure (milliseconds)	RetryInterval	The length of time that the adapter waits between attempts to reestablish connection after an error during inbound operations.
Return dummy business object for RetrieveSP	ReturnDummyBOForSP	Specifies whether to return output parameters when the result set is empty.
SQL query to verify the connection	PingQuery	The SQL query used to test the reliability of the connection to the database.
Stop the adapter when an error is encountered while polling	StopPollingOnError	Specifies whether the adapter stops polling for events when it encounters an error during polling.
Stored procedure to run after polling	SPAfterPoll	The name of the stored procedure that you want to be run after each poll cycle
Stored procedure to run before polling	SPBeforePoll	The name of the stored procedure that you want to be run before the actual poll query is called
Type of delivery	DeliveryType	Determines the order in which events are delivered by the adapter to the export.
User name	UserName	The database user name to use for inbound events
"Database connection information (ConnectionType)" on page 175	ConnectionType	Specifies how the adapter establishes connection to the database.

Adapter Instance for event filtering (AdapterInstanceEventFilter)

This property controls whether the adapter instance processes specific events in the event store.

Table 66. Adapter Instance for event filtering details

Required	No
Default	`null`
Property type	String

Table 66. Adapter Instance for event filtering details (continued)

Usage	This property helps you migrate from WebSphere Business Integration Adapter for Oracle Applications to WebSphere Adapter for Oracle E-Business Suite. WebSphere Business Integration Adapter for Oracle Applications allows you to perform load balancing on high-volume event types by allowing multiple adapter instances to process events of the same type. When load balancing is not required, a single adapter instance processes all events of a given type. This property is to enable seamless migration for WBIA customers to JCA for customers who are currently taking advantage of the connectorID filtering. WebSphere Adapter for Oracle E-Business Suite typically does not require load balancing in this way, but supports it so that you can migrate without modifying the database triggers or other mechanisms that write events to the event store. The AdapterInstanceEventFilter property corresponds to the ConnectorID property of the WebSphere Business Integration Adapter for Oracle Applications. To use this feature, the database triggers or other mechanisms that create events in the event store must assign the appropriate value to the connectorId column. Table 67 shows the interaction between the AdapterInstanceEventFilter property and the value in the connectorId column in the event store. If the EventTypeFilter and AdapterInstanceEventFilter properties are both set, the adapter processes only events that meet both criteria. That is, it processes only those events whose type is specified in the EventTypeFilter property and whose connectorId column matches the AdapterInstanceEventFilter property.
Example	See Table 67.
Globalized	Yes
Bidi supported	Yes

Table 67. Interaction of the AdapterInstanceEventFilter property with the connectorId column in the event store

AdapterInstanceEventFilter property	connectorId column of an event	Result
null	null	The adapter processes the event.
null	Instance1	The adapter processes the event, because the connectorId column is not checked.
Instance1	Instance1	The adapter processes the event.
Instance1	Instance2	The adapter does not process the event, because the instance IDs do not match.
Instance1	null	The adapter does not process the event, because the instance IDs do not match.

Additional JDBC driver connection properties [name:value;name:value] (JDBCDriverConnectionProperties)

This property contains additional information for connecting to the database using the JDBC driver.

Table 68. Additional JDBC driver connection properties details

Row	Explanation
Required	No

Table 68. Additional JDBC driver connection properties details (continued)

Row	Explanation
Possible values	Database connection properties are database-specific.
Default	No default value
Property type	String
Usage	These connection properties are used in addition to the UserName and Password properties to customize the database connection used by the adapter. Specify the connection properties as one or more *name:value* pairs separated by the semicolon character (;).
Example	The following value of this property specifies a login timeout interval: `loginTimeout:20; ConnectionRetryCount:5; ConnectionRetryDelay:5`
Globalized	Yes
Bidi supported	No

Business object namespace (BusinessObjectNameSpace)

This property specifies the namespace for the business object definitions.

Table 69. Business object namespace property characteristics

Required	No
Default	`http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle`
Property type	String
Usage	This value is added as a prefix to the business object name to keep business object names logically separated.
Example	The following example shows the `Schema1Customer` business object with the default namespace: `http://www.ibm.com/xmlns/prod/websphere/j2ca/oracle/Schema1Customer`
Bidi supported	No

User-defined delete query (CustomDeleteQuery)

Use this property to specify the SQL statement, stored procedure, or stored function to run after each event is processed to delete records that can be deleted after the event is delivered.

Table 70. Custom delete query details

Required	No
Default	No default value
Property type	String
Usage	Use this property to specify the SQL statement, stored procedure, or stored function to run when the EventQueryType property is set to `Dynamic`.
Globalized	Yes
Bidi supported	Yes

User-defined update query for failed event delivery (CustomUpdateQueryForFailedEvent)

Use this property to specify the SQL statement, stored procedure, or stored function to run to update event status when an event is not delivered successfully.

Table 71. Custom update query for failed event details

Required	No
Default	No default value
Property type	String
Usage	Use this property to specify the SQL statement, stored procedure, or stored function to run when the EventQueryType property is set to Dynamic. If an event is delivered successfully, the property "User-defined update query (CustomUpdateQuery)" on page 162 is used to update event status. If an event is not delivered successfully, the property "User-defined update query for failed event delivery (CustomUpdateQueryForFailedEvent)" is used to update event status.
Globalized	Yes
Bidi supported	Yes

User-defined event query (CustomEventQuery)

Use this property to specify the SQL statement, stored procedure, or stored function to run to poll for events in custom event processing.

Table 72. Custom event query details

Required	No
Default	No default value
Property type	String
Usage	Use this property to specify the SQL statement, stored procedure, or stored function to run during each poll cycle when the EventQueryType property is set to Dynamic.

Examples	In the following example, the custom event query runs an SQL statement that returns the event ID, object key, and object name of every record in the MY_EVENT_TABLE event store whose status column has the value 0:
	
	`select event_id, object_key, object_name from MY_EVENT_TABLE where status = 0`
	
	The following example limits the returned event records to the value of the PollQuantity property:
	
	`select event_id, object_key, object_name from MY_EVENT_TABLEwhere status = 0`
	`and rownum < POLL QUANTITY`
	
	The following example runs a stored procedure with two parameters:
	
	`CALL MY_EVENT_STORED_PROC (?,?)`
	
	The following example runs a stored function with one parameter and one return value:
	
	`? = CALL MY_EVENT_FUNCTION(?)`

Globalized	Yes
Bidi supported	Yes

User-defined update query (CustomUpdateQuery)

Use this property to specify the SQL statement, stored procedure, or stored function to run after each event is processed so that the same event does not get picked up for processing in the subsequent event cycle.

Table 73. Custom update query details

Required	No
Default	No default value
Property type	String
Usage	Use this property to specify the SQL statement, stored procedure, or stored function to run when the EventQueryType property is set to `Dynamic`.
Globalized	Yes
Bidi supported	Yes

Data source JNDI name (DataSourceJNDIName)

This property specifies the name of the JNDI data source to use to establish a connection to the database.

Table 74. Data source JNDI name details

Required	No
Default	No default value
Property type	String

Usage	Use this property to specify the JNDI name of a data source in WebSphere Application Server that specifies connection information for the target database.
	
	To improve the performance of inbound or outbound operations, specify the name of a data source that is enabled for prepared statement caching.
	
	If the UserName and Password properties are also set, they override the user name and password in the data source.
	
	Do not confuse the data source JNDI name property with the JNDI name of a managed connection factory or activation specification on the server. The following list highlights important differences between the types of JNDI names:
	• Data source JNDI name
	  – Specifies a connection to a database
	  – Is used instead of saving user name and password in adapter properties
	  – Is saved as an adapter property
	• JNDI name of the managed connection factory or activation specification
	  – Specifies a connection to a managed connection factory or activation specification on the server
	  – Is used instead of specifying the value of each managed connection factory or activation specification property in the wizard
	  – Is saved as the connection target in the import file

Globalized	Yes

Table 74. Data source JNDI name details (continued)

Bidi supported	No

Database URL (DatabaseURL)

This property specifies the JDBC driver-specific URL for creating a connection to the database.

Table 75. Database URL details

Required	Yes, when the connection is of the type LocalConnectionProps or ConnectionProps.
Default	No default value
Property type	String
Usage	In the J2C Bean wizard, compose the database URL by entering in database-specific fields. For example, the database URL for an Oracle database is composed of the system ID (SID), server host name, and database port number. In the administrative console, type the complete database URL value. If your database server supports IPv6, you can specify the host name portion of the database URL in IPv6 format. If you specify the host name as an IP address in IPv6 format, enclose the IP address in square brackets ([]).
Examples	The following example illustrates typical database URL values for a common database:`jdbc:oracle:thin:@9.26.248.148:1521:dev`
Globalized	Yes
Bidi supported	Yes

Database vendor (DatabaseVendor)

This property specifies the type of database that is used. The type is determined by the database vendor name.

Table 76. Database vendor details

Required	Yes
Possible values	`Oracle`
Default	`ORACLE`
Property type	String
Usage	Some SQL statements require special processing, which varies by database type. For example, the Struct and Array data types in Oracle require special processing. This property specifies the RDBMS that is used, which determines the database type. For other databases, the adapter does not perform any special processing. Make sure that the correct driver is specified in the JDBCDriverClass property. Specify `Oracle` as the value that corresponds to your database vendor.
Globalized	No

Table 76. Database vendor details (continued)

Bidi supported	No

Delivery type (DeliveryType)

This property specifies the order in which events are delivered by the adapter to the export.

Table 77. Delivery type details

Required	No
Possible values	ORDERED UNORDERED
Default	ORDERED
Property type	String
Usage	The following values are supported: • ORDERED: The adapter delivers events to the export one at a time. • UNORDERED: The adapter delivers all events to the export at once.
Globalized	No
Bidi supported	No

Do not process events that have a timestamp in the future (FilterFutureEvents)

This property specifies whether the adapter filters out future events by comparing the timestamp on each event with the system time.

Table 78. Do not process events that have a timestamp in the future details

Required	Yes
Possible values	True False
Default	False
Property type	Boolean
Usage	If set to True, the adapter compares the time of each event to the system time. If the event time is later than the system time, the event is not delivered. If set to False, the adapter delivers all events.
Globalized	No
Bidi supported	No

Ensure once-only event delivery (AssuredOnceDelivery)

This property specifies whether to provide ensure once-only event delivery for inbound events.

Table 79. Ensure once-only event delivery details

Required	Yes
Possible values	True False

Table 79. Ensure once-only event delivery details (continued)

Default	True
Property type	Boolean
Usage	When this property is set to True, the adapter provides assured once event delivery. This means that each event is delivered once and only once. A value of False does not provide assured once event delivery, but provides better performance. When this property is set to True, the adapter attempts to store transaction (XID) information in the event store. If it is set to False, the adapter does not attempt to store the information. This property is used only if the export component is transactional. If it is not, no transaction can be used, regardless of the value of this property.
Globalized	No
Bidi supported	No

Event order by (EventOrderBy)

The order in which events are retrieved and processed.

Table 80. Event order by details

Required	No
Possible values	A comma-separated (,) list of column names in the event store, and the order attributes asc and desc
Default	event_time, event_priority
Property type	String
Usage	Specify a comma-separated list of column names from the event store, with the optional attributes for ascending or descending order.
Examples	To present events ordered first by time and then by priority, specify: event_time, event_priority To present events ordered first by object name in ascending order and then by event time in descending order, specify: object_name asc, event_time desc
Globalized	Yes
Bidi supported	Yes

Event query type (EventQueryType)

This property specifies whether to use standard or custom query processing.

Table 81. Event query type details

Required	Yes
Possible values	Standard Dynamic
Default	Standard
Property type	String

Table 81. Event query type details (continued)

Usage	The valid values are `Standard`, for standard event processing, and `Dynamic`, for custom event processing. If this property is set to `Dynamic`, the CustomEventQuery, CustomUpdateQuery, and CustomDeleteQuery properties are used. If this property is set to `Standard`, those properties are ignored.
Globalized	No
Bidi supported	No

Event table name (EventTableName)

This property specifies the name of the table in the target database that contains the event store, which is used for inbound processing.

Table 82. Event table name details

Required	Yes
Default	`WBIA_Oracle_EventStore`
Property type	String
Usage	Create the event store before starting to configure the adapter. For standard event processing, the event is generated by the database through a trigger or other mechanism. For custom query processing, the adapter saves events in the event store as it receives the result of the custom queries.
Globalized	Yes
Bidi supported	Yes

Event types to process (EventTypeFilter)

This property contains a delimited list of event types that indicates to the adapter which events it should deliver.

Table 83. Event types to process details

Required	No
Possible values	A comma-delimited (`,`) list of business object types
Default	`null`
Property type	String

Table 83. Event types to process details (continued)

Usage	Events are filtered by business object type and operations. If the property is set, the adapter delivers only those events that are in the list. A value of `null` or `*` indicates that no filter will be applied and that all events will be delivered to the export.
	When the default operation set (Create/Update/Delete) is not preferred, the adapter will provide an operation-based event polling capability.
	For example:
	If you select `Create|Update` operation, the adapter will poll only those events with Create or Update operation for all the BOs. The default string generated for eventTypeFilter is `*: Create|Update`.
	Syntax: `BOName:Operation1|Operation2, BOName:Operation1|Operation2|Operation3`
	• "," is used for separating the business objects. The objects that are not in this list are ignored.
	• ":" is used for separating the business object name and the operation name.
	• "|" is used for separating the supported operations, such as Create, Delete, and Update.
Example	To receive events related to the Customer and Order business objects only, specify this value: `Customer,Order`
	• To receive events related to the Customer and Order business objects regardless of operations, specify this value: `Customer,Order` or specify this value for all supported operations: `Customer:Create|Update|Delete,Order:Create|Update|Delete`.
	• To receive all events of the Customer business object and the Create and Delete events of the Order business object, specify this value: `Customer,Order:Create|Delete`
	• To receive the Create and Delete events of all the business objects, specify this value: `*:Create|Delete`.
	Note: In this property, you can use only those operations that are chosen while selecting business objects and services.
	If the EventTypeFilter and AdapterInstanceEventFilter properties are both set, the adapter processes only events that meet both criteria. That is, it processes only those events whose type is specified in the EventTypeFilter property and whose connectorId column matches the AdapterInstanceEventFilter property.
Globalized	No
Bidi supported	No

Retry limit for failed events (FailedEventRetryLimit)

This property specifies the number of times that the adapter attempts to redeliver an event before marking the event as failed.

Table 84. Retry limit for failed events details

Required	No
Possible values	Integers
Default	5
Property type	Integer

Table 84. Retry limit for failed events details (continued)

Usage	Use this property to control how many times the adapter tries to send an event before marking it as failed. It accepts the following values:
	Default If this property is not set, the adapter tries five additional times before marking the event as failed.
	0 The adapter tries to deliver the event an infinite number of times. When the property is set to 0, the event remains in the event store and the event is never marked as failed.
	> 0 For integers greater than zero, the adapter retries the specified number of times before marking the event as failed.
	< 0 For negative integers, the adapter does not retry failed events.
Globalized	No
Bidi supported	No

JDBC driver class (JDBCDriverClass)

This property specifies the class name of the JDBC driver that is used to connect to the database.

Table 85. JDBC driver class details

Row	Explanation
Required	Yes, when the connection is of the type LocalConnectionProps or ConnectionProps.
Possible values	The value specified when you set the connection properties to connect to the database instance.
Default	No default value
Property type	String
Usage	In the J2C Bean wizard, the JDBC driver class is displayed if you select a common database software and driver combination, such as type 4 drivers for recent versions of Oracle. For type 2 drivers , you must type the database class name. If you select a type 2 driver or a generic driver, you must type the JDBC driver class name. In the administrative console, type the database-specific name of the driver.
Examples	Values for the JDBC driver class is displayed in both the J2C Bean wizard and the administrative console. The following examples depict the JDBC driver class properties in both the J2C Bean wizard and the administrative console.**In the J2C Bean wizard:** • To connect to an Oracle 10 database using the type 4 driver, select `Oracle Thin Driver`. **In the administrative console:** **Oracle Thin JDBC driver** `oracle.jdbc.driver.OracleDriver`
Globalized	No

Table 85. JDBC driver class details (continued)

Bidi supported	No

Maximum connections (MaximumConnections)

This property specifies the maximum number of connections that the adapter can use for inbound event delivery.

Table 86. Maximum connections details

Required	No
Default	1
Property type	Integer
Usage	Only positive values are valid. Any value less than 1 is treated as 1 by the adapter.
Globalized	No
Bidi supported	No

Minimum connections (MinimumConnections)

This property specifies the minimum number of connections that the adapter can use for inbound event delivery.

Table 87. Minimum connections details

Required	No
Default	1
Property type	Integer
Usage	Only positive values are valid. Any value less than 1 is treated as 1 by the adapter.
Globalized	No
Bidi supported	No

Password (Password)

This property specifies the password for the database user name.

Table 88. Password details

Required	No. For inbound processing, if you set the Authentication alias or the DataSourceJNDIName, the password is not mandatory. However, if you set the DataSourceJNDIName, and the Password field, the value specified for the Password takes precedence. For outbound processing, if you set the Authentication alias, XADataSourceJNDIName or PoolDataSourceJNDIName property, the password is not mandatory. However, if you set the XADataSourceJNDIName or PoolDataSourceJNDIName property, and the Password field, the value specified for the Password takes precedence.
Default	No default value
Property type	String

Table 88. Password details (continued)

Usage	For inbound processing, setting this property overrides the password specified for the data source on the server using the Authentication alias or DataSourceJNDIName property.
	For outbound processing, setting this property overrides the password specified for the data source on the server using the Authentication alias, or XADataSourceJNDIName or PoolDataSourceJNDIName property.
	If you specify JAAS as the security credential, the authentication alias will override this property.
Globalized	Yes
Bidi supported	Yes

SQL query to verify the connection (PingQuery)

This property specifies the SQL query that is used to test the reliability of the connection to the database.

Table 89. Ping query details

Required	No
Property type	String
Default	No default value
Usage	This property contains the SQL query statement that you want to run to determine whether the adapter can connect to the database.
	The adapter runs the ping query every time it receives a SQLException exception while performing the outbound operation.
	The adapter does not try to recover the connection. If the ping query indicates that the connection to the database is no longer valid, the adapter notifies the container. It is the responsibility of the connection pool manager to remove the stale connection from the pool, which allows subsequent outbound requests to be processed.
Globalized	No
Bidi supported	No

Interval between polling periods (PollPeriod)

This property specifies the length of time that the adapter waits between polling periods.

Table 90. Interval between polling periods details

Required	Yes
Possible values	Integers greater than or equal to 0.
Default	2000
Unit of measure	Milliseconds
Property type	Integer

Table 90. Interval between polling periods details (continued)

Usage	The poll period is established at a fixed rate, which means that if running the poll cycle is delayed for any reason (for example, if a prior poll cycle takes longer than expected to complete) the next poll cycle occurs immediately to make up for the lost time caused by the delay.
Globalized	No
Bidi supported	No

Maximum events in polling period (PollQuantity)

This property specifies the number of events that the adapter delivers to the export during each poll period.

Table 91. Maximum events in polling period details

Required	Yes
Default	10
Property type	Integer
Usage	The value must be greater than 0. If this value is increased, more events are processed per polling period and the adapter may perform less efficiently. If this value is decreased, fewer events are processed per polling period and the adapter's performance might improve slightly.
Globalized	No
Bidi supported	No

Query timeout (seconds) (QueryTimeOut)

This property specifies the maximum number of seconds a query can take to run all SQL statements.

Table 92. Query timeout details

Required	No
Default	No default value
Unit of measure	Seconds
Property type	Integer
Usage	If the query takes longer than the number of seconds specified, the database generates an SQL exception that is captured. The associated message is logged in the log file. If a value is not specified, no timeout is set on the query.
Globalized	No
Bidi supported	No

Time between retries in case of system connection failure (RetryInterval)

When the adapter encounters an error related to the inbound connection, this property specifies the length of time the adapter waits before trying to reestablish a connection.

Table 93. Retry interval details

Required	Yes
Default	2000
Unit of measure	Milliseconds
Property type	Integer
Usage	Only positive values are valid. When the adapter encounters an error related to the inbound connection, this property specifies the length of time the adapter waits before trying to establish a new connection.
Globalized	No
Bidi supported	No

Maximum number of retries in case of system connection failure (RetryLimit)

This property specifies the number of times the adapter tries to reestablish an inbound connection.

Table 94. Maximum number of retries in case of system connection failure

Required	No
Possible values	0 and positive integers
Default	0
Property type	Integer
Usage	This property controls how many times the adapter retries the connection if the adapter cannot connect to the Oracle database to perform inbound processing. A value of 0 indicates an infinite number of retries. To control whether the adapter retries if it cannot connect to the Oracle database when it is first started, use the RetryConnectionOnStartup property.
Globalized	No
Bidi supported	No

Retry EIS connection on startup (RetryConnectionOnStartup)

This property controls whether the adapter attempts to connect again to the Oracle database if it cannot connect at startup.

Table 95. Retry EIS connection on startup details

Required	No
Possible values	True False
Default	False
Property type	Boolean

Table 95. Retry EIS connection on startup details (continued)

Usage	This property indicates whether the adapter should retry the connection to the Oracle database if the connection cannot be made when the adapter is started: • Set the property to `False` when you want immediate feedback about whether the adapter can establish a connection to the Oracle database, for example, when you are building and testing the application that receives events from the adapter. If the adapter cannot connect, the adapter writes log and trace information and stops. The administrative console shows the application status as `Stopped`. After you resolve the connection problem, start the adapter manually. • Set the property to `True` if you do not need immediate feedback about the connection. If the adapter cannot connect during startup, it writes log and trace information, and then attempts to reconnect, using the RetryInterval property to determine how frequently to retry and the value of the RetryLimit property to retry multiple times until that value is reached. The administrative console shows the application status as `Started`.
Globalized	No
Bidi supported	No

Return business object even when the stored procedure result set is empty (ReturnDummyBOForSP)

This property specifies whether to return output parameters when the result set is empty.

Table 96. Return business object even when the stored procedure result set is empty details

Required	No
Possible values	`True` `False`
Default	`False`
Property type	Boolean
Usage	The Retrieve Stored Procedure (RetrieveSP) operation returns a result set. If the result set is empty and the ReturnDummyBOForSP property is set to `False`, no business objects are created, and the output parameters returned by the procedure call cannot be retrieved. However, if the ReturnDummyBOForSP property is set to `True`, a dummy business object is created and the returned parameters (including the out and inout parameters) from the stored procedure are populated into the corresponding attributes
Globalized	No
Bidi supported	No

Stop the adapter when an error is encountered while polling (StopPollingOnError)

This property specifies whether the adapter will stop polling for events when it encounters an error during polling.

Table 97. Stop the adapter when an error is encountered while polling details

Required	No

Possible values	True False
Default	False
Property type	Boolean
Usage	If this property is set to True, the adapter stops polling when it encounters an error. If this property is set to False, the adapter logs an exception when it encounters an error during polling and continues polling.
Globalized	No
Bidi supported	No

Stored procedure to run after polling (SPAfterPoll)

This property specifies the name of the stored procedure or stored function to run after each polling cycle.

Table 98. Stored procedure to run after poll details

Required	No
Default	No default value
Property type	String
Usage	The stored procedure takes one parameter for poll quantity.
Globalized	Yes
Bidi supported	Yes

Stored procedure to run before polling (SPBeforePoll)

This property specifies the name of any stored procedure or stored function to run before the actual poll query is called.

Table 99. Stored procedure to run before poll details

Required	No
Default	No default value
Property type	String
Usage	The stored procedure takes one parameter for poll quantity.
Globalized	Yes
Bidi supported	Yes

User name (UserName)

This property specifies the database user name that is used to access the database.

Table 100. User name details

Required	No. For inbound processing, if you set the Authentication alias or the DataSourceJNDIName, the User name property is not mandatory. However, if you set the DataSourceJNDIName, and User name field, the value specified for the User name takes precedence.
	For outbound processing, if you set the Authentication alias, XADataSourceJNDIName or PoolDataSourceJNDIName property, the User name is not mandatory. However, if you set the XADataSourceJNDIName or PoolDataSourceJNDIName property, and the User name field, the value specified for the User name takes precedence.
Default	No default value
Property type	String
Usage	For inbound processing setting this property overrides the user name specified for the data source on the server using the DataSourceJNDIName property or the Authentication alias.
	For outbound processing, setting this property overrides the user name specified for the data source on the server using the XADataSourceJNDIName property, PoolDataSourceJNDIName property or the Authentication alias.
	If you specify JAAS as the security credential, the authentication alias will override this property.
Globalized	Yes
Bidi supported	Yes

Database connection information (ConnectionType)

This property specifies how the adapter establishes connection to the database.

Table 101. Database connection information

Required	Yes
Possible values	ConnectionProps or DataSourceJNDI
Default	ConnectionProps
Property type	String

Table 101. Database connection information (continued)

Usage	This property specifies how adapter establish database connection at run time. This property can have the following values:
	• ConnectionProps - Indicates that the database connection is established using the DatabaseURL and JDBCDriverClass properties.
	• DataSourceJNDI - Indicates that the database connection is established using the DataSourceJNDIName property corresponding to predefined datasource.
	For a new application, this property is automatically set by J2C Bean wizard.
	• If the DataSourceJNDIName property is set, the value of this property is set to DataSourceJNDI.
	• If the DataSourceName is set, the value of this property is set to ConnectionProps.
	If this property is not set, adapter uses backward compatibility mode to establish database connection. In backward compatibility mode, the properties for connecting to the database are used in the following order:
	1. If the DataSourceJNDIName property is set, the adapter uses it to establish the connection to the database.
	2. If the DataSourceJNDIName is not set, then the adapter uses the DatabaseURL, JDBCDriverClass, UserName and Password properties to establish the connection.
Globalized	No
Bidi supported	No

Globalization

WebSphere Adapter for Oracle E-Business Suite is a globalized application that can be used in multiple linguistic and cultural environments. Based on character set support and the locale of the host server, the adapter delivers message text in the appropriate language. The adapter supports bidirectional script data transformation between integration components.

Globalization and bidirectional transformation

The adapter is globalized to support single and multibyte character sets and deliver message text in the specified language. The adapter also performs bidirectional transformation, which refers to the task of processing data that contains both right-to-left (Hebrew or Arabic, for example) and left-to-right (a URL or file path, for example) semantic content within the same file.

Globalization

Globalized software applications are designed and developed for use within multiple linguistic and cultural environments rather than a single environment. WebSphere Adapters, Rational Application Developer for WebSphere Software, WebSphere Application Server are written in Java. The Java run time environment within the Java virtual machine (JVM) represents data in the Unicode character code set. Unicode contains encoding for characters in most known character code sets (both single- and multi-byte). Therefore, when data is transferred between these integration system components, there is no need for character conversion.

To log error and informational messages in the appropriate language and for the appropriate country or region, the adapter uses the locale of the system on which it is running.

Bidirectional transformation

Languages such as Arabic and Hebrew are written from right to left, yet they contain embedded segments of text that are written left to right, resulting in bidirectional script. There are multiple ways that a software application might display and process bidirectional script. WebSphere Application Server uses the Windows standard format, but an enterprise information system exchanging data with WebSphere Application Server can use a different format. WebSphere Adapters transform bidirectional script data is passed between the two systems so that it is accurately processed and displayed on both sides of a transaction.

Bidirectional format

WebSphere Application Server uses the bidirectional format of ILYNN (implicit, left-to-right, on, off, nominal). This is the format used by Windows. If an enterprise information system uses a different format, the adapter converts the format prior to introducing the data to WebSphere Application Server.

The bidirectional format consists of five attributes. When you set bidirectional properties, you assign values for each of these attributes. The attributes and settings are listed in the following table.

Table 102. Bidirectional format attributes

Letter position	Purpose	Values	Description	Default setting
1	Order schema	I	Implicit (Logical)	I
		V	Visual	
2	Direction	L	Left-to-Right	L
		R	Right-to-Left	
		C	Contextual Left-to-Right	
		D	Contextual Right-to-Left	
3	Symmetric Swapping	Y	Symmetric swapping is on	Y
		N	Symmetric swapping is off	
4	Text Shaping	S	Text is shaped	N
		N	Text is not shaped (Nominal)	
		I	Initial shaping	
		M	Middle shaping	
		F	Final shaping	
		B	Isolated shaping	
5	Numeric Shaping	H	National (Hindi)	N
		C	Contextual shaping	
		N	Numbers are not shaped (Nominal)	

The adapter transforms data into a logical, left-to-right format before sending the data to WebSphere Application Server.

Using bidirectional properties

You can use multiple bidirectional properties to control the transformation of both content data and metadata. You can set special bidirectional properties to exclude either content data or metadata from bidirectional transformation, or to identify data that requires special treatment during a transformation.

The following table describes the types of bidirectional properties.

Table 103. Bidirectional property types

Property type	Data transformations
EIS	Controls the format for content data, or data that is sent by the enterprise information system, which is, the database.
Metadata	Controls the format for metadata, or data that provides information about the content data.
Skip	Identifies content or metadata to exclude from transformation.
Special Format	Identifies certain text, such as file paths or URLs, which require different treatment during the transformation process. Can be set for either content data or metadata.

You can set properties that control bidirectional transformation in the following areas:

- **Resource adapter properties:** These properties store default configuration settings, including the TurnBiDiOff property, which controls whether the adapter instance performs bidirectional transformation or not. Use the administrative console of the server to configure these properties.
- **Managed connection factory properties:** These properties are used at run time to create an outbound connection instance with an enterprise information system. After the managed connection factory properties are created, they are stored in the deployment descriptor.
- **Activation specification properties:** These properties hold the inbound event processing configuration information for a message endpoint. Set them when you use the J2C Bean wizard, or use the administrative console of the server.

Property scope and lookup mechanism

After you set values for bidirectional properties for an adapter, the adapter performs bidirectional transformations. It does so by using logic that relies on a hierarchical inheritance of property settings and a lookup mechanism.

Properties defined within the resource adapter are at the top of the hierarchy, while those defined within other areas or annotated within a business object are at lower levels of the hierarchy. So for example, if you set values for EIS-type bidirectional properties only for the resource adapter, those values are inherited and used by transformations that require a defined EIS-type bidirectional property, whether they arise from an inbound (activation specification) transaction or an outbound (managed connection factory) transaction.

However, if you set values for EIS-type bidirectional properties for both the resource adapter and the activation specification, a transformation arising from an inbound transaction uses the values set for the activation specification.

The processing logic uses a lookup mechanism to search for bidirectional property values to use during a transformation. The lookup mechanism begins its search at the level where the transformation arises and searches upward through the hierarchy for defined values of the appropriate property type. It uses the first valid value it finds. It searches the hierarchy from child to parent only; siblings are not considered in the search.

Properties enabled for bidirectional data transformation

WebSphere Adapter for Oracle E-Business Suite has several configuration properties that are enabled for bidirectional data transformation.

The adapter enables the exchange of bidirectional data between a client application and the database, even if the data in the database is in a different bidirectional format than is used by the runtime environment. You can use bidirectional characters when configuring the adapter and in the application-specific information of your business objects. The following sets of properties and application-specific information are enabled for bidirectional support:

- Configuration properties
  - Activation specification properties
  - Connection properties for the J2C Bean wizard
  - Managed connection factory properties
- Application-specific information
  - Business object level ASI
  - Operation level ASI
  - Attribute level ASI

The sections which follow list the specific configuration properties and application-specific information that are enabled for bidirectional transformation.

Activation specification properties

The following activation specification properties are enabled for bidirectional script data transformation:

- Custom delete query
- Custom event query
- Custom update query
- Additional JDBC driver connection properties
- Database URL
- Event order by
- Event table name
- Password
- Stored procedure to run before polling
- Stored procedure to run after polling
- User name

Connection properties used in the wizard

The following connection properties for the J2C Bean wizard are enabled for bidirectional script data transformation:

- User name
- Password

Managed connection factory properties

The following managed connection properties are enabled for bidirectional script data transformation:

- Additional JDBC driver connection properties
- Database URL
- Password
- User name

Business object application-specific information

The following business object application-specific information parameters are enabled for bidirectional script data transformation:

- TableName
- StatusColumnName
- SPName
- SelectStatement

Operation application-specific information

The following operation application-specific information parameters are enabled for bidirectional script data transformation:

- StoredProcedureName
- PropertyName in Parameters

Attribute application-specific information

The following attribute application-specific information parameters are enabled for bidirectional script data transformation:

- ColumnName

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Department 2Z4A/SOM1
294 Route 100
Somers, NY 10589-0100
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: (c) (your company name) (year). Portions of

this code are derived from IBM Corp. Sample Programs. (c) Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming interface information

Programming interface information, if provided, is intended to help you create application software using this program.

General-use programming interfaces allow you to write application software that obtain the services of this program's tools.

However, this information may also contain diagnosis, modification, and tuning information. Diagnosis, modification and tuning information is provided to help you debug your application software.

Warning:

Do not use this diagnosis, modification, and tuning information as a programming interface because it is subject to change.

Trademarks and service marks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol ($^{\circledR}$ or $^{\texttrademark}$), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A complete and current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

This product includes software developed by the Eclipse Project (http://www.eclipse.org).

Index

Special characters

, adding 79

A

activation specification 87, 88
activation specification properties
  list of 155
  setting in administrative console 96, 99
Active-Passive 40
adapter
  package files 91
  project, create 44
adapter application
  starting 112
  stopping 113
Adapter for Oracle E-Business Suite module
  exporting as EAR file 84
  installing EAR file on server 84
  starting 112
  stopping 113
adapter implementation
  security 35
administrative console 86, 87, 88
after-image 3
application-specific information 125
  adding to object 49, 66
  for attributes of type child business object 123
  for simple attributes 117
assured once delivery 12
attribute properties 115
attribute type, business object 116
authentication
  description 36
  in the wizard 36
  runtime environment 36
authentication alias
  J2C 36

B

Batch Processing 40
business object information 115
business object structure
  for query business objects 20
  for stored procedure business objects 20
  for table or view business objects 19
business objects 18, 63, 125
  attribute types 116
  attributes 115
  cardinality 21
  composite keys 63, 76
  how to view 50, 67
  multiple parents 63, 76
  naming conventions 128
  query 32

business objects *(continued)*
  stored procedure 25

C

cardinality 21, 115
cluster level 87, 88
clustered environment 86, 87, 88
  adapters version conflict 40
  deployment 40
  inbound process 40
  inbound processes 41
  load balancing 40
  outbound process 40
  outbound processes 41
compatibility matrix 1
complex data types 26
confidential data, disguising 35
confidential tracing 35
configuration properties
  inbound 150
configuring
  logging properties 91
configuring JDBC adapter for XA 62
connector project 44
Create operation 4
custom properties
  activation specification 96, 99
  managed connection factory 95, 98
  resource adapter 94, 97
custom queries
  standard SQL 13
  stored function 14
  stored procedure 14

D

data types
  complex 26
debugging
  self-help resources 103, 114
Delete operation 9
delta 3
deploy module 86, 87
deployment 84
  production environment 82
  test environment 79
deployment environment 79
deployment options 37
disguising confidential data 35
distributed transactions, see *XA* transactions 143

E

EAR file
  exporting 84
  installing on server 84
embedded adapter 86

embedded adapter *(continued)*
  activation specification properties, setting 96
  considerations for using 38
  managed connection factory properties, setting 95
  resource adapter properties, setting 94
  usage considerations 37
embedded adapters 86, 87
  changing configuration properties 94
  setting activation specification properties 96
  setting managed (J2C) connection factory properties 95
  setting resource adapter properties 94
embedded deployment 80
enableHASupport property 41
endpoint applicaiton
  troubleshoot 106
event delivery 164
event processing
  custom 3, 13
  standard 2, 12
event store 14, 15
event store set up 43
Execute operation 9
Exist operation
  limitations 10
Exists operation 10
  for database table business objects 10
  for database views business objects 10
  for nicknames business objects 10
  for synonyms business objects 10
exporting module as EAR file 84

F

fdxconnector dependencies, adding 80
FFDC (first-failure data capture) 105
files
  SystemOut.log log file 93
  trace.log trace file 93
first-failure data capture (FFDC) 105
flat business objects 21
foreign key 115

G

global transaction 62
  with XA DataSource JNDI name 62
  with XADataSourceName and XADatabaseName 62

H

HA Active-Active 40
hardware and software requirements 1

© Copyright IBM Corp. 2006, 2012

185

IBM ®

Printed in USA