



IBM Software Group

## IBM WebSphere® Adapters V6.0.2

### *WebSphere Adapter for JDBC V6.0.2*



@business on demand.

© 2007 IBM Corporation  
Updated March 12, 2007

This presentation provides specific details of WebSphere Adapter for JDBC V6.0.2

## Section

# ***Overview and installation***

This section provides an overview and installation preparation steps of the WebSphere Adapter for JDBC. Note that the installation and deployment of the WebSphere Adapter for JDBC is also covered in a separate presentation common for all WebSphere Adapters.

## Overview: WebSphere Adapter for JDBC

- IBM WebSphere Adapter for JDBC implements the Java™ 2 Enterprise Edition (J2EE) Connector Architecture (JCA), version 1.5 specification
- Enables bi-directional connectivity for integration with database applications
  - ▶ Requires database with JDBC driver that supports the JDBC 2.0 or higher specification
- Uses standard JDBC APIs to retrieve and update information in the database
  - ▶ Executes SQL statements or Stored Procedures
  - ▶ Specified in the Business Object
- Supports both Local Transactions and Global (XA) as supported by the database
  - ▶ Configure as XA Data source through WebSphere Process Server data source
  - ▶ DataSourceJNDIName is specified as part of Managed Connection Factory
  - ▶ Supports all databases - no longer restricted to DB2® and Oracle

The IBM WebSphere Adapter for JDBC implements the JCA 1.5 specification and enables bi-directional connectivity, both inbound and outbound, with those Enterprise Information System business applications that communicate with database applications. The JDBC driver supports the JDBC 2.0 or higher specification and the adapter uses JDBC APIs to retrieve and update information in the database by running SQL statements or stored procedures, depending on what you specify. This is stored and specified in the business object. A stored procedure is a group of SQL statements that forms a logical unit and performs 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 JDBC Adapter supports both local transactions and global XA transactions. This version of the adapter provides Global XA support for outbound connections, and can configure as XA Data Source through WebSphere Process Server Data Source. DataSourceJNDIName is a new property in the Managed Connection Factory where the data source name is specified. In this case, it supports all databases and is no longer restricted to just DB2 and Oracle.

## Advanced features of JDBC adapter

- UID Support (auto generated column)
  - ▶ Sequences and Identity Columns
- Order By
- Fixed Length Strings
- Stored Procedures
- Query Timeout

The adapter supports a set of advanced features listed here. An automatically generated column (UID) can be tied to a sequence or be defined as an identity column (numeric value). Sequences can be defined for DB2 and Oracle only and identity columns can be defined for SQL Server and DB2 databases. Ascending or descending order can be specified using the Order By feature. The adapter also supports processing with fixed length strings. If the attribute is of type fixed length when columns are CHAR, the adapter will pad the field with blanks. Support for stored procedures and query timeout complete the list of advanced features.

## Overview: New V6.0.2 functions



- CLOB and BLOB SQL data type support
- User customizable event processing
  - ▶ Accommodate additional event processing scenarios while minimizing the need for customization and changes to the sources database
    - User-specified event store poll queries
    - Optionally deliver contents from event store
    - User-specified event store update after delivery
    - User-specified SQL before poll/after poll
- New event manager support
  - ▶ Uses of EDT for Assured once delivery – deprecated
  - ▶ Implants a XAResource to keeps track of transaction IDs in the Event Table
- Stored Functions Support



The WebSphere Adapter for JDBC version 6.0.2 supports a set of new features and functions listed here.

Version 6.0.2 supports SQL data types Character Large Object (CLOB) and Binary Large Object (BLOB) along with the other supported types.

User-customizable event processing is supported to accommodate additional event processing scenarios beyond the earlier design while minimizing the need for customization and changes to the source databases. The enhancement enable user-specified event store poll queries, optional delivery of contents from the event store, user-specified event store update after delivery, and user-specified SQL before or after poll.

The Event Distribution Table is no longer supported by the adapter. Instead, the adapter now implements the new EventStoreWithXid interface. The event table contains an additional XID field, which the adapter will query and update as per functionality implemented by adapter foundation classes.

The adapter now supports Stored Functions and it will check for the ReturnValue ASI, or Application Specific Information. If it exists, then it will run the stored function.

JDBC Enterprise Metadata Discovery will allow you to associate one or more Stored function on a business object which is based on a table or view.

## Overview: New V6.0.2 functions (cont.)



- User specified SQL processing
- Allow selection of Stored Procedures using filter in EMD
  - ▶ Specify option to use a filter to narrow the list of stored procedures in EMD User specified SQL processing support
- Store Procedure independent of operation
  - ▶ Support stored procedures that are not tied to Create/Delete/Update/Retrieve/RetrieveAll operations
- Data source for JCA JDBC Adapter support
- Performance Enhancement
  - ▶ Remove extra retrieval of top level Business Object

6

The adapter provides support for user specified SQL processing. It will run the query based on the select query Application Specific Information, or ASI, and the JDBC where clause. The adapter will generate a container business object and populate that Business Object in the container with the returned records. You need to specify one select statement, and set other information during Enterprise Metadata Discovery. This includes business object name, the type and dummy value for all parameters in the 'where' clause. Then one query business object will be generated based on this select statement. The JDBC Adapter can then process this query business object once all necessary ASIs have been set properly.

Another new function is using a filter to narrow the list of Stored Procedures in EMD when associating stored procedure with a business object. You can specify a valid string as filter criteria and that valid values of the Stored Procedure will only list those stored procedures whose names conform to this filter criteria.

Supports Stored Procedures that are not tied to Create/Delete/Update/Retrieve/RetrieveAll operations. It will run the Stored Procedure and generate Business Objects for the output parameters. Complex type struct and array, nested struct and array are now supported.

The adapter supports establishing database connections through the WASDataSource instance. Both local and global XA transactions are supported by the Data Source connections.

For performance enhancement, the adapter was doing an extra retrieve on the top level Business Object for the RetrieveAll verb. That extra retrieval of the top level business object has been corrected.

## Installation Preparation Steps

- WebSphere Adapter for JDBC RAR file is included in Resource Adapter directory of WebSphere Integration Developer
- JDBC 2.0 Driver dependent jar file for Database
  - ▶ Dependent jars need to be in the classpath of the WebSphere Integration Developer project after importing the Resource Adapter Archive (RAR) file
- Event Store table – need to create
  - ▶ Required for inbound processing
  - ▶ Sample scripts provided for DB2, Oracle and SQL Server
  - ▶ When data changes in application tables, appropriate events are inserted into the event store
    - Use triggers defined on respective tables, or other means to capture changed data, such as Oracle Change Data Capture for Oracle databases

The JCA adapters are packaged as resource adapter archive, or RAR, files, and the WebSphere Adapter for JDBC RAR file is now included in Resource Adapter directory of WebSphere Integration Developer. Any JDBC 2.0 driver for the database that you will be using must be added to the class path. You also need the database and tables names used by the adapter.

The event table, also referred to as the event store, is required by the adapter for inbound processing. You must create this table and specify the database, table, and driver used to connect.

As data changes in the particular application table, appropriate events are inserted into the event store table through the use of triggers you have created and defined on the application tables. You can also use whatever other means your database supports to capture changed data. For example, with Oracle, the Oracle Change Data Capture can be used.

## Section

# ***Enterprise service discovery***

This section provides an overview of the Enterprise Service Discovery of the WebSphere Adapter for JDBC.



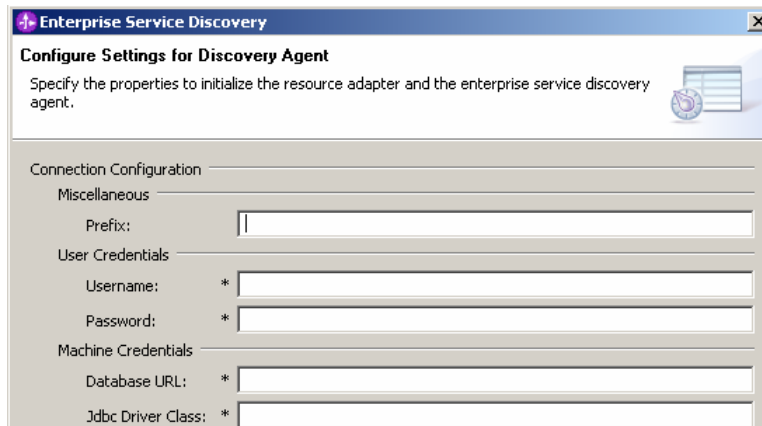
## Enterprise service discovery

- Discovery of objects in a database
  - ▶ Objects can be created from tables, views, stored procedures, synonyms or nicknames, custom query business objects
  - ▶ Generate a business object corresponding to a database object
  - ▶ Generate properties in the business object corresponding to properties in the database object
  - ▶ Set Application Specific Information on the business object
  - ▶ Provide Service Descriptions (inbound and outbound) used to generate the SCA-related files
    - Imports, Exports, and WSDL files

The Enterprise Service Discovery allows for the discovery of objects in a database, generates Business Objects from the selected objects, and generates the service constructs that enable the adapter to run as an SCA component. Objects from which business objects can be created include tables, views, stored procedures, synonyms or nicknames, and custom query Business Objects.

## Enterprise service discovery


- Configure Settings for Discovery Agent
  - ▶ Properties: UserName, Password, Database URL, JDBCDriverClass, Prefix



Here is screen capture of the 'Configure Settings for Discovery Agent' panel. User and Machine Credentials are required properties. They include User name, Password, Database URL, and JDBCDriverClass.

## Enterprise service discovery

- Metadata discovery system capabilities
  - ▶ Analyze database to identify schemas
  - ▶ Set filter properties to limit number of objects displayed and whether to set ASI values for the business objects



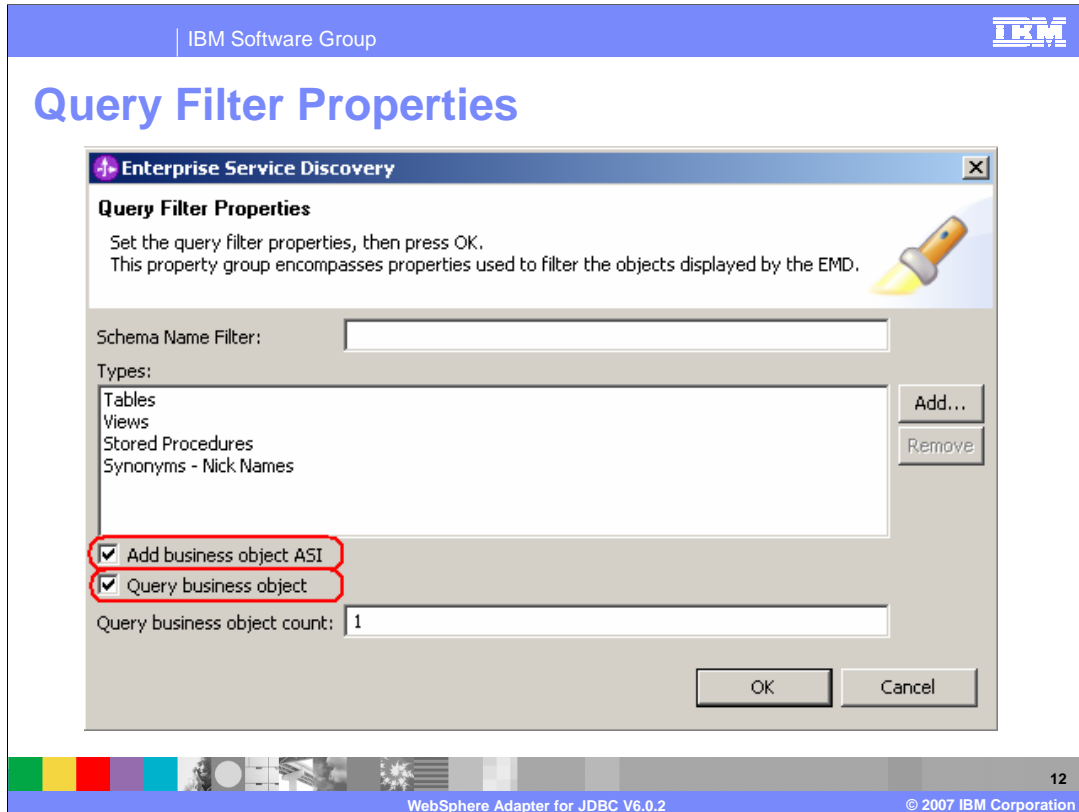
11

WebSphere Adapter for JDBC V6.0.2

© 2007 IBM Corporation

Metadata Discovery System Capabilities Schemas will be displayed as top-level nodes in a tree. Nodes, labeled Tables, Views, Stored Procedures, and Synonyms/Nicknames for that schema are found under each schema. These nodes are selectable. Before displaying the tree, the SchemaNameFilter property will be used to filter the list of schemas displayed. If the SchemaNameFilter property is not set, all schemas will be displayed. The Types property will then be used to determine which type nodes to add under the schemas listed. Upon expanding a node, the ObjectNameFilter property will be used to determine which database objects to display.

Edit query can be selected to set query filters based on schema name, type. It also allows to check the “add Business object ASI” and “query business objects” boxes to select stored procedures or set custom query business objects.



From the query filter properties panel, you can filter on schema name or types and optionally select “add business object ASI” to specify stored procedures. Or you can select “Query business object” to set query business objects. The query business count, which represents the maximum integer value of Query business objects, can be generated at one time. This property will be enabled only when the value of QueryBO property is ‘true’. Its default value is 1.

## Business object application specific information

- **StatusColumnName**
  - Name of the database column used to perform logical deletes
- **StatusValue**
  - Value that specifies whether a business object is inactive or deleted
- **StoredProcedureName**
  - Associate with the Stored Procedure Type

The screenshot shows a dialog box titled "Enterprise Service Discovery" with a sub-title "Configuration Parameters for ADDRESS (1 of 2)". The dialog contains the following fields and controls:

- Status Column Name: ADDRID
- Status Value: (empty)
- VerbASL: CreateSP
- Schema Name Filter: (empty)
- Schema: ADCUSER
- SP Name Filter: (empty)
- Stored Procedure: CREATEADDRESS
- Stored Procedure Parameters:
  - VAR0: addrid
  - VAR1: custid
  - VAR2: city
  - VAR3: zipcode

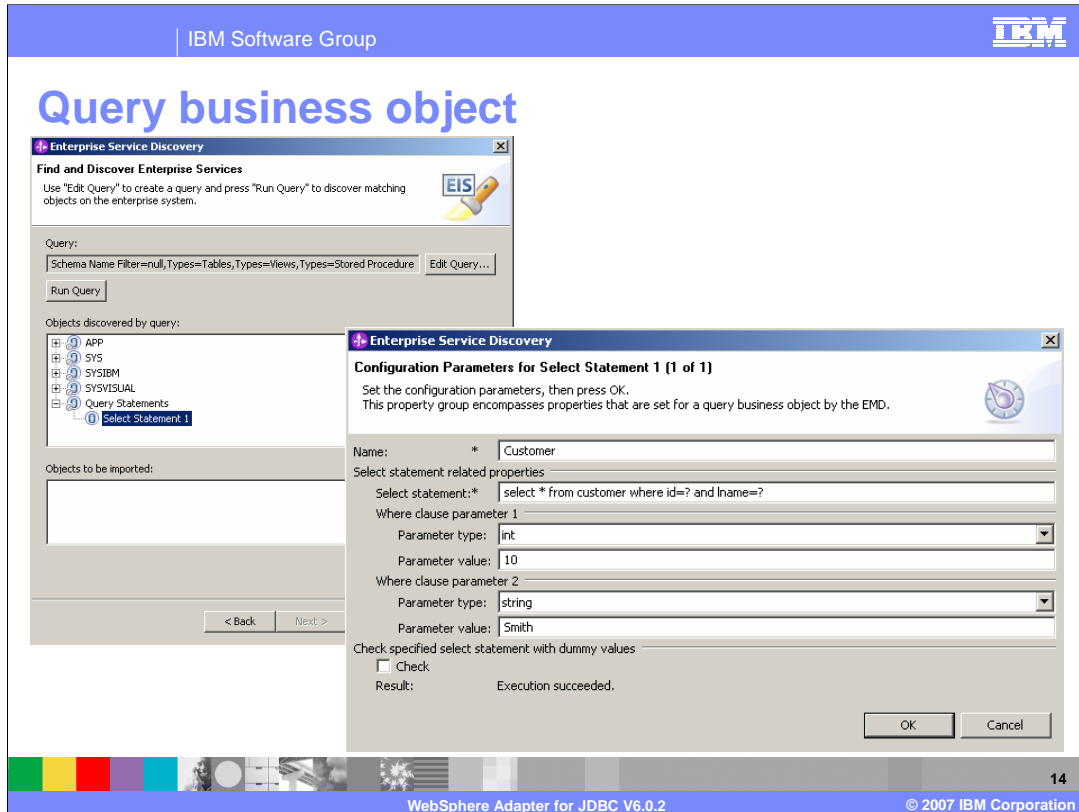
Buttons for "Add...", "Remove", "OK", and "Cancel" are visible.

13

WebSphere Adapter for JDBC V6.0.2

© 2007 IBM Corporation

You can set configuration properties such as StatusColumnName and StatusValue, which come into play when you want to use logical deletes from the database. This is also where you associate the stored procedure type.



In the Configuration Parameters for Select Statement panel, you can select “Check” to see if the specified select statement with dummy values runs correctly. If the statement is valid, you will see the message indicating success. Otherwise, you will need to double check the select statement or configurations for the ‘where’ clause parameters.

## Enterprise service discovery - Inbound

- Activation Specification Properties
- Delivery Mode and Polling Info, Credentials, and Miscellaneous

Properties for Activation Spec

BO Namespace: \*

Delivery Mode and Polling Information

Delivery Type:

Poll period [Int]:

Poll quantity [Int]:

Retry interval [Integer]:

Retry limit [Integer]:

Stop polling on error [Boolean]

Assured once delivery [Boolean]

Filter future events [Boolean]

Event Type filter [String]:

User Credentials

Username: \*

Password: \*

Machine Credentials

Data Source JNDI Name:

Database URL: \*

Jdbc Driver Class: \*

Miscellaneous

Event Order By:

Event Table Name: \*

SP Before Poll:

SP After Poll:

Event Query Type:

15

WebSphere Adapter for JDBC V6.0.2 © 2007 IBM Corporation

Once the query has been run and the selection of business objects has been done, you will then specify Inbound or Outbound service type. Here you can see an example of the set of properties specified with an Inbound service type.

## Enterprise service discovery - Outbound

- J2C Authentication Data Entry and other Managed Connection Factory Properties
- Resource Adapter Properties and Miscellaneous

Deploy connector with module

Specify the connection properties which will be used to connect to the Enterprise Information System at runtime:

Use connection properties specified on server

Use discovered connection properties

J2C Authentication Data Entry:

Managed Connection Properties

User Credentials

Username:

Password:

Machine Credentials

Auto Commit

XA Data Source Name:

XA Database Name:

Data Source JNDI Name:

Database URL: \*

Jdbc Driver Class: \*

Resource Adapter Properties

Logging and Tracing

Adapter ID (String): \*

Log file size (Integer):

Log file name (String):

Log Files (Integer):

Trace file size (Integer):

Trace file name (String):

Trace files (Integer):

Miscellaneous

Ping Query:

Database Vendor: \*

QueryTimeout:

Return Dummy BO For SPI:

Here you can see an example of the set of properties specified with an outbound service type including Miscellaneous



## Section

# ***Configuration properties***

This section provides details of configurations properties of the WebSphere Adapter for JDBC.

## Configuration properties

- Configuration Properties used by the adapter to communicate with the specific database application
- Types of Properties
  - ▶ Enterprise Service Discovery Connection
    - Required properties
  - ▶ Activation Specification
  - ▶ Managed (J2C) Connection Factory
  - ▶ Resource Adapter

This group of configuration properties contains attributes used by the adapter to set up a communication channel to a specific database application. There are various types of properties used such as enterprise service discovery connection which is required for both inbound and outbound processing. For specific adapter configuration properties, they include resource adapter, managed (J2C) connection factory, and activation specification properties.

## Enterprise service discovery connection properties

Property	Description
UserName	▪ User ID to connect to the database
Password	▪ Password for corresponding user ID
Database URL	▪ DatabaseURL used to connect to the database
Prefix	▪ The prefix to be added to the name of the business object.
JDBCDriverClass	▪ JDBC Driver class for the driver used to connect to the database

Shown here is a list of connection properties for Enterprise Service Discovery, which are defined in the Adapter User Guide and Information Center. These properties are required for both Inbounds and Outbound services. These properties are similar to previous version.

## Activation specification properties

Property	Description
SPBeforePoll	<ul style="list-style-type: none"> <li>Any stored procedure that you want to run before the actual poll query is called. It will take one input parameter for poll quantity.</li> </ul>
SPAfterPoll	<ul style="list-style-type: none"> <li>Any stored procedure that you want to run after each poll cycle. It will take one input parameter for poll quantity.</li> </ul>
Retry Interval	<ul style="list-style-type: none"> <li>Time in milliseconds between retries in case of EIS connection failure</li> </ul>
Retry Limit	<ul style="list-style-type: none"> <li>Number of times to attempt to retry the inbound connection</li> </ul>
Stop Polling on Error	<ul style="list-style-type: none"> <li>Stop the adapter when an error is encountered while polling. The default value is False.</li> </ul>
Assured Once Delivery	<ul style="list-style-type: none"> <li>This option provides once-and-only-once event delivery. Turning it off may provide a performance benefit.</li> </ul>
EventFilterType	<ul style="list-style-type: none"> <li>This option will allow to filter the events to be processed by business object type. It has a comma delimited list of business object types, and only the types specified in the property are picked up for processing.</li> </ul>

Here are just highlights of some new activation specification properties available in V6.0.2. These properties hold the inbound event processing configuration information for a message endpoint. They can be set through either the enterprise service discovery wizard or the WebSphere Application Server or WebSphere Enterprise Service Bus administrative console.

More information on a complete list of activation specification properties can be found in the WebSphere Adapter Information Center. The link to the Information Center is provided at the end of this presentation.

## Activation specification properties (cont.)

Property	Description
Filter Future Events	<ul style="list-style-type: none"> <li>If the value is true, the adapter will not process events that have a timestamps in the future.</li> </ul>
EventQueryType	<ul style="list-style-type: none"> <li>This property values will decide whether to use standard event store or custom query. The valid values are Standard (for the standard event store) and Dynamic (for custom event processing)</li> </ul>
CustomEventQuery	<ul style="list-style-type: none"> <li>The SQL query, stored procedure or stored function for custom event processing. This will be run during each poll cycle when the EventQueryType is set to Dynamic.</li> </ul>
CustomUpdateQuery	<ul style="list-style-type: none"> <li>The custom update query that will be run after each event is processed so that the same event does not get picked up for processing in the subsequent event cycles.</li> </ul>
CustomDeleteQuery	<ul style="list-style-type: none"> <li>The custom delete query that will be run after each event is processed.</li> </ul>
DataSourceJNDIName	<ul style="list-style-type: none"> <li>This is the JNDI data source name to be used by the adapter to establish connection to the database.</li> </ul>

New additional inbound properties are added to support user-customizable event processing. These include properties for custom event query, custom update query, custom delete query, Stored Procedure before and after poll. An event query type property will determine if custom query or standard event store should be used for event processing.

## Managed (J2C) connection factory properties

Property	Description
UserName	▪ User ID to connect to the EIS
Password	▪ Associated password for User ID
DatabaseURL	▪ DatabaseURL used to connect to the database
DataSourceJNDIName	▪ JNDI name of the data source used to get the JDBC connection
AutoCommit	▪ Autocommit to be set on the connection
JDBCDriverClass	▪ JDBC Driver class for the driver used to connect to the database
XADataSourceName	▪ XA Data Source name to be used to establish an XA connection to the database
XADatabaseName	▪ Database name to be used for the XA connection

Here are highlights of some of managed connection factory configuration properties. They are used at run time to create an outbound connection instance with an enterprise information system. Once the Managed connection factory properties are created, they are stored in the deployment descriptor.

A J2C connection factory manages connection pooling and provides configuration information for outbound connectivity to a single JDBC application instance from an application by way of the adapter.

## Resource adapter properties

Property	Description
PingQuery	<ul style="list-style-type: none"> <li>SQL Query is test valid connection to the database</li> </ul>
DatabaseVendor	<ul style="list-style-type: none"> <li>Specifies which RDBMS the adapter uses for special processing (DB2, Oracle, or SQL Server, Other for Cloudscape)</li> </ul>
enableHASupport	<ul style="list-style-type: none"> <li>When the enableHASupport property is set to true, only one of the replicated adapter instances actively polls for events while other instances are in standby mode. If the enableHASupport property is set to false, all of the adapter instances replicated on cluster members actively poll for events. This may result in event duplication. Do <i>not</i> change the value of enableHASupport to false for a single server environment.</li> </ul>
QueryTimeOut	<ul style="list-style-type: none"> <li>Sets the QueryTimeOut for all SQL statements to the number of seconds specified. SQL exception is captured for timeout</li> </ul>
ReturnDummyBOForSP	<ul style="list-style-type: none"> <li>Used to return output parameters even when the result set is empty (Dummy business object with values from output/input parameters will be returned)</li> </ul>

Resource adapter properties consist of logging and tracing and activities specific to the adapter, such as the default configuration properties of the adapter. One new property in V6.0.2 is enableHASupport which supports multiple adapter instances in clustered environment. So, if enableHASupport is set to true, only one of the replicated adapter instances actively polls for events while other instances are in standby mode. And if set to false, all of the adapter instances replicated on cluster members actively poll for events. This capability improves adapter performance and availability. You can configure these properties using the enterprise service discovery wizard or the administrative console of the server.

## Section

# ***Summary and references***

This section provides a summary of the WebSphere Adapter for JDBC.



## Summary

- WebSphere Adapter for JDBC enables integration with SCA Applications and JDBC 2.0 Databases
  - ▶ Inbound and Outbound support
- Provided support for Enterprise Metadata Discovery for discovering services
- New and enhanced features support
- Looked at configuration properties for inbound and outbound services

To summarize this presentation, the WebSphere Adapter for JDBC enables integration with SCA business integration applications and Enterprise Information System database applications. The adapter supports both inbound and outbound interaction. Enterprise service discovery is used for discovery of services and creating the service description. It is also used to specify values for custom adapter properties and discovery of business objects. New and enhanced features are now supported in V6.0.2. They include BLOB/CLOB data type, customizing event processing, new event manager support, stored functions support, selection of stored procedures, and user specified SQL processing. Lastly, configuration properties for inbound and outbound services are also examined.

## Feedback

### Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

[Click to send e-mail feedback](#)



You can help improve the quality of IBM Education Assistant content by providing feedback.

## Trademarks, copyrights, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

DB2 IBM WebSphere

J2EE, Java, JDBC, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. 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 in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2007. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.