

IBM WebSphere Business Integration Adapters



Mainframe Adapter Suite User Guide

Adapter Version 24.x

IBM WebSphere Business Integration Adapters



Mainframe Adapter Suite User Guide

Adapter Version 24.x

Note!

Before using this information and the product it supports, read the information in "Notices" on page 179.

19December2003

This edition of this document applies to the Mainframe Adapter Suite version 2.4.x, and to all subsequent releases and modification until otherwise indicated in new editions.

To send us your comments about this document, e-mail doc-comments@us.ibm.com. We look forward to hearing from you.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 2001, 2003. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	v
About this document	vii
Audience	vii
Prerequisites for this document.	vii
Related documents.	vii
Typographic conventions	viii
New in this release.	ix
New in release 2.4.x.	ix
New in release 2.3.x.	ix
New in release 2.2.x.	ix
New in release 2.1.x	x
New in release 1.9.x	x
New in release 1.8.x.	xi
New in release 1.7.x.	xi
New in release 1.6.x.	xi
Chapter 1. Overview	1
Terminology	1
Mainframe Adapter Suite	1
Connector components	2
How the connector works	2
Chapter 2. Installing and configuring the connector.	7
Adapter environment	7
Prerequisites	8
Installing the adapters and related files	9
Installing the MAS drivers.	9
Installed file structure for the ADABAS connector	10
Installed file structure for the CICS/TS connector.	11
Installed file structure for the DB2 connector	12
Installed file structure for the IMS/DB connector.	13
Installed file structure for the IMS/TM connector.	15
Installed file structure for the VSAM connector	16
Installed file structure for the Natural connector	17
Installed file structure for the IDMS Database connector	19
Chapter 3. Configuring driver data sources	21
MAS drivers overview.	21
Configuring the ODBC driver	22
Configuring the JDBC driver	34
Testing the MAS driver connections	35
Tracing	37
Chapter 4. Configuring the connector	45
Enabling the application for the connector	45
Configuring the connector	47
Creating multiple instances of a connector	58
Starting the connector	59
Stopping the connector	61
Chapter 5. Developing business objects for the Mainframe Adapter Suite	63
Business object and attribute naming conventions	63

Business object structure	63
Business object verb processing	68
Business object attribute properties	83
Business object application-specific information	85
Chapter 6. Troubleshooting and error handling	97
Startup problems	97
Event processing	97
Mapping (ICS integration broker only)	97
Error handling and logging	97
Appendix A. Standard configuration properties for connectors	99
New and deleted properties	99
Configuring standard connector properties	99
Summary of standard properties	100
Standard configuration properties	104
Appendix B. Connector Configurator	115
Overview of Connector Configurator	115
Starting Connector Configurator	116
Running Configurator from System Manager	117
Creating a connector-specific property template	117
Creating a new configuration file	119
Using an existing file	120
Completing a configuration file	121
Setting the configuration file properties	122
Saving your configuration file	127
Changing a configuration file	128
Completing the configuration	128
Using Connector Configurator in a globalized environment	128
Appendix C. MAS driver keywords	131
Setting a driver keyword	131
Driver keyword descriptions	133
Driver keyword settings	150
Appendix D. Business object samples.	153
DB2 Connector	153
VSAM Connector	154
NATURAL Connector	156
CICS Connector	161
IDMS(DB) Connector	167
ADABAS Connector	169
IMS_TM Connector	172
Appendix E. Support for null and blank values	177
Pass and fail scenarios	177
Functionality	178
Notices	179
Programming interface information	180
Trademarks and service marks	180

Figures

1. Business object request architecture	2	15. Client debug information	41
2. ODBC Data Source Administrator dialog box	24	16. Select Trace File	42
3. Create new data source dialog box	25	17. Typical single-cardinality relationship	65
4. MAS debug configuration dialog box	25	18. Multiple-cardinality business object relationship	67
5. ODBC data source type	26	19. Single-cardinality business object with relationship stored in the child	67
6. TCP/IP connection type	26	20. Example of relationships among business objects	90
7. LU 6.2 Connection type	27	21. ODBC Data Source Administrator screen	132
8. WebSphere MQ connection type	28	22. Data source name screen	132
9. WebSphere MQ connection information	28	23. Advanced Information Screen	133
10. Client Advanced Information	29	24. Optional Information Screen	133
11. Optional settings	31		
12. Select data source dialog box	36		
13. Debug configuration	40		
14. Advanced information	40		

About this document

The IBM[®] WebSphere[®] Business Integration Adapter portfolio supplies integration connectivity for leading e-business technologies, enterprise applications, and legacy and mainframe systems. The product set includes tools and templates for customizing, creating, and managing components for business process integration.

This document describes the installation, configuration, business object development, and troubleshooting for the IBM WebSphere Business Integration Mainframe Adapter Suite.

Audience

This document is for consultants, developers, and system administrators who use the WebSphere business integration system at customer sites.

Prerequisites for this document

You should be familiar with the:

- IBM WebSphere business integration system (if you are using WebSphere InterChange Server as your integration broker)
- IBM WebSphere MQ Integrator Broker (if you are using MQ Integrator Broker as your integration broker)
- OS/390-based mainframe legacy application that the MAS integrates with the WebSphere business integration adapter: CICS/TS, DB2, VSAM, ADABAS, IMS/TM, IMS/DB, Natural, or IDMS Database.

Related documents

The complete set of documentation available with this product describes the features and components common to all WebSphere Business Integration Adapters installations, and includes reference material on specific components.

You can install related documentation from the following sites:

- For general adapter information; for using adapters with WebSphere message brokers (WebSphere MQ Integrator, WebSphere MQ Integrator Broker, WebSphere Business Integration Message Broker); and for using adapters with WebSphere Application Server, see the IBM WebSphere Business Integration Adapters InfoCenter:
<http://www.ibm.com/websphere/integration/wbiadapters/infocenter>
- For using adapters with WebSphere InterChange Server, see the IBM WebSphere InterChange Server InfoCenters:
<http://www.ibm.com/websphere/integration/wicserver/infocenter>
<http://www.ibm.com/websphere/integration/wbicollaborations/infocenter>
- For more information about WebSphere message brokers:
<http://www.ibm.com/software/integration/mqfamily/library/manualsa/>
- For more information about WebSphere Application Server:
<http://www.ibm.com/software/webservers/appserv/library.html>

These sites contain simple directions for downloading, installing, and viewing the documentation.

Typographic conventions

This document uses the following conventions:

<code>courier font</code>	Indicates a literal value, such as a command name, filename, information that you type, or information that the system prints on the screen.
bold	Indicates a new term the first time that it appears.
<i>italic, italic</i>	Indicates a variable name or a cross-reference.
<i>blue outline</i>	A blue outline, which is visible only when you view the manual online, indicates a cross-reference hyperlink. Click inside the outline to jump to the object of the reference.
<i>ProductDir</i>	Represents the directory where the IBM WebSphere Business Integration Adapters product is installed. The default product directory is WebSphereAdapters.
{ }	In a syntax line, curly braces surround a set of options from which you must choose one and only one.
[]	In a syntax line, square brackets surround an optional parameter.
...	In a syntax line, ellipses indicate a repetition of the previous parameter. For example, <code>option[,...]</code> means that you can enter multiple, comma-separated options.
< >	In a naming convention, angle brackets surround individual elements of a name to distinguish them from each other, as in <code><server_name><connector_name>tmp.log</code> .
/, \	In this document, backslashes (\) are used as the convention for directory paths. For UNIX installations, substitute slashes (/) for backslashes. All WebSphere business integration system product pathnames are relative to the directory where the product is installed on your system.
UNIX/Windows:	Paragraphs beginning with either of these indicate notes listing operating system differences.
u	This symbol indicates the end of a UNIX/Windows paragraph; it can also indicate the end of a multiparagraph note.
<code>%text%</code> and <code>\$text</code>	Text within percent (%) signs indicates the value of the Windows text system variable or user variable. The equivalent notation in a UNIX environment is <code>\$text</code> , indicating the value of the <code>text</code> UNIX environment variable.

New in this release

New in release 2.4.x

Updated in December 2003. The release of this document for adapter version 2.4.x contains the following new or corrected information:

- Information formerly in Chapter 2 about installing the connector has been removed. See chapter 2 for the new location.
- Also in Chapter 2, in Tables 3 and 4, the Subdirectory names have been revised.
- In Chapter 3, under “Syntax for CICS/TS and IMS/TM CALL statements,” the CICS/TS CALL Syntax and the IMS/TM CALL Syntax have been revised. And the statement for making SQL requests for IMS has been revised.
- Under “Configuring the connector” in Chapter 4, the connector-specific configuration properties of ApplicationPassword and ApplicationUserName are not required when trusted authentication is being used.
- Under “Business object verb processing ” in Chapter 5, the instructions for DeltaUpdate operations have been added and the instructions for Delete operations have been revised.
- Under “Business object application-specific information” in Chapter 5, information has been added to describe how a CLOB datatype would be defined.
- Appendix C, “Connector feature list” has been removed.
- In Appendix E, the business object samples have been replaced with new samples.

New in release 2.3.x

Updated in July 2003. The release of this document for adapter version 2.3.x contains the following new or corrected information:

- The adapter can now use WebSphere Application Server as an integration broker. For further information, see “Broker compatibility” on page 7.
- The adapter now runs on the following platforms:
 - HP-UX11i
 - AIX 5.x
 - Solaris 8
- Support for return of result set from Oracle stored procedures has been added.
- Support for CLOB data types has been added.
- Grandparent access support for copy attributes has been added. Copy attributes can now be accessed from the parent, which allows propagation of attributes down the business object hierarchy.
- The restriction has been removed that an eventId must be a numeric data type.

New in release 2.2.x

Updated in March 2003. The “CrossWorlds” name is no longer used to describe an entire system or to modify the names of components or tools, which are otherwise mostly the same as before. For example “CrossWorlds System Manager” is now “System Manager,” and “CrossWorlds InterChange Server” is now “WebSphere InterChange Server.”

The release of this document for connector version 2.2.x contains the following new or corrected information:

- Support has been added for the following:
 - Wrapper objects at the top-level of business objects
 - LIKE operator
 - Hex/binary data
 - Stored procedures for the RetrieveUpdate verb
 - Verb Application Specific Information for RetrieveByContent
 - Verb Application Specific Information in the WHERE clause when the WHERE clause length is 0 in RetrieveByContent
- The ConnectorID property has been changed from an int to a String to allow for a more descriptive name.
- The DRIVERLIB variable has been added to point to the native libraries used by the custom JDBC drivers.
- More functionality has been added to check for loss of database connectivity during object processing.

New in release 2.1.x

The IBM WebSphere Business Integration Adapters in the Mainframe Connector Suite have been released with the same functionality as in previous releases.

New in release 1.9.x

The IBM WebSphere Business Integration Adapters in the Mainframe Connector Suite includes connectors for the following adapters:

- WebSphere Business Integration Adapter for IMS Database Manager
- WebSphere Business Integration Adapter for IMS Transaction Manager
- WebSphere Business Integration Adapter for DB2 Databases
- WebSphere Business Integration Adapter for CICS
- WebSphere Business Integration Adapter for VSAM
- WebSphere Business Integration Adapter for ADABAS

These adapters operate with both the InterChange Server (ICS) and WebSphere MQ Integrator integration brokers. An integration broker, which is an application that performs integration of heterogeneous sets of applications, provides services that include data routing. The adapters include:

- An application component specific to the adapter
- Sample business objects
- IBM WebSphere Adapter Framework, which consists of:
 - Connector Framework
 - Development tools (including Business Object Designer and Connector Configurator)
 - APIs (including CDK)

This manual provides information about using this adapter with both integration brokers: ICS and MQIntegrator.

Important: Because the connector has not been internationalized, do not run it against ICS version 4.1.1 if you cannot guarantee that only ISO Latin-1 data will be processed.

New in release 1.8.x

The Mainframe Connector Suite has been enabled for IBM CrossWorlds 4.1.x.

New in release 1.7.x

CrossWorlds installation now provides the IBM branded JDBC driver for the MS SQL Server to replace the WebLogic JDBC driver, and continues to provide the Oracle thin driver.

New in release 1.6.x

Version 1.6.x marks the first public release of the CrossWorlds Mainframe Connectivity Suite. The connectors include the following features:

- Full support for configuring data sources on UNIX and Windows. See Chapter 3, “Configuring driver data sources,” on page 21.
- A utility to test the MCS driver connections. See “Testing the MAS driver connections” on page 35.
- Support for stored procedures. See “Using stored procedures” on page 76.

Chapter 1. Overview

This chapter documents the components of the WebSphere Business Integration Adapters in the Mainframe Adapter Suite (MAS) and presents a high-level description of how the connectors work. It contains the following sections:

- “Terminology”
- “Mainframe Adapter Suite”
- “Connector components” on page 2
- “How the connector works” on page 2

The connectors for MAS are runtime components of the WebSphere Business Integration Adapters in the MAS. The connectors allow the WebSphere integration broker to exchange business objects with MAS legacy-enabled business processes.

Terminology

The Mainframe Adapter Suite incorporates, and is bundled with, third party software and documentation. The terminology used in the third party documentation differs from that used in WebSphere documentation as shown in Table 1:

Table 1. MainFrame Adapter Suite terminology

WebSphere usage	Third party usage
Mainframe Agent	Shadow Server
MAS Drivers	Neon Client -JDBC/ODBC Drivers

Note: *Legacy application* and *application* are used interchangeably.

Mainframe Adapter Suite

The Mainframe Adapter Suite (MAS) is a collection of eight connectors, an MAS driver component, and a Mainframe Agent. The MAS integrates the following IBM OS/390-based mainframe legacy applications with the WebSphere business integration environment:

- CICS/TS
- DB2
- VSAM
- ADABAS
- IMS/TM
- IMS/DB
- Natural
- IDMS Database

The MAS driver component and Mainframe Agent provide a single technology interface for all of the MAS connectors. Like all of the MAS connectors, the MAS drivers must be configured to support the legacy application you wish to integrate. See Chapter 3, “Configuring driver data sources,” on page 21, for more on configuring data sources for the drivers.

Connector components

Connectors consist of an application-specific component and the connector framework. The application-specific component contains code tailored to a particular application. The connector framework, whose code is common to all connectors, acts as an intermediary between the integration broker and the application-specific component. The connector framework provides the following services between the integration broker and the application-specific component:

- Receives and sends business objects
- Manages the exchange of startup and administrative messages

The connectors allow the integration broker to exchange business objects with legacy applications built on the IBM OS/390 mainframe system. This section presents a high-level description of the connector's architecture.

When the connector is started, it establishes a connection pool with the legacy application. It uses connections from this pool for all transaction processing with the database/application. On termination of the connector, all connections in the pool are closed.

Connector architecture

Figure 1 shows the MAS components and their relationships within the business integration system.

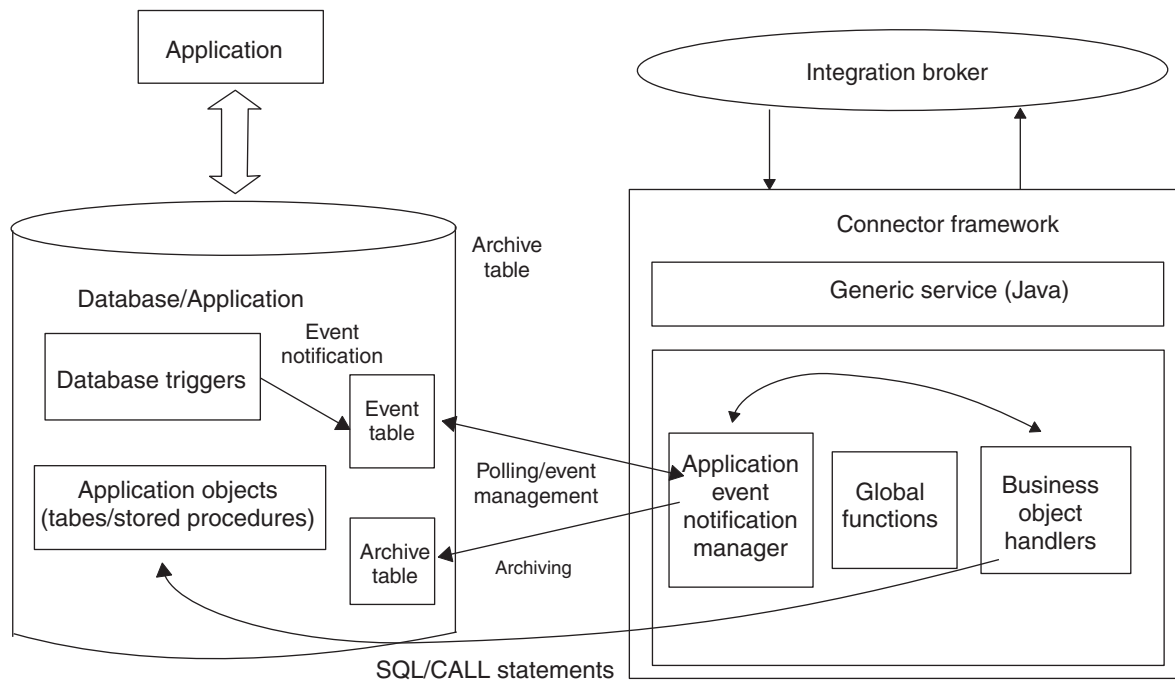


Figure 1. Business object request architecture

How the connector works

This section describes how meta-data enhances the connector's flexibility and presents a high-level description of business object processing and event notification.

The connector and meta-data

The connector is meta-data-driven. **Meta-data**, in the WebSphere environment, is application-specific data that is stored in business objects and that assists the connector in its interaction with the application. A meta-data-driven connector handles each business object that it supports based on meta-data encoded in the business object definition rather than on instructions hardcoded in the connector.

Business object meta-data includes the structure of a business object, the settings of its attribute properties, and the content of its application-specific information. Because the connector is meta-data driven, it can handle new or modified business objects without requiring modifications to the connector code.

The connector executes SQL statements or stored procedures to retrieve or change data in the legacy application. To build dynamic SQL statements or stored procedures, the connector uses application-specific meta-data. These SQL statements and stored procedures perform the required retrieval from or changes to the legacy application for the business object and for the verb that the connector is processing. For details about using application-specific information, see Chapter 5, “Developing business objects for the Mainframe Adapter Suite,” on page 63.

Business object processing

This section provides an overview of how the connector processes business object requests and application events. For more detailed information, see “Business object verb processing” on page 68.

Processing business object requests

When the connector receives a request to perform an application operation, the connector processes hierarchical business objects recursively; that is, it performs the same steps for each child business object until it has processed all individual business objects. The order in which the connector processes child business objects and the top-level business object depends on whether the child business objects are contained with or without ownership and whether they are contained with single cardinality or multiple cardinality.

Note: 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 a single business object, independent of any child business objects it might contain or that contain it. The term **top-level** business object refers to the individual business object at the top of the hierarchy that does not itself have a parent business object.

Business object retrieval: When a business object request asks the connector to retrieve a hierarchical business object from the legacy application, the connector attempts to return a business object that exactly matches the current application representation of that business object. In other words, all simple attributes of each individual business object returned match the value of the corresponding field in the application. Also, the number of individual business objects in each array contained by the returned business object match the number of children in the application for that array.

To perform such a retrieval, the connector uses the primary key values in the top-level business object received from the integration broker to recursively descend through the corresponding data in the application.

Business object RetrievalByContent: When a business object request asks the connector to retrieve a hierarchical business object based on values in non-key attributes in the top-level business object, the connector uses the value of all non-null attributes as the criteria for retrieving the data.

Business object creation: When a business object request asks the connector to create a hierarchical business object in the application, the connector performs the following steps:

1. Recursively creates each single-cardinality child business object contained with ownership into the application. Processes each single-cardinality child business object contained without ownership.

Note: All single cardinality child business objects are processed based on occurrence in the business object and before the parent business object is processed. Child object ownership and non-ownership do not determine the processing sequence, but do determine the type of processing.

2. Creates the top-level business object in the application.
3. Creates each single-cardinality child business object that stores the parent/child relationship in the child.
4. Creates each multiple-cardinality child business object.

Business object modification: When a business object request asks the connector to update a hierarchical business object in the application, the connector performs the following steps:

1. Uses the primary key values of the source business object to retrieve the corresponding entity from the application.
2. Recursively updates all single-cardinality children of the top-level business object.
3. Updates all simple attributes of the retrieved business object except those whose corresponding attribute in the source business object contain the value CxIgnore.
4. Processes all arrays of the retrieved business object.

Business object deletion: When a business object request asks the connector to delete a hierarchical business object from the application, the connector performs the following steps:

1. Deletes the single-cardinality children.
2. Deletes the multiple-cardinality children.
3. Deletes the top-level business object.

Processing application events

An event table should be used to store events associated with business object operations. The connector is notified of application events such as create, update and delete by means of a database trigger or some other suitable mechanism. The event table must be populated with these application events. The event table resides in the application database.

The connector handles the Create, Update, and Delete events generated by the application in the manner described below.

Note: The CICS/TS, Natural, IDMS Database, VSAM, and IMS/TM legacy applications may not support all of the features described below.

Create notification: When the connector encounters a Create event in the event table, it creates a business object of the type specified by the event, sets the key values for the business object (using the keys specified in the event table), and retrieves the business object from the application. After it retrieves the business object, the connector sends it with the Create verb to the integration broker.

Update notification: When the connector encounters an Update event in the event table, it creates a business object of the type specified by the event, sets the key values for the business object (using the keys specified in the event table), and retrieves the business object from the application. After it retrieves the business object, the connector sends it with the Update verb.

Delete notification: When the connector encounters a Delete event in the event table, it creates a business object of the type specified by the event, sets the key values for the business object (using the keys specified in the event table), and sends it with the Delete verb to the integration broker. All values other than the key values are set to CxIgnore. If any of the non-key fields are significant at your site, modify the value of the fields.

The connector handles logical and physical Delete operations that are triggered by its application. In the case of physical deletes, the SmartFiltering mechanism removes all of the business object's unprocessed events (such as Create or Update) before inserting the Delete event into the event table. In the case of logical deletes, the connector inserts a Delete event in the event table without removing other events for the business object.

Retrieving business objects for event processing: A Retrieve can be done in two ways on a business object for event processing. The first is a Retrieve based on key attributes in a business object. The second is a Retrieve based on both key and non-key attributes. In this case, the business object needs to support the RetrieveByContent verb and must use name_value pair for the object keys.

Note: If the object key does not use name_value pair, the keys in the object key field should follow the same order as the keys in the business object.

Event processing

The connector's event detection mechanism uses an event table, an archive table, stored procedures, and database triggers. Because there are potential failure points associated with the processing of events, the event management process does not delete an event from the event table until it has been inserted into the archive table.

The database triggers populate an event table whenever an event of interest occurs in the application database. The connector polls this table at a regular, configurable interval, retrieves the events, and processes the events first by priority and then sequentially. When the connector has processed an event, the event's status is updated. You must add the triggers to the application database as part of the installation procedure.

The setting of its ArchiveProcessed property determines whether the connector archives an event into the archive table after updating its status. For more information on the ArchiveProcessed property, see Chapter 4, "Configuring the connector," on page 45.

Table 2 on page 6 illustrates the archiving behavior depending on the setting of the ArchiveProcessed property.

Table 2. Archiving behavior

Archive processed setting	Reason deleted from event table	Connector behavior
true or no value	Successfully processed	Archived with status of Sent to InterChange
	Unsuccessfully processed	Archived with status of Error
false	No subscription for business object	Archived with status of Unsubscribed
	Successfully processed	Not archived and deleted from event table
	Unsuccessfully processed	Remains in event table with status of s
	No subscription for business object	Remains in event table with status of Unsubscribed

WebSphere has implemented SmartFiltering, a mechanism within the database triggers that minimizes the amount of processing the integration broker performs. For example, if an application has updated the Contract business object 15 times since the connector last polled for events, the SmartFiltering stores those changes as a single Update event.

Handling lost database connections

There are numerous reasons for losing a connection to the application or database. If this occurs, the connector terminates. The JDBC specification does not provide a mechanism for detecting lost connections. The "PingQuery" on page 54 property is provided to handle this detection. If a failure occurs during a service call request, the connector executes this PingQuery to confirm that the failure was not due to a lost connection to an application/database. If the PingQuery fails and the AutoCommit property is set to false, the connector will attempt to create a new connection to the database. If it succeeds in creating a new connection to the database it will continue processing, otherwise the connector returns an APPRESPONSETIMEOUT, which results in the termination of the connector.

The "PingQuery" on page 54 is executed if a failure occurs when accessing an application for any type of transaction. For example:

- While accessing the event and archive tables
- While retrieving the business object that is related to the event
- While creating or updating a record pertaining to a business object

Chapter 2. Installing and configuring the connector

This chapter describes how to install an MAS connector. Subsequent chapters describe how to configure the drivers and connector.

This chapter contains the following sections:

- “Adapter environment”
- “Prerequisites” on page 8
- “Installing the MAS drivers” on page 9
- “Installed file structure for the ADABAS connector” on page 10
- “Installed file structure for the CICS/TS connector” on page 11
- “Installed file structure for the DB2 connector” on page 12
- “Installed file structure for the IMS/DB connector” on page 13
- “Installed file structure for the IMS/TM connector” on page 15
- “Installed file structure for the VSAM connector” on page 16
- “Installed file structure for the Natural connector” on page 17
- “Installed file structure for the IDMS Database connector” on page 19

Adapter environment

Before installing, configuring, and using the adapter, you must understand its environment requirements. They are listed in the following section.

- “Broker compatibility”
- “Adapter platforms” on page 8
- “Globalization” on page 8

Broker compatibility

The adapter framework that an adapter uses must be compatible with the version of the integration broker (or brokers) with which the adapter is communicating. The 2.3.x version of the adapter for MAS is supported on the following adapter framework and integration brokers:

- **Adapter framework:**
WebSphere Business Integration Adapter Framework versions 2.3.x and 2.4.
- **Integration brokers:**
 - WebSphere InterChange Server, versions 4.1.1, 4.2, 4.2.1, 4.2.2
 - WebSphere MQ Integrator, version 2.1.0
 - WebSphere MQ Integrator Broker, version 2.1.0
 - WebSphere Business Integration Message Broker, version 5.0
 - WebSphere Application Server Enterprise, version 5.0.2, with WebSphere Studio Application Developer Integration Edition, version 5.0.1

See *Release Notes* for any exceptions.

Note: For instructions on installing your integration broker and its prerequisites, see the following guides.

For WebSphere InterChange Server (ICS), see *IBM WebSphere InterChange Server System Installation Guide for UNIX or for Windows*.

For WebSphere message brokers, see *Implementing Adapters with WebSphere*

Message Brokers.

For WebSphere Application Server, see *Implementing Adapters with WebSphere Application Server*.

Adapter platforms

The adapter is supported on the following software.

Operating systems:

- AIX 5.1, AIX 5.2
- Solaris 7.0, Solaris 8.0
- HP UX 11.0, HP UX 11i
- Windows 2000

Globalization

This adapter is DBCS (double-byte character set)-enabled.

Prerequisites

Before you use an MAS connector, you must do the following:

- Install WebSphere business integration system software version 4.2.0 or later or WebSphere Business Integration Adapter Framework version 2.0 or later, including the Adapter Development Kit (ADK).

If the connector runs on a different machine from WebSphere InterChangeServer (ICS) or WebSphere MQ Integrator Broker, install the ADK that is compatible with the ICS or WebSphere MQ Integrator Broker version.

- Install JDK 1.3.1 software. See the *IBM WebSphere InterChange Server System Installation Guide for Windows* or *System Installation Guide for UNIX*.
- Install Neon Client version 3.8.788
- Verify that the Mainframe Agent¹ has been installed
- Verify the existence of a user account in the application

For the connector to process data in the application, with which it talks directly, it must have access to a user account and password that is valid for the application. The user account must have the privileges to retrieve, insert, update, and delete data from the application's database. If you do not already have such an account, you must create one.

MAS Drivers Prerequisites

The Mainframe Agent must be installed and you should be on one of the following platforms currently supported by the MAS ODBC and JDBC drivers.

- Windows
 - Windows 3.1
 - Windows 95
 - Windows 98
 - Windows 2000
- UNIX
 - Solaris 2.4 and above
 - HP-UX 11i

1. Shadow Server

Installing the adapters and related files

For information on installing WebSphere Business Integration adapter products, refer to the *Installation Guide for WebSphere Business Integration Adapters*, located in the WebSphere Business Integration Adapters Infocenter at the following site:

<http://www.ibm.com/websphere/integration/wbiadapters/infocenter>

Installing the MAS drivers

The MAS drivers are required for every MAS connector. The following subsections describe the contents of the MAS drivers on UNIX and Windows systems.

The MAS drivers on a UNIX system

Table 3 shows the file structure of the MAS drivers for a UNIX system.

Table 3. Installed UNIX file structure for the MAS drivers

Subdirectory of <i>\$ProductDir</i>	Description
IBM Shadow Client/Shadow	Contains Readme12, a text file describing any changes in file location, class path environment variables, and other information.
IBM Shadow Client/Shadow	Contains the MCS driver <i>dsa</i> , <i>neontrace.set</i> , <i>odbc (driver)</i> , <i>scjdlog</i> , and <i>scodbcdm</i> files.
IBM Shadow Client/Shadow	Contains the <i>jdbc2_0-stdext.jar</i> , <i>scjd12.jar</i> , and <i>scjd12ts.jar</i> files as well as the Java Class files <i>scjds01</i> and <i>scjds02</i> .
IBM Shadow Client/Shadow	Contains the <i>libscryp.so</i> , <i>libscjd12.so</i> , <i>libscjd12ts.so</i> , <i>libscodbc_r.so.1</i> , <i>libscodbcts_r.so.1</i> , and <i>libscssl.so</i> files.

The MAS drivers on a Windows system

Table 4 shows the file structure of the MAS drivers for a Windows system.

Table 4. Installed Windows file structure for the MAS drivers

Subdirectory of <i>%ProductDir%</i>	Description
IBM Shadow Client/Shadow	Contains <i>Dsa32.exe</i> , <i>scjdlog.exe</i> , <i>scod32dm.exe</i> , <i>scod32dm.map</i> , and <i>Vbdemo.32.exe</i> .
IBM Shadow Client/Shadow	Contains <i>scjd12.jar</i> and <i>scjdts12.jar</i> files as well as the Java Class file <i>scjds01</i> and <i>Readme11</i> , a text file describing any changes in file location, class path environment variables, and other information. The <i>scjd12.dll</i> and <i>scjd12ts.dll</i> files have also been added.
IBM Shadow Client/Shadow	Contains the <i>.dll</i> files <i>Scod32.dll</i> , <i>Scod32r.dll</i> , <i>Scod32tr.dll</i> , and <i>Scod32ts.dll</i>

Installed file structure for the ADABAS connector

The following subsections describe the installed file structure of the ADABAS connector on a UNIX or Windows system.

Installed file structure on a UNIX system

Table 5 describes the UNIX file structure used by the connector.

Table 5. Installed UNIX file Structure for the MAS adapter for ADABAS

Subdirectory of <i>\$ProductDir</i>	Description
connectors/ADABAS	Contains the connector CWADABAS.jar and the start_ADABAS.sh files. The startup script for the connector. The script is called from the generic connector manager script. When you click Install from Connector Configurator (WebSphere MQ Integration Broker as the integration broker) or the Connector Configuration screen of System Manager (ICS as the integration broker), the installer creates a customized wrapper for this connector manager script. When the connector works with ICS, use this customized wrapper only to start and stop the connector. When the connector works with WebSphere MQ Integration Broker, use this customized wrapper only to start the connector; use mqsiremotestopadapter to stop the connector.
connectors/messages	Contains the ADABASConnector.txt file.
repository/ADABAS	Contains the CN_ADABAS.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for UNIX* (when ICS is used as integration broker)
- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45.

Installed file structure on a Windows system

Table 6 describes the Windows file structure used by the connector.

Table 6. Installed Windows file structure for the MAS connector for ADABAS

Subdirectory of <i>%ProductDir%</i>	Description
connectors\ADABAS	Contains the connector CWADABAS.jar and the start_ADABAS.bat files.
connectors\messages	Contains the ADABASConnector.txt file.
repository\ADABAS	Contains the CN_ADABAS.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for Windows* (when ICS is used as integration broker)
- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45.

Installed file structure for the CICS/TS connector

The following subsections describe the installed file structure of the CICS/TS connector on a UNIX or Windows system.

Installed file structure on a UNIX system

Table 7 describes the UNIX file structure used by the connector.

Table 7. Installed UNIX file structure for the MAS connector for CICS/TS

Subdirectory of <code>\$ProductDir</code>	Description
connectors/Cics	Contains the connector <code>CWCICs.jar</code> and the <code>start_CICS.sh</code> files. The startup script for the connector. The script is called from the generic connector manager script. When you click Install from Connector Configurator (WebSphere MQ Integration Broker as the integration broker) or the Connector Configuration screen of System Manager (ICS as the integration broker), the installer creates a customized wrapper for this connector manager script. When the connector works with ICS, use this customized wrapper only to start and stop the connector. When the connector works with WebSphere MQ Integration Broker, use this customized wrapper only to start the connector; use <code>mqsiremotestopadapter</code> to stop the connector.
connectors/messages	Contains the <code>CICSConnector.txt</code> file.
repository/Cics	Contains the <code>CN_CICS.txt</code> file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for UNIX* (when ICS is used as integration broker)
- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45.

Installed file structure on a Windows system

Table 8 describes the Windows file structure used by the connector.

Table 8. Installed Windows file structure for the MAS connector for CICS/TS

Subdirectory of %ProductDir%	Description
connectors\Cics	Contains the connector CWCICS.jar and the start_CICS.bat files.
connectors\messages	Contains the CICSConnector.txt file.
repository\Cics	Contains the CN_CICS.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for Windows* (when ICS is used as integration broker)
- *IBM WebSphere Business Integration Adapter Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45.

Installed file structure for the DB2 connector

The following subsections describe the installed file structure of the DB2 connector on a UNIX or Windows system.

Installed file structure on a UNIX system

Table 9 describes the UNIX file structure used by the connector.

Table 9. Installed UNIX file structure for the MAS connector for DB2

Subdirectory of \$ProductDir	Description
connectors/Db2	Contains the connector CWDB2.jar and the start_DB2.sh files. The startup script for the connector. The script is called from the generic connector manager script. When you click Install from Connector Configurator (WebSphere MQ Integration Broker as the integration broker) or the Connector Configuration screen of System Manager (ICS as the integration broker), the installer creates a customized wrapper for this connector manager script. When the connector works with ICS, use this customized wrapper only to start and stop the connector. When the connector works with WebSphere MQ Integration Broker, use this customized wrapper only to start the connector; use mqsi remotestopadapter to stop the connector.
connectors/messages	Contains the DB2Connector.txt file.
repository/Db2	Contains the CN_DB2.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for UNIX* (when ICS is used as integration broker)
- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as the integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45.

Installed file structure on a Windows system

Table 10 describes the Windows file structure used by the connector.

Table 10. Installed Windows file structure for the MAS connector for DB2

Subdirectory of %ProductDir%	Description
connectors\Db2	Contains the connector CWDB2.jar and the start_DB2.bat files.
connectors\messages	Contains the DB2Connector.txt file.
repository\Db2	Contains the CN_DB2.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for Windows* (when ICS is used as integration broker)
- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45.

Installed file structure for the IMS/DB connector

The following subsections describe the installed file structure of the IMS/DB connector on a UNIX or Windows system.

Installed file structure on a UNIX system

Table 11 describes the UNIX file structure used by the connector.

Table 11. Installed UNIX file structure for the MAS connector for IMS/DB

Subdirectory of <i>\$ProductDir</i>	Description
connectors/Ims_db	Contains the connector CWIMS_DB.jar and the start_IMS_DB.sh files. The startup script for the connector. The script is called from the generic connector manager script. When you click Install from Connector Configurator (WebSphere MQ Integration Broker as the integration broker) or the Connector Configuration screen of System Manager (ICS as the integration broker), the installer creates a customized wrapper for this connector manager script. When the connector works with ICS, use this customized wrapper only to start and stop the connector. When the connector works with WebSphere MQ Integration Broker, use this customized wrapper only to start the connector; use mqsiremotestopadapter to stop the connector.
connectors/messages	Contains the IMS_DBConnector.txt file.
repository/Ims_db	Contains the CN_IMS_DB.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for UNIX* (when ICS is used as integration broker)
- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45.

Installed file structure on a Windows system

Table 12 describes the Windows file structure used by the connector.

Table 12. Installed Windows file structure for the MAS connector for IMS/DB

Subdirectory of <i>%ProductDir%</i>	Description
connectors\Ims_db	Contains the connector CWIMS_DB.jar and the start_IMS_DB.bat files.
connectors\messages	Contains the IMS_DBConnector.txt file.
repository\Ims_db	Contains the CN_IMS_DB.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange SErver System Installation Guide for Windows* (when ICS is used as integration broker)

- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45.

Installed file structure for the IMS/TM connector

The following subsections describe the installed file structure of the IMS/TM connector on a UNIX or Windows system.

Installed file structure on a UNIX system

Table 13 describes the UNIX file structure used by the connector.

Table 13. Installed UNIX file structure for the MAS connector for IMS/TM

Subdirectory of <i>\$ProductDir</i>	Description
connectors/Im _s _tm	Contains the connector CWIMS_TM.jar and the start_IMS_TM.sh files. The startup script for the connector. The script is called from the generic connector manager script. When you click Install from Connector Configurator (WebSphere MQ Integration Broker as the integration broker) or the Connector Configuration screen of System Manager (ICS as the integration broker), the installer creates a customized wrapper for this connector manager script. When the connector works with ICS, use this customized wrapper only to start and stop the connector. When the connector works with WebSphere MQ Integration Broker, use this customized wrapper only to start the connector; use mqsi remotestopadapter to stop the connector.
connectors/messages	Contains the IMS_TMConnector.txt file.
repository/Im _s _tm	Contains the CN_IMS_TM.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for UNIX* (when ICS is used as integration broker)
- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45.

Installed file structure on a Windows system

To install the connector on a Windows system, run IBM WebSphere InterChange Server Installer and select the connector. IBM WebSphere InterChange Server Installer installs standard files associated with the connector. Table 14 on page 16 describes the Windows file structure used by the connector.

Table 14. Installed Windows file structure for the MAS connector for IMS/TM

Subdirectory of %ProductDir%	Description
connectors\Ims_tm	Contains the connector CWIMS_TM.jar and the start_IMS_TM.bat files.
connectors\messages	Contains the IMS_TMConnector.txt file.
repository\Ims_tm	Contains the CN_IMS_TM.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for Windows* (when ICS is used as integration broker)
- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45.

Installed file structure for the VSAM connector

The following subsections describe the installed file structure of the VSAM connector on a UNIX or Windows system.

Installed file structure on a UNIX system

Table 15 describes the UNIX file structure used by the connector.

Table 15. Installed UNIX file structure for the MAS connector for VSAM

Subdirectory of \$ProductDir	Description
connectors/Vsam	Contains the connector CWVSAM.jar and the start_VSAM.sh files. The startup script for the connector. The script is called from the generic connector manager script. When you click Install from Connector Configurator (WebSphere MQ Integration Broker as the integration broker) or the Connector Configuration screen of System Manager (ICS as the integration broker), the installer creates a customized wrapper for this connector manager script. When the connector works with ICS, use this customized wrapper only to start and stop the connector. When the connector works with WebSphere MQ Integration Broker, use this customized wrapper only to start the connector; use mqsiremotestopadapter to stop the connector.
connectors/messages	Contains the VSAMConnector.txt file.
repository/Vsam	Contains the CN_VSAM.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for UNIX* (when ICS is used as integration broker)

- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45

Installed file structure on a Windows system

Table 16 describes the Windows file structure used by the connector.

Table 16. Installed Windows file structure for the MAS connector for VSAM

Subdirectory of %ProductDir%	Description
connectors\Vsam	Contains the connector CWVSAM.jar and the start_VSAM.bat files.
connectors\messages	Contains the VSAMConnector.txt file.
repository\Vsam	Contains the CN_VSAM.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for Windows* (when ICS is used as integration broker)
- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45.

Installed file structure for the Natural connector

The following subsections describe the installed file structure of the Natural connector on a UNIX or Windows system.

Installed file structure on a UNIX system

Table 17 describes the UNIX file structure used by the connector.

Table 17. Installed UNIX file structure for the MAS connector for Natural

Subdirectory of <i>\$ProductDir</i>	Description
connectors/Natural	Contains the connector BIA_NATURAL.jar and the start_NATURAL.sh files. The startup script for the connector. The script is called from the generic connector manager script. When you click Install from Connector Configurator (WebSphere MQ Integration Broker as the integration broker) or the Connector Configuration screen of System Manager (ICS as the integration broker), the installer creates a customized wrapper for this connector manager script. When the connector works with ICS, use this customized wrapper only to start and stop the connector. When the connector works with WebSphere MQ Integration Broker, use this customized wrapper only to start the connector; use mqsi remotestopadapter to stop the connector.
connectors/messages	Contains the BIA_NATURALConnector.txt file.
repository/Natural	Contains the BIA_CN_NATURAL.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for UNIX* (when ICS is used as integration broker)
- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45

Installed file structure on a Windows system

Table 18 describes the Windows file structure used by the connector.

Table 18. Installed Windows file structure for the MAS connector for Natural

Subdirectory of <i>%ProductDir%</i>	Description
connectors\Natural	Contains the connector BIA_NATURAL.jar and the start_NATURAL.bat files.
connectors\messages	Contains the BIA_NATURALConnector.txt file.
repository\NATURAL	Contains the BIA_CN_NATURAL.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for Windows* (when ICS is used as integration broker)

- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45.

Installed file structure for the IDMS Database connector

The following subsections describe the installed file structure of the IDMS Database connector on a UNIX or Windows system.

Installed file structure on a UNIX system

Table 19 describes the UNIX file structure used by the connector.

Table 19. Installed UNIX file structure for the MAS connector for IDMS Database

Subdirectory of <i>\$ProductDir</i>	Description
connectors/IDMS_DB	Contains the connector BIA_IDMS_DB.jar and the start_IDMS_DB.sh files. The startup script for the connector. The script is called from the generic connector manager script. When you click Install from Connector Configurator (WebSphere MQ Integration Broker as the integration broker) or the Connector Configuration screen of System Manager (ICS as the integration broker), the installer creates a customized wrapper for this connector manager script. When the connector works with ICS, use this customized wrapper only to start and stop the connector. When the connector works with WebSphere MQ Integration Broker, use this customized wrapper only to start the connector; use mqsiremotestopadapter to stop the connector.
connectors/messages	Contains the BIA_IDMS_DBConnector.txt file.
repository/Idms_db	Contains the BIA_CN_IDMS_DB.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for UNIX* (when ICS is used as integration broker)
- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45

Installed file structure on a Windows system

Table 20 describes the Windows file structure used by the connector.

Table 20. Installed Windows file structure for the MAS connector for IDMS Database

Subdirectory of %ProductDir%	Description
connectors\Idms_db	Contains the connector BIA_IDMS_DB.jar and the start_IDMS_DB.bat files.
connectors\messages	Contains the BIA_IDMS_DBConnector.txt file.
repository\Idms_db	Contains the BIA_CN_IDMS_DB.txt file.

For more information on installing the connector component, refer to one of the following guides, depending on the integration broker you are using:

- *IBM WebSphere InterChange Server System Installation Guide for Windows* (when ICS is used as integration broker)
- *IBM WebSphere Business Integration Adapters Implementation Guide for WebSphere MQ Integrator Broker* (when MQ Integrator Broker is used as integration broker)

Note: Before you can use the connector, you must configure it. See Chapter 4, “Configuring the connector,” on page 45.

Chapter 3. Configuring driver data sources

This chapter describes how to configure data sources for the Mainframe Adapter Suite drivers. It contains the following sections:

- “MAS drivers overview”
- “Configuring the ODBC driver” on page 22
- “Configuring the JDBC driver” on page 34
- “Testing the MAS driver connections” on page 35
- “Tracing” on page 37

MAS drivers overview

The MAS drivers provide connectivity to mainframe systems with both an ODBC (Open DataBase Connectivity) and a JDBC (Java Database Connectivity) API. The ODBC and JDBC drivers provide the following:

- Access to DB2, CICS, IMS/TM, VSAM, ADABAS, Natural, IDMS Database, and other sources
- Multiple data sources with one connection
- Tight integration with OS/390 security (RACF, ACF2, and Top Secret) for authentication and message handling
- Work over TCP/IP, SNA, or WebSphere MQ
- DBCS support

Many Web-based applications that access DB2 UDB for OS/390 require support for the stored procedures in DB2. The MAS Mainframe Agent and MAS drivers provide full support for DB2's stored procedures, which enable JDBC applications to call COBOL, C/C++, PL/I, or Assembler programs that access DB2 as well as CICS/TS, IMS/TM, IMS/DB, ADABAS, VSAM, Natural, IDMS Database, and more.

Windows versus UNIX environments

In the Windows environment, the ODBC application interfaces directly with the driver manager provided freely by Microsoft. This driver manager is responsible for actually loading the ODBC driver and invoking the ODBC APIs; the application does not contain any static reference to the driver.

In the UNIX environment, ODBC driver managers are generally not available unless purchased from third-party vendors. Because most UNIX applications tend to be larger in scope and complexity, an application rarely needs to switch from one DBMS to another to take advantage of the standardized ODBC API interface. In most cases, the application is designed to connect to one DBMS type. For this reason the driver manager is unnecessary and often introduces another point of failure in an already complex client/server design.

The MAS ODBC drivers are designed to allow the application to bypass the driver manager and be linked with the driver libraries directly. However, in rare cases when a driver manager is required, a few simple steps can be taken to allow it to work.

MAS ODBC driver

The MAS ODBC driver implements Level 2 of the ODBC specification and supports many of the current Level 3 applications. It also provides ODBC access to non-relational OS/390 data sources such as those that access existing IMS transactions, CICS programs, and VSAM and ADABAS data, or via customer-written OS/390-resident programs. The MAS ODBC driver is used with client/server applications, 3-tier Web servers, EIS, ad hoc query tools, Enterprise Resource Planning (ERP) solutions, EAI, and data warehouses that are ODBC compliant.

MAS JDBC Driver

The MAS JDBC driver enables Java applications to integrate Multiple Virtual Storage (MVS) data and transactional sources through the JDBC API. JDBC is designed for use by Java database applications.

The JDBC driver is JDBC 2.0 compliant and supports JDK 1.2.3 (J2EE) and Java servlets. Specifically, it supports access and integration of OS/390 resident data and transactions using the standard JDBC API for access to data in DB2 Universal Database (UDB) for OS/390, ADABAS, VSAM, IMS/DB, IMS/TM, CICS/TS, and more.

Configuring the ODBC driver

The MAS ODBC Client provides one basic connectivity option: MAS ODBC Direct. This option and its features are described in the sections below.

MAS ODBC direct

This is the recommended option. It allows Windows NT and 2000, OS/2, UNIX, and Linux client systems to access many types of data including DB2-OS/390, IMS/DB, IMS/TM, CICS/TS, ADABAS, VSAM, PDSs, flat files, and many more. There are drivers available for TCP/IP users and LU 6.2 users. The MAS Mainframe Agent will interact with any application that uses the ODBC API.

If you select ODBC for your client connectivity, you will need to know your DBMS type, which is the means by which you will be accessing data. Initially, the DB2 type is displayed in the ODBC Data Source Administrator, the tool you use to configure the ODBC driver data source. See "Using the ODBC data source administrator" on page 24.

Controlling the ODBC driver

After you configure the MAS ODBC driver, you can control it using driver keywords. These keywords are automatically given default values that can be changed depending on the connector you are using. For further information, see Appendix C, "MAS driver keywords," on page 131.

Data sources

The data source you will be defining stores information about how to connect to the indicated data provider.

You will be configuring the data source on the ODBC Administrator, which was automatically installed. The ODBC Administrator lets you configure multiple data sources of any type or types on the same workstation according to your requirements.

Data sources for MAS are desktop. All connection information is stored locally in each user's ODBC.INI file. If you make any changes to the connection information, the new ODBC.INI file must be propagated to everyone's PC. If the data source moves to another computer, everything must be reconfigured.

Using support for multi-threaded applications

The MAS ODBC driver is enhanced to work with multi-threaded applications, which include many popular server products that use concurrent threads or thread-based connection pooling models to invoke ODBC operations. The MAS ODBC driver is fully thread-safe and does not require thread affinity.

By default, the driver does not run in the thread-safe mode; it must be enabled. The default is set to remove the performance penalty caused by the requirement to control access to shared resources by concurrent threads and because most applications do not take advantage of the multithreading programming model.

Running the driver in multi-threaded mode

To enable the driver to run in thread-safe mode when called by a multi-threaded application, you can set the NEONTHREADSAFE variable to YES or 1 in one of the following ways:

- Modifying the registry key (Windows only). To enable the driver to run in thread-safe mode, you can modify the registry key by setting the NEONTHREADSAFE registry key to YES or 1, as follows:

```
NEONTHREADSAFE = YES
```

Modifying the registry can enable the driver to run in thread-safe mode when called by all applications, regardless of whether they are designed to be single-threaded or multi-threaded. Setting the NEONTHREADSAFE variable via the registry key to enable thread-safe mode is a machine-wide setting, available only via Microsoft Windows platforms

- Setting the environment variable (Windows or UNIX). To enable the driver to run in thread-safe mode by setting the NEONTHREADSAFE environment variable, add the following line to the application startup file (.bat or.sh):

```
NEONTHREADSAFE = YES
```

Using an environment variable to enable thread-safe mode will only affect applications that inherit the environment settings when started.

Note: If you manually edit the ODBC.INI registry keys or file and add this keyword or if you append this keyword in the connection string, it will not produce thread-safe behavior.

Enabling debug tracing in multi-threaded mode

When you use the NEONTRACE environment variable, its value controls the trace settings for that session. This means any values set by the NEONTRACE ODBC keyword are ignored. To enable debug tracing and thread-safe behavior, set all trace keywords with the NEONTRACE environment variable.

For example, to force thread-safeness and generate an INFO level trace file to C:\NEONLOG.TXT on Windows systems:

- For Windows 95 and higher: Add the following line to your AUTOEXEC.BAT:

```
NEONTRACE=LOCKTHREAD INFO LOG c:\neonlog.txt
```

- For Windows NT: Open the **Control Panel**, select the **System** icon, and then add the following line to the user's environment variable set:

```
NEONTRACE=LOCKTHREAD INFO LOG=C:\neonlog
```

For UNIX systems:

- If you are using the debugging version of the MAS connector, you can use the trace facility. To enable debugging you must edit the `odbc.ini` file, specifying values for the `NEONTRACE` variable and a path to a trace log file. Here is an example from a sample `odbc.ini` file:

```
NEONTRACE=INFO BUFFERTRACE LOG=/home/ai38aas/shadow/bin/neonlog.txt
```

Note: For further information on tracing, see “Tracing” on page 37.

Configuring the ODBC driver data source for Windows

You configure the data source on Windows systems using the MAS ODBC Data Source Administrator, a graphical user interface (GUI). The sections below describe how to do this.

Using the ODBC data source administrator

When MAS was installed, the icon for the MAS ODBC Data Source Administrator automatically placed on your PC's **Control Panel**. Using the **Control Panel** is the easiest way to access the driver, although you can access it using the path specified during the installation process.

To configure the MAS ODBC driver data source:

1. Open the **Start** menu and select **Settings/Control Panel** (if you are using Windows NT 3.0 or higher) or **Settings/Control Panel/Administrative Tools** (if you are using Windows 95 or higher).
2. Double-click the **ODBC Data Sources** icon (or item, if viewing by List or Details).

The system displays the **ODBC Data Source Administrator** dialog box.

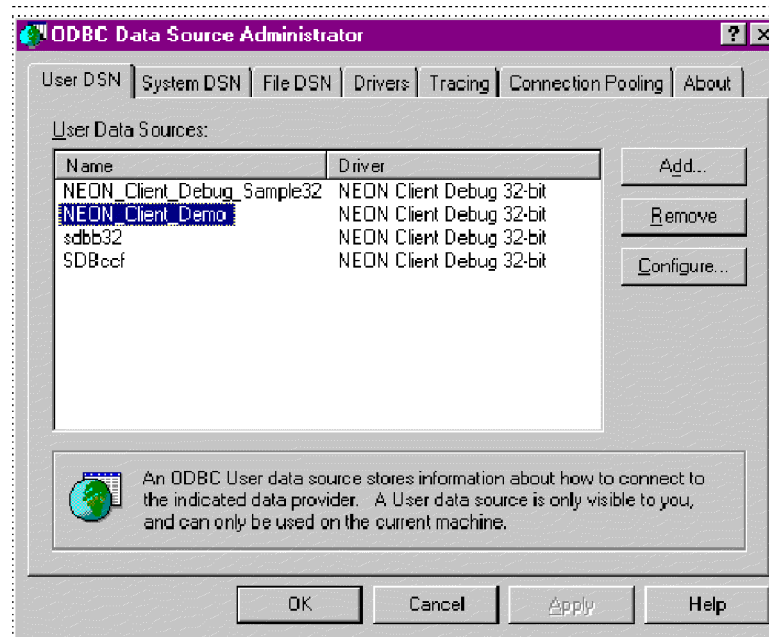


Figure 2. ODBC Data Source Administrator dialog box

3. Select the **User DSN** tab.
4. Click **Add**.

This displays the Create New Data Source box as shown in Figure 3.

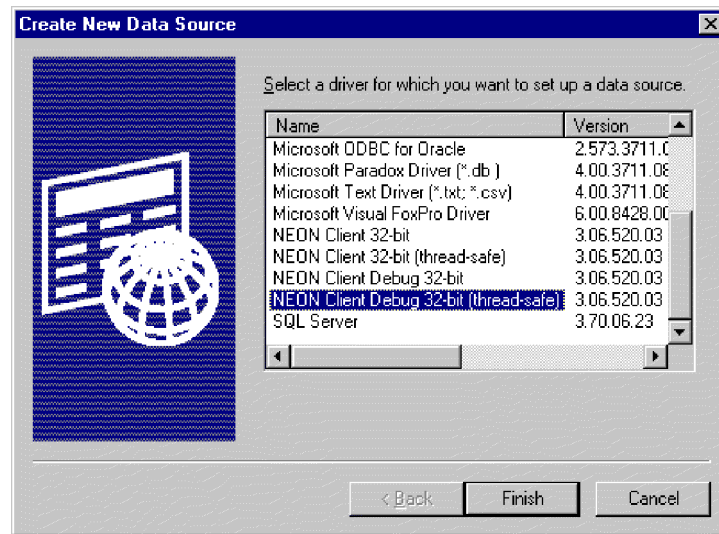


Figure 3. Create new data source dialog box

5. Select the Neon Client Debug 32-bit (thread-safe) or Neon Client 32-bit (thread-safe) driver and click **Finish**.

Note: Only thread-safe is advisable.

The system will display a configuration dialog box, as shown in Figure 4.

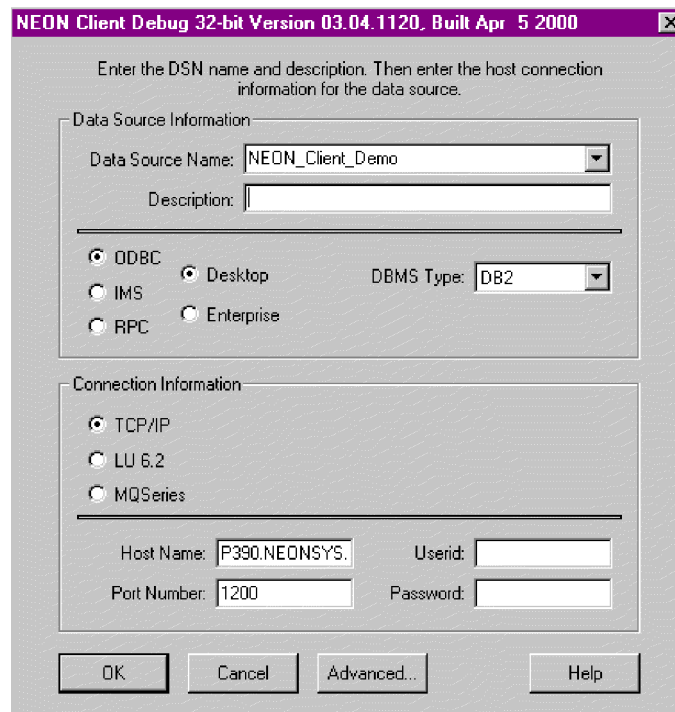


Figure 4. MAS debug configuration dialog box

6. Enter the **Data Source Name** and **Description**.

7. Select the data source type: **ODBC**. The dialog box options will change depending on the data source type.

ODBC:

- If you select **ODBC**, the system displays the options shown in Figure 5.

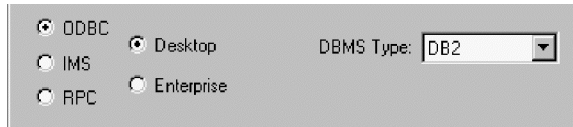


Figure 5. ODBC data source type

- Enter the **DBMS Type**, which is the means by which you will be accessing data. In this case, it is DB2 by default. See Table 21.

Table 21. Data source DBMS type

Data source	Enter in DBMS type field
DB2	DB2
VSAM	VSAM
VSAMCICS	VSAMCICS
ADABAS	ADABAS
IMSDB	IMSDB
IMS/TM	DB2 (see note)
CICS/TS	DB2 (see note)
Natural	DB2 (see note)
IDMS Database	DB2 (see note)

Note: You can select DB2 as your DBMS type when you are accessing IMS/TM, CICS/TS, Natural, and IDMS Database. In these cases, all creates, modifications, retrieves, and deletes will be processed using CALL statements. See “Syntax for CICS/TS and IMS/TM CALL statements” on page 31.

8. Select the **Connection Information** (shown in Figure 4 on page 25): **TCP/IP**, **LU 6.2**, or **WebSphere MQ**. Like the data source type, this portion of the dialog box changes depending on the option you select.

TCP/IP:

- If you select **TCP/IP**, the system displays the options shown in Figure 6.

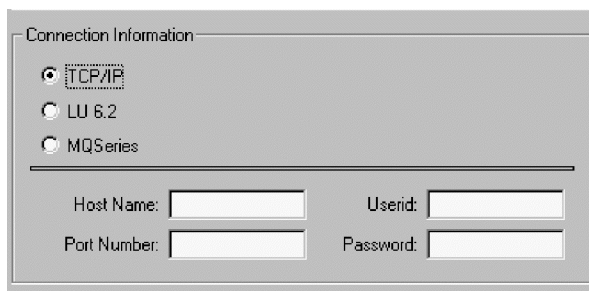


Figure 6. TCP/IP connection type

- Enter the following information:

Host Name (required): The TCP/IP address or host name for the connection. This field can contain a hostname string or a TCP/IP address in dot-notation format.

Port Number (required): The port number on which the MAS Mainframe Agent is listening. You can find this number by examining Shadow Server. It is often set to 1200.

Userid: Your mainframe userid. If you don't enter a userid, you are prompted for one at connection time.

Password: Your mainframe password. This is automatically encrypted (both on the screen and in the ODBC.INI file). If you don't enter a password, you are prompted for one at connection time.

LU 6.2:

- If you select **LU 6.2**, the system displays the options shown in Figure 7.

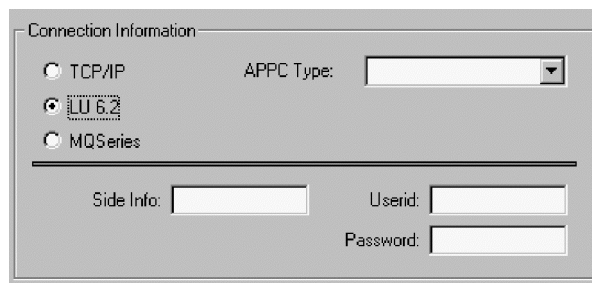


Figure 7. LU 6.2 Connection type

- Enter the following information:

APPC Type: The APPC type from the drop-down list provided.

Side Info (required): An entry in the *Side Info* table, which contains the partner LU name, conversation mode, and transaction program name for each side info name value. These values are used to establish an LU 6.2 connection with the host.

Userid: Your mainframe userid. If you don't enter a userid, you are prompted for one at connection time.

Password: Your mainframe password. This is automatically encrypted (both on the screen and in the ODBC.INI file). If you don't enter a password, you are prompted for one at connection time.

WebSphere MQ:

- If you select **WebSphere MQ**, the system displays the options shown in Figure 8 on page 28.

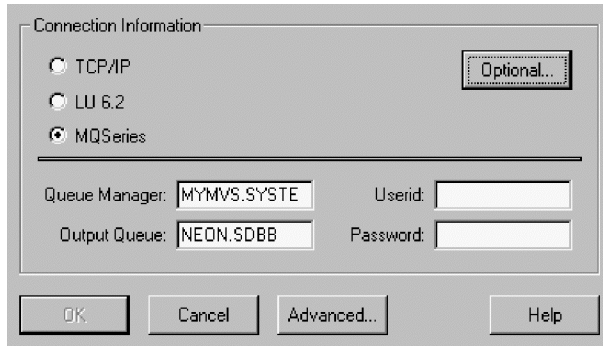


Figure 8. WebSphere MQ connection type

- a. Enter the following information:

Queue Manager (required): There must be a Queue Manager on both the client and host. Figure 8 shows the client side. If the client, such as Windows 3.1 or Windows for Workgroups, can't support the Queue Manager, you must click **Optional** to define the **WebSphere MQ Client** and the **MQSERVER** (see 8b).

Output Queue (required): The location where the request is stored until it is sent. There must be an Output Queue. If the Queue Manager is placed on the server, then the Output Queue must be defined on the server.

Userid: Your mainframe userid. If you don't enter a userid, you are prompted for one at connection time.

Password: Your mainframe password. This is automatically encrypted (both on the screen and in the ODBC.INI file). If you don't enter a password, you are prompted for one at connection time.

- b. If you need to define the **WebSphere MQ Client**, click the **Optional** button. The system displays the **Optional WebSphere MQ Connection Information** dialog box (Figure 9).

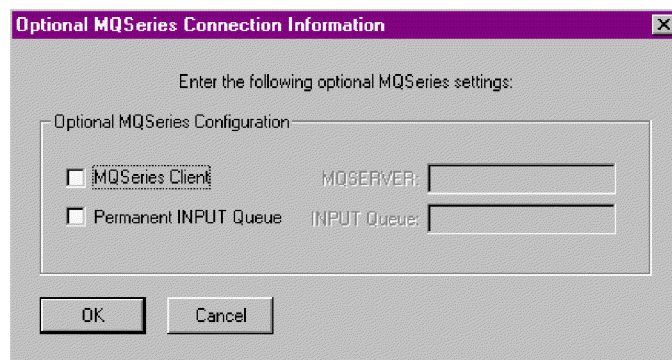


Figure 9. WebSphere MQ connection information

- c. Define the WebSphere MQ configuration as appropriate:

WebSphere MQ Client: Allows the client to accept MQI calls from an application and communicate with a Queue Manager on a server system.

MQSERVER: The server where the Queue Manager is placed.

- d. Click **OK** to return to the previous screen (the configuration dialog box shown in Figure 4 on page 25).

9. Click **Advanced**. The system displays the **Advanced Information** dialog box, shown in Figure 10 (based on the data source type you selected, some of the fields shown in the figure may be missing.)

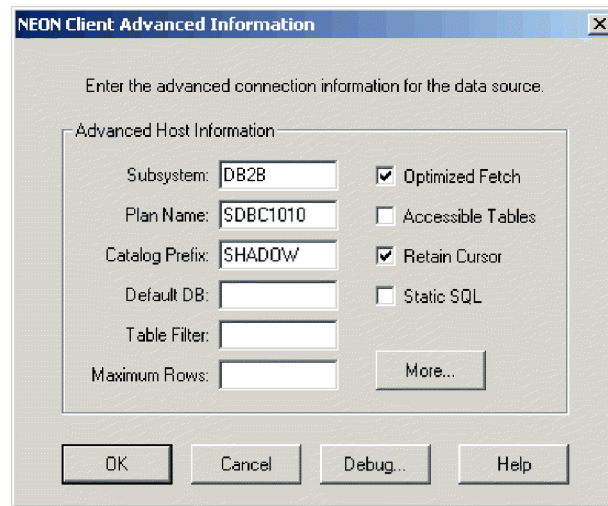


Figure 10. Client Advanced Information

10. If you selected ODBC as your data source type, enter or select the following information:
- **Subsystem:** The copy of DB2 on the host to which you connect. The default is the first four characters of the data source name. If you do not want to open a connection to DB2, specify NONE. This is applicable to customers who use the MAS connector for CICS/TS, for IMS/TM, or for ADABAS, and other types of access that do not require a connection to DB2.
 - **Plan Name:** MAS uses plan names on the host to make a connection to DB2. The first two letters of the name are always "SD." The third letter is always the last letter of the Mainframe Agent subsystem name, which is "B" by default. The fourth letter is either "C," if the plan was bound with cursor stability, or "R," if the plan was bound with repeatable read. The final four characters are always a version number (such as, 1010 or 1040). For example, SDBC1010. This default is release-dependent.
- Note:** If the 4th character of the plan name is an R, the MAS ODBC driver assumes that your application is using a plan where the plan was bound using an isolation value of repeatable read. If you are not using repeatable read, please ensure that your plan name does NOT have an R in the 4th character of the plan name as does the default plan SDBR1010. If the 4th character is any character other than an R, it is assumed the plan was bound with an isolation value of cursor stability.
- Tip:** If the client sets Subsystem to DFLT, then the DEFAULTDB2SUBSYS parameter that is set on the Mainframe Agent will be used for that connection. If the client sets Plan Name to either DFLT or DEFAULT, then the DEFAULTDB2PLANNAM parameter that is set on the Mainframe Agent will be used for that connection
- **Catalog Prefix:** Enter a value *only* if you are using the optimized catalog feature. The catalog prefix should contain the table owner (in DB2 terminology, the authorization ID) for a set of DB2 tables that are optimized

to support ODBC catalog queries. If no value is specified for this field, the standard system catalogs (SYSIBM.SYS*) are used.

Note: WebSphere considers optimized catalogs to be essential for smooth operation of pre-packaged ODBC applications with MAS connectors.

- **Default DB:** If you need to create tables in DB2 using your ODBC application, you must enter the name of the database in which your userid has authority to create tables. This value is added to CREATE TABLE statements, which do not specify a database in which to create the table. If the statement did not specify the "IN DATABASE" clause, the driver then appends this clause. This is **required** for users who do not have authority to create tables in DB2's default database, which includes most users.
- **Table Filter:** This field, along with the accessible tables feature described later, is used to restrict the information returned by catalog inquiries (SQLTables, SQLTablePrivileges, and SQLColumns). When your ODBC application requests a list of all the tables in your RDBMS instance, MAS constructs this list based on the value of the table filter. Only tables that fall under the first-level name (username) are returned. This filter can include the two SQL wildcard characters '%' and '_'. Here, '%' substitutes for zero or more characters and '_' for exactly one character. This field is useful because the RDBMS typically sends thousands of rows if it is not limited.

Tip: If the Table Filter field is left blank, the filter defaults to the userid value, which means MAS only returns tables that you own or created under your userid. If you are sure that no filter is needed, enter "" in the field. For example: "ai38%" will return ai38og.staff, ai38pds.staff, and ai38pds.xyz. More than one table filter can be specified by comma delimiting the entries.

- **Maximum Rows:** This field contains an integer that limits the number of rows returned from a single query. If no value is entered, no restriction is placed on the number of rows returned.
- **Optimized Fetch:** This box should be checked if you want to use the block fetch feature of the driver to speed up your queries. The increase in performance is significant.

Note: There is one restriction in using the block fetch feature: when the block fetch feature is enabled, cursor-based deletes (DELETE WHERE CURSOR OF) cannot be used.

- **Accessible Tables:** Many applications require information about the tables in the connected data source. However in DB2's case, this list of tables is unmanageably large unless you provide restrictions on the size. To do this, select the **Accessible Tables** box. Only those tables for which you have SELECT authority are returned. This includes those in PUBLIC and PUBLIC AT ALL LOCATIONS but not those where you have SELECT authority by virtue of a secondary authorization ID.
- **Retain Cursor:** This box controls how DB2 COMMIT operations affect cursors in the data source. **Retain Cursor** should be selected to preserve cursors in the same position as before the COMMIT operation took place. This enables the application to execute or fetch without preparing the statement again. The default is set to ON. If you check the box (this sets the value to OFF), the application will close and delete cursors after a COMMIT operation, meaning you must prepare and execute the next statement from scratch. If you are not sure about this value, set it to ON (checked).

Note: The Retain Cursor option also causes the Mainframe Agent to issue COMMITs after every SELECT statement as long as Auto-Commit is enabled by the application.

- **Static SQL:** This box should be checked to place the MAS connector into STATIC SQL mode. In this mode, every query passed to the driver will be converted and stored as static SQL on the host. Before using this option, you must convert your application with the Dynamic-to-Static Analyzer.

11. Click **More**. The system will display the **Optional Information** dialog box, as shown in Figure 11.

You can use this dialog box to set keyword values.

Note: Some of these keyword values must be set to non-default values, depending on the interface you are using. These keywords and their values are covered in Appendix C, “MAS driver keywords,” on page 131.

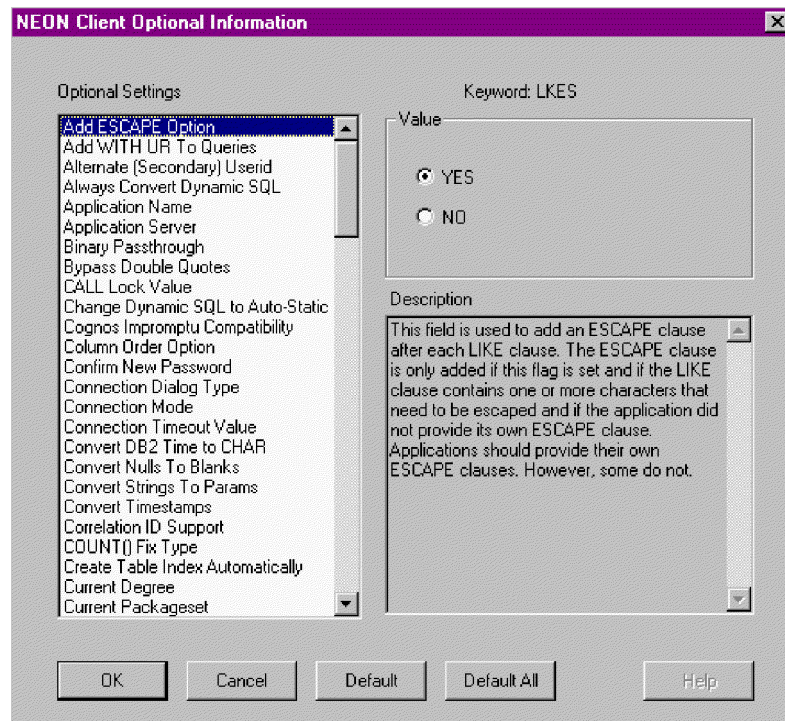


Figure 11. Optional settings

12. Click OK until you exit you from the ODBC Data Source Administrator to write the entered values into the ODBC.INI file, which will be used as your default the next time you connect to the data source.

Syntax for CICS/TS and IMS/TM CALL statements

This section discusses syntax for CICS/TX and IMS/TM CALL statements.

CICS/TS CALL Syntax: CICS/TS statements use the following syntax:

```
CALL <CICS Stored procedure> (<parameters>)
```

an example would be:

```
CALL CICSEX.NESPSEL ('logonid')
```


it will return one row with four columns and it will look like this:

CA_IN_LID	CA_OUT_NAME	CA_OUT_PHONE	CA_OUT_EMAIL
SCHFS		Henry Sobieski	
8005056366		HSOBIESKI@NEONSYS.COM	

IMS/TM CALLS**Syntax:** IMS/TM statements use the following syntax:

CALL <IMS Stored procedure> (<parameters>)

an example would be:

```
call ims.imssp('','','s','Sobieski', 'Henry','F','','','','','','')
```

A record will be retrieved by specifying the first name, last name and middle initial in parm 5 (D05751), parm 6 (D05757), and parm 7 (D05752).

Note that the parm 4 (REVFLD) is populated with an 's'. The resultset will look like this:

	ADDFLD	CHGFLD	REMFLD	
REVFLD	D05751	D05757	D05752	
D05754	D05755	D05756	D15333	
D15334	D07786		F00028	
SOBIESKI	HENRY	F	FAIRFAX	VA
	IBM TEST	SDBB000	0SDBB00001246581	

ITINERARY RETRIEVED

Configuring the ODBC data source for UNIX

The ODBC drivers for UNIX are provided as UNIX shareable libraries but, unlike the Windows environment, the UNIX drivers do not require the presence of an ODBC driver manager. Instead, you need only edit the .INI file. However, if your application requires a third-party ODBC manager, MAS provides a procedure. The sections below describe:

- The MAS UNIX ODBC file structure
- How to configure the ODBC driver data source by editing the .INI file
- How to enable a third-party ODBC driver manager
- A sample `odbc.ini` file
- How to enable tracing
- Programming notes

ODBC driver file structure

The installation process creates a directory tree with the root directory named `shadow`. The installed files are `untar`'ed. The directory structure is as follows:

```
SHADOW - +--- SAMPLES (scodbcsm.c)
|
|   scimbmsm.c
|   scimccsm.c
|   schorpsm.c,
|   Readme files,
|   makefile)
|
+--- INCLUDE (header files for UNIX)
|
+--- LIB (shareable library)
|
|   libscodbc.so.1 (production version)
|   libscodbcts.so.1 (debug version)
```

```
+--- BIN (Dynamic-to-static)
|      dsa                (dynamic-to-static)
|      scodbcdm executable (sample program - See
|                          appendix.)
```

Editing the odbc.ini file

To configure data sources on the SunOS/Solaris platform, you specify a search path and use a text editor to modify the `odbc.ini` file:

1. Specify the Search Path.

Make sure the search path to the MAS run-time library is specified. For example, using `csh` under SunOS and Solaris, you would specify the directory name where the libraries are located by modifying the `LD_LIBRARY_PATH` environment variable:

```
setenv LD_LIBRARY_PATH ${LD_LIBRARY_PATH}:/shadow/lib
```

2. Configure the `.odbc.ini` or `odbc.ini` file.

All ODBC data connection information must be stored in a configuration file. MAS supports two naming conventions: `.odbc.ini` (used by many UNIX ODBC driver manager vendors) and `odbc.ini`. This configuration file must be stored in the directory pointed to by the `$HOME` environment variable. The driver searches for the `.odbc.ini` file first, then the `odbc.ini`.

Note: The file name for the configuration file must be in lower case. For example, `.odbc.ini` is a valid file name, but `.ODBC.INI` is not.

Tip: When the ODBC application runs as a daemon process or is spawned by a master process, the `$HOME` environment variable might not be set in the process context calling the driver. You can either set the variable programatically using `putenv(3C)` or by setting the environment variable in the context of the parent process.

Using a UNIX-based ODBC application linked with third-party ODBC driver manager

If an application requires the presence of a third-party driver manager and the run-time library dependency cannot be altered, then you must do the following to allow the application to use MAS drivers:

1. Ensure that the MAS ODBC driver is setup correctly.
2. Create two new symbolic links from the standard MAS driver library with the driver manager library's file names. These are usually named `libodbc.so` and `libodbcinst.so`. For example:

```
ln -s shadow/lib/libscodbcts.so shadow/lib/libodbc.so
ln -s shadow/lib/libscodbcts.so shadow/lib/libodbcinst.so
```

3. Ensure these two new symbolic links can be searched ahead of the real ODBC driver manager libraries by modifying the `LD_LIBRARY_PATH` variable.

Sample odbc.ini file

The DSN keyword is used to identify a valid keyword section in the `odbi.ini` file.

```
[ODBC Data Sources]
DB2.DEMO=Shadow Direct UNIX Debug Driver
DB2.TEST=Shadow Direct UNIX Debug Driver
VSAM.TEST=Shadow Direct UNIX Debug Driver
```

```
[DB2.DEMO]
Driver=<put the load path of the library here...>
Description=<your description here...>
APPL=ODBC
MR=0
```

```

HD=NO
CPFX=SHADOW
LINK=TCPIP
UID=<user id>
PWD=<password>
HOST=<hostname or IP address>
PORT=<server listening port>
CD=NO
PLAN=SDBC1010
SUBSYS=DSN2
# this is a comment line
# this is a comment line
NEONTRACE=INFO log=/tmp/neonlog.txt

```

Search order for UNIX odbc.ini file

The following is the search order for the odbc.ini file:

1. /usr/local/lib/neon/odbc.ini
2. \$ODBC_INI
3. \$ODBC_INI (e.g., setenv ODBC_INI "/home/somebody/.my_odbc.ini")
4. \$HOME/.odbc.ini
5. \$HOME/.ODBC.INI
6. \$HOME/odbc.ini
7. \$HOME/ODBC.INI

The driver requires only the READ permission on this file.

Configuring the JDBC driver

The MAS JDBC driver can operate in two different modes:

- **Type 1:** In this mode, the MAS JDBC driver interacts with the ODBC driver manager.
- **Type 2:** In this mode, the MAS JDBC driver bypasses the ODBC driver manager and interacts directly with the MAS ODBC driver.

You can specify the mode you want to use in the NEONTRACE environment variable by entering the keyword TYPE1 or TYPE2. **If you do not specify a mode, TYPE1 will be used on the Windows NT platform and TYPE2 will be used on the SunOS platform.**

Formatting the MAS JDBC driver connection string

Ensure that the MAS JDBC driver connection string has been formatted. The format for the connection string passed to driver.connect() is:

```

jdbc:neon:<data-source-name>;
<attr1>=<value1>;...;<attrN>=<valueN>

```

All of the attributes that are valid for the ODBC connection string are also valid for the JDBC connection string, with the addition of the following JDBC-specific attributes:

- **DRPM** controls whether or not the driver should prompt for missing logon information. This attribute can have one of the following values:
 - **NOPROMPT** (default) corresponds to SQL_DRIVER_NOPROMPT.
 - **COMPLETE** corresponds to SQL_DRIVER_COMPLETE.
 - **PROMPT** corresponds to SQL_DRIVER_PROMPT.

An example of this attribute is:

```
DRPM=COMPLETE
```

DBCS support

The MAS JDBC driver supports DBCS on Windows NT, and SUNOS. DBCS is not supported on Windows 95/98.

Testing the MAS driver connections

VB Demo is a Windows user interface provided with the MAS connector. This interface is used to access data structures. The following sections describe how to use VB Demo:

- Overview
- Planning considerations
- Administering VB Demo
 - Starting VB Demo
 - Making a SQL request
 - Closing VB Demo

An overview of VB demo

VB Demo was installed onto your computer during the MAS installation (see Chapter 2, “Installing and configuring the connector,” on page 7). By default, it is stored in the Systems folder, located on your C:\ drive as follows:

```
IBM Shadow Client\Shadow\Bin32\VBdemo32.exe
```

With VB Demo, you can use Microsoft Access, Microsoft Excel, Crystal Reports, and many other tools to access DB2, IMS, CICS, and ADABAS data. You can also use standard SQL statement calls to access data from IMS, ADABAS, and VSAM.

You can also use VB Demo to verify installation of the Shadow Direct Interfaces for ADABAS, DB2, IMS, CICS, and VSAM.

Planning considerations

Before you begin, ensure that the following actions have been completed:

- Microsoft Visual Basic has been installed.
- Microsoft Data Access Component (latest version) has been installed.
- MAS components have both been installed.
- The MAS drivers have been configured and connected to the data source.

Administering VB demo

Starting VB demo

To run VB Demo, perform the following steps:

1. Open the Start Menu and select Programs/Neon Systems/VB Demo 32-bit application.

The system will display the **VB Demo** window.

2. From the **Connections** menu, select **Add New Connection**.

Note: You can have more than one connection open at the same time.

The system will display the **Select Data Source** dialog box, as shown in Figure 12.

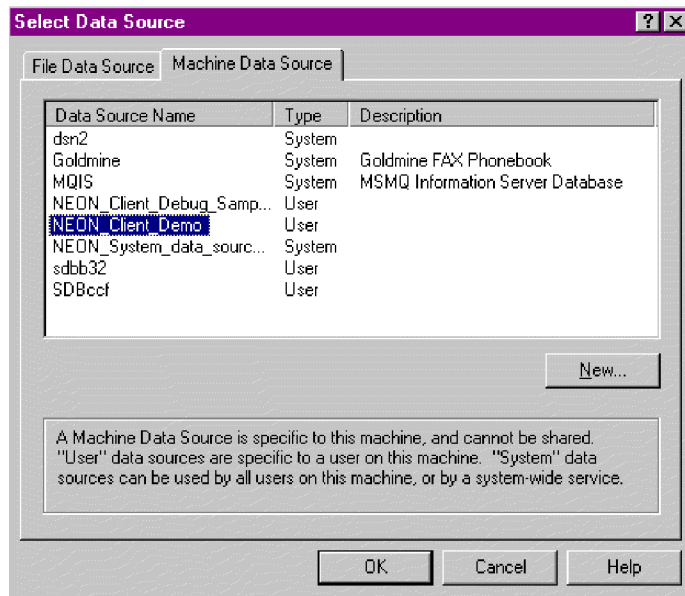


Figure 12. Select data source dialog box

3. Select the **Machine Data Source** tab, and then click **NEON_Client_Demo**.
4. Click **OK**. The system will display the **User Authentication** dialog box.
5. Enter your mainframe userid in the **Userid** field, and enter your mainframe password in the **Password** field.
6. Click **OK**. The system issues a message indicating that the connection was established.
7. Click **OK** to return to the **VB Demo** window.

Making SQL requests

Perform the following steps to make an SQL request:

1. Enter one of the following statements, depending on the type of data you are accessing:

For DB2-OS/390:

```
select * from q.staff
```

Note: If you select the Tables button, only the tables defined in the table filter are shown. The table filter was defined during data source configuration (see “Configuring the ODBC driver” on page 22 and “Configuring the JDBC driver” on page 34). If you did not define a table filter, the system will display only those tables that you own or that you created under your user id.

For ADABAS:

```
('select first_name birth sex from employees where last_name = "jones")
```

For CICS:

```
CALL CICSEX.NESPSEL
```

For IMS:

IMS/TM statements use the following syntax:

```
call ims.imssp('','','s','Sobieski', 'Henry','F','','','','','','')
```

A record will be retrieved by specifying the first name, last name and middle initial in parm 5 (D05751), parm 6 (D05757), and parm 7 (D05752). Note that parm 4 (REVFLD) is populated with an 's'.

The resultset will look like this:

	ADDFLD	CHGFLD	REMFLD
REVFLD	D05751	D05757	D05752
D05754	D05755	D05756	D15333
D15334	D07786		F00028
SOBIESKI	HENRY	F	FAIRFAX VA
IBM TEST		SDBB000	0SDBB00001246581

ITINERARY RETRIEVED

2. ClickQuery. The system displays the following information, depending on the type of data you requested.

Closing VB demo

To close VB Demo, perform the following steps

1. To exit the VB Demo program, click **Disconnect**. The system will display a message indicating that the connection was closed.
2. Close the VB Demo window.

Tracing

The MAS Trace Facility can be very useful in solving problems that can occur during installation and use. You can enable the client side trace for both the MAS ODBC Driver and the MAS JDBC Driver. The sections below cover the following topics:

- An overview
- Planning considerations
 - Configuring a data source
 - Performance issues
 - Controlling trace log files
- Enabling the MAS Trace Facility
 - Controlling the MAS Trace Facility

An overview of the MAS trace facility

The MAS Trace Facility traces selected events. It creates a file named NEONLOG.TXT in a selected directory and records the traced events in that file. If a file by that name is already present, the trace facility adds any newly recorded events to it.

The format of the trace output data always includes a date, time, and pertinent information about each event.

Example of trace output data

```

Fri Oct 01 22:19:30 1993  pcbColName      = 0x0a5f:759a
Fri Oct 01 22:19:30 1993  pfSqlType     = 0x0a5f:759e
Fri Oct 01 22:19:30 1993  pcbColDef     = 0x0a5f:75a0
Fri Oct 01 22:19:30 1993  pibScale     = 0x0a5f:7594
Fri Oct 01 22:19:30 1993  pfNullable   = 0x0a5f:759c
Fri Oct 01 22:19:30 1993  SQLDescribeCol exiting - return = SQL_SUCCESS(0)
Fri Oct 01 22:19:30 1993  szColName    = 'REMARKS'
Fri Oct 01 22:19:30 1993  *pcbColName  = 7

```

```

Fri Oct 01 22:19:30 1993 *pfSqlType      = SQL_CHAR(1)
Fri Oct 01 22:19:30 1993 *pcbColDef     = 64
Fri Oct 01 22:19:30 1993 *pibScale      = 9999
Fri Oct 01 22:19:30 1993 *pfNullable    = SQL_NULLABLE_UNKNOWN(2)
Fri Oct 01 22:19:30 1993 SQLFetch entered
Fri Oct 01 22:19:30 1993 lpstmt         = 0x095f:0000
Fri Oct 01 22:19:30 1993 internal error detected: file scodbcre.c line 1228 rc =
    0 from scclxlal
Fri Oct 01 22:19:30 1993 SQLFetch exiting - return = SQL_SUCCESS(0)
Fri Oct 01 22:19:30 1993 SQLGetData entered
Fri Oct 01 22:19:30 1993 lpstmt         = 0x095f:0000

```

Planning considerations

Before using the MAS Trace Facility, there are some important factors to consider. These include:

- Configuring a data source
- Performance issues
- Controlling trace log files

Configuring data sources

Make sure you have configured a dynamic data source. (See “Configuring the ODBC driver” on page 22).

Performance issues

The MAS Trace Facility is a debugging tool that can slow the execution of your ODBC application. Because of this, you may want to turn it off once you have the driver and your applications operating perfectly. This can be done by setting up the DOS environment variable SET NEONTRACE=NONE.

Controlling trace logfiles

If a fully qualified pathname (including a directory) for the NEONLOG.TXT file has not been specified, the MAS Mainframe Agent creates the NEONLOG.TXT file in the current directory. Over time, this could result in several NEONLOG.TXT files in several different directories. These files can become very large, especially if you are using a low severity level trigger to trace events. These files should be deleted as necessary.

Enabling the MAS ODBC driver trace for Windows

1. Open the **ODBC Data Source Administrator**.
2. Select the **Tracing** Tab.
3. Click **Start Tracing Now**.
4. Click **OK**.

Enabling the MAS JDBC driver trace for Windows

1. Follow steps 1 through 3 for “Enabling the MAS ODBC driver trace for Windows” on page 38.
2. Add the following item to the NEONTRACE environment variable in DOS:
JDBCLOG=xyz
where xyz is the full pathname of a log file. For example, on Windows NT, you may have:
THREADID INFO BUFFER JDBCLOG=C:\temp\neonjdbc.txt LOG=C:\temp\neonlog.txt

Note: The JDBCLOG keyword must precede the LOG keyword.

Two separate files will be generated: one for the ODBC driver and one for the JDBC driver.

3. Click **OK**.

Using the JDBC trace on Windows NT

To perform the trace in the Windows NT environment if you are running in TYPE1 mode (default), ensure that the trace file specified in the data source has the same file name as the one specified in the LOG keyword. If you omit the LOG keyword from the NEONTRACE parameter, the trace will be written to the trace file specified in the data source.

Controlling the MAS Trace Facility

The MAS Trace Facility's default values provide adequate tracing in most cases. However, you might need to change the MAS Trace Facility's severity level trigger or tracing options. You can do this in one of the following ways:

- Selecting the options using the **Debug Information** dialog box.
- Setting the appropriate NEONTRACE environment variable in DOS.

Using the debug dialog box

The most commonly used method for controlling the MAS Trace Facility is to use the **Debug Information** dialog box. This lists all of the options available to you and ensures that the correct keyword is changed in the ODBC.INI file.

1. Open the **Start** menu and select **Settings/Control Panel** (if you are using Windows 95 or higher) or **Settings/Control Panel/Administrative Tools** (if you are using Windows NT 3.0 or higher).
2. Double-click the **ODBC Data Sources** icon (or item, if viewing by List or Details).

The system displays the **ODBC Data Source Administrator** dialog box.

3. From the **User Data Sources** list, select the appropriate data source.
4. Click **Configure**. The system displays a configuration dialog box, as shown in Figure 13 on page 40.

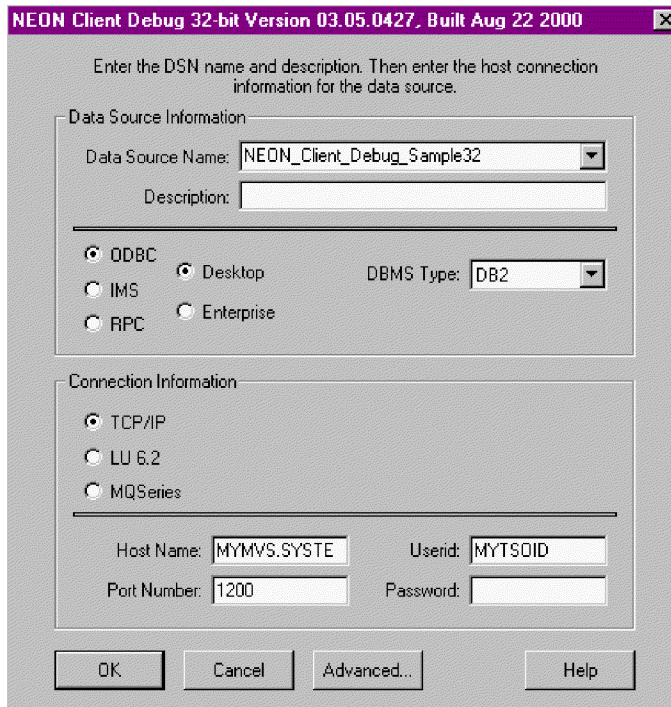


Figure 13. Debug configuration

5. Click **Advanced**. The system displays the **Advanced Information** dialog box (Figure 14).

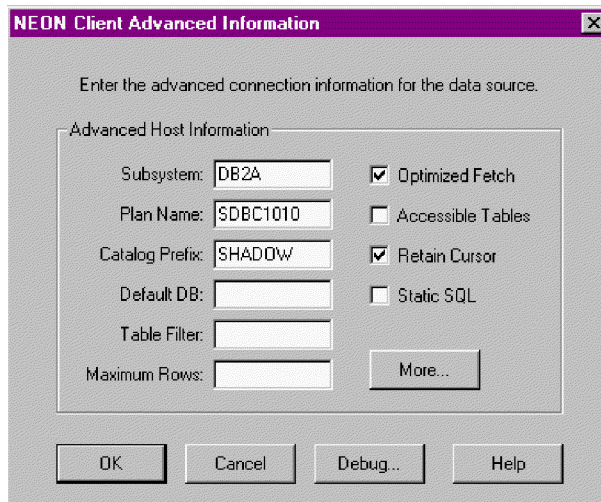


Figure 14. Advanced information

6. Click **Debug**. The system displays the **Debug Information** dialog box (Figure 15 on page 41).

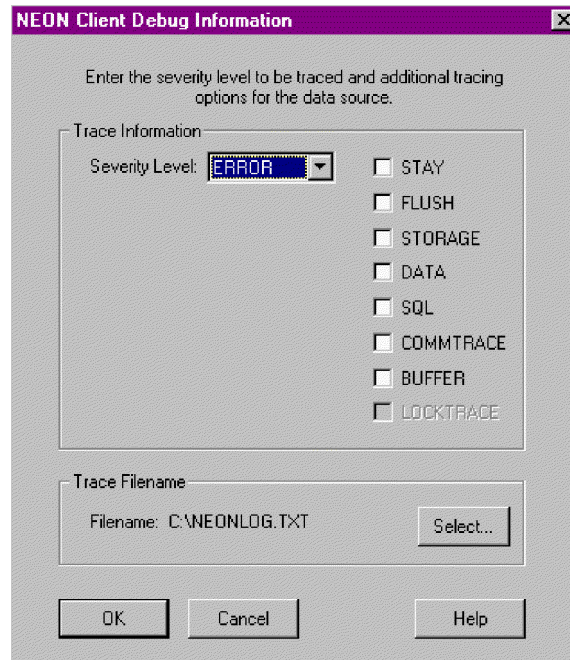


Figure 15. Client debug information

Note: If you had selected IMS as your client type from the configuration dialog box shown in Figure 13 on page 40, the Debug button would have appeared on the Advanced Information dialog box, shown in Figure 14 on page 40

7. Select the appropriate severity level trigger of the events to be traced. When a severity level is selected, all events at and above that severity level will be traced and recorded. The severity options are listed below from the highest to the lowest level:
 - **NONE:** With this option, no messages are traced or recorded. No trace file is created or opened. This option effectively turns off the MAS Trace Facility.
 - **FATAL:** This option, the highest severity level, causes only program termination messages to be traced and recorded.
 - **SEVERE:** (Default selection) This option causes only events of a severe or higher level to be traced and recorded.
 - **ERROR:** This option causes all error messages to be traced and recorded. These messages usually include a description of what went wrong, a record of where the error was detected, any relevant return code, and a detailed error text message.
 - **WARNING:** This option allows any warning messages issued by the MAS ODBC Driver to be traced and recorded.
 - **INFO:** This option allows all informational events to be traced. Note that values passed to and returned from the MAS ODBC Driver are included here. This option will cause a large quantity of information to be traced and recorded and is not recommended for general use. However, this option can be useful for application debugging.
 - **DETAIL:** This option, the lowest severity level, causes all events, including function call events, to be traced. This value will cause huge amounts of information to be traced and recorded, and it is not recommended for general use.

8. Select additional options by clicking the appropriate boxes. The additional options include the following:
 - **STAY:** This option keeps the trace file open. This can improve performance if you are using a disk file to store trace messages.
 - **FLUSH:** Use this option with STAY. It helps to ensure that messages are not lost if your machine fails.
 - **STORAGE:** This option causes all storage GET and FREE operations to be traced. The trace record shows the name of the file requesting or freeing the storage area and the line number within the file. The record also indicates the size of the data area being obtained or freed, the address of the data area, and other important information. The final storage report is generated in the Windows environment when the DLL containing the storage manager functions is unloaded.
 - **DATA:** This option causes all conversion operations to and from DB2, ODBC C, and ODBC SQL data types to be traced. In each case, the input and output data is displayed in hexadecimal and character format. In addition, information is provided that designates which row and column of the table is currently being processed.
 - **SQL:** This option causes all SQL statements passed to the host to be traced, even if the trace severity level is higher than INFO. It also causes the column information for SELECT statements to be displayed, including the column number, name, and data type. For SQL_NUMERIC and SQL_DECIMAL data types, the precision and scale values are also displayed.
9. Click **Select**. The system displays the **Select Trace File** dialog box (Figure 16).

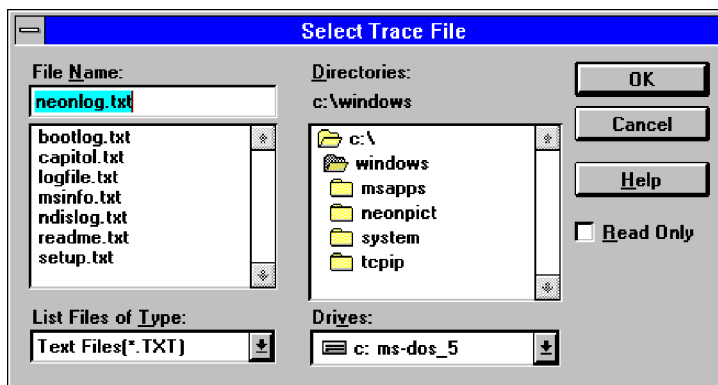


Figure 16. Select Trace File

10. Select a name and directory for the trace file.
11. Click **OK** to return to the **Debug Information** dialog box (shown in Figure 15 on page 41).
12. Click **OK**.

This writes a NEONTRACE keyword in the appropriate section of the ODBC.INI file. The NEONTRACE keyword will reflect the options you selected with the data source that you selected from the **ODBC Data Source Administrator**. The contents of the ODBC.INI file augment and override any DOS environment variable values set earlier using the SET NEONTRACE=[string] command.

Setting the trace environment variable in DOS

The other approach for controlling the MAS Trace Facility is to set a trace environment variable in DOS. To set this variable, use the following form:

SET NEONTRACE=[string]

where [string] consists of the options you select, delimited with spaces.

Set this variable in one of the following ways:

- Manually enter the SET NEONTRACE=[string] command from the DOS command line.
- Add the SET NEONTRACE=[string] command to the AUTOEXEC.BAT and CONFIG.SYS files. Reboot your computer for the command to take effect.

Note: If you prefer to enter the command manually, be sure to do so within DOS itself and not through a DOS window while Windows is still running. Setting the trace environment variable from a DOS window (which uses a second-level command processor) will have no effect at all on the MAS Trace Facility.

Chapter 4. Configuring the connector

This chapter describes how to configure the MAS connector and how to configure applications to work with the connector. It contains the following sections:

- “Enabling the application for the connector”
- “Configuring the connector” on page 47
- “Creating multiple instances of a connector” on page 58
- “Starting the connector” on page 59
- “Stopping the connector” on page 61

Enabling the application for the connector

You must set up the event notification mechanism in the application/database before the connector can process event delivery. To do this, you must complete the following tasks:

- Create the event and archive tables in the legacy application database.
- Install database triggers on the application’s tables to support the business objects needed by the collaborations running at your site. WebSphere assumes that you develop your own database triggers.
- Optionally, install a counter table. Perform this step only if you require the connector to generate a unique ID when creating a business object. For more information on generating unique IDs, see the `UID=CW.uidcolumnname[=UseIfMissing]` parameter.

The sections that follow provide information on creating and configuring the event and archive tables.

Event and archive tables

The connector uses the event table to queue events for pickup. If you have set the `ArchiveProcessed` property to `true` or to no value, the connector uses the archive table to store events after updating their status in the event table.

For each event, the connector gets the business object’s name, verb, and key from the event table. The connector uses this information to retrieve the entire entity from the application. If the entity was changed after the event was first logged, the connector gets the initial event and all subsequent changes. In other words, if an entity is created and updated before the connector gets it from the event table, the connector gets both data changes in the single retrieval.

The following three outcomes are possible for each event processed by a connector:

- Event was processed successfully
- Event was not processed successfully
- Event was not subscribed to

If events are not deleted from the event table after the connector picks them up, they occupy unnecessary space there. However, if they are deleted, all events that are not processed are lost and you cannot audit the event processing. Therefore, IBM recommends that you also create an archive table and keep the

ArchiveProcessed property set to true. Whenever an event is deleted from the event table, the connector inserts it into the archive table.

Note: If problems accessing the application database cause the connector to fail while deleting an event from the event table or inserting an event into the archive table, the connector returns APPRESPONSETIMEOUT and is terminated.

Configuring event and archive processing

To configure event and archive processing, you must use configuration properties to specify the following information:

- The name of the event table (EventTableName). You need not specify a value for this property if you use the connector only to process business object requests.
- The interval frequency (PollFrequency).
- The number of events for each polling interval (PollQuantity).
- The name of the archive table (ArchiveTableName).
- Whether the connector archives unsubscribed and unprocessed events (ArchiveProcessed).
- The unique ID of the connector, which is important when multiple connectors poll the same table (ConnectorID).

You can also specify a value for the EventKeyDel property to specify the order of events to be processed. For information on these and other configuration properties, see Table 23 on page 48.

Note: Creation of the event and archive tables is optional. However, if you specify a value for EventTableName but do not use the connector to poll for events and do not create an event table, the connector times out. To prevent such time-out, leave the value of EventTableName as null (as a string).

By default, the name of the event queue table is xworlds_events, and the name of the archive queue table is xworlds_archive_events.

Note: If your site will not archive events into the archive table, set the value of ArchiveProcessed to false.

To use the connector only for request processing, use the -fno option when starting it and set the value of EventTableName to null (as a string).

If the driver being used does not support Java class DatabaseMetaData and you want the connector to avoid checking for the existence of event and archive tables, then disable the CheckForEventTableInInit by setting it to false. By default, this property is set to true.

Event and archive table schema

Table 22 describes the columns in the event and archive tables. Use these as guides for creating and installing these tables, taking special note of the order and data type in table columns.

Table 22. Event and archive table schema

Name	Description	Type	Constraint
event_id	Internal identifier of the event	INTEGER	Primary key

Table 22. Event and archive table schema (continued)

Name	Description	Type	Constraint
connector_id	Unique ID of the connector for which the event is destined. This value is important when multiple connectors poll the same table.	VARCHAR	
object_key	Primary key of the business object. Multiple keys can be concatenated with a colon or other configurable delimiter, for example, 1000065:10056:2333. See the DateFormat property for more information.	VARCHAR	Not null
object_name	Name of the business object	VARCHAR	Not null
object_verb	Verb associated with the event	VARCHAR	Not null
event_priority	Event priority (0 is highest, n is lowest), which the connector uses to get events on a priority basis. The connector does not use this value to lower or raise priorities.	INTEGER	Not null
event_time	Date and time the event occurred	DATETIME	Default current date/time (for archive table, actual event time)
archive_time	Date and time the event was archived (applies only to the archive table)	DATETIME	Archive date/time
event_status	-2 (Error sending event to InterChange server) -1 (Error processing event) 0 (Ready for poll) 1 (Sent to InterChange Server) 2 (No Subscriptions for the business object) 3 (In Progress). This status is used only in the event table and not in the archive table.	INTEGER	Not null
event_comment	Description of the event or error string	VARCHAR	

Configuring the connector

Connectors have two types of configuration properties: standard configuration properties and connector-specific configuration properties. You must set the values of these properties before running the connector. Use one of the following tools to set a connector's configuration properties:

- Connector Configurator (if ICS is the integration broker)—Access to this tool is from the System Manager.
- Connector Configurator (if WebSphere MQ Integration Broker is the integration broker)—Access this tool from the WebSphere Business Integration Adapter program folder. For more information see Appendix B, "Connector Configurator," on page 115.

Standard connector properties

Standard configuration properties provide information that all connectors use. See Appendix A, "Standard configuration properties for connectors," on page 99, for documentation of these properties.

Note: Because this connector supports both the ICS and WebSphere MQ Integration Broker integration broker, configuration properties for both brokers relevant to the connector.

Connector-specific properties

Connector-specific configuration properties provide information needed by the connector at runtime. Connector-specific properties also provide a way of changing static information or logic within the connector without having to recode and rebuild the agent.

Table 23 lists the connector-specific configuration properties for the connector. See the sections that follow for explanations of the properties.

Table 23. Connector-specific configuration properties

Name	Possible values	Default value	Required
ApplicationPassword	<i>Password for WebSphere user account</i>		Yes*
ApplicationUserName	<i>Name of WebSphere user account</i>		Yes*
ArchiveProcessed	true or false	true	No
ArchiveTableName	<i>Name of archive queue table</i>	xworlds_archive_events	Yes if Archive Processed is true
AutoCommit	true or false	false	No
CheckforEventTableInInit	true or false	true	No
ChildUpdatePhyDelete	true or false	false	No
CloseDBConnection	true or false	false	No
ConnectorID	<i>Unique ID for the connector</i>	null	No
DatabaseURL	<i>Name of the database server</i>		Yes
DateFormat	<i>A time pattern String</i>	MM/dd/yyyy HH:mm:ss	No
DriverConnectionProperties	<i>Additional driver connection properties</i>		No
EventKeyDel	<i>Delimiter character for event key column of event table</i>	semicolon (;) colon (:) equal and semicolon(=;) equal and colon (=:)	No
EventOrderBy	none, <i>ColumnName</i> [, <i>ColumnName</i> , ...]		No
EventQueryType	Fixed or Dynamic	Fixed	No
EventTableName	<i>Name of event queue table</i>	xworlds_events	Yes, if polling is required; null (as a string) if polling is not required
JDBCdriverClass	driver classname		No
MaximumDatabaseConnections	<i>Number of simultaneous database connections</i>	5	Yes
PingQuery	SELECT 1 FROM <tablename>		No
PollQuantity	Values are 1 to 500	1	No
PreserveUIDSeq	true or false	true	No
RDBMS.initsession	<i>SQL statement that initializes every database session</i>		No
RDBMSVendor	Others		No

Table 23. Connector-specific configuration properties (continued)

Name	Possible values	Default value	Required
ReplaceAllStr	true or false	false	No
ReplaceStrList	A set composed of a single character, a character delimiter, and the character's substitution string. Also, multiple such sets with a termination delimiter between them.	Q,DSQNote: In System Manager, these characters represent a single quotation mark, followed by a comma, followed by two single quotation marks.	No
RetryCountAndInterval	<i>Count, interval in seconds</i>	3,20	No
SPBeforePollCall	Name of the stored procedure to be executed for each poll call.		No
SPBeforePollCall	The character and termination delimiters used in the ReplaceStrList property	;	No
TimingStats	0, 1, 2	0	No
UniqueIDTableName	<i>Name of table used for generation of IDs</i>	xworlds_uid	No
UseDefaults	true or false	false	Yes
UseDefaultsForCreatingChildBOs	true or false	false	No
UseDefaultsForRetrieve	true or false	false	No

*ApplicationPassword and ApplicationUserNames are not required if you are using trusted authentication.

ApplicationPassword

Password for WebSphere user account.

There is no default value.

ApplicationUserName

Name of WebSphere user account.

There is no default value.

ArchiveProcessed

Specifies whether the connector archives events for which there are no current subscriptions.

Set this property to true or false to cause events to be inserted into the archive table after they are deleted from the event table.

Set this property to true or false to cause the connector not to perform archive processing. In this case, it does not check the value of the ArchiveTableName property. If ArchiveProcessed is set to false, the connector performs the following behavior:

- If the event is successfully processed, the connector deletes it from the event table and does not archive it.
- If the connector does not subscribe to the event's business object, the connector leaves the event in the event table and changes its event status to Unsubscribed.
- If the business object encounters a problem while being processed, the connector leaves the event in the event table with event status of Error.

If this property is set to false and the poll quantity is low, the connector appears to be polling the event table, but it is simply picking up the same events repeatedly.

If this property has no value, the connector assumes the value to be true. If the ArchiveTableName property also has no value, the connector assumes the archive table's name is xworlds_archive_events.

The default value is true.

ArchiveTableName

Name of archive queue table.

If the ArchiveProcessed property is set to false, it is unnecessary to set a value for this property.

The default name is xworlds_archive_events.

AutoCommit

This property makes the AutoCommit setting configurable. When set to true, all transactions are automatically committed. Some databases (such as Sybase) require AutoCommit to be set to true. If set to false, stored procedures on Sybase to fail.

If the database connection is lost, the connector will attempt to create a new connection to restart the complete processing as long as AutoCommit is set to false. If the new connection is invalid, or if AutoCommit is set to true, the connector returns APPRESPONSETIMEOUT, which results in the termination of the connector.

The default value is false.

CheckforEventTableInInit

Setting this connector property to false prevents the connector from checking for the existence of the event and archive tables during connector initialization. WebSphere recommends that you always set it to true unless the JDBC driver you are using does not support the JDBC class DatabaseMetaData.

When the property is set to false, although the connector does not check for the existence of EventTable and ArchiveTable, the event and archive tables should always exist because the connector uses them during the initialization process. To prevent the connector from using the event and archive tables during initialization, set the property EventTableName to null.

The default value is true.

ChildUpdatePhyDelete

During an update operation, specifies how the connector handles data represented by a child business object that is missing from the incoming business object but exists in the database.

Set this property to true to cause the connector to physically delete the data record from the database.

Set this property to false to cause the connector to logically delete the data record from the database by setting the status column to the appropriate value. The connector obtains the name of the status column and its value from the StatusColumnValue (SCN) parameter specified in its business-object level

application-specific information. For more information, see “Application-specific information at the business-object level” on page 86.

Default value is false.

CloseDBConnection

This property makes the closing of the database connection configurable. When set to true, for every service call request and poll call, the database connection is closed. Setting this property to true impairs performance and is not advisable.

The default value is false.

ConnectorID

A unique ID for the connector. This ID is useful to retrieve events for a particular instance of the connector.

Default value is null.

DatabaseURL

Name of the database server to which the connector should connect. The name specified when configuring the MAS ODBC driver data source.

In the example:

```
jdbc:neon:mydatasource
```

mydatasource is the name of the database server, and hence the value of DatabaseURL.

You must provide this value for the connector to process successfully. For further information on configuring the data source name, see “Configuring the ODBC driver data source for Windows” on page 24 or “Configuring the ODBC data source for UNIX” on page 32.

DateFormat

Specifies the date format that the connector expects to receive and return. This property supports any format that is based on the syntax as contained in Table 24 on page 52.

Table 24 on page 52 Defines the time format syntax using a time pattern string. In this pattern, all ASCII letters are reserved as pattern letters.

Table 24. Time format syntax

Symbol	Meaning	Presentation	Example
G	era designator	(Text)	AD
Y	year	(Number)	1996
M	month in year	(Text & Number)	July & 07
D	day in month	(Number)	10
h	hour in am/pm(1-12)	(Number)	12
H	hour in day(0-23)	(Number)	0
M	minute in hour	(Number)	30
S	second in minute	(Number)	55
S	millisecond	(Number)	978
E	day in week	(Text)	Tuesday
D	day in year	(Number)	189
F	day in week in month	(Number)	2 (2nd Wed in July)
W	week in year	(Number)	27
W	week in month	(Number)	2
a	am/pm marker	(Text)	PM
k	hour in day(1-24)	(Number)	24
K	hour in am/pm(0-11)	(Number)	0
Z	time zone	(Text)	Pacific Std Time
'	escape for text	(Delimiter)	
"	single quote	(Literal)	'

Table 25. Examples using the US standard

Format pattern	Result
"yyy.MM.dd G 'at'hh:mm:ss z'	1996.07.10 AD at 15:08:56 PDT
"EEE, MMM d, 'yy"	Wed, July 10, '96
"h:mm a"	12:08 PM
"hh 'o'clock' a, zzzz"	12 o'clock PM, Pacific Daylight Time
"K:mm a, z"	0:00 PM, PST
"yyyy.MMMMM.dd GGG hh:mm aaa"	1996.July.10 AD 12:08 PM

DriverConnectionProperties

Besides the user name and password, the JDBC driver might need additional properties or information. The `DriverConnectionProperties` connector property will take additional properties that a JDBC driver needs, as name-value pairs. The properties should be specified as follows:

```
property1=value1[:property2=value2...]
```

The properties must be given as name value pairs, separated by semi-colons. The property is separated from its value by an equals sign (with no extra spaces).

For example, assume the JDBC driver needs license information and port number. The property name it expects for license information is `MyLicense` and the value is `ab23jk5`. The property name it expects for port number is `PortNumber` and value is `1200`. The `DriverConnectionProperties` should be set to the value `MyLicense=ab23jk5;PortNumber=1200`.

EventKeyDel

Specifies the delimiter when the `object_key` column of the event table contains multiple attribute values.

There are two ways to retrieve the business object that has been created, updated, or deleted in the triggering application.

- The first is to populate the `object_key` column with values for attributes that are keys in a business object. Set the `EventKeyDel` configuration property to a single character that is not part of the key field. For **example**, if the delimiter is specified as “;”, then the `object_key` will be as follows: `xxx;123`
- The second is to populate the `object_key` column with values for any attribute in a business object. These values should be represented as `name_value` pair. The first delimiter will be for the `name_value` and the second is for the keys. For example, if the delimiter is specified as “=;”, then the `object_key` will be as follows: `CustomerName=xxx;CustomerId=123;`

If the delimiter is specified as “=:”, then the `object_key` will be as follows:
`CustomerName=xxx:CustomerId=123:`

Note: The order that the key values are defined should follow the same order as the key attributes in a business object.

Important: If you use Date attribute data, avoid using a colon (:) delimiter, because it may be included in the attribute’s data.

The default value is a semicolon (;), which is based on keys, not `name_value` pairs.

EventOrderBy

Specifies whether to turn off the ordering of events, or specifies an order of event processing that is different from the default order.

By default, at each poll the connector pulls only the number of events specified in its `PollQuantity` property, and orders event processing by the values in the `event_time` and `event_priority` columns of the Event table.

To cause the connector not to order events, set the value of this property to none.

To cause the connector to order by different columns in the Event table, specify the names of those columns. Separate column names with a comma (,). Specifying a value for this property overwrites the default value.

There is no default value for this property.

EventQueryType

The `EventQueryType` property is used to indicate whether the connector should dynamically generate a query to retrieve events from the event table or use its built in query. For the dynamically generated query, the connector maps its event structure to the columns in the event table. The order of the data in the table columns is very important. Please refer to the “Event and archive table schema” on page 46 to view the correct order.

If the value in the `EventQueryType` is Fixed (as a string), the default query is executed. If set to Dynamic (as a string), a new query is built by getting the column names from the table that is specified in the “`EventTableName`” on page 54 property.

The event table column names can change but the order and data type of the columns must remain the same as specified in the event table creation section. "EventOrderBy" on page 53 will be appended to either the default or the dynamically generated query.

If the EventQueryType property is not added or it contains no value, it is defaulted to Fixed.

Default value is Fixed (as a string).

EventTableName

Name of event queue table, which is used by the connector's polling mechanism.

The default name is xworlds_events.

Set this to null (as a string) when polling is turned off for the connector. This prevents validation of the existence of the event and archive tables.

For more information, see Table 22 on page 46.

No default value is provided.

For a user defined event table, ensure that the event_id maps to one of the following JDBC types: INTEGER, BIGINT, NUMERIC, VARCHAR.

JDBCDriverClass

Specifies the class name of a driver. For MAS connectors, specify:

```
com.neon.jdbc.Driver
```

MaximumDatabaseConnections

Specifies the maximum number of simultaneous database connections allowed. At runtime, the number of open database connections is the sum of this value plus 1.

If the "PreserveUIDSeq" on page 55 property is set to false, at runtime, the number of open database connections is the sum of this value plus 2.

The default value is 5.

PingQuery

Specifies the SQL statement or stored procedure that the connector executes to check database connectivity.

The following is an example of an SQL statement used as a ping query:

```
SELECT 1 FROM <tablename>
```

The following is an example of a stored procedure call (sampleSP) used as a ping query with an Oracle or DB2 database:

```
call sampleSP( )
```

Note that stored procedure calls cannot have output parameters. If an input parameter is required by the database, the input value must be specified as part of the ping query. For example:

```
Call checkproc(2)
```

There is no default value. For more information, see "Handling lost database connections" on page 6.

Note: PingQuery cannot be supported for CICS and IMS applications.

PollQuantity

Number of rows in the database table that the connector retrieves per polling interval. Allowable values are 1 to 500.

The default value is 1.

PreserveUIDSeq

Specifies whether or not the incoming unique ID sequence will be preserved in the unique identifier table.

If set to true, the unique ID is not committed until the business object is successfully processed in the destination application. All other processes attempting to access the unique identifier table must wait until the transaction is committed.

If set to false, the unique ID is committed when the business object requests it. The business object processing and the unique ID processing each have their own transaction block (internal to the connector). This is only possible if the transaction relating to the unique identifier table has its own connection.

Note: If this property is not added to the connector configuration, the default behavior is the same as if this property were added and set to true. Also, if “AutoCommit” on page 50 is set to true, the connector executes the same behavior as if PreserveUIDSeq is set to false.

If the “PreserveUIDSeq” property is set to false, at runtime, the number of open database connections is the sum of this value plus 2.

The default value is true.

RDBMS.initsession

SQL statement that initializes every session with the database. The connector takes a query and executes it at startup. There should not be a return value for this query. The property name is required, but a value is not.

There is no default value.

RDBMSVendor

Specifies which RDBMS the connector uses for special processing. You need not define this property for MAS connectors.

No default value is provided.

ReplaceAllStr

Specifies whether the connector replaces all instances of each character identified in the ReplaceStrList property with the substitution string specified in that property. The connector evaluates ReplaceAllStr only if the ESC=[true|false] parameter of each attribute’s AppSpecificInfo property does not contain a value. In other words, if the ESC parameter has been specified, its value takes precedence over the value set for the ReplaceAllStr property. To cause the connector to use the value of ReplaceAllStr, verify that the ESC parameter has not been specified.

The default value of ReplaceAllStr is false.

ReplaceStrList

Specifies one or more substitution sets, each composed of an individual character to be replaced, a character delimiter, and a substitution string. The connector performs this substitution on an attribute's value only if a value has been specified for the ESC=[true|false] parameter of the attribute's AppSpecificInfo property or for the connector's ReplaceAllStr property.

The syntax for this attribute is:

```
single_char1,substitution_str1  
[:single_char2,substitution_str2[:...]]
```

where:

single_char A character to be replaced.

substitution_str The substitution string that the connector uses to replace the character.

,

The character delimiter, which separates the character to be replaced from the string that replaces it. By default, the character delimiter is a comma (.). You can configure this delimiter by setting the first delimiter in the StrDelimiter property.

:

The termination delimiter, which separates substitution sets (each of which is composed of the character to be replaced, a character delimiter, and the substitution string). By default, the termination delimiter is a colon (:). You can configure this delimiter by setting the second delimiter in the StrDelimiter property.

For example, assume you want to replace a single percent sign (%) with two percent signs (%%), and a caret (^) with a backslash and a caret (\^). By default, StrDelimiter specifies a comma (,) as the character delimiter, and a colon (:) as the termination delimiter. If you keep the default delimiters, use the following string as the value of ReplaceStrList:

```
%,%:^^,\^
```

Note: A restriction of System Manager prevents entering single quotation marks. Therefore, WebSphere requires you to represent a single quotation with the character Q, and two single quotations with the characters DSQ. In the above example, if you also want to substitute a single quotation mark (') with two single quotation marks (''), use the following notation: Q,DSQ:%,%%:^^,\^

RetryCountAndInterval

Specifies the number of attempts and the interval in seconds that the connector should use when it is unable to lock data while performing an update operation.

Before it performs an update, the connector locks rows related to the update and attempts to retrieve current data. If the connector cannot lock the rows, it tries again to get the lock for the count and interval specified in this configuration property. The connector eventually times out if the lock is not obtainable within the values specified here.

Specify the value in the format: *count, interval in seconds*. For example, a value of 3,20 specifies three retries with an interval of 20 seconds in between.

The default is 3,20.

SPBeforePollCall

This property names the stored procedure that is executed for every poll call. If the property `SPBeforePollCall` has a value (the name of a stored procedure), then at the start of each poll call, the connector calls the stored procedure, passing it the values of the connector properties `ConnectorID` and `PollQuantity`. The procedure will update `PollQuantity` number of rows, setting the connector-id column to `ConnectorID` where `status=0` and connector-id is null. This enables load balancing in the connector.

Note: In the case where a poll call fails prematurely (the database is down, or the connection is lost), the connector-id remains set. This may result in records being skipped during polling. It is therefore recommended that periodically, the connector-id is reset back to null for all records in the event table with a status of 0.

StrDelimiter

Specifies the character and termination delimiters for use in the `ReplaceStrList` property.

- The character delimiter separates the character to be replaced from the string that replaces it. The character delimiter occupies the first (left-hand) position of this property's values and defaults to a comma (,).
- The termination delimiter separates substitution sets (each of which is composed of the character to be replaced, a character delimiter, and the substitution string). The termination delimiter occupies the second (right-hand) position of this property's values and defaults to a colon (:).

You can specify your own value for either or both of these delimiters. If you do so, do not specify a space or other character between them.

Default value is a comma followed immediately by a colon (, :)

TimingStats

Allows you to time each verb operation of the connector to look for problems. Available settings are:

- 0 (No timing statistics)
- 1 (Timing displayed at entry and exit of the verb operation for an entire hierarchical business object).
- 2 (Timing displayed at entry and exit of each verb operation for each individual business object in a hierarchical business object).

Timing messages are log messages rather than trace messages. They can be turned on and off, independent of trace levels.

The default value is 0.

UniqueIDTableName

Specifies the table that contains the latest value used for generation of a unique ID. By default, the table has one column (`id`). You can customize the table to add a column for each attribute that requires generation of a UID.

The default value is `xworlds_uid`.

UseDefaults

If UseDefaults is set to true or is not set, the connector checks whether a valid value or a default value is provided for each required business object attribute. If a value is provided, the Create succeeds; otherwise, it fails.

If UseDefaults is set to false, the connector checks only whether a valid value is provided for each required business object attribute; the Create operation fails if a valid value is not provided.

The default value is false.

UseDefaultsForCreatingChildBOs

If UseDefaultsForCreatingChildBOs is set to true or is not set, the connector checks whether a valid value or a default value is provided for each required business object attribute. If a value is provided, the Create succeeds; otherwise, it fails.

If UseDefaultsForCreatingChildBOs is set to false, the connector checks only whether a valid value is provided for each required business object attribute; the Create operation fails if a valid value is not provided.

UseDefaultsForRetrieve

For polling: If UseDefaultsForRetrieve is not defined and set to true, the default values will be set in the BO before it is retrieved from the database and dispatched to the server.

If UseDefaultsForRetrieve is defined and set to false, the default values will not be set in the BO before it is retrieved from the database and dispatched to the server.

For request processing: If UseDefaultsForRetrieve is not defined and set to false, the default values will not be set in the BO before it is retrieved from the database and dispatched to the server.

If UseDefaultsForRetrieve is defined and set to true, the default values will be set in the BO before it is retrieved from the database and dispatched to the server.

Creating multiple instances of a connector

Creating multiple instances of a connector is in many ways the same as creating a custom connector. You can set your system up to create and run multiple instances of a connector by following the steps below. You must:

- Create a new directory for the connector instance
- Make sure you have the requisite business object definitions
- Create a new connector definition file
- Create a new start-up script

Create a new directory

You must create a connector directory for each connector instance. This connector directory should be named:

```
ProductDir\connectors\connectorInstance
```

where connectorInstance uniquely identifies the connector instance.

If the connector has any connector-specific meta-objects, you must create a meta-object for the connector instance. If you save the meta-object as a file, create this directory and store the file here:

`ProductDir\repository\connectorInstance`

Create business object definitions

If the business object definitions for each connector instance do not already exist within the project, you must create them.

1. If you need to modify business object definitions that are associated with the initial connector, copy the appropriate files and use Business Object Designer to import them. You can copy any of the files for the initial connector. Just rename them if you make changes to them.
2. Files for the initial connector should reside in the following directory:

`ProductDir\repository\initialConnectorInstance`

Any additional files you create should be in the appropriate `connectorInstance` subdirectory of `ProductDir\repository`.

Create a connector definition

You create a configuration file (connector definition) for the connector instance in Connector Configurator. To do so:

1. Copy the initial connector's configuration file (connector definition) and rename it.
2. Make sure each connector instance correctly lists its supported business objects (and any associated meta-objects).
3. Customize any connector properties as appropriate.

Create a start-up script

To create a startup script:

1. Copy the initial connector's startup script and name it to include the name of the connector directory:
`dirname`
2. Put this startup script in the connector directory you created in "Create a new directory" on page 58.
3. Create a startup script shortcut (Windows only).
4. Copy the initial connector's shortcut text and change the name of the initial connector (in the command line) to match the name of the new connector instance.

You can now run both instances of the connector on your integration server at the same time.

For more information on creating custom connectors, refer to the *Connector Development Guide for C++ or for Java*.

Starting the connector

A connector must be explicitly started using its **connector start-up script**. The startup script should reside in the connector's runtime directory:

`ProductDir\connectors\connName`

where `connName` identifies the connector. The name of the startup script depends on the operating-system platform, as Table 26 shows.

Table 26. Startup scripts for a connector

Operating system	Startup script
UNIX-based systems	connector_manager_connName
Windows	start_connName.bat

You can invoke the connector startup script in any of the following ways:

- On Windows systems, from the **Start** menu
 Select **Programs>IBM WebSphere Business Integration Adapters>Adapters>Connectors**. By default, the program name is “IBM WebSphere Business Integration Adapters”. However, it can be customized. Alternatively, you can create a desktop shortcut to your connector.
 - From the command line
 - On Windows systems:
`start_connName connName brokerName [-cconfigFile]`
 - On UNIX-based systems:
`connector_manager_connName -start`
 where *connName* is the name of the connector and *brokerName* identifies your integration broker, as follows:
 - For WebSphere InterChange Server, specify for *brokerName* the name of the ICS instance.
 - For WebSphere message brokers (WebSphere MQ Integrator, WebSphere MQ Integrator Broker, or WebSphere Business Integration Message Broker) or WebSphere Application Server, specify for *brokerName* a string that identifies the broker.
- Note:** For a WebSphere message broker or WebSphere Application Server on a Windows system, you *must* include the `-c` option followed by the name of the connector configuration file. For ICS, the `-c` is optional.
- From Adapter Monitor (WebSphere Business Integration Adapters product only), which is launched when you start System Manager
 You can load, activate, deactivate, pause, shutdown or delete a connector using this tool.
 - From System Monitor (WebSphere InterChange Server product only)
 You can load, activate, deactivate, pause, shutdown or delete a connector using this tool.
 - On Windows systems, you can configure the connector to start as a Windows service. In this case, the connector starts when the Windows system boots (for an Auto service) or when you start the service through the Windows Services window (for a Manual service).

For more information on how to start a connector, including the command-line startup options, refer to one of the following documents:

- For WebSphere InterChange Server, refer to the *System Administration Guide*.
- For WebSphere message brokers, refer to *Implementing Adapters with WebSphere Message Brokers*.
- For WebSphere Application Server, refer to *Implementing Adapters with WebSphere Application Server*.

Stopping the connector

The way to stop a connector depends on the way that the connector was started, as follows:

- If you started the connector from the command line, with its connector startup script:
 - On Windows systems, invoking the startup script creates a separate “console” window for the connector. In this window, type “Q” and press Enter to stop the connector.
 - On UNIX-based systems, connectors run in the background so they have no separate window. Instead, run the following command to stop the connector:
`connector_manager_connName -stop`
where *connName* is the name of the connector.
- From Adapter Monitor (WebSphere Business Integration Adapters product only), which is launched when you start System Manager
You can load, activate, deactivate, pause, shutdown or delete a connector using this tool.
- From System Monitor (WebSphere InterChange Server product only)
You can load, activate, deactivate, pause, shutdown or delete a connector using this tool.
- On Windows systems, you can configure the connector to start as a Windows service. In this case, the connector stops when the Windows system shuts down.

Chapter 5. Developing business objects for the Mainframe Adapter Suite

This chapter describes how the MAS connector processes business objects and describes the assumptions the connector makes when retrieving and modifying data. It contains the following sections:

- “Business object and attribute naming conventions”
- “Business object structure”
- “Business object verb processing” on page 68
- “Business object attribute properties” on page 83
- “Business object application-specific information” on page 85

You can use this information as a guide to modifying existing business objects or as suggestions for implementing new ones.

The connector makes assumptions about the structure of its supported business objects, the relationships between parent and child business objects, the format of the application-specific information, and the application representation of the business object. Therefore, when you create or modify a business object that will be processed by the connector, your modifications must conform to the rules the connector is designed to follow. If they do not, the connector cannot process new or modified business objects correctly.

Business object and attribute naming conventions

The name of a business object used by the connector can consist only of alphanumeric characters or the underscore character. Business object attribute names also can consist only of alphanumeric characters or the underscore character.

Business object structure

In most cases, the connector assumes that every individual business object is represented by one application database table or view unless the business object is based on a stored procedure. In the case of IMS/TM or CICS/TS systems, the business object is based on SHADOW calls to the data source. **Although this chapter refers to business objects that are defined as table- or view-based, the general architecture applies to stored procedure-based business objects, including those for IMS/TM or CICS/TS.** The only difference is that IMS/TM or CICS/TS objects use SHADOW calls to the data source instead of SQL statements.

The connector assumes that each **simple attribute** (that is, an attribute that represents a single value, such as a String or Integer or Date) within the object is represented by a column in that table or view. Thus, attributes within the same individual business object cannot be stored in different application database tables. However, the following situations are possible:

- The application database table might have more columns than the corresponding individual business object has simple attributes (that is, some columns in the database are not represented in the business object). Include in your design only those columns needed for the business object processing.

- The individual business object might have more simple attributes than the corresponding application database table has columns (that is, some attributes in the business object are not represented in the application database). The attributes that do not have a representation in the application database either have no application-specific information or are set with a default value or specify stored procedures.
- The individual business object can represent a view that spans multiple database tables. The connector can use such a business object when processing Create, Retrieve, Update, and Delete events triggered in the application. However, when processing business object requests, the connector can use such a business object only for Retrieve requests.
- The individual business object can represent a wrapper object that is used as a container for unrelated business objects. The wrapper object is not represented by a database table or view. Wrapper objects may not be used as children of other objects.

Note: If a business object is based on a stored procedure, each simple attribute (other than the special SP attributes) may or may not have application-specific information. For more information, see “Using stored procedures” on page 76.

WebSphere business objects can be flat or hierarchical. All the attributes of a **flat** business object are simple and represent a single value.

A hierarchical business object has attributes that represent a child business object, an array of child business objects, or a combination of both. In turn, each child business object can contain a child business object or an array of business objects, and so on. A **single-cardinality relationship** occurs when an attribute in a parent business object represents a single child business object. In this case, the attribute is of the same type as the child business object.

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 an array of the same type as the child business objects.

Note: 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 a single business object, independent of any child business objects it might contain or that contain it. The term **top-level** business object refers to the individual business object at the top of the hierarchy that does not itself have a parent business object.

The connector supports the following relationships among business objects:

- “Single-cardinality relationships” on page 65
- “Single-cardinality relationships and data without ownership” on page 65
- “Multiple-cardinality relationships” on page 66
- “Wrapper objects” on page 68

In each type of cardinality, the relationship between the parent and child business objects is described by the application-specific information of the key attribute of the business object storing the relationship. For more details about this application-specific information, see "FK=[fk_object_name.]fk_attribute_name" on page 88 in Table 28 on page 88.

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 single-cardinality child of a given type.

Figure 17 illustrates a typical single-cardinality relationship. In the example, `fk1` is the simple attribute that contains the child's primary key, and `child[1]` is the attribute that represents the child business object.

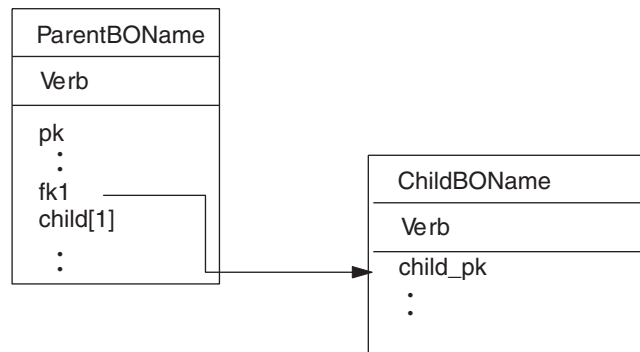


Figure 17. Typical single-cardinality relationship

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 a single Address business object, when a new customer is created, a new row is inserted into both the customer and the 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, there are situations where multiple hierarchical business objects contain the same data, which none of them owns. For example, assume that an Address business object has a `StateProvince[1]` attribute that represents the StateProvince lookup table with single cardinality. Because the lookup table is rarely updated and is maintained independently of the address data, creation or modification of address data does not affect the data in the lookup table. The connector either finds an existing state or province name or fails. It does not add or change values in the lookup table.

When multiple business objects contain the same single-cardinality child business object, the foreign-key attribute in each parent business object must specify the relationship as `NO_OWNERSHIP`. When a business object request sends the connector a hierarchical business object with a Create, Delete, or Update request, the connector ignores single-cardinality children contained without ownership. The connector performs only retrieves on these business objects. If the connector fails to retrieve such a single-cardinality business object, it returns an error and stops processing.

For information on how to specify the relationship without ownership, see “Attributes that represent a single-cardinality child business Object” on page 92. For more information on specifying foreign key relationships, see “Specifying an attribute’s foreign key” on page 89.

Denormalized data and data without ownership

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

Synchronizing from normalized to denormalized data: Specifying a relationship as `NO_OWNERSHIP` allows you to create or change data when you synchronize from a normalized application to a denormalized one. For example, assume that your normalized source application stores data in two tables, A and B. Assume further that your denormalized destination application stores all the data in a single table such that each entity A redundantly stores B data.

In this example, to synchronize a change in table B data from your source application to your destination application, you must trigger a table A event whenever table B data changes. Moreover, 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.

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

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

Multiple-cardinality relationships

Typically, a business object that contains an array of child business objects has only one attribute that represents the relationship. The type of the attribute is an array of the same type as the child business objects. In order for a parent to contain more than one child, the foreign keys that establish the relationship are stored in each 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.

Figure 18 illustrates a multiple-cardinality relationship. In the example, `parentId` is the simple attribute that contains the parent’s primary key, and `child[n]` is the attribute that represents the array of child business objects.

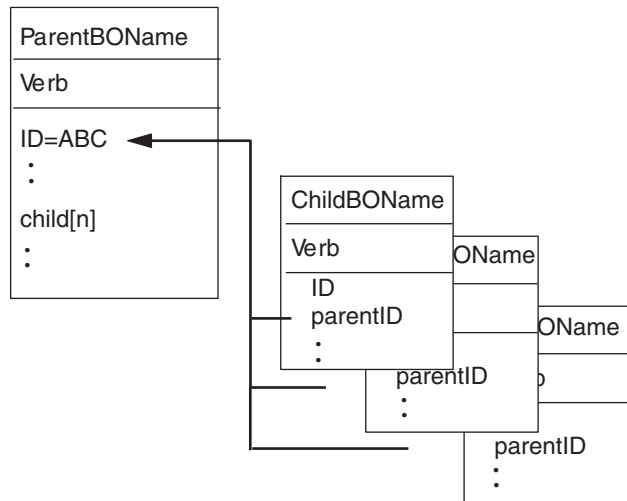


Figure 18. Multiple-cardinality business object relationship

Single-cardinality relationships that store the relationship in the child

Some applications store a single child entity so that the 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.

Figure 19 illustrates this special type of single-cardinality relationship.

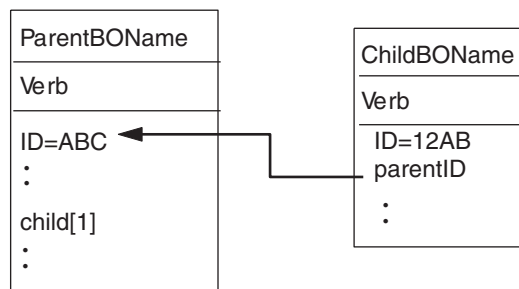


Figure 19. Single-cardinality business object with relationship stored in the child

Applications use this type of single-cardinality relationship when child data does not exist independently of its parent and can be accessed only through its parent. Such child data is never owned by more than one parent, and requires that the parent and its primary-key value exist before the child and its foreign-key value can be created.

To accommodate such applications, the connector also supports hierarchical business objects that contain a child with single cardinality but store the relationship in the child rather than in the parent.

To specify that a parent business object contains a single-cardinality child in this special way, when you specify the application-specific information of the attribute

that contains the child, do not include the CONTAINMENT parameter. For more information, see “Attributes that represent a single-cardinality child business Object” on page 92.

Wrapper objects

The wrapper object is a top-level business object that does not correspond to any database table or view. The wrapper object is denoted by the top-level business object property of WRAPPER with a value of true. The wrapper object is a dummy parent that is used as a container for unrelated children. In processing the wrapper object, the connector ignores the top-level business object and processes only the children. The wrapper object may contain N cardinality or N-1 cardinality entities or both.

A N cardinality entity should have at least one unique attribute marked as a primary key and at least one attribute marked as a foreign key. This foreign key will then be added as a primary key in the wrapper object. The entity’s foreign key will reference the wrapper object’s primary key that was just added.

In the case of a N-1 cardinality entity, the primary key should be marked as both a primary key and a foreign key, referencing the primary key in the wrapper, which is the same as the primary key in the N-1 entity.

Business object verb processing

This section describes the following aspects of processing a business object’s verbs:

- “Verb determination,” which explains how the connector determines the verb to use for each individual, source business object
- “Afterimages and deltas,” which defines the terms and explains how the connector works with afterimages
- “Verb processing” on page 70, which explains the steps the connector takes when creating, retrieving, updating, or deleting a business object
- “Using SQL statements” on page 76, which explains how the connector uses stored procedures
- “Using stored procedures” on page 76, which describes when the connector uses them and how to specify them
- “Processing business objects using stored procedures or simple SQL statements” on page 79, which explains how the connector processes stored procedures and SQL statements
- “Transaction commit and rollback” on page 83, which briefly explains how the connector uses transaction blocks

Verb determination

A top-level business object and each of its individual child business objects can contain their own verbs. Therefore, a business object request can pass a business object that has different verbs for parent and child business objects to the connector. When this occurs, the connector uses the verb of the top-level parent business object to determine how to process the entire business object. For more information, see “Verb processing” on page 70.

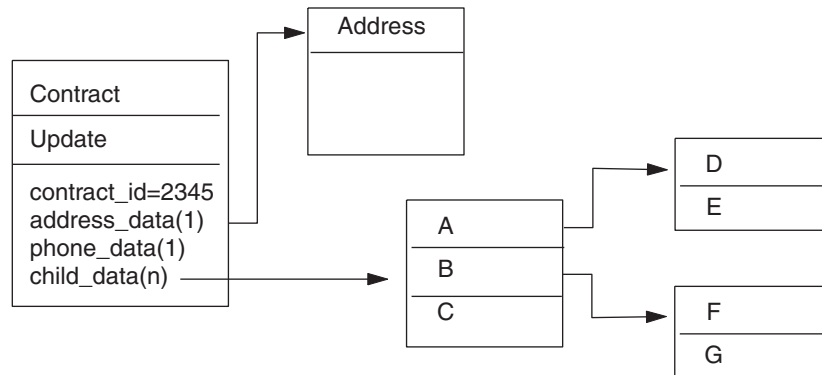
Afterimages and deltas

An **afterimage** is the state of a business object after all changes have been made to it. A **delta** is a business object used in an update operation that contains only key values and the data to be changed. Because the connector supports only

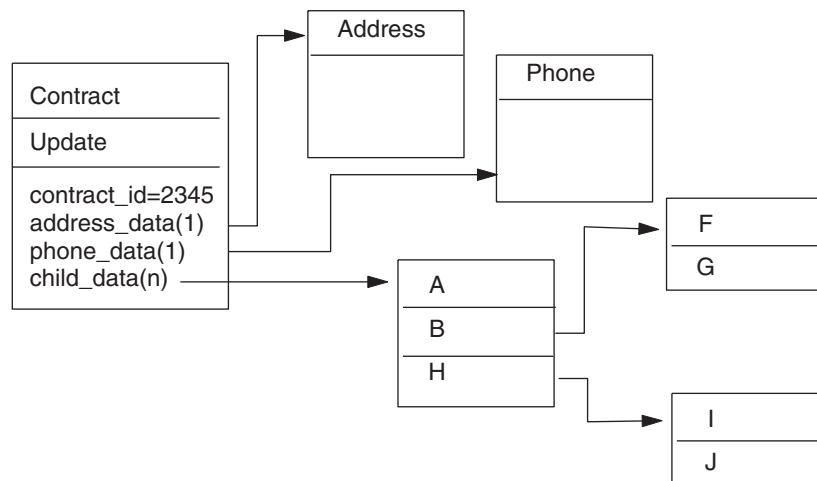
afterimages, when it receives a business object for update, the connector assumes that the business object represents the desired state of the data after update.

Therefore, a connector receives a business object with the Update verb, the connector changes the current representation of the business object in the application so that it exactly matches the source business object. To do this, the connector changes simple attribute values and adds or removes child business objects.

For example, assume the current state of Contract 2345 in the database is as follows:



Assume further that the connector receives the following business object:



To process the update, the connector applies the following changes to the database:

- Update the simple attributes in the top-level Contract and Address business objects
- Create the Phone business object
- Update the simple attributes in the child business objects A, B, F and G
- Delete the child business objects C, D and E
- Create the child business objects H, I and J

Because the connector assumes that each business object it receives represents an afterimage, it is important that developers ensure that each business object sent to such a connector for updating contains valid existing child business objects. Even if

none of a child business object's simple attributes have changed, the child business object must be included in the source business object.

There is a way, however, that you can prevent some connectors from deleting missing child business objects during an update operation. You can use the application-specific information for the attribute that represents the child or array of children to instruct the connector to keep child business objects that are not included in the source business object. To do so, set `KEEP_RELATIONSHIP` to true. For more information, see "Specifying an attribute's foreign key" on page 89.

Verb processing

This section outlines the steps the connector takes when creating, retrieving, updating, or deleting a business object. The connector processes hierarchical business objects recursively; that is, it performs the same steps for each child business object until it has processed all individual business objects.

Note: A top-level business object that is a wrapper supports the create, retrieve, update and delete verbs. The only difference in processing a wrapper object is that the wrapper object is not processed, only the objects that it contains are processed.

Business object comparison

At various points in the processing outlined below, the connector compares two business objects to see if they are the same. For example, during an update operation, the connector determines whether a particular business object exists in an array of business objects. To perform the check, the connector compares the business object to each business object within the array. For two business objects to be identical, the following two conditions must be satisfied:

- The type of the business objects being compared must be the same. For example, a Customer business object is never considered identical to a Contact business object even if all of their attributes are exactly the same.
- All corresponding key attributes in the two business objects must contain identical values. If a key attribute is set to `CxIgnore` in both business objects, the connector considers them identical. However, if a key attribute is set to `CxIgnore` in one business object but not in the other, the business objects are not identical.

Create operations

When creating a business object, the connector returns a status either of `VALCHANGE` if the operation was successful (regardless of whether the operation caused changes to the business object), or `FAIL` if the operation failed.

The connector performs the following steps when creating a hierarchical business object:

1. Recursively inserts each single-cardinality child business object contained with ownership in the application. In other words, the connector 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 connector ignores the attribute. However, if the business object definition requires that attribute to represent a child and it does not, the connector returns an error and stops processing.

2. Processes each single-cardinality child business object contained without ownership as follows:

- a. Recursively attempts to retrieve the child from the database using the key values passed in by the business object request.
- b. If the retrieve is unsuccessful, indicating that the child does not currently exist in the application, the connector returns an error and stops processing. If the retrieve is successful, the connector recursively updates the child business object.

Note: For this approach to work correctly when the child business object already exists in the application, developers must ensure that primary key attributes in child business objects are cross-referenced correctly on create operations. If the child business object does not already exist in the application database, set the primary key attributes to CxBlank.

3. Inserts the top-level business object in the database as follows:
 - a. Sets each of its foreign-key values 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 connector 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's application-specific information. If an attribute has an associated database sequence or counter, the value generated by the connector overwrites any value passed in by the request. For more information on specifying a database sequence or counter, see UID=AUTO in "Application-specific information for simple attributes" on page 87.
 - c. Copies the value of an attribute to the value of another attribute as specified by the CA (CopyAttribute) parameter of the attribute's application-specific information. For more information on using the CA parameter, see CA=set_attr_name in "Application-specific information for simple attributes" on page 87.
 - d. Inserts the top-level business object into the database.

Note: A top-level business object that is a wrapper will not be inserted into the database.

4. Processes each of its single-cardinality child business objects that stores the parent/child relationship in the child, as follows:
 - a. Sets the foreign-key values in the child to reference the value in the corresponding primary-key attributes in the parent. Because the parent's primary-key values may have been generated during the creation of the parent, this ensures that the foreign-key values in each child are correct before the connector inserts the child in the application.
 - b. Inserts the child in the application.
5. 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's primary-key values may have been generated during the creation of the parent, this ensures that the foreign-key values in each child are correct before the connector inserts the child in the application.
 - b. Inserts each of its multiple-cardinality child business objects into the database.

Retrieve operations

The connector performs the following steps when retrieving a hierarchical business object:

1. Removes all child business objects from the top-level business object that it received from the business object request.
2. Retrieves the top-level business object from the application.
 - If the retrieval returns 1 row, the connector continues processing.
 - If the retrieval returns no rows, indicating that the top-level business object does not exist in the database, the connector returns `BO_DOES_NOT_EXIST`.
 - If the retrieval returns more than one row, the connector returns `FAIL`.

Note: A business object can have attributes that do not map to any database column, such as placeholder attributes. During retrieval, the connector does not change such attributes in the top-level business object; they remain set to the values received from the business object request. In child business objects, the connector sets such attributes to their default values during retrieval.

Note: A top-level business object that is a wrapper must contain any attribute values from the objects at the level immediately below the wrapper object, which would be necessary to retrieve the objects, including keys and placeholder attributes. The wrapper object must have all keys and placeholder attributes populated. Simple attributes in the wrapper object that will be used as foreign keys in the objects one level below the wrapper should be marked as keys in the wrapper object.

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

Note: The connector does not enforce uniqueness when populating an array of business objects. It is the database's responsibility to ensure uniqueness. If the application returns duplicate child business objects, the connector 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 occurrence in the business object and before the parent business object is processed. Child object ownership and non-ownership do not determine the processing sequence, but do determine the type of processing.

RetrieveByContent operations

A `RetrieveByContent` verb is applicable only for the top-level business object, because the connector performs a retrieval based on attributes only in the top-level business object.

If a top-level business object uses the `RetrieveByContent` verb, all of the attributes (including non-key attributes) that are not null are used as retrieval criteria.

If more than one row is returned, the connector uses the first row as the result row and returns `MULTIPLE_HITS`.

Note: A `RetrieveByContent` verb is not applicable for a top-level business object that is a wrapper.

Update operations

When updating a business object, the connector returns a status either of VALCHANGE if the operation was successful (regardless of whether the operation caused changes to the business object), or FAIL if the operation failed.

The connector performs the following steps when updating a hierarchical business object:

1. Uses the primary-key values of the source business object to retrieve the corresponding entity from the application. 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 application, the connector returns BO_DOES_NOT_EXIST and the update fails.

Note: A top-level business object that is a wrapper does not have to exist in the database. However, it must contain any attribute values from the objects at the level immediately below the wrapper object, which would be necessary to retrieve the objects including keys and placeholder attributes. The wrapper object must have all keys and placeholder attributes populated. Simple attributes in the wrapper object that will be used as foreign keys in the objects one level below the wrapper should be marked as keys in the wrapper object.

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

If all the simple attributes in the top-level business object represent keys, the connector cannot generate an update query for the top-level business object. In this case, the connector logs a warning and continues to step 2.

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

If the business object definition requires that an attribute represent a child business object, the child must exist in both the source business object and the retrieved business object. If it does not, the update fails, and the connector returns an error.

The connector 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 already existing child in the database, the connector deletes the existing child and creates the new child.
- If the child is present in the source business object but not in the retrieved business object, the connector recursively creates it in the database.
- If the child is present in the retrieved business object but not in the source business object, the connector recursively deletes it from the database. The type of delete, physical or logical, depends on the value of its ChildUpdatePhyDelete property.

For single-cardinality children contained without ownership, the connector attempts to retrieve every child from the application that is present in the source business object. If it successfully retrieves the child, the connector

populates the child business object, but does not update it; single-cardinality child objects that are contained without ownership are never modified by the connector.

3. For single-cardinality child business objects that store the relationship in the parent, sets each foreign-key value in the parent to the value of the primary key in the corresponding single-cardinality child business object. This step is necessary because single-cardinality children may have been added to the database during previous steps, resulting in the generation of new unique IDs.
4. Updates all simple attributes of the retrieved business object except those whose corresponding attribute in the source business object contain the value CxIgnore.

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

5. Sets all foreign-key values in each child that stores the parent/child relationship in the child (both multiple-cardinality and single-cardinality) to the primary-key value of its corresponding parent business object. Typically, when ICS is the integration broker, these values have been cross-referenced during data mapping. However, this step is important to ensure that the foreign-key values of new children that store the relationship in the child are correct before the connector updates those children.
6. Processes each multiple-cardinality child of the retrieved business object in one of the following ways:
 - If the child is present in both the source and the retrieved business objects, instead of updating the already existing child in the database, the connector deletes the existing child and creates the new child.
 - If the child exists in the source array but not in the retrieved business object's array, the connector recursively creates it in the application.
 - If the child exists in the retrieved business object's array but not in the source array, the connector recursively deletes it from the application unless the application-specific information for the attribute that represents the child in the parent has KEEP_RELATIONSHIP set to true. In this case, the connector does not delete the child from the application. For more information, see "Specifying an attribute's foreign key" on page 89. The type of delete, physical or logical, depends on the value of its ChildUpdatePhyDelete property.

Note: The developer's code must ensure that business objects contained with multiple cardinality in the source business object are unique (that is, that an array does not contain two or more copies of the same business object). If the connector receives duplicates of a business object in a source array, it processes the business object twice, with possibly unpredictable results.

DeltaUpdate operations

DeltaUpdate verb processing is different from update verb processing as follows:

1. On a DeltaUpdate no retrieve is done before updating, as is done in update verb processing.
2. No comparisons are made between the incoming business object and the business object in the database.
3. All children will be processed based on the verb set in each child object. If a child doesn't have a verb set in it, the connector will return an error.

When delta updating a business object, the connector returns a status of either VALCHANGE if the operation was successful (regardless of whether the operation caused changes to the business object) or FAIL if the operation failed.

The connector performs the following steps when delta updating a hierarchical business object:

1. Recursively processes all single-cardinality children of the parent object. If a child is marked as IsRequired in the business object specification, it must be present in the inbound object. If not, the delta update will fail and the connector will return an error.
2. Sets all foreign key values in the parent that reference attributes in single-cardinality children to their corresponding child values. This is necessary because single-cardinality children may have been added to the database during the previous steps, resulting in the generation of new sequence values.
3. Updates the current object being processed via an SQL UPDATE statement or a stored procedure. All simple attributes of the individual business object are updated, except those attributes set to IsIgnore in the inbound business object. The connector does not compare the inbound object to the current object on an attribute level to determine which attributes need to be added to the update statement; they are all updated. Since the object being updated should be unique, the connector checks to make sure that only one row is processed as a result. An error is returned if more than one row is processed.
4. Sets all foreign key values in all cardinality N children of the current object that reference parent attributes to the corresponding parent values. Usually these values will already be cross-referenced during data mapping; however, this may not be the case for new children in cardinality N containers. This ensures that the foreign key values in all cardinality N children are correct before those children are updated.
5. Updates all cardinality N containers of the current object.

When the child objects are processed, each child's verb is taken and the appropriate operation is done. The allowed verbs on a child in DeltaUpdate are Create, Delete and DeltaUpdate.

- If a Create verb is found in the child, the child gets created in the database if it is an ownership child. Non-ownership children are retrieved to validate their existence in the database.
- If a Delete verb is found in the child, that child gets deleted.
- If a DeltaUpdate verb is found in the child, the child gets updated in the database.

Delete operations

When deleting a business object, the connector returns a status of SUCCESS if the operation was successful or FAIL if the operation failed. The parent business object is first retrieved and then the adapter recursively deletes all single-cardinality children that have an ownership relationship to the parent, then the parent business object itself, and finally all cardinality N children. Single-cardinality no-ownership children are never deleted. If the business object does not exist, the connector returns a FAIL.

The connector supports logical and physical deletes, depending on the Status Column Name (SCN) value in the object's application-specific information. If the SCN value is defined, the connector performs a logical delete. If the SCN value is not defined, the connector performs a physical delete.

Physical deletes: The connector performs the following steps when physically deleting a hierarchical business object:

1. Recursively deletes all single-cardinality child business objects contained with ownership.
2. Deletes the top-level business object.
3. Recursively deletes all multiple-cardinality child business objects.

Note: A top-level business object that is a wrapper does not have a corresponding database table, hence it will not be deleted from the database. Any simple attribute values for a wrapper will be ignored.

Logical deletes: When logically deleting a business object, the connector performs the following steps:

1. Issues an UPDATE that sets the business object's status attribute to the value specified by the business object's application-specific information. The connector ensures that only one database row is updated as a result, and it returns an error if this is not the case.
2. Recursively logically deletes all single-cardinality children contained with ownership and all multiple-cardinality children. The connector does not delete single-cardinality children contained without ownership.

Using SQL statements

The connector can use simple SQL statements for select, update, retrieve or delete operations. The column names for SQL statements are derived from an attribute's `AppSpecificInfo` property. Each query spans one table only, unless posted to a view.

Using stored procedures

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 connector to execute on an object in a database server. An attribute's `AppSpecificInfo` property is used to pass parameters to stored procedures.

The connector calls stored procedures in the following circumstances:

- Before processing a business object, to perform preparatory operational processes
- After processing a business object, to perform post-operational processes
- To perform a set of operations on a business object, instead of using a simple INSERT, RETRIEVE, UPDATE, or DELETE statement

When it processes a hierarchical business object, the connector can use a stored procedure to process the top-level business object or any of its child business objects. However, each business object or array of business objects must have its own stored procedure.

Specifying a stored procedure

This section describes the steps you must perform to cause the connector to use a stored procedure for a business object. It contains the following sections:

- "Adding attributes to the business object" on page 77
- "Syntax of a stored procedure" on page 77
- "Examples of stored procedures" on page 78
- "Specifying the stored procedure" on page 78

Adding attributes to the business object: You must add a special kind of attribute to the business object for each type of stored procedure that the connector processes. These attributes represent only the stored procedure's type and the application-specific information that defines it. These attributes do not use the application-specific information parameters available for a standard simple attribute.

Name the attribute according to the type of stored procedure to be used. For example, to cause the connector to use AfterUpdate and BeforeRetrieve stored procedures, add the AfterUpdateSP and BeforeRetrieveSP attributes.

The connector recognizes the following business object attribute names:

```
BeforeCreateSP
AfterCreateSP
CreateSP
BeforeUpdateSP
AfterUpdateSP
UpdateSP
BeforeDeleteSP
AfterDeleteSP
DeleteSP
BeforeRetrieveSP
AfterRetrieveSP
RetrieveSP
BeforeRetrieveByContentSP
AfterRetrieveByContentSP
RetrieveByContentSP
BeforeRetrieveUpdateSP
AfterRetrieveUpdateSP
RetrieveUpdateSP
```

Note: Create an attribute only for those stored procedures that you want the connector to execute. For example, when ICS is the integration broker, use the application-specific information or mapping to specify values for these attributes before the business object is sent to the connector. The connector must be restarted to recognize changes to these values for subsequent calls on a business object.

Syntax of a stored procedure: The syntax for specifying a stored procedure is:

```
SPN=StoredProcedureName;RS=[true|false]
[;IP=Attribute_Name1[:Attribute_Name2
[:...]]];OP=Attribute_Name1|RS[:Attribute_Name2|RS
[:...]]];IO=Attribute_Name1[:Attribute_Name2[:...]]]
```

where:

StoredProcedureName

The name of the stored procedure.

RS

Result Set: If true, the stored procedure returns a result set, else false if no result set is returned. The default is false. If the value is true, the ColumnName property in an attribute's application-specific information points to the appropriate column in the result set. If RS is part of the output parameter list, then that particular parameter returns a result set. Only one result set OUT parameter is supported. If more than one result set is returned as an OUT parameter, only the first result set is returned and all others are ignored. Currently, this feature is supported for Oracle 8i

and above, stored procedures that use the Oracle JDBC Driver. For the stored procedure in the database, the corresponding parameter should return a REF CURSOR type.

IP	Input Parameters: The list of business object attributes whose values the connector should use as input values when executing the stored procedure.
OP	Output Parameters: The list of business object attributes to which the connector should return values after executing the stored procedure. See RS for a description of the result set.
IO	InputOutput Parameters: The list of business object attributes whose values the connector should use as input values and to which the connector should return values after executing the stored procedure.

Note: The order of the SPN and RS parameters is important; the order of the other parameters is not syntactically important. In other words, it makes no difference to the connector if the stored procedure groups all parameters of each type or intersperses the types of parameters.

When multiple parameters of the same type are grouped together, separate the values with a colon delimiter; you need not repeat the parameter's name for each value. Separate parameters of different types with a semicolon delimiter. When specifying parameter values, do not put a blank space on either side of the equal sign (=).

Examples of stored procedures: The following examples use stored procedures named CustomerInsert and VendorInsert that get values from two input attributes, and return values to four output attributes. The examples illustrate different structures for stored procedures.

- Parameters of the same type are grouped together (IP, IP, OP, OP, OP, IO):
SPN=CustomerInsert;RS=false;IP=LastName:FirstName;
OP=CustomerName:CustomerID:ErrorStatus:ErrorMessage;
IO=VendorID
- Parameters of the same type are interspersed (IP, OP, OP, OP, IP, IO, OP):
SPN=VendorInsert;RS=false;IP=LastName;OP=CustomerName:CustomerID:
ErrorStatus;IP=FirstName;IO=VendorID;OP=ErrorMessage

The connector supports only the simple data types supported by the JDBC driver.

Specifying the stored procedure: There are two ways to specify the stored procedure name and its parameter values:

- Attribute's AppSpecificInfo property
If the length of the text that specifies the stored procedure is less than or equal to 4000 bytes, you can specify the value in the attribute's AppSpecificInfo property. You can use this property to specify the stored procedure regardless of whether the connector has polled for the business object (that is, the business object represents an application event) or has received the business object as a request.

The following example illustrates specification of the stored procedure in application-specific information. In this case, the value specified for the MaxLength property is not important to the stored procedure.

```
[Attribute]
Name = BeforeCreateSP
Type = String
MaxLength = 15
IsKey = false
IsRequired = false
AppSpecificInfo =SPN=ContactInsert;IP=LastName:FirstName;OP=CustomerName:
CustomerID:ErrorStatus:ErrorMessage
```

[End]

- **Attribute's value**

If the length of the text that specifies the stored procedure is more than 4000 bytes, you must specify the stored procedure. When ICS is the integration broker, you can use mapping to specify the stored procedure, but only if the business object represents a request. In other words, you cannot use an attribute's value to specify a stored procedure when the connector is polling for events.

If the text of the stored procedure is longer than 4000 bytes and you use mapping to specify it, remember to expand the value of the MaxLength property to accommodate the full text.

Note: If a stored procedure that handles a create, update, or delete operation is executed on a hierarchical business object containing an array of child business objects, the connector processes each child business object individually. For example, if the connector executes a BeforeCreate stored procedure, it does not process the array as a unit but processes each member in the array. When it processes a BeforeRetrieve stored procedure, the connector operates on a single business object. When it processes an AfterRetrieve stored procedure, the connector operates on all business objects returned by the retrieval.

Processing business objects using stored procedures or simple SQL statements

The following sections explain how the connector processes stored procedures and SQL statements:

- "Business object create operations"
- "Business object update operations" on page 80
- "Business object delete operations" on page 80
- "Business object retrieve operations" on page 81
- "Business object RetrieveByContent operations" on page 82
- "Business object Retrieve-for-Update operations" on page 82

Business object create operations

A Create stored procedure usually returns values that the connector uses to populate the simple attributes in the top-level business object. The connector performs the following steps when processing a business object Create operation using SQL statements and stored procedures (BeforeCreate, Create, AfterCreate):

1. Checks whether the business object contains a BeforeCreateSP attribute. If it does, calls the BeforeCreate stored procedure.
2. If the stored procedure returns values through output parameters, uses the values to set the value of simple attributes in the business object.
3. Creates the single-cardinality child business objects.
4. Sets each of the top-level business object's foreign key values to the primary-key value of each single-cardinality child business object.

5. Checks whether the business object contains a CreateSP attribute. If it does, calls the Create stored procedure to create the top-level business object. If it does not, builds and executes an INSERT statement to create the top-level business object.
6. If the Create stored procedure returns values through output parameters, uses the values to set the value of simple attributes in the business object.
7. Sets the foreign-key value in each multiple-cardinality child to the value of its parent's primary-key attribute.
8. Creates the multiple-cardinality child business objects.
9. Checks whether the business object contains an AfterCreateSP attribute. If it does, calls the AfterCreate stored procedure.
10. If the stored procedure returns values through output parameters, uses the values to set the values of simple attributes in the business object.

The connector can use values returned in step 10 to change the values of a business object that it created in steps 3 or 5.

Business object update operations

An Update stored procedure usually returns values that the connector uses to populate the simple attributes in the top-level business object. The connector performs the following steps when processing the business object Update operation using SQL statements and stored procedures (BeforeUpdate, Update, AfterUpdate):

1. Checks whether the business object contains a BeforeUpdateSP attribute. If it does, calls the BeforeUpdate stored procedure.
2. If the BeforeUpdate stored procedure returns values through output parameters, uses the values to set the value of simple attributes in the business object.
3. Updates the single-cardinality child business objects.
4. Sets each of the top-level business object's foreign-key values to the primary-key value of each child business object contained with single cardinality.
5. Checks whether the business object contains an UpdateSP attribute. If it does, calls the Update stored procedure to update the top-level business object. If it does not, builds and executes an UPDATE statement to update the top-level business object.
6. If the Update stored procedure returns values through output parameters, uses the values to set the value of simple attributes in the business object.
7. Sets foreign-key values in the multiple-cardinality children to reference the value in the corresponding primary-key attributes in the parent.
8. Updates the multiple-cardinality child business objects.
9. Checks whether the business object contains an AfterUpdateSP attribute. If it does, calls the AfterUpdate stored procedure.
10. If the stored procedure returns values through output parameters, uses the values to set the value of simple attributes in the business object.

Business object delete operations

A Delete stored procedure does not return values to the connector. The connector performs the following steps when processing the business object Delete operations using simple SQL statements and stored procedures (BeforeDelete, Delete, AfterDelete):

1. Checks whether the business object contains a BeforeDeleteSP attribute. If it does, calls the BeforeDelete stored procedure.

2. Deletes the single-cardinality child business objects.
3. Deletes the multiple-cardinality child business objects.
4. Checks whether the business object contains a DeleteSP attribute. If it does, calls the Delete stored procedure to delete the top-level business object. If it does not, builds and executes a DELETE statement.
5. Checks whether the business object contains an AfterDeleteSP attribute. If it does, calls the AfterDelete stored procedure.

Business object retrieve operations

For simple RETRIEVE operations, stored procedures can be used for top-level business object, single cardinality children, as well as multiple cardinality children. The order of the procedures is as follows:

- BeforeRetrieve
- Retrieve
- AfterRetrieve

The connector creates a temporary object to retrieve a single cardinality child business object or a multiple cardinality child business object. The connector applies the BeforeRetrieve stored procedure to the temporary business object. The AfterRetrieve stored procedure is applied to each of the child objects retrieved for the container.

The connector executes the AfterRetrieve stored procedure after it executes a Retrieve query generated dynamically from the business object meta-data or stored procedure on the business object.

According to the JDBC specification there are three types of StoredProcedure calls as follows:

- {call <spName>(?,?,?)}
- {call <spName>}
- {?= call <spName>(?,?,?)}

The connector supports the first two types. It will process the ResultSet that is returned from StoredProcedure.

In the stored procedure syntax, if RS=true, the result set from the stored procedure is processed. If RS=false, the result set is not processed. By default the value of RS is false. After the result set values are processed, the stored procedure output variables are processed. If RS=true, multiple cardinality children cannot specify the output variables in the related stored procedure.

Note: Result set processing is supported only for Retrieve verb operations and for RetrieveSP only.

Processing result set returned from retrieve stored procedure (RetrieveSP):

ResultSetMetaData is obtained for the result set returned from the stored procedure. Values of all the columns in the result set are obtained and set on the corresponding attribute of the business object. The ColumnName property of an attribute's application-specific information should contain the ResultSet column name to match the attribute to the column.

For single cardinality objects, the corresponding result set should consist of only one row. If multiple rows are returned in the result set, an error is reported.

For multiple cardinality children, multiple rows can be returned through the result set. For each row returned, a new object is created and added to the container. The container is then added to the parent object at the required attribute index.

Business object RetrieveByContent operations

For simple RetrieveByContent operations, stored procedures can be used only for the top-level business object and its single-cardinality children; that is, they cannot be used to return a result set or multiple rows. The order of the procedures is as follows:

- BeforeRetrieveByContent
- RetrieveByContent
- AfterRetrieveByContent

The connector creates a temporary object to retrieve a single cardinality child business object or a multiple cardinality child business object. For multiple cardinality business objects, the connector applies the BeforeRetrieveByContent stored procedure to the temporary business object. The AfterRetrieveByContent stored procedure is applied to each of the child objects retrieved for the container.

The connector executes the AfterRetrieveByContent stored procedure after it executes a RetrieveByContent query generated dynamically from the business object meta-data or stored procedure on the business object. In this case, even though the retrieval of a hierarchical business object also retrieves its child business objects, the connector executes the AfterRetrieveByContent stored procedure on every business object present in the array.

Business object Retrieve-for-Update operations

The following stored procedures are called on the top-level business object and retrieve all child business objects in the same way as the simple Retrieve.

The order of the procedures is as follows:

- BeforeRetrieveUpdate
- RetrieveUpdate
- AfterRetrieveUpdate

These stored procedures perform the same operations as BeforeRetrieve and AfterRetrieve. They have distinguishing names so that you can create separate attributes to cause the connector to perform both BeforeRetrieve and BeforeRetrieveUpdate operations, as well as AfterRetrieve and AfterRetrieveUpdate operations.

The connector creates a temporary object to retrieve a single cardinality child business object or a multiple cardinality child business object. For multiple cardinality business objects, the connector applies the BeforeRetrieveUpdate stored procedure to the temporary business object. The AfterRetrieveUpdate stored procedure is applied to each of the child objects retrieved for the container.

The connector executes the AfterRetrieveUpdate stored procedure after it executes a RETRIEVE query generated dynamically from the business object meta-data or stored procedure on the business object. In this case, even though the retrieval of a hierarchical business object also retrieves its child business objects, the connector executes the AfterRetrieveUpdate stored procedure on every business object present in the array.

Transaction commit and rollback

Whenever the connector receives a business object for processing, it begins a transaction block. All SQL statements that the connector executes while processing that business object are encapsulated within the transaction block. When the connector finishes processing the business object, it commits the transaction block if the processing was successful, or rolls back the transaction if it encountered an error.

Business object attribute properties

Business object architecture defines various properties that apply to attributes. This section describes how the connector interprets several of these properties and describes how to set them when modifying a business object.

Name property

Each business object attribute must have a unique name.

Type property

Each business object attribute must have a type, such as Integer, String, or the type of a child business object. When the connector encounters an attribute of type Date, Long Text, or String, the connector wraps the value in quotation marks and handles the value as character data.

Cardinality property

Each business object attribute that represents a child or array of child business objects has the value of 1 or n, respectively, in this attribute. All attributes that represent child business objects also have a `ContainedObjectVersion` property (which specifies the child's version number) and a `Relationship` property (which specifies the value Containment).

Max length property

If the attribute is of type String, this property specifies the maximum length allowed for the attribute's value.

Key property

At least one simple attribute in each business object must be specified as the key. To define an attribute as a key, set this property to Yes. If the business object attribute is of type String, WebSphere recommends that the data type in the database is of type Varchar instead of char.

Note: The connector does not support specifying an attribute that represents a child business object or an array of child business objects as a key attribute.

If the key property is set to true for a simple attribute, the connector adds that attribute to the WHERE clause of SELECT, UPDATE, RETRIEVE, and DELETE SQL statements that it generates while processing the business object.

If the key property is set to true for an attribute in a child that stores the parent/child relationship in the child (both multiple-cardinality and single-cardinality), the connector uses the parent's primary keys in the WHERE clause of the SELECT statement, and it does not use the Key property. For information on

specifying the name of business object attributes whose values are used to set the child's foreign-key attributes, see "Application-specific information at the attribute level" on page 87.

Foreign key property

The connector uses this property to determine whether an attribute is a foreign key.

Required property

The Required property specifies whether an attribute must contain a value.

If this property is specified for an attribute that represents a single-cardinality child business object, the connector requires the parent business object to contain a child business object for this attribute.

When the connector receives a business object with a Create request, the connector causes the Create operation to fail if both of the following conditions are true:

- The business object does not have a valid value or a default value for a required attribute.
- Application-specific information does not specify that the connector generate the unique ID.

When the connector receives a business object with a Retrieve request and the business object does not have a valid value or a default value for a required attribute, the connector causes the retrieval operation to fail.

The connector does not use this property for attributes that contain an array of child business objects.

Note: If the key attribute uses a sequence or counter or is populated by the database (UID=AUTO), it should not be marked as Required.

AppSpecificInfo

For information on this property, see "Application-specific information at the attribute level" on page 87.

Default value property

This property specifies a default value that the connector uses to populate a simple attribute if it is not populated with a value from the database table. The connector does not evaluate this property for attributes that represent a child business object or an array of child business objects.

The connector evaluates this property only if the UseDefaults configuration property is set to true. For more information, see Table 23 on page 48.

Special attribute value

Simple attributes in business object can have the special value, CxIgnore. When it receives a business object, the connector ignores all attributes with a value of CxIgnore. It is as if those attributes were invisible to the connector.

When the connector retrieves data from the database and the SELECT statement returns a null value for an attribute, the connector sets the value of that attribute to

CxIgnore by default. If a value has been specified for the UNVL parameter of the attribute's application-specific information, the connector uses that value to represent the null.

Because the connector requires every business object to have at least one primary-key attribute, developers of collaborations and maps should ensure that business objects passed to the connector have at least one primary key that is not set to CxIgnore. The only exception to this requirement is a business object whose primary key is to be generated by the connector using a counter or sequence, or is generated by the database.

When the connector inserts data into the database and a business object attribute has no value specified, it uses the value specified by the attribute's UseNullValue property. For more information about UseNullValue, see UNVL=value in Table 28 on page 88.

Business object application-specific information

Application-specific information in business object definitions provides the connector with application-dependent instructions on how to process business objects. The connector parses the application-specific information from the attributes or verb of a business object or from the business object itself to generate queries for create, update, retrieve, and delete operations.

The connector stores some of the business object's application-specific information in cache and uses this information to build queries for all the verbs.

If you extend or modify an application-specific business object, you must make sure that the application-specific information in the business object definition matches the syntax that the connector expects.

This section provides information on the object-level, attribute, and verb application-specific information format for business objects supported by the connector.

Table 27 provides an overview of the functionality available in business object application-specific information.

Table 27. Overview of application-specific information in supported business objects

Scope of application-specific information	Functionality
Entire business object	Specifies: <ul style="list-style-type: none">• The name of the corresponding database table.• Defines the column whose value the connector uses in the WHERE clause to perform a logical (or soft) delete.• That the top-level business object is a wrapper.

Table 27. Overview of application-specific information in supported business objects (continued)

Scope of application-specific information	Functionality
Simple attributes	<p>Specifies:</p> <ul style="list-style-type: none"> • The database column name for an attribute. • The foreign key relationship between an attribute in the current business object and a parent or child business object. • Automatic generation of unique identifier values. • The name of another attribute within the same business object whose value the connector must use to set the value of the current attribute. • Whether to use the current attribute when sorting a retrieval. • The value to use when the value of the current attribute is null. • String substitution behavior. • Whether to use the LIKE operator or = operator when comparing strings. • The value to use as the wildcard position when the LIKE operator is used.
Attributes that contain a child or an array of child business objects	Specifies whether a single-cardinality child is owned by the parent. Specifies whether the connector deletes child data during an update operation if the data is not represented in the source business object.
Business object verb	Used only for the Retrieve verb, this text specifies the attributes to be included in the WHERE clause for a retrieval. You can also specify operators and attribute values.

The following sections discuss this functionality in more detail.

Application-specific information at the business-object level

Application-specific information at the business-object level allows you to:

- Specify the name of the corresponding application database table.
- Provide the information necessary to perform a physical or logical delete.
- Specify that the top-level business object is a wrapper object.

At the business-object level, application-specific information format consists of parameters separated by colon (:) or semicolon (;) delimiters:

```
TN=TableName; SCN=StatusColumnName:StatusValue
```

where TableName identifies the database table, StatusColumnName is the name of the application database column used to perform logical deletes, and StatusValue is the value that signifies that a business object is inactive or deleted.

For example, assume that a Customer business object has the following value specified for its business object application-specific information:

```
TN=CUSTOMER; SCN=CUSTSTATUS:DELETED
```

Assume also that the connector receives a request to delete the customer. Such a value causes the connector to issue the following SQL statement:

```
UPDATE CUSTOMER SET CUSTSTATUS = 'DELETED' WHERE CUSTOMER_ID = 2345
```

If the SCN parameter is not included or no value is specified for it, the connector physically deletes the business object from the application database. In other words, if the business object with the Delete verb includes the SCN parameter in its application-specific information, the connector performs a logical delete. If the business object with the Delete verb does not include the SCN parameter in its application-specific information, the connector performs a physical delete.

Both update and delete operations may use the value of the SCN property:

- When performing an update, the connector uses the value of its `ChildUpdatePhyDelete` property to determine whether to physically or logically delete missing child data. If logically deleting the child data, it uses the value of its SCN parameter to obtain the name of the status column and the text of the status value. For more information, see “Update operations” on page 73.
- When performing a delete, the connector uses the value of its SCN parameter to determine whether to physically or logically delete the entire business object. If the SCN parameter contains a value, the connector performs a logical delete. If the SCN parameter does not contain a value, the connector performs a physical delete. For more information, see “Delete operations” on page 75.

At the business-object level, application-specific information may be used to specify a wrapper:

```
WRAPPER=true|false
```

If the wrapper parameter is set to true, the top-level business object is a wrapper object. The wrapper object is not represented by a database table or view. A wrapper is used as a container for unrelated business objects. The connector ignores the top-level object and processes only the children. The wrapper object may contain N cardinality or N-1 cardinality entities or both.

Application-specific information at the attribute level

The application-specific information for attributes differs depending on whether the attribute is a simple attribute or an attribute that represents a child or an array of child business objects. The application-specific information for an attribute that represents a child also differs depending on whether the parent/child relationship is stored in the child or in the parent. For information on application-specific information for attributes that represent a child or array of child business objects, see “Specifying an attribute’s foreign key” on page 89.

Application-specific information for simple attributes

For simple attributes, application-specific information format consists of eleven name-value parameters, each of which includes the parameter name and its value. Each parameter set is separated from the next by a colon (:) delimiter.

The format of attribute application-specific information is shown below. Square brackets ([]) surround an optional parameter. A vertical bar (|) separates the members of a set of options. Reserve the colon as a delimiter.

```
CN=col_name:[FK=[fk_object_name.]fk_attribute_name]:  
[UID=[AUTO|uid_name] schema_name.uid_name [=UseIfMissing]|CW.uidcolumnname  
[=UseIfMissing]]]:  
[CA=set_attr_name|.set_attr_name]:[OB=[ASC|DESC]]:[UNVL=value]:  
[ESC=true|false]:[FIXEDCHAR=true|false]:  
[BYTEARRAY=true|false]:[USE_LIKE=true|false]:  
[WILDCARD_POSITION=non-negative number|NONE|BEGIN|END|BOTH]]:  
[CLOB=true]
```

The only required parameter for a simple attribute that you want the connector to process is the column name. For example, to specify only the column name, use the following format:

CN=customer_id

Table 28 describes each name-value parameter.

Table 28. Name-value parameters in attribute application-specific information

Parameter	Description
CN= <i>col_name</i>	The name of the application database column for this attribute.
FK=[<i>fk_object_name.</i> <i>fk_attribute_name</i>]	The value of this property depends on whether the parent/child relationship is stored in the parent business object or the child. If an attribute is not a foreign key, do not include this parameter in the application-specific information. For more information, see “Specifying an attribute’s foreign key” on page 89.
UID=AUTO	The connector uses this parameter to generate the unique ID for the business object. If an attribute does not require generation of a unique ID, do not include this parameter in the application-specific information. See the “PreserveUIDSeq” on page 55 property description for details on preserving the unique ID during business object processing. For more information, see “Generating a business object’s unique identifier” on page 92.
UID= <i>uid_name</i> <i>schema_name.uid_name</i> [=UseIfMissing]	
UID=CW. <i>uidcolumnname</i> [=UseIfMissing]	
CA= <i>set_attr_name</i> .. <i>set_attr_name</i>	<p>If <i>set_attr_name</i> is set to the name of another attribute within the current individual business object, the connector uses the value of the specified attribute to set the value of this attribute before it adds the business object to the application database during a Create operation.</p> <p>The value of <i>set_attr_name</i> cannot reference an attribute in a parent or child business object, but it can reference an attribute in the parent business object if there if <i>set_attr_name</i> is preceded by the two periods.</p> <p>If you do not include this parameter in the application-specific information, the connector uses the value of the current attribute without copying the attribute’s value (CA) from another attribute.</p>
OB=[ASC DESC]	<p>If a value is specified for this parameter and the attribute is in a child business object, the connector uses the value of the attribute in the ORDER BY clause of retrieval queries.</p> <p>The connector can retrieve child business objects in ascending order or descending order.</p> <p>Use ASC to specify retrieval in ascending order.</p> <p>Use DESC to specify retrieval in descending order.</p> <p>If you do not include this parameter in the application-specific information, the connector does not use this attribute when specifying retrieval order.</p>
UNVL= <i>value</i>	Specifies the value the connector uses to represent a null when it retrieves a business object with null-valued attributes. If you do not include this parameter in the application-specific information, the connector inserts a CxIgnore for the attribute’s value.
ESC=[true false]	Determines whether the connector replaces all instances of each character identified in the ReplaceAllStr property with the substitution strings also specified in the ReplaceStrList property. If this parameter does not contain a value, the connector uses the value of the ReplaceStrList property to make this determination.

Table 28. Name-value parameters in attribute application-specific information (continued)

Parameter	Description
FIXEDCHAR=true false	Specifies whether the attribute is of fixed length when the columns in the table are of type CHAR, not VARCHAR. For example, if a particular attribute is linked to a column that is of type CHAR, the connector expects FIXEDCHAR in length; for the application specific information of that attribute specify FIXEDCHAR=true. Ensure that the MaxLength property of the attribute is of the CHAR length, which is specified in the application database. By default, FIXEDCHAR=false.
BYTEARRAY=true false	If BYTEARRAY=true, the connector will read and write binary data to the database and will send that data as a string to ICS/WebSphere MQ Integration Broker. BYTEARRAY=false is the default.
USE_LIKE=true false	Specifies whether the connector compares strings using the = operator or the LIKE operator. If USE_LIKE is set to true, wildcard queries can be performed by setting WILDCARD_POSITION. If USE_LIKE is set to false, the =operator will be used.
WILDCARD_POSITION=non-negative number NONE BEGIN END BOTH	If USE_LIKE is true, the WILDCARD_POSITION is used to specify the position that is the wildcard. This value can be any non-negative number, NONE, BEGIN, END, or BOTH. For example, using BEGIN will place the wildcard character in the first position of the string (%string). Using END will place the wildcard character in the last position of the string (string%). Using BOTH will place wildcard characters in both the first and last position in the string (%string%).
CLOB=true	<p>Only applicable for String Attribute Type. Specifies that the database column that corresponds to this attribute is a CLOBdatatype.</p> <p>Note: ACLOB datatype is defined as follows:</p> <ul style="list-style-type: none"> • The CLOB attribute has a String Type whose length is used to define the length of the CLOB • The CLOB attribute has ASI=CN=xyz; CLOB=true • Any other attribute type with reference to CLOB in the ASI would result in an error • CLOB=false would result in an error <p>A regular String Type would be the same and with no reference to CLOB in the ASI. CLOB datatypes of 4k and larger can be inserted or updated. But they can be used only with Oracle and require the latest thin driver with CLOB support. Using any other driver may cause errors.</p>

Note: If none of the application-specific information in any of a business object's attributes cause the connector to build or execute a query, the connector logs a warning and continues operating. It does not throw an exception or return a failure.

Specifying an attribute's foreign key: The value of this property depends on whether the parent/child relationship is stored in the parent business object or the child:

- Stored in the parent—set the value to include both the type of the child business object and the name of the attribute in the child to be used as the foreign key.
- 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 the value of *fk_object_name* does not match the type of the child business object, and the value of *fk_attribute_name* does not match the name of the attribute in the parent or child (as applicable), the connector cannot process this attribute as a foreign key. The case of the business object's name and the attribute's name is significant.

For example, assume that the Customer business object contains the Addr[1] attribute, which represents the Address child business object, and the AID attribute, which stores the primary key of the child business object as a foreign key. In this case, the application-specific information of the parent's foreign key attribute must contain the type of the child business object (Address) as well as the name of its primary key attribute (ID). In this example, the application-specific information of the AID attribute would include `FK=Address.ID`.

Naming a foreign key attribute: Multiple parent business objects can contain the same child business object, regardless of whether the child is stored with single cardinality or multiple cardinality, and regardless of whether the parent/child relationship is stored on the parent or on the child. However, all parent business objects that store the parent/child relationship must use identically named attributes to contain the child's primary key. Moreover, all child business objects that store the parent/child relationship must use identically named attributes to contain the parent's primary key. Figure 20 illustrates these relationships.

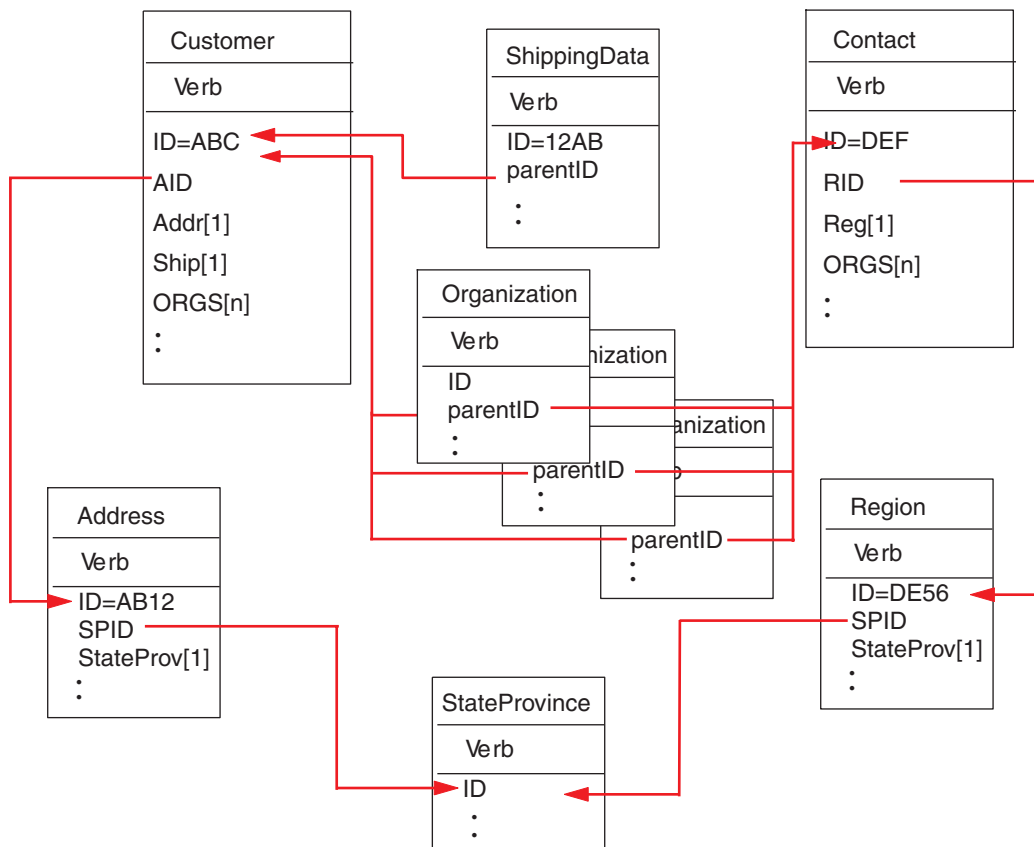


Figure 20. Example of relationships among business objects

Figure 20 illustrates the following relationships:

- The ORGS[n] attribute of Customer ABC and Contact DEF represents an array of Organization business objects. The foreign key value for each business object in the array of Organizations corresponds to the primary key value in the ID attribute in the Customer and Contact business objects. In this case, each business object in the array is contained by multiple parents.

The application-specific information for the ORGS attribute might be:

`KEEP_RELATIONSHIP=true`

For more information on the KEEP_RELATIONSHIP parameter, see “Application-specific information for attributes that represent children” on page 92.

The application-specific information for the parentID attribute of each child in the array of Organizations contains the name of the column in the application database that corresponds to the current attribute, and specifies the current attribute’s foreign key by containing the name of the parent’s primary key attribute; for example:

```
CN=ORG_ID:FK=ID
```

Note: For multiple business objects to contain the same child (where the parent/child relationship is stored in the child), all parent business objects must use an identically named attribute to contain the foreign key for the child. The foreign key parameter of that child’s application-specific information identifies only the attribute’s name and not the type of the parent business object. The connector assumes that the direct parent is the owner of each child.

- The Addr[1] attribute of Customer represents the Address business object with ownership. The AID attribute of Customer identifies the primary key of the Address business object as a foreign key in the parent. In this case, the parent’s foreign key attribute must contain the type of the child business object as well as the name of its primary key attribute. The single-cardinality child, Address, is contained by only one parent.

The application-specific information for the Addr attribute is:

```
CONTAINMENT=OWNERSHIP
```

The application-specific information for the AID attribute contains the name of the column in the application database that corresponds to the current attribute, and specifies the current attribute’s foreign key by containing the type of the child business object and the name of its primary key attribute; for example:

```
CN=FK_AD:FK=Address.ID
```

The application-specific information for the child’s primary-key attribute is

```
CN=pk
```

- The StateProv[1] attributes of the Address and Region business objects represent the StateProvince business object without ownership. The SPID attributes of the Address and Region business objects contain the type of the child business object (StateProvince) and the name of its primary key attribute, which serve as the parent’s foreign key. The same single-cardinality child, StateProvince, is contained by multiple parents.

The application-specific information for the SPID attribute is:

```
CONTAINMENT=NOOWNERSHIP
```

For more information on the CONTAINMENT parameter, see “Application-specific information for attributes that represent children” on page 92.

The application-specific information for the Address SPID attribute contains the name of the column in the database that corresponds to the current attribute, and specifies the current attribute’s foreign key by containing the type of the child business object and the name of its primary key attribute; for example:

```
CN=FK_SP:FK=StateProvince.ID
```

The application-specific information for the child’s primary key attribute is:

```
CN=SP_ID
```

Note: For multiple business objects (that store the parent/child relationship in the parent) to contain the same child, all child business objects must use an identically named attribute to contain the foreign key for the parent.

- The Ship[1] attribute of Customer represents a ShippingData business object that contains the customer's shipping information. The ID attribute of Customer functions as the foreign key for the shipping data. In this case, because ShippingData cannot exist independently of its parent and is created only after its parent is created, the parent/child relationship is stored in the child.

The application-specific information for the child's parentID attribute contains the name of the column in the database that corresponds to the current attribute, and specifies the current attribute's foreign key by containing the name of its parent's primary key attribute; for example:

```
CN=SD_ID:FK=ID
```

Generating a business object's unique identifier: The connector uses the UID parameter to generate the unique ID for the business object. The connector generates unique IDs by using sequences, or counters (which are structured as tables), and then issues the INSERT statement.

The connector uses a sequence or counter to generate the ID value and then issues the INSERT statement:

- If UID=CW.*uidcolumnname*, the connector uses a WebSphere counter table to generate a unique ID for the attribute. The table, whose name is configurable, is created with a single column named id. You can customize the table to add a column for each attribute that requires generation of a UID. Use the *uidcolumnname* parameter to specify the name of the column for the connector to use when generating the unique ID. Note that the connector supports only the numeric data type for columns that require generation of a UID.

If the WebSphere counter table is used to generate a UID for one attribute only, then create a table with one column of type integer (choose database-appropriate data types). Configure the connector property *UniqueIdTableName* with an appropriate table name.

If the WebSphere counter table is used to generate UIDs for more than one attribute, then create a table with as many columns as attributes of type integer only.

- If UID=CW.*uidcolumnname*=UseIfMissing and if the value of the attribute is not CxIgnore, the connector uses the attribute's value rather than generating a unique ID. The =UseIfMissing parameter cannot contain blanks and is case-insensitive.

See the "PreserveUIDSeq" on page 55 property for information on preserving the unique ID sequence during processing.

Application-specific information for attributes that represent children

Attributes that represent a single-cardinality child business object can specify whether the child is owned by the parent or shared among multiple parents.

Attributes that represent a single-cardinality child or an array of child business objects can specify the connector's behavior when updating the parent and a subset of the children.

Attributes that represent a single-cardinality child business Object: The format of the application-specific information for attributes that represent a single-cardinality child business object is:

```
CONTAINMENT= [OWNERSHIP|NO_OWNERSHIP]
```

Set CONTAINMENT to OWNERSHIP to represent a single-cardinality relationship where the parent owns the child business object. Set CONTAINMENT to NO_OWNERSHIP to represent a single-cardinality relationship where the parent shares the child business object. Do not include the CONTAINMENT parameter when you represent a single-cardinality relationship that stores the relationship in the child rather than in the parent.

For more information, see “Single-cardinality relationships and data without ownership” on page 65 and “Wrapper objects” on page 68.

Attributes that represent a child that stores the parent’s Key: For Update operations on an array of business objects that store the parent/child relationship in the child, there is a special value for the attribute that represents the child: you can set KEEP_RELATIONSHIP to true to prevent the connector from deleting existing child data that is not represented in the source business object.

For example, assume an existing contract is associated with an existing site, such as New York. Assume further that the connector receives a request to update a Contract business object that contains a single child business object that associates San Francisco as the site. If KEEP_RELATIONSHIP evaluates to true for the attribute that represents the site data, the connector updates the contract to add its association with San Francisco and does not delete its association with New York.

However, if KEEP_RELATIONSHIP evaluates to false, the connector deletes all existing child data that is not contained in the source business object. In such a case, the contract is associated only with San Francisco.

The format for this application-specific information is:

```
KEEP_RELATIONSHIP=[true|false]
```

Case is ignored in checking for this application-specific information.

Working with Binary Data: If BYTEARRAY=true, the connector will read and write binary data to the database. Since there is no support for binary data in the current version of the WebSphere business integration system framework, the binary data is converted to a String and then sent to the integration broker. The format of this string is a hexadecimal number with 2 characters per byte. For example, if the binary data in the database is 3 bytes with the (decimal) values (1, 65, 255), the string will be "0141ff".

Application-specific information format for verbs

The connector uses verb application-specific information for the Retrieve and RetrieveByContent verbs. This text allows you to specify the attributes to be included in the WHERE clause for a retrieval. You can also specify operators and attribute values.

The syntax for application-specific information for the Retrieve and RetrieveByContent verbs is shown below:

```
[condition_variable conditional_operator @ [...]:[.]attribute_name [, ...]]
```

where:

<i>condition_variable</i>	The name of the application database column.
<i>conditonal_operator</i>	The operator supported by the application database, for example =, >, OR, AND, and IN (<i>value1</i> , <i>value2</i>).

@	A variable that is substituted with the value retrieved by <code>getAttrValue(attribute_name)</code> . The substitution is positional; that is, the connector substitutes the first @ with the value of the first <code>attribute_name</code> variable specified after the <code>:</code> delimiter.
..	The attribute specified in the <code>attribute_name</code> variable belongs to the immediate parent business object; if this value is missing, the attribute belongs in the current business object.
attribute_name	The name of the attribute whose value the connector substitutes for @.

To understand the syntax of this property, assume that an Item business object has an `item_id` attribute whose value is XY45 and a `Color` attribute whose value is RED. Assume further that you specify the Retrieve verb's `AppSpecificInfo` property as:

```
Color='RED'
```

The above application-specific information value causes the connector to build the following WHERE clause for a retrieval:

```
where item_id=XY45 and Color = 'RED'
```

For a more complicated example, assume that the Customer business object has a `customer_id` attribute whose value is 1234 and a `creation_date` attribute whose value is 01/01/90. Assume also that this business object's parent has a `quantity` attribute whose value is 20.

Assume further that you specify the Retrieve verb's `AppSpecificInfo` property as:

```
creation_date > @ OR quantity = @ AND customer_status  
IN ('GOLD', 'PLATINUM') : creation_date, ..quantity
```

The above application-specific information value causes the connector to build the following WHERE clause for a retrieval:

```
where customer_id=1234 and creation_date > '01/01/90'  
OR quantity = 20 AND customer_status IN ('GOLD', 'PLATINUM')
```

The connector gets the date value ('01/01/90') from the `creation_date` attribute in the current business object. It gets the quantity value (20) from the `quantity` attribute in the parent business object (as indicated by `..quantity` in the application-specific information).

After the connector parses the application-specific information for the Retrieve verb, it adds the text to the WHERE clause of the RETRIEVE statement that it constructs from the business object's primary or foreign keys. The connector adds the leading AND to the WHERE clause. The value of the application-specific information must be valid SQL syntax. In the case of `RetrieveByContent`, the application-specific information is added to the WHERE clause of the RETRIEVE statement that it constructs from the business object's attributes that have their values populated.

The WHERE clause can also refer to placeholder attributes instead of the actual attributes in the parent business object. These placeholders do not have any application-specific information. An attribute can be a placeholder if it satisfies one of the following conditions for its ASI:

1. Simple attribute with `ASI=null` or ''

2. Simple attribute with ASI=PH=TRUE

For example: An Order business object contains a multiple cardinality line item business object, and retrieval of only specific line items is needed. This retrieval can be handled through a placeholder attribute in the Order business object. This placeholder is required in the parent object because the child objects are all pruned. When ICS is the integration broker, the placeholder attribute can be populated at runtime by a map with a list of the specific line items, separated by a comma (,).

For this example, you would add the following information to the WHERE clause for the retrieve verb on the child line item business object:

```
line_item_id in(@)..placeholder
```

Where `line_item_id` is the ID in the child business object, `placeholder` is the attribute in the parent. If `placeholder` contains the values 12,13,14 the query would select the following from the WHERE clause:

```
line_item_id in(12,13,14)
```

Where `SELECT:..FROM:..WHERE x in (1,2,3)` is a standard database SQL syntax.

In the `RetrieveByContent` verb, if the length of the WHERE clause is 0, the connector will use the application-specific information in the WHERE clause of the RETRIEVE statement. With this feature, the user can send a business object with no attribute values populated and specify verb application-specific information for `RetrieveByContent`, and the connector will build the WHERE clause based on what was specified in the verb application-specific information alone.

Chapter 6. Troubleshooting and error handling

The chapter describes problems that you may encounter when starting up or running an MAS connector, and contains the following sections:

- “Startup problems”
- “Event processing”
- “Mapping (ICS integration broker only)”
- “Error handling and logging”

Startup problems

If you encounter difficulties when trying to start the connector, check to make sure that InterChange Server is up and running.

Event processing

If there are events in the event table, and they are not being processed while the connector is running, ensure that:

- The relevant business object request is running.
- The name of the business object in the event table matches the name of the business object specified for the relevant port.

Mapping (ICS integration broker only)

If the business objects are not being mapped or mapping is not being invoked, check to make sure the maps have been installed in the correct directory.

Error handling and logging

The connector logs an error message whenever it encounters a condition that causes its current processing of a business object and verb to fail. When such an error occurs, the connector also prints a textual representation of the failed business object as it was received. It writes the text to the connector log file or the standard output stream, depending on its configuration. You can use the text as an aid in determining the source of the error.

Error types

Table 29 describes the types of tracing messages that the connector outputs at each trace level. These messages are in addition to any tracing messages output by the WebSphere architecture.

Table 29. Connector tracing messages

Tracing level	Tracing messages
Level 0	Message that identifies the connector version. No other tracing is done at this level. This is the default value.

Table 29. Connector tracing messages (continued)

Tracing level	Tracing messages
Level 1	<ul style="list-style-type: none"> • Status messages • Messages that provide identifying (key) information for each business object processed • Messages delivered each time the pollForEvents method is executed
Level 2	<ul style="list-style-type: none"> • Business object handler messages that contain information such as the arrays and child business objects that the connector encounters or retrieves during the processing of a business object • Messages logged each time a business object is posted to InterChange Server, either from gotAppIEvent() or executeCollaboration() • Messages that indicate that a business object has been received as a request
Level 3	<ul style="list-style-type: none"> • Foreign key processing messages that contain such information as when the connector has found or has set a foreign key in a business object • Messages that provide information about business object processing. For example, these messages are delivered when the connector finds a match between business objects, or finds a business object in an array of child business objects
Level 4	<ul style="list-style-type: none"> • Application-specific text messages, for example, messages showing the values returned by the functions that parse the business object's application-specific information fields • Messages that identify when the connector enters or exits a function, which helps trace the process flow of the connector • All thread-specific messages. If the connector spawns multiple threads, a message appears for the creation of each new thread
Level 5	<ul style="list-style-type: none"> • Messages that indicate connector initialization, for example, messages showing the value of each configuration property retrieved from the integration broker. • Messages that include statements executed in the application. At this trace level, the connector log file contains all statements executed in the destination application and the value of any variables that are substituted. • Messages that comprise a representation of a business object before the connector begins processing it (displaying its state as the connector receives it) and after the connector has completed its processing (displaying its state as the connector returns it) • Messages that comprise a business object dump • Messages that indicate the status of each thread the connector spawns while it is running

Appendix A. Standard configuration properties for connectors

This appendix describes the standard configuration properties for the connector component of WebSphere Business Integration adapters. The information covers connectors running on the following integration brokers:

- WebSphere InterChange Server (ICS)
- WebSphere MQ Integrator, WebSphere MQ Integrator Broker, and WebSphere Business Integration Message Broker, collectively referred to as the WebSphere Message Brokers (WMQI).
- WebSphere Application Server (WAS)

Not every connector makes use of all these standard properties. When you select an integration broker from Connector Configurator, you will see a list of the standard properties that you need to configure for your adapter running with that broker.

For information about properties specific to the connector, see the relevant adapter user guide.

Note: In this document, backslashes (\) are used as the convention for directory paths. For UNIX installations, substitute slashes (/) for backslashes and follow the conventions for each operating system.

New and deleted properties

These standard properties have been added in this release.

New properties

- XMLNamespaceFormat

Deleted properties

- RestartCount

Configuring standard connector properties

Adapter connectors have two types of configuration properties:

- Standard configuration properties
- Connector-specific configuration properties

This section describes the standard configuration properties. For information on configuration properties specific to a connector, see its adapter user guide.

Using Connector Configurator

You configure connector properties from Connector Configurator, which you access from System Manager. For more information on using Connector Configurator, refer to the Connector Configurator appendix.

Note: Connector Configurator and System Manager run only on the Windows system. If you are running the connector on a UNIX system, you must have a Windows machine with these tools installed. To set connector properties

for a connector that runs on UNIX, you must start up System Manager on the Windows machine, connect to the UNIX integration broker, and bring up Connector Configurator for the connector.

Setting and updating property values

The default length of a property field is 255 characters.

The connector uses the following order to determine a property's value (where the highest number overrides other values):

1. Default
2. Repository (only if WebSphere InterChange Server is the integration broker)
3. Local configuration file
4. Command line

A connector obtains its configuration values at startup. If you change the value of one or more connector properties during a run-time session, the property's **Update Method** determines how the change takes effect. There are four different update methods for standard connector properties:

- **Dynamic**
The change takes effect immediately after it is saved in System Manager. If the connector is working in stand-alone mode (independently of System Manager), for example with one of the WebSphere message brokers, you can only change properties through the configuration file. In this case, a dynamic update is not possible.
- **Component restart**
The change takes effect only after the connector is stopped and then restarted in System Manager. You do not need to stop and restart the application-specific component or the integration broker.
- **Server restart**
The change takes effect only after you stop and restart the application-specific component and the integration broker.
- **Agent restart (ICS only)**
The change takes effect only after you stop and restart the application-specific component.

To determine how a specific property is updated, refer to the **Update Method** column in the Connector Configurator window, or see the Update Method column in the Property Summary table below.

Summary of standard properties

Table 30 on page 101 provides a quick reference to the standard connector configuration properties. Not all the connectors make use of all these properties, and property settings may differ from integration broker to integration broker, as standard property dependencies are based on RepositoryDirectory.

You must set the values of some of these properties before running the connector. See the following section for an explanation of each property.

Table 30. Summary of standard configuration properties

Property name	Possible values	Default value	Update method	Notes
AdminInQueue	Valid JMS queue name	<i>CONNECTORNAME</i> /ADMININQUEUE	Component restart	Delivery Transport is JMS
AdminOutQueue	Valid JMS queue name	<i>CONNECTORNAME</i> /ADMINOUTQUEUE	Component restart	Delivery Transport is JMS
AgentConnections	1-4	1	Component restart	Delivery Transport is MQ or IDL: Repository directory is <REMOTE>
AgentTraceLevel	0-5	0	Dynamic	
ApplicationName	Application name	Value specified for the connector application name	Component restart	
BrokerType	ICS, WMQI, WAS			
CharacterEncoding	ascii7, ascii8, SJIS, Cp949, GBK, Big5, Cp297, Cp273, Cp280, Cp284, Cp037, Cp437 Note: This is a subset of supported values.	ascii7	Component restart	
ConcurrentEventTriggeredFlows	1 to 32,767	1	Component restart	Repository directory is <REMOTE>
ContainerManagedEvents	No value or JMS	No value	Component restart	Delivery Transport is JMS
ControllerStoreAndForwardMode	true or false	True	Dynamic	Repository directory is <REMOTE>
ControllerTraceLevel	0-5	0	Dynamic	Repository directory is <REMOTE>
DeliveryQueue		<i>CONNECTORNAME</i> /DELIVERYQUEUE	Component restart	JMS transport only
DeliveryTransport	MQ, IDL, or JMS	JMS	Component restart	If Repository directory is local, then value is JMS only
DuplicateEventElimination	True or False	False	Component restart	JMS transport only: Container Managed Events must be <NONE>
FaultQueue		<i>CONNECTORNAME</i> /FAULTQUEUE	Component restart	JMS transport only

Table 30. Summary of standard configuration properties (continued)

Property name	Possible values	Default value	Update method	Notes
jms.FactoryClassName	CxCommon.Messaging.jms.IBMMQSeriesFactory or CxCommon.Messaging.jms.SonicMQFactory or any Java class name	CxCommon.Messaging.jms.IBMMQSeriesFactory	Component restart	JMS transport only
jms.MessageBrokerName	If FactoryClassName is IBM, use crossworlds.queue.manager. If FactoryClassName is Sonic, use localhost:2506.	crossworlds.queue.manager	Component restart	JMS transport only
jms.NumConcurrentRequests	Positive integer	10	Component restart	JMS transport only
jms.Password	Any valid password		Component restart	JMS transport only
jms.UserName	Any valid name		Component restart	JMS transport only
JvmMaxHeapSize	Heap size in megabytes	128m	Component restart	Repository directory is <REMOTE>
JvmMaxNativeStackSize	Size of stack in kilobytes	128k	Component restart	Repository directory is <REMOTE>
JvmMinHeapSize	Heap size in megabytes	1m	Component restart	Repository directory is <REMOTE>
ListenerConcurrency	1- 100	1	Component restart	Delivery Transport must be MQ
Locale	en_US, ja_JP, ko_KR, zh_CN, zh_TW, fr_FR, de_DE, it_IT, es_ES, pt_BR Note: This is a subset of the supported locales.	en_US	Component restart	
LogAtInterchangeEnd	True or False	False	Component restart	Repository Directory must be <REMOTE>
MaxEventCapacity	1-2147483647	2147483647	Dynamic	Repository Directory must be <REMOTE>
MessageFileName	Path or filename	InterchangeSystem.txt	Component restart	
MonitorQueue	Any valid queue name	CONNECTORNAME/MONITORQUEUE	Component restart	JMS transport only: DuplicateEvent Elimination must be True
OADAutoRestartAgent	True or False	False	Dynamic	Repository Directory must be <REMOTE>

Table 30. Summary of standard configuration properties (continued)

Property name	Possible values	Default value	Update method	Notes
OADMaxNumRetry	A positive number	1000	Dynamic	Repository Directory must be <REMOTE>
OADRetryTimeInterval	A positive number in minutes	10	Dynamic	Repository Directory must be <REMOTE>
PollEndTime	HH:MM	HH:MM	Component restart	
PollFrequency	A positive integer in milliseconds no (to disable polling) key (to poll only when the letter p is entered in the connector's Command Prompt window)	10000	Dynamic	
PollQuantity	1-500	1	Agent restart	JMS transport only: Container Managed Events is specified
PollStartTime	HH:MM(HH is 0-23, MM is 0-59)	HH:MM	Component restart	
RepositoryDirectory	Location of metadata repository		Agent restart	For ICS: set to <REMOTE> For WebSphere MQ message brokers and WAS: set to C:\crossworlds\repository
RequestQueue	Valid JMS queue name	CONNECTORNAME/REQUESTQUEUE	Component restart	Delivery Transport is JMS
ResponseQueue	Valid JMS queue name	CONNECTORNAME/RESPONSEQUEUE	Component restart	Delivery Transport is JMS: required only if Repository directory is <REMOTE>
RestartRetryCount	0-99	3	Dynamic	
RestartRetryInterval	A sensible positive value in minutes: 1 - 2147483547	1	Dynamic	
RHF2MessageDomain	mrm, xml	mrm	Component restart	Only if Delivery Transport is JMS and WireFormat is CwXML.

Table 30. Summary of standard configuration properties (continued)

Property name	Possible values	Default value	Update method	Notes
SourceQueue	Valid WebSphere MQ name	CONNECTORNAME/SOURCEQUEUE	Agent restart	Only if Delivery Transport is JMS and Container Managed Events is specified
SynchronousRequestQueue		CONNECTORNAME/ SYNCHRONOUSREQUESTQUEUE	Component restart	Delivery Transport is JMS
SynchronousRequestTimeout	0 - any number (milliseconds)	0	Component restart	Delivery Transport is JMS
SynchronousResponseQueue		CONNECTORNAME/ SYNCHRONOUSRESPONSEQUEUE	Component restart	Delivery Transport is JMS
WireFormat	CwXML, CwBO	CwXML	Agent restart	CwXML if Repository Directory is not <REMOTE>; CwBO if Repository Directory is <REMOTE>
WsifSynchronousRequest Timeout	0 - any number (milliseconds)	0	Component restart	WAS only
XMLNamespaceFormat	short, long	short	Agent restart	WebSphere MQ message brokers and WAS only

Standard configuration properties

This section lists and defines each of the standard connector configuration properties.

AdminInQueue

The queue that is used by the integration broker to send administrative messages to the connector.

The default value is CONNECTORNAME/ADMININQUEUE.

AdminOutQueue

The queue that is used by the connector to send administrative messages to the integration broker.

The default value is CONNECTORNAME/ADMINOUTQUEUE.

AgentConnections

Applicable only if RepositoryDirectory is <REMOTE>.

The AgentConnections property controls the number of ORB connections opened by `orb.init[]`.

By default, the value of this property is set to 1. There is no need to change this default.

AgentTraceLevel

Level of trace messages for the application-specific component. The default is 0. The connector delivers all trace messages applicable at the tracing level set or lower.

ApplicationName

Name that uniquely identifies the connector's application. This name is used by the system administrator to monitor the WebSphere business integration system environment. This property must have a value before you can run the connector.

BrokerType

Identifies the integration broker type that you are using. The options are ICS, WebSphere message brokers (WMQI, WMQIB or WBIMB) or WAS.

CharacterEncoding

Specifies the character code set used to map from a character (such as a letter of the alphabet, a numeric representation, or a punctuation mark) to a numeric value.

Note: Java-based connectors do not use this property. A C++ connector currently uses the value `ascii7` for this property.

By default, a subset of supported character encodings only is displayed in the drop list. To add other supported values to the drop list, you must manually modify the `\Data\Std\stdConnProps.xml` file in the product directory. For more information, see the appendix on Connector Configurator.

ConcurrentEventTriggeredFlows

Applicable only if `RepositoryDirectory` is `<REMOTE>`.

Determines how many business objects can be concurrently processed by the connector for event delivery. Set the value of this attribute to the number of business objects you want concurrently mapped and delivered. For example, set the value of this property to 5 to cause five business objects to be concurrently processed. The default value is 1.

Setting this property to a value greater than 1 allows a connector for a source application to map multiple event business objects at the same time and deliver them to multiple collaboration instances simultaneously. This speeds delivery of business objects to the integration broker, particularly if the business objects use complex maps. Increasing the arrival rate of business objects to collaborations can improve overall performance in the system.

To implement concurrent processing for an entire flow (from a source application to a destination application), you must:

- Configure the collaboration to use multiple threads by setting its `Maximum number of concurrent events` property high enough to use multiple threads.
- Ensure that the destination application's application-specific component can process requests concurrently. That is, it must be multi-threaded, or be able to use connector agent parallelism and be configured for multiple processes. Set the `Parallel Process Degree` configuration property to a value greater than 1.

The `ConcurrentEventTriggeredFlows` property has no effect on connector polling, which is single-threaded and performed serially.

ContainerManagedEvents

This property allows a JMS-enabled connector with a JMS event store to provide guaranteed event delivery, in which an event is removed from the source queue and placed on the destination queue as a single JMS transaction.

The default value is `No value`.

When `ContainerManagedEvents` is set to `JMS`, you must configure the following properties to enable guaranteed event delivery:

- `PollQuantity` = 1 to 500
- `SourceQueue` = `CONNECTORNAME/SOURCEQUEUE`

You must also configure a data handler with the `MimeType`, `DHClass`, and `DataHandlerConfigMOName` (optional) properties. To set those values, use the **Data Handler** tab in Connector Configurator. The fields for the values under the Data Handler tab will be displayed only if you have set `ContainerManagedEvents` to `JMS`.

Note: When `ContainerManagedEvents` is set to `JMS`, the connector does *not* call its `pollForEvents()` method, thereby disabling that method's functionality.

This property only appears if the `DeliveryTransport` property is set to the value `JMS`.

ControllerStoreAndForwardMode

Applicable only if `RepositoryDirectory` is `<REMOTE>`.

Sets the behavior of the connector controller after it detects that the destination application-specific component is unavailable.

If this property is set to `true` and the destination application-specific component is unavailable when an event reaches ICS, the connector controller blocks the request to the application-specific component. When the application-specific component becomes operational, the controller forwards the request to it.

However, if the destination application's application-specific component becomes unavailable **after** the connector controller forwards a service call request to it, the connector controller fails the request.

If this property is set to `false`, the connector controller begins failing all service call requests as soon as it detects that the destination application-specific component is unavailable.

The default is `true`.

ControllerTraceLevel

Applicable only if `RepositoryDirectory` is `<REMOTE>`.

Level of trace messages for the connector controller. The default is `0`.

DeliveryQueue

Applicable only if DeliveryTransport is JMS.

The queue that is used by the connector to send business objects to the integration broker.

The default value is CONNECTORNAME/DELIVERYQUEUE.

DeliveryTransport

Specifies the transport mechanism for the delivery of events. Possible values are MQ for WebSphere MQ, IDL for CORBA IIOP, or JMS for Java Messaging Service.

- If ICS is the broker type, the value of the DeliveryTransport property can be MQ, IDL, or JMS, and the default is IDL.
- If the RepositoryDirectory is a local directory, the value may only be JMS.

The connector sends service call requests and administrative messages over CORBA IIOP if the value configured for the DeliveryTransport property is MQ or IDL.

WebSphere MQ and IDL

Use WebSphere MQ rather than IDL for event delivery transport, unless you must have only one product. WebSphere MQ offers the following advantages over IDL:

- Asynchronous communication:
WebSphere MQ allows the application-specific component to poll and persistently store events even when the server is not available.
- Server side performance:
WebSphere MQ provides faster performance on the server side. In optimized mode, WebSphere MQ stores only the pointer to an event in the repository database, while the actual event remains in the WebSphere MQ queue. This saves having to write potentially large events to the repository database.
- Agent side performance:
WebSphere MQ provides faster performance on the application-specific component side. Using WebSphere MQ, the connector's polling thread picks up an event, places it in the connector's queue, then picks up the next event. This is faster than IDL, which requires the connector's polling thread to pick up an event, go over the network into the server process, store the event persistently in the repository database, then pick up the next event.

JMS

Enables communication between the connector and client connector framework using Java Messaging Service (JMS).

If you select JMS as the delivery transport, additional JMS properties such as `jms.MessageBrokerName`, `jms.FactoryClassName`, `jms.Password`, and `jms.UserName`, appear in Connector Configurator. The first two of these properties are required for this transport.

Important: There may be a memory limitation if you use the JMS transport mechanism for a connector in the following environment:

- AIX 5.0
- WebSphere MQ 5.3.0.1
- When ICS is the integration broker

In this environment, you may experience difficulty starting both the connector controller (on the server side) and the connector (on the client side) due to memory use within the WebSphere MQ client. If your installation uses less than 768M of process heap size, IBM recommends that you set:

- The `LDR_CNTRL` environment variable in the `CWSharedEnv.sh` script.
This script resides in the `\bin` directory below the product directory. With a text editor, add the following line as the first line in the `CWSharedEnv.sh` script:

```
export LDR_CNTRL=MAXDATA=0x30000000
```


This line restricts heap memory usage to a maximum of 768 MB (3 segments * 256 MB). If the process memory grows more than this limit, page swapping can occur, which can adversely affect the performance of your system.
- The `IPCCBaseAddress` property to a value of 11 or 12. For more information on this property, see the *System Installation Guide for UNIX*.

DuplicateEventElimination

When you set this property to true, a JMS-enabled connector can ensure that duplicate events are not delivered to the delivery queue. To use this feature, the connector must have a unique event identifier set as the business object's **ObjectEventId** attribute in the application-specific code. This is done during connector development.

This property can also be set to false.

Note: When `DuplicateEventElimination` is set to true, you must also configure the `MonitorQueue` property to enable guaranteed event delivery.

FaultQueue

If the connector experiences an error while processing a message then the connector moves the message to the queue specified in this property, along with a status indicator and a description of the problem.

The default value is `CONNECTORNAME/FAULTQUEUE`.

JvmMaxHeapSize

The maximum heap size for the agent (in megabytes). This property is applicable only if the `RepositoryDirectory` value is `<REMOTE>`.

The default value is 128m.

JvmMaxNativeStackSize

The maximum native stack size for the agent (in kilobytes). This property is applicable only if the `RepositoryDirectory` value is `<REMOTE>`.

The default value is 128k.

JvmMinHeapSize

The minimum heap size for the agent (in megabytes). This property is applicable only if the `RepositoryDirectory` value is `<REMOTE>`.

The default value is 1m.

jms.FactoryClassName

Specifies the class name to instantiate for a JMS provider. You *must* set this connector property when you choose JMS as your delivery transport mechanism (DeliveryTransport).

The default is `CxCommon.Messaging.jms.IBMMQSeriesFactory`.

jms.MessageBrokerName

Specifies the broker name to use for the JMS provider. You *must* set this connector property when you choose JMS as your delivery transport mechanism (DeliveryTransport).

The default is `crossworlds.queue.manager`.

jms.NumConcurrentRequests

Specifies the maximum number of concurrent service call requests that can be sent to a connector at the same time. Once that maximum is reached, new service calls block and wait for another request to complete before proceeding.

The default value is 10.

jms.Password

Specifies the password for the JMS provider. A value for this property is optional.

There is no default.

jms.UserName

Specifies the user name for the JMS provider. A value for this property is optional.

There is no default.

ListenerConcurrency

This property supports multi-threading in MQ Listener when ICS is the integration broker. It enables batch writing of multiple events to the database, thus improving system performance. The default value is 1.

This property applies only to connectors using MQ transport. The `DeliveryTransport` property must be set to MQ.

Locale

Specifies the language code, country or territory, and, optionally, the associated character code set. The value of this property determines such cultural conventions as collation and sort order of data, date and time formats, and the symbols used in monetary specifications.

A locale name has the following format:

ll_TT.codeset

where:

ll a two-character language code (usually in lower case)

<i>TT</i>	a two-letter country or territory code (usually in upper case)
<i>codeset</i>	the name of the associated character code set; this portion of the name is often optional.

By default, only a subset of supported locales appears in the drop list. To add other supported values to the drop list, you must manually modify the `\Data\Std\stdConnProps.xml` file in the product directory. For more information, see the appendix on Connector Configurator.

The default value is `en_US`. If the connector has not been globalized, the only valid value for this property is `en_US`. To determine whether a specific connector has been globalized, see the connector version list on these websites:

<http://www.ibm.com/software/websphere/wbiadapters/infocenter>, or
<http://www.ibm.com/websphere/integration/wicsserver/infocenter>

LogAtInterchangeEnd

Applicable only if `RepositoryDirectory` is `<REMOTE>`.

Specifies whether to log errors to the integration broker's log destination. Logging to the broker's log destination also turns on e-mail notification, which generates e-mail messages for the `MESSAGE_RECIPIENT` specified in the `InterchangeSystem.cfg` file when errors or fatal errors occur.

For example, when a connector loses its connection to its application, if `LogAtInterChangeEnd` is set to `true`, an e-mail message is sent to the specified message recipient. The default is `false`.

MaxEventCapacity

The maximum number of events in the controller buffer. This property is used by flow control and is applicable only if the value of the `RepositoryDirectory` property is `<REMOTE>`.

The value can be a positive integer between 1 and 2147483647. The default value is 2147483647.

MessageFileName

The name of the connector message file. The standard location for the message file is `\connectors\messages`. Specify the message filename in an absolute path if the message file is not located in the standard location.

If a connector message file does not exist, the connector uses `InterchangeSystem.txt` as the message file. This file is located in the product directory.

Note: To determine whether a specific connector has its own message file, see the individual adapter user guide.

MonitorQueue

The logical queue that the connector uses to monitor duplicate events. It is used only if the `DeliveryTransport` property value is `JMS` and `DuplicateEventElimination` is set to `TRUE`.

The default value is CONNECTORNAME/MONITORQUEUE

OADAutoRestartAgent

Valid only when the RepositoryDirectory is <REMOTE>.

Specifies whether the connector uses the automatic and remote restart feature. This feature uses the MQ-triggered Object Activation Daemon (OAD) to restart the connector after an abnormal shutdown, or to start a remote connector from System Monitor.

This property must be set to true to enable the automatic and remote restart feature. For information on how to configure the MQ-triggered OAD feature, see the *Installation Guide for Windows or for UNIX*.

The default value is false.

OADMaxNumRetry

Valid only when the RepositoryDirectory is <REMOTE>.

Specifies the maximum number of times that the MQ-triggered OAD automatically attempts to restart the connector after an abnormal shutdown. The OADAutoRestartAgent property must be set to true for this property to take effect.

The default value is 1000.

OADRetryTimeInterval

Valid only when the RepositoryDirectory is <REMOTE>.

Specifies the number of minutes in the retry-time interval for the MQ-triggered OAD. If the connector agent does not restart within this retry-time interval, the connector controller asks the OAD to restart the connector agent again. The OAD repeats this retry process as many times as specified by the OADMaxNumRetry property. The OADAutoRestartAgent property must be set to true for this property to take effect.

The default is 10.

PollEndTime

Time to stop polling the event queue. The format is HH:MM, where *HH* represents 0-23 hours, and *MM* represents 0-59 seconds.

You must provide a valid value for this property. The default value is HH:MM, but must be changed.

PollFrequency

The amount of time between polling actions. Set PollFrequency to one of the following values:

- The number of milliseconds between polling actions.
- The word *key*, which causes the connector to poll only when you type the letter *p* in the connector's Command Prompt window. Enter the word in lowercase.
- The word *no*, which causes the connector not to poll. Enter the word in lowercase.

The default is 10000.

Important: Some connectors have restrictions on the use of this property. To determine whether a specific connector does, see the installing and configuring chapter of its adapter guide.

PollQuantity

Designates the number of items from the application that the connector should poll for. If the adapter has a connector-specific property for setting the poll quantity, the value set in the connector-specific property will override the standard property value.

PollStartTime

The time to start polling the event queue. The format is *HH:MM*, where *HH* represents 0-23 hours, and *MM* represents 0-59 seconds.

You must provide a valid value for this property. The default value is *HH:MM*, but must be changed.

RequestQueue

The queue that is used by the integration broker to send business objects to the connector.

The default value is `CONNECTOR/REQUESTQUEUE`.

RepositoryDirectory

The location of the repository from which the connector reads the XML schema documents that store the meta-data for business object definitions.

When the integration broker is ICS, this value must be set to `<REMOTE>` because the connector obtains this information from the InterChange Server repository.

When the integration broker is a WebSphere message broker or WAS, this value must be set to `<local directory>`.

ResponseQueue

Applicable only if `DeliveryTransport` is JMS and required only if `RepositoryDirectory` is `<REMOTE>`.

Designates the JMS response queue, which delivers a response message from the connector framework to the integration broker. When the integration broker is ICS, the server sends the request and waits for a response message in the JMS response queue.

RestartRetryCount

Specifies the number of times the connector attempts to restart itself. When used for a parallel connector, specifies the number of times the master connector application-specific component attempts to restart the slave connector application-specific component.

The default is 3.

RestartRetryInterval

Specifies the interval in minutes at which the connector attempts to restart itself. When used for a parallel connector, specifies the interval at which the master connector application-specific component attempts to restart the slave connector application-specific component. Possible values ranges from 1 to 2147483647.

The default is 1.

RHF2MessageDomain

WebSphere message brokers and WAS only.

This property allows you to configure the value of the field domain name in the JMS header. When data is sent to WMQI over JMS transport, the adapter framework writes JMS header information, with a domain name and a fixed value of `mrm`. A configurable domain name enables users to track how the WMQI broker processes the message data.

A sample header would look like this:

```
<mcd><Msd>mrm</Msd><Set>3</Set><Type>
Retek_POPhyDesc</Type><Fmt>CwXML</Fmt></mcd>
```

The default value is `mrm`, but it may also be set to `xml`. This property only appears when `DeliveryTransport` is set to `JMSand` and `WireFormat` is set to `CwXML`.

SourceQueue

Applicable only if `DeliveryTransport` is `JMS` and `ContainerManagedEvents` is specified.

Designates the JMS source queue for the connector framework in support of guaranteed event delivery for JMS-enabled connectors that use a JMS event store. For further information, see “`ContainerManagedEvents`” on page 106.

The default value is `CONNECTOR/SOURCEQUEUE`.

SynchronousRequestQueue

Applicable only if `DeliveryTransport` is `JMS`.

Delivers request messages that require a synchronous response from the connector framework to the broker. This queue is necessary only if the connector uses synchronous execution. With synchronous execution, the connector framework sends a message to the `SynchronousRequestQueue` and waits for a response back from the broker on the `SynchronousResponseQueue`. The response message sent to the connector bears a correlation ID that matches the ID of the original message.

The default is `CONNECTORNAME/SYNCHRONOUSREQUESTQUEUE`

SynchronousResponseQueue

Applicable only if `DeliveryTransport` is `JMS`.

Delivers response messages sent in reply to a synchronous request from the broker to the connector framework. This queue is necessary only if the connector uses synchronous execution.

The default is CONNECTORNAME/SYNCHRONOUSRESPONSEQUEUE

SynchronousRequestTimeout

Applicable only if DeliveryTransport is JMS.

Specifies the time in minutes that the connector waits for a response to a synchronous request. If the response is not received within the specified time, then the connector moves the original synchronous request message into the fault queue along with an error message.

The default value is 0.

WireFormat

Message format on the transport.

- If the RepositoryDirectory is a local directory, the setting is CwXML.
- If the value of RepositoryDirectory is <REMOTE>, the setting is CwB0.

WsifSynchronousRequest Timeout

WAS integration broker only.

Specifies the time in minutes that the connector waits for a response to a synchronous request. If the response is not received within the specified, time then the connector moves the original synchronous request message into the fault queue along with an error message.

The default value is 0.

XMLNameSpaceFormat

WebSphere message brokers and WAS integration broker only.

A strong property that allows the user to specify short and long name spaces in the XML format of business object definitions.

The default value is short.

Appendix B. Connector Configurator

This appendix describes how to use Connector Configurator to set configuration property values for your adapter.

You use Connector Configurator to:

- Create a connector-specific property template for configuring your connector
- Create a configuration file
- Set properties in a configuration file

Note:

In this document, backslashes (\) are used as the convention for directory paths. For UNIX installations, substitute slashes (/) for backslashes and follow the conventions for each operating system.

The topics covered in this appendix are:

- “Overview of Connector Configurator” on page 115
- “Starting Connector Configurator” on page 116
- “Creating a connector-specific property template” on page 117
- “Creating a new configuration file” on page 119
- “Setting the configuration file properties” on page 122
- “Using Connector Configurator in a globalized environment” on page 128

Overview of Connector Configurator

Connector Configurator allows you to configure the connector component of your adapter for use with these integration brokers:

- WebSphere InterChange Server (ICS)
- WebSphere MQ Integrator, WebSphere MQ Integrator Broker, and WebSphere Business Integration Message Broker, collectively referred to as the WebSphere Message Brokers (WMQI)
- WebSphere Application Server (WAS)

You use Connector Configurator to:

- Create a **connector-specific property template** for configuring your connector.
- Create a **connector configuration file**; you must create one configuration file for each connector you install.
- Set properties in a configuration file.

You may need to modify the default values that are set for properties in the connector templates. You must also designate supported business object definitions and, with ICS, maps for use with collaborations as well as specify messaging, logging and tracing, and data handler parameters, as required.

The mode in which you run Connector Configurator, and the configuration file type you use, may differ according to which integration broker you are running. For example, if WMQI is your broker, you run Connector Configurator directly, and not from within System Manager (see “Running Configurator in stand-alone mode” on page 116).

Connector configuration properties include both standard configuration properties (the properties that all connectors have) and connector-specific properties (properties that are needed by the connector for a specific application or technology).

Because **standard properties** are used by all connectors, you do not need to define those properties from scratch; Connector Configurator incorporates them into your configuration file as soon as you create the file. However, you do need to set the value of each standard property in Connector Configurator.

The range of standard properties may not be the same for all brokers and all configurations. Some properties are available only if other properties are given a specific value. The Standard Properties window in Connector Configurator will show the properties available for your particular configuration.

For **connector-specific properties**, however, you need first to define the properties and then set their values. You do this by creating a connector-specific property template for your particular adapter. There may already be a template set up in your system, in which case, you simply use that. If not, follow the steps in “Creating a new template” on page 117 to set up a new one.

Note: Connector Configurator runs only in a Windows environment. If you are running the connector in a UNIX environment, use Connector Configurator in Windows to modify the configuration file and then copy the file to your UNIX environment.

Starting Connector Configurator

You can start and run Connector Configurator in either of two modes:

- Independently, in stand-alone mode
- From System Manager

Running Configurator in stand-alone mode

You can run Connector Configurator independently and work with connector configuration files, irrespective of your broker.

To do so:

- From **Start>Programs**, click **IBM WebSphere InterChange Server>IBM WebSphere Business Integration Toolset>Development>Connector Configurator**.
- Select **File>New>Configuration File**.
- When you click the pull-down menu next to **System Connectivity Integration Broker**, you can select ICS, WebSphere Message Brokers or WAS, depending on your broker.

You may choose to run Connector Configurator independently to generate the file, and then connect to System Manager to save it in a System Manager project (see “Completing a configuration file” on page 121.)

Running Configurator from System Manager

You can run Connector Configurator from System Manager.

To run Connector Configurator:

1. Open the System Manager.
2. In the System Manager window, expand the **Integration Component Libraries** icon and highlight **Connectors**.
3. From the System Manager menu bar, click **Tools>Connector Configurator**. The Connector Configurator window opens and displays a **New Connector** dialog box.
4. When you click the pull-down menu next to **System Connectivity Integration Broker**, you can select ICS, WebSphere Message Brokers or WAS, depending on your broker.

To edit an existing configuration file:

1. In the System Manager window, select any of the configuration files listed in the Connector folder and right-click on it. Connector Configurator opens and displays the configuration file with the integration broker type and file name at the top.
2. Click the Standard Properties tab to see which properties are included in this configuration file.

Creating a connector-specific property template

To create a configuration file for your connector, you need a connector-specific property template as well as the system-supplied standard properties.

You can create a brand-new template for the connector-specific properties of your connector, or you can use an existing file as the template.

- To create a new template, see “Creating a new template” on page 117.
- To use an existing file, simply modify an existing template and save it under the new name.

Creating a new template

This section describes how you create properties in the template, define general characteristics and values for those properties, and specify any dependencies between the properties. Then you save the template and use it as the base for creating a new connector configuration file.

To create a template:

1. Click **File>New>Connector-Specific Property Template**.
2. The **Connector-Specific Property Template** dialog box appears, with the following fields:
 - **Template**, and **Name**
Enter a unique name that identifies the connector, or type of connector, for which this template will be used. You will see this name again when you open the dialog box for creating a new configuration file from a template.
 - **Old Template**, and **Select the Existing Template to Modify**
The names of all currently available templates are displayed in the Template Name display.

- To see the connector-specific property definitions in any template, select that template's name in the **Template Name** display. A list of the property definitions contained in that template will appear in the **Template Preview** display. You can use an existing template whose property definitions are similar to those required by your connector as a starting point for your template.
3. Select a template from the **Template Name** display, enter that template name in the **Find Name** field (or highlight your selection in **Template Name**), and click **Next**.

If you do not see any template that displays the connector-specific properties used by your connector, you will need to create one.

Specifying general characteristics

When you click **Next** to select a template, the **Properties - Connector-Specific Property Template** dialog box appears. The dialog box has tabs for General characteristics of the defined properties and for Value restrictions. The General display has the following fields:

- **General:**
 - Property Type
 - Updated Method
 - Description
- **Flags**
 - Standard flags
- **Custom Flag**
 - Flag

After you have made selections for the general characteristics of the property, click the **Value** tab.

Specifying values

The **Value** tab enables you to set the maximum length, the maximum multiple values, a default value, or a value range for the property. It also allows editable values. To do so:

1. Click the **Value** tab. The display panel for Value replaces the display panel for General.
2. Select the name of the property in the **Edit properties** display.
3. In the fields for **Max Length** and **Max Multiple Values**, make any changes. The changes will not be accepted unless you also open the **Property Value** dialog box for the property, described in the next step.
4. Right-click the box in the top left-hand corner of the value table and click **Add**. A **Property Value** dialog box appears. Depending on the property type, the dialog box allows you to enter either a value, or both a value and range. Enter the appropriate value or range, and click **OK**.
5. The **Value** panel refreshes to display any changes you made in **Max Length** and **Max Multiple Values**. It displays a table with three columns:
 - The **Value** column shows the value that you entered in the **Property Value** dialog box, and any previous values that you created.
 - The **Default Value** column allows you to designate any of the values as the default.
 - The **Value Range** shows the range that you entered in the **Property Value** dialog box.

After a value has been created and appears in the grid, it can be edited from within the table display. To make a change in an existing value in the table, select an entire row by clicking on the row number. Then right-click in the **Value** field and click **Edit Value**.

Setting dependencies

When you have made your changes to the **General** and **Value** tabs, click **Next**. The **Dependencies - Connector-Specific Property Template** dialog box appears.

A dependent property is a property that is included in the template and used in the configuration file *only if* the value of another property meets a specific condition. For example, `Pol1Quantity` appears in the template only if JMS is the transport mechanism and `DuplicateEventElimination` is set to `True`.

To designate a property as dependent and to set the condition upon which it depends, do this:

1. In the **Available Properties** display, select the property that will be made dependent.
2. In the **Select Property** field, use the drop-down menu to select the property that will hold the conditional value.
3. In the **Condition Operator** field, select one of the following:
 - == (equal to)
 - != (not equal to)
 - > (greater than)
 - < (less than)
 - >= (greater than or equal to)
 - <=(less than or equal to)
4. In the **Conditional Value** field, enter the value that is required in order for the dependent property to be included in the template.
5. With the dependent property highlighted in the **Available Properties** display, click an arrow to move it to the **Dependent Property** display.
6. Click **Finish**. Connector Configurator stores the information you have entered as an XML document, under `\data\app` in the `\bin` directory where you have installed Connector Configurator.

Creating a new configuration file

When you create a new configuration file, your first step is to select an integration broker. The broker you select determines the properties that will appear in the configuration file.

To select a broker:

- In the Connector Configurator home menu, click **File>New>Connector Configuration**. The **New Connector** dialog box appears.
- In the **Integration Broker** field, select ICS, WebSphere Message Brokers or WAS connectivity.
- Complete the remaining fields in the **New Connector** window, as described later in this chapter.

You can also do this:

- In the System Manager window, right-click on the **Connectors** folder and select **Create New Connector**. Connector Configurator opens and displays the **New Connector** dialog box.

Creating a configuration file from a connector-specific template

Once a connector-specific template has been created, you can use it to create a configuration file:

1. Click **File>New>Connector Configuration**.
2. The **New Connector** dialog box appears, with the following fields:
 - **Name**

Enter the name of the connector. Names are case-sensitive. The name you enter must be unique, and must be consistent with the file name for a connector that is installed on the system.

Important: Connector Configurator does not check the spelling of the name that you enter. You must ensure that the name is correct.
 - **System Connectivity**

Click ICS or WebSphere Message Brokers or WAS.
 - **Select Connector-Specific Property Template**

Type the name of the template that has been designed for your connector. The available templates are shown in the **Template Name** display. When you select a name in the Template Name display, the **Property Template Preview** display shows the connector-specific properties that have been defined in that template.

Select the template you want to use and click **OK**.
3. A configuration screen appears for the connector that you are configuring. The title bar shows the integration broker and connector names. You can fill in all the field values to complete the definition now, or you can save the file and complete the fields later.
4. To save the file, click **File>Save>To File** or **File>Save>To Project**. To save to a project, System Manager must be running.

If you save as a file, the **Save File Connector** dialog box appears. Choose *.cfg as the file type, verify in the File Name field that the name is spelled correctly and has the correct case, navigate to the directory where you want to locate the file, and click **Save**. The status display in the message panel of Connector Configurator indicates that the configuration file was successfully created.

Important: The directory path and name that you establish here must match the connector configuration file path and name that you supply in the startup file for the connector.
5. To complete the connector definition, enter values in the fields for each of the tabs of the Connector Configurator window, as described later in this chapter.

Using an existing file

You may have an existing file available in one or more of the following formats:

- A connector definition file.

This is a text file that lists properties and applicable default values for a specific connector. Some connectors include such a file in a \repository directory in their delivery package (the file typically has the extension .txt; for example, CN_XML.txt for the XML connector).
- An ICS repository file.

Definitions used in a previous ICS implementation of the connector may be available to you in a repository file that was used in the configuration of that connector. Such a file typically has the extension .in or .out.

- A previous configuration file for the connector.
Such a file typically has the extension *.cfg.

Although any of these file sources may contain most or all of the connector-specific properties for your connector, the connector configuration file will not be complete until you have opened the file and set properties, as described later in this chapter.

To use an existing file to configure a connector, you must open the file in Connector Configurator, revise the configuration, and then resave the file.

Follow these steps to open a *.txt, *.cfg, or *.in file from a directory:

1. In Connector Configurator, click **File>Open>From File**.
2. In the **Open File Connector** dialog box, select one of the following file types to see the available files:
 - Configuration (*.cfg)
 - ICS Repository (*.in, *.out)
Choose this option if a repository file was used to configure the connector in an ICS environment. A repository file may include multiple connector definitions, all of which will appear when you open the file.
 - All files (*.*)
Choose this option if a *.txt file was delivered in the adapter package for the connector, or if a definition file is available under another extension.
3. In the directory display, navigate to the appropriate connector definition file, select it, and click **Open**.

Follow these steps to open a connector configuration from a System Manager project:

1. Start System Manager. A configuration can be opened from or saved to System Manager only if System Manager has been started.
2. Start Connector Configurator.
3. Click **File>Open>From Project**.

Completing a configuration file

When you open a configuration file or a connector from a project, the Connector Configurator window displays the configuration screen, with the current attributes and values.

The title of the configuration screen displays the integration broker and connector name as specified in the file. Make sure you have the correct broker. If not, change the broker value before you configure the connector. To do so:

1. Under the **Standard Properties** tab, select the value field for the BrokerType property. In the drop-down menu, select the value ICS, WMQI, or WAS.
2. The Standard Properties tab will display the properties associated with the selected broker. You can save the file now or complete the remaining configuration fields, as described in “Specifying supported business object definitions” on page 124..
3. When you have finished your configuration, click **File>Save>To Project** or **File>Save>To File**.

If you are saving to file, select *.cfg as the extension, select the correct location for the file and click **Save**.

If multiple connector configurations are open, click **Save All to File** to save all of the configurations to file, or click **Save All to Project** to save all connector configurations to a System Manager project.

Before it saves the file, Connector Configurator checks that values have been set for all required standard properties. If a required standard property is missing a value, Connector Configurator displays a message that the validation failed. You must supply a value for the property in order to save the configuration file.

Setting the configuration file properties

When you create and name a new connector configuration file, or when you open an existing connector configuration file, Connector Configurator displays a configuration screen with tabs for the categories of required configuration values.

Connector Configurator requires values for properties in these categories for connectors running on all brokers:

- Standard Properties
- Connector-specific Properties
- Supported Business Objects
- Trace/Log File values
- Data Handler (applicable for connectors that use JMS messaging with guaranteed event delivery)

Note: For connectors that use JMS messaging, an additional category may display, for configuration of data handlers that convert the data to business objects.

For connectors running on **ICS**, values for these properties are also required:

- Associated Maps
- Resources
- Messaging (where applicable)

Important: Connector Configurator accepts property values in either English or non-English character sets. However, the names of both standard and connector-specific properties, and the names of supported business objects, must use the English character set only.

Standard properties differ from connector-specific properties as follows:

- Standard properties of a connector are shared by both the application-specific component of a connector and its broker component. All connectors have the same set of standard properties. These properties are described in Appendix A of each adapter guide. You can change some but not all of these values.
- Application-specific properties apply only to the application-specific component of a connector, that is, the component that interacts directly with the application. Each connector has application-specific properties that are unique to its application. Some of these properties provide default values and some do not; you can modify some of the default values. The installation and configuration chapters of each adapter guide describe the application-specific properties and the recommended values.

The fields for **Standard Properties** and **Connector-Specific Properties** are color-coded to show which are configurable:

- A field with a grey background indicates a standard property. You can change the value but cannot change the name or remove the property.
- A field with a white background indicates an application-specific property. These properties vary according to the specific needs of the application or connector. You can change the value and delete these properties.
- Value fields are configurable.
- The **Update Method** field is informational and not configurable. This field specifies the action required to activate a property whose value has changed.

Setting standard connector properties

To change the value of a standard property:

1. Click in the field whose value you want to set.
2. Either enter a value, or select one from the drop-down menu if it appears.
3. After entering all the values for the standard properties, you can do one of the following:
 - To discard the changes, preserve the original values, and exit Connector Configurator, click **File>Exit** (or close the window), and click **No** when prompted to save changes.
 - To enter values for other categories in Connector Configurator, select the tab for the category. The values you enter for **Standard Properties** (or any other category) are retained when you move to the next category. When you close the window, you are prompted to either save or discard the values that you entered in all the categories as a whole.
 - To save the revised values, click **File>Exit** (or close the window) and click **Yes** when prompted to save changes. Alternatively, click **Save>To File** from either the File menu or the toolbar.

Setting application-specific configuration properties

For application-specific configuration properties, you can add or change property names, configure values, delete a property, and encrypt a property. The default property length is 255 characters.

1. Right-click in the top left portion of the grid. A pop-up menu bar will appear. Click **Add** to add a property. To add a child property, right-click on the parent row number and click **Add child**.
2. Enter a value for the property or child property.
3. To encrypt a property, select the **Encrypt** box.
4. Choose to save or discard changes, as described for “Setting standard connector properties.”

The Update Method displayed for each property indicates whether a component or agent restart is necessary to activate changed values.

Important: Changing a preset application-specific connector property name may cause a connector to fail. Certain property names may be needed by the connector to connect to an application or to run properly.

Encryption for connector properties

Application-specific properties can be encrypted by selecting the **Encrypt** check box in the **Edit Property** window. To decrypt a value, click to clear the **Encrypt** check box, enter the correct value in the **Verification** dialog box, and click **OK**. If the entered value is correct, the value is decrypted and displays.

The adapter user guide for each connector contains a list and description of each property and its default value.

If a property has multiple values, the **Encrypt** check box will appear for the first value of the property. When you select **Encrypt**, all values of the property will be encrypted. To decrypt multiple values of a property, click to clear the **Encrypt** check box for the first value of the property, and then enter the new value in the **Verification** dialog box. If the input value is a match, all multiple values will decrypt.

Update method

Refer to the descriptions of update methods found in the *Standard configuration properties for connectors* appendix, under “Setting and updating property values” on page 100.

Specifying supported business object definitions

Use the **Supported Business Objects** tab in Connector Configurator to specify the business objects that the connector will use. You must specify both generic business objects and application-specific business objects, and you must specify associations for the maps between the business objects.

Note: Some connectors require that certain business objects be specified as supported in order to perform event notification or additional configuration (using meta-objects) with their applications. For more information, see the *Connector Development Guide for C++* or the *Connector Development Guide for Java*.

If ICS is your broker

To specify that a business object definition is supported by the connector, or to change the support settings for an existing business object definition, click the **Supported Business Objects** tab and use the following fields.

Business object name: To designate that a business object definition is supported by the connector, with System Manager running:

1. Click an empty field in the **Business Object Name** list. A drop-down list displays, showing all the business object definitions that exist in the System Manager project.
2. Click on a business object to add it.
3. Set the **Agent Support** (described below) for the business object.
4. In the File menu of the Connector Configurator window, click **Save to Project**. The revised connector definition, including designated support for the added business object definition, is saved to the project in System Manager.

To delete a business object from the supported list:

1. To select a business object field, click the number to the left of the business object.
2. From the **Edit** menu of the Connector Configurator window, click **Delete Row**. The business object is removed from the list display.
3. From the **File** menu, click **Save to Project**.

Deleting a business object from the supported list changes the connector definition and makes the deleted business object unavailable for use in this implementation of this connector. It does not affect the connector code, nor does it remove the business object definition itself from System Manager.

Agent support: If a business object has Agent Support, the system will attempt to use that business object for delivering data to an application via the connector agent.

Typically, application-specific business objects for a connector are supported by that connector's agent, but generic business objects are not.

To indicate that the business object is supported by the connector agent, check the **Agent Support** box. The Connector Configurator window does not validate your Agent Support selections.

Maximum transaction level: The maximum transaction level for a connector is the highest transaction level that the connector supports.

For most connectors, Best Effort is the only possible choice.

You must restart the server for changes in transaction level to take effect.

If a WebSphere Message Broker is your broker

If you are working in stand-alone mode (not connected to System Manager), you must enter the business name manually.

If you have System Manager running, you can select the empty box under the **Business Object Name** column in the **Supported Business Objects** tab. A combo box appears with a list of the business object available from the Integration Component Library project to which the connector belongs. Select the business object you want from the list.

The **Message Set ID** is an optional field for WebSphere Business Integration Message Broker 5.0, and need not be unique if supplied. However, for WebSphere MQ Integrator and Integrator Broker 2.1, you must supply a unique **ID**.

If WAS is your broker

When WebSphere Application Server is selected as your broker type, Connector Configurator does not require message set IDs. The **Supported Business Objects** tab shows a **Business Object Name** column only for supported business objects.

If you are working in stand-alone mode (not connected to System Manager), you must enter the business object name manually.

If you have System Manager running, you can select the empty box under the Business Object Name column in the Supported Business Objects tab. A combo box appears with a list of the business objects available from the Integration Component Library project to which the connector belongs. Select the business object you want from this list.

Associated maps (ICS only)

Each connector supports a list of business object definitions and their associated maps that are currently active in WebSphere InterChange Server. This list appears when you select the **Associated Maps** tab.

The list of business objects contains the application-specific business object which the agent supports and the corresponding generic object that the controller sends to the subscribing collaboration. The association of a map determines which map

will be used to transform the application-specific business object to the generic business object or the generic business object to the application-specific business object.

If you are using maps that are uniquely defined for specific source and destination business objects, the maps will already be associated with their appropriate business objects when you open the display, and you will not need (or be able) to change them.

If more than one map is available for use by a supported business object, you will need to explicitly bind the business object with the map that it should use.

The **Associated Maps** tab displays the following fields:

- **Business Object Name**

These are the business objects supported by this connector, as designated in the **Supported Business Objects** tab. If you designate additional business objects under the Supported Business Objects tab, they will be reflected in this list after you save the changes by choosing **Save to Project** from the **File** menu of the Connector Configurator window.

- **Associated Maps**

The display shows all the maps that have been installed to the system for use with the supported business objects of the connector. The source business object for each map is shown to the left of the map name, in the **Business Object Name** display.

- **Explicit**

In some cases, you may need to explicitly bind an associated map.

Explicit binding is required only when more than one map exists for a particular supported business object. When ICS boots, it tries to automatically bind a map to each supported business object for each connector. If more than one map takes as its input the same business object, the server attempts to locate and bind one map that is the superset of the others.

If there is no map that is the superset of the others, the server will not be able to bind the business object to a single map, and you will need to set the binding explicitly.

To explicitly bind a map:

1. In the **Explicit** column, place a check in the check box for the map you want to bind.
2. Select the map that you intend to associate with the business object.
3. In the **File** menu of the Connector Configurator window, click **Save to Project**.
4. Deploy the project to ICS.
5. Reboot the server for the changes to take effect.

Resources (ICS)

The **Resource** tab allows you to set a value that determines whether and to what extent the connector agent will handle multiple processes concurrently, using connector agent parallelism.

Not all connectors support this feature. If you are running a connector agent that was designed in Java to be multi-threaded, you are advised not to use this feature, since it is usually more efficient to use multiple threads than multiple processes.

Messaging (ICS)

The messaging properties are available only if you have set MQ as the value of the `DeliveryTransport` standard property and ICS as the broker type. These properties affect how your connector will use queues.

Setting trace/log file values

When you open a connector configuration file or a connector definition file, Connector Configurator uses the logging and tracing values of that file as default values. You can change those values in Connector Configurator.

To change the logging and tracing values:

1. Click the **Trace/Log Files** tab.
2. For either logging or tracing, you can choose to write messages to one or both of the following:
 - To console (STDOUT):
Writes logging or tracing messages to the STDOUT display.

Note: You can only use the STDOUT option from the **Trace/Log Files** tab for connectors running on the Windows platform.

- To File:
Writes logging or tracing messages to a file that you specify. To specify the file, click the directory button (ellipsis), navigate to the preferred location, provide a file name, and click **Save**. Logging or tracing message are written to the file and location that you specify.

Note: Both logging and tracing files are simple text files. You can use the file extension that you prefer when you set their file names. For tracing files, however, it is advisable to use the extension `.trace` rather than `.trc`, to avoid confusion with other files that might reside on the system. For logging files, `.log` and `.txt` are typical file extensions.

Data handlers

The data handlers section is available for configuration only if you have designated a value of JMS for `DeliveryTransport` and a value of JMS for `ContainerManagedEvents`. Not all adapters make use of data handlers.

See the descriptions under `ContainerManagedEvents` in Appendix A, *Standard Properties*, for values to use for these properties. For additional details, see the *Connector Development Guide for C++* or the *Connector Development Guide for Java*.

Saving your configuration file

When you have finished configuring your connector, save the connector configuration file. Connector Configurator saves the file in the broker mode that you selected during configuration. The title bar of Connector Configurator always displays the broker mode (ICS, WMQI or WAS) that it is currently using.

The file is saved as an XML document. You can save the XML document in three ways:

- From System Manager, as a file with a `*.con` extension in an Integration Component Library, or
- In a directory that you specify.
- In stand-alone mode, as a file with a `*.cfg` extension in a directory folder.

For details about using projects in System Manager, and for further information about deployment, see the following implementation guides:

- For ICS: *Implementation Guide for WebSphere InterChange Server*
- For WebSphere Message Brokers: *Implementing Adapters with WebSphere Message Brokers*
- For WAS: *Implementing Adapters with WebSphere Application Server*

Changing a configuration file

You can change the integration broker setting for an existing configuration file. This enables you to use the file as a template for creating a new configuration file, which can be used with a different broker.

Note: You will need to change other configuration properties as well as the broker mode property if you switch integration brokers.

To change your broker selection within an existing configuration file (optional):

- Open the existing configuration file in Connector Configurator.
- Select the **Standard Properties** tab.
- In the **BrokerType** field of the Standard Properties tab, select the value that is appropriate for your broker.
When you change the current value, the available tabs and field selections on the properties screen will immediately change, to show only those tabs and fields that pertain to the new broker you have selected.

Completing the configuration

After you have created a configuration file for a connector and modified it, make sure that the connector can locate the configuration file when the connector starts up.

To do so, open the startup file used for the connector, and verify that the location and file name used for the connector configuration file match exactly the name you have given the file and the directory or path where you have placed it.

Using Connector Configurator in a globalized environment

Connector Configurator is globalized and can handle character conversion between the configuration file and the integration broker. Connector Configurator uses native encoding. When it writes to the configuration file, it uses UTF-8 encoding.

Connector Configurator supports non-English characters in:

- All value fields
- Log file and trace file path (specified in the **Trace/Log files** tab)

The drop list for the CharacterEncoding and Locale standard configuration properties displays only a subset of supported values. To add other values to the drop list, you must manually modify the `\Data\Std\stdConnProps.xml` file in the product directory.

For example, to add the locale `en_GB` to the list of values for the Locale property, open the `stdConnProps.xml` file and add the line in boldface type below:

```
<Property name="Locale"
isRequired="true"
updateMethod="component restart">
  <ValidType>String</ValidType>
  <ValidValues>
    <Value>ja_JP</Value>
    <Value>ko_KR</Value>
    <Value>zh_CN</Value>
    <Value>zh_TW</Value>
    <Value>fr_FR</Value>
    <Value>de_DE</Value>
    <Value>it_IT</Value>
    <Value>es_ES</Value>
    <Value>pt_BR</Value>
    <Value>en_US</Value>
    <Value>en_GB</Value>
  </ValidValues>
  <DefaultValue>en_US</DefaultValue>
</Property>
```

Appendix C. MAS driver keywords

The MAS ODBC driver is controlled using certain driver keywords. These keywords are automatically given default values that can be changed depending on the connector you are using. This chapter will show you:

- How to change a keyword setting.
- Details about each keyword.

Setting a driver keyword

You can set driver keywords using the following methods, or a combination of the following:

- ODBC connection string (passed by an application program).
- ODBC.INI file.
- ODBC Datasource Administrator.
- SQLConnect function
- SQLSetConnectOption function
- SQLSetStmtOption function

Using the ODBC connection string

To specify a setting using an ODBC connection string, code the keyword in your program, as shown in the following example:

```
rc = SQLDriver Connect (hdbc, NULL, (u_char far *)
    "NOAS=no;"
    "UID=ai00xyz;PWD=abcdef;"
    "PORT=1200;HOST=129.4.0.160;"
    "LGID=dan;"
    "DSN=Tu1d;")
```

Using the ODBC.INI File

The following example shows how to set a keyword value using the ODBC.INI file:

```
Driver=C:\WINDOWS\SYSTEM\scodbc.dll
Description=Sample Shadow Direct Driver
APPL=ODBC
LINK=TCPIP
LGID=ita
UID=ai00xyz
PWD=abcdef
HOST=129.4.0.160
```

Using the ODBC datasource administrator

Keywords can also be set using the MAS ODBC Datasource Administrator. This method can be used after the MAS ODBC driver has been installed and configured. To use this method:

1. Open the Control Panel, and double-click the ODBC icon.
The system displays the following screen:

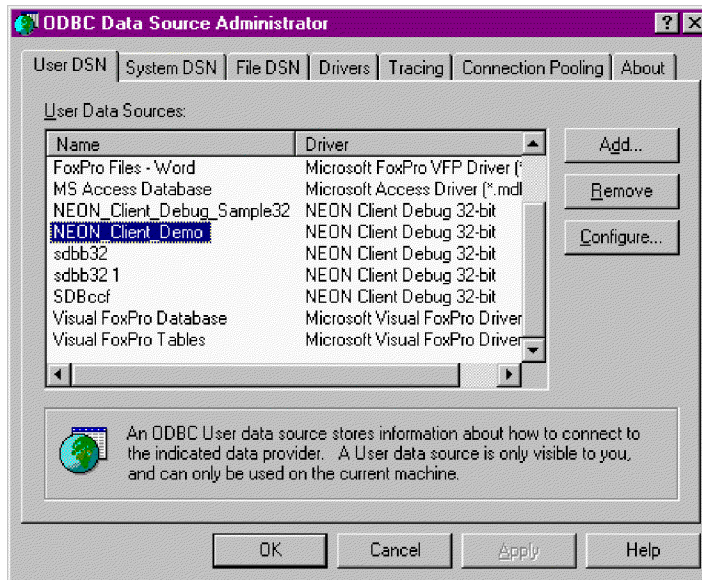


Figure 21. ODBC Data Source Administrator screen

2. Click **Configure**.

The system displays the following screen:

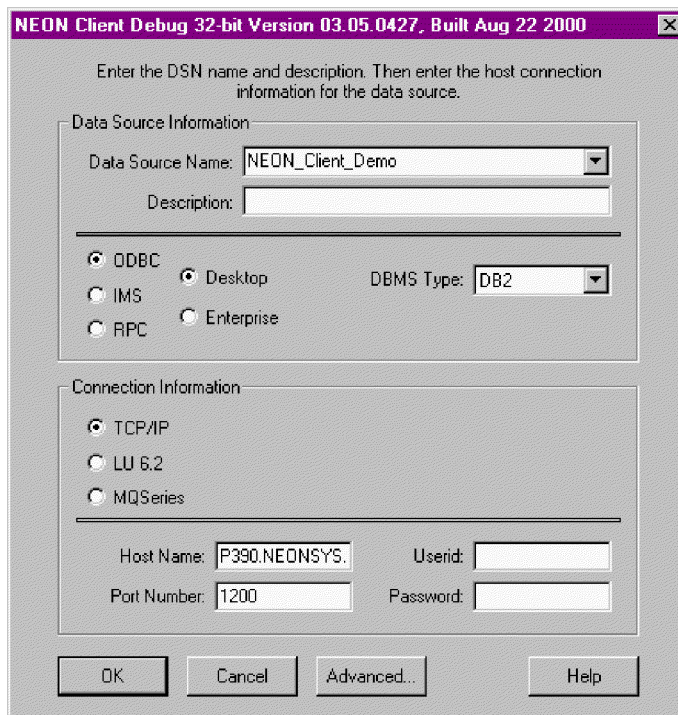


Figure 22. Data source name screen

3. Click **Advanced**.

The system displays the following screen:

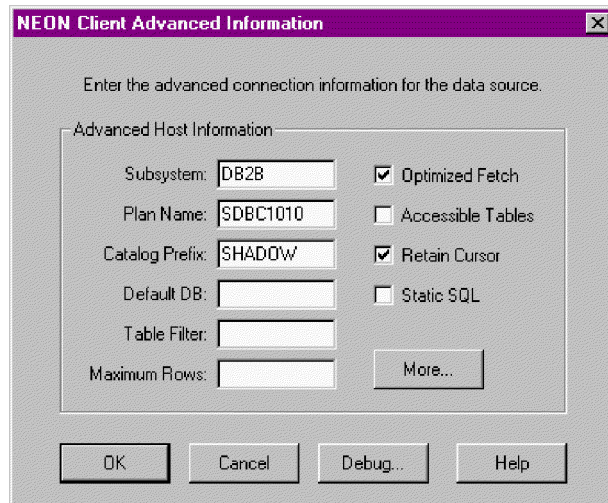


Figure 23. Advanced Information Screen

4. Click **More**.

The system displays the following screen:

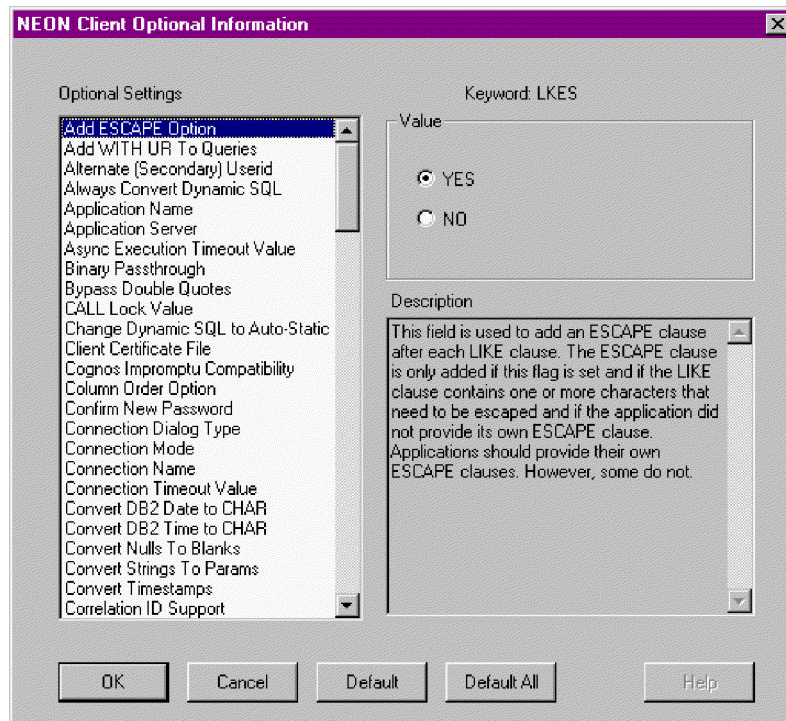


Figure 24. Optional Information Screen

5. Select the Keyword, in this example it is LKES. You can change the setting from Yes to No.

Driver keyword descriptions

Table 31 on page 134 lists the driver keywords that can be set, along with a brief explanation of each and default settings. All of the keywords are not necessarily able to be set in all the methods described above (see Table 32 on page 150).

Table 31. Driver keywords

Description	Keyword	(Format) Explanation	Default
Add ESCAPE Option	LKES	This keyword is used to add an ESCAPE clause after each LIKE clause. The ESCAPE clause is only added if: <ul style="list-style-type: none"> • This keyword is set. • The LIKE clause contains one or more characters that need to be escaped. • The application did not provide its own ESCAPE clause. 	YES
Add WITH UR to Queries	WIUR	This keyword controls whether or not a WITH UR clause should be added to queries. The WITH UR clause reduces the amount of CPU time needed for some queries. There is also the possibility it could change the results of these queries by allowing uncommitted data to be read. If this keyword is set to YES, a WITH UR clause will be added to queries. If this keyword is set to NO, a WITH UR clause will not be added to queries.	NO
Alternate (Secondary) Userid	ALUS	This keyword is used to set the host secondary userid for client applications. This keyword must be eight or fewer characters long. If this keyword is set to a non-blank, non-null value, a SET CURRENT SQLID statement is issued after DSNALI OPEN processing is completed.	
Always Convert Dynamic SQL	ALCD	This keyword controls what happens if dynamic SQL can not be converted to static SQL. If this keyword is set to YES, and dynamic SQL cannot be converted to static, an error is reported to the application. If this keyword is set to NO, the dynamic SQL is sent to the host for processing. Note: This keyword is not even tested unless the primary Dynamic-To-Static SQL keyword has been set to YES. This keyword must be set to NO if dynamic SQL and static SQL are going to be used together.	YES
Application Name	APNA	This keyword is the Application name. The Application Name is sent to the host as part of the logon information. The Application Name is normally used to group SQL statements within a plan. If the Application Name is not set, all of the SQL associated with a plan will be considered to be part of one large group. If the SQL used with one plan must be divided into subgroups (for conversion to static SQL), the Application Name must be set.	
Application Server	SECN	This keyword is used to select the initial Application Server value on the host. It must be sixteen or fewer characters long. If this keyword is set to a non-blank, non-null value, the current application will be connected to the specified Application Server after DSNALI OPEN processing is completed.	
Async Execution Timeout Value	ASTM	This keyword controls the default wait time for an ODBC function that is executing asynchronously. This keyword is measured in seconds, and must be an integer.	2
Binary Passthrough	BIPA	This keyword controls whether or not host binary data should be returned to client applications without being converted to hexadecimal. If this keyword is set to YES, host binary data is passed through to client applications unchanged. If this keyword is set to NO, host binary data is converted to hexadecimal (in accordance with the ODBC specification).	NO
Bypass Double Quote	BYDB	This keyword is set to YES if double quotes should be left alone. This keyword is provided to fix certain application bugs.	NO

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
CALL Lock Value	CALK	This keyword controls the type of lock associated with CALL statements on the host. If this keyword is set to NONE, then the host code will assume that CALL statements do not obtain any host database locks. The other possible values are SHARE, UPDATE, and EXCLUSIVE.	EXCLUSIVE
Change Dynamic SQL to Static	AUST	This keyword is used to control whether or not Shadow Direct should try to convert dynamic SQL to Auto-Static SQL. Dynamic SQL can only be converted to Auto-Static SQL if this keyword is set to YES and if the host supports Auto-Static SQL. Note: If this keyword is set to YES, all dynamic SQL statements will be looked up in a lookaside buffer and converted to static SQL, if possible. If this keyword is to NO, normal dynamic SQL processing will be performed.	YES
Client Certificate File	CLCT	No description is available.	
Cognos Impromptu Compatibility	CGFX	This keyword is set to YES to resolve certain problems with Cognos Impromptu. This keyword should not be set for any other reason.	YES
Column Order Option	CLOR	This keyword shows how column names should be returned by ODBC catalog functions. Some functions explicitly specify an order. Other functions don't. This keyword is used by the ODBC catalog functions (SQLColumns) that don't specify a column name order.	NUMBER
Confirm New Password	CMNP	If this keyword is set, the user will be asked to confirm the new password string if the new password is set in the ODBC PWD keyword.	NO
Connection Dialog Type	CNDG	This keyword shows what type of connection dialog should be displayed. This keyword can be set using the ODBC.INI file or using a connection string. If the dynamic dialog is specified, either the simple or the detailed dialog is displayed depending on how much information is needed.	DYNAMIC
Connection Mode	CNMD	This keyword controls the connection mode used by ODBC applications. The connection mode determines how long each physical connection (session or conversation) lasts and if SQL operations are blocked together. The default is to use a permanent connection and to send each SQL operation standalone to the server. In BLOCK mode, the session is permanent. However, SQL operations are blocked and sent together. In TRANSACTION mode each SQL operation is sent standalone but the session is terminated at the end of each Logical Unit Of Work (LUOW). In TRANSBLOCK mode the session is terminated at the end of each LUOW, and SQL operations are blocked and sent together. In MESSAGE mode SQL operations are blocked and sent together using messages. In MESSAGE mode, no session is ever maintained.	PERMANENT
Connection Name	CNNA	This keyword is used to specify the connection name. The use of the connection name is application specific. The name can be up to eight (8) byte long, however, the application may or may not use all of the bytes. The connection name is padded on the right with blanks.	

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
Connection Timeout Value	CNTM	This keyword controls how many seconds the ODBC client will wait for a connection to the host server to complete. If this keyword is set to zero, then the system default value will be used. Otherwise, the specified value will be used. This value should never be negative, but can be zero. This value has no effect in UDP messaging. Note: The use of this keyword may cause some unpredictable results in many environments. This option is only available in the TCP/IP and the MQ link types, but is applicable to both the permanent and the messaging modes.	0
Convert DB2 Date to CHAR	DTCH	When this keyword is set, the DB2 date keyword will be converted to SQL_CHAR.	NO
Convert DB2 Time to CHAR	TSCH	When this keyword is set, the DB2 timestamp and time keyword will be converted to SQL_CHAR. If the keyword AF (MS Access Compatibility) is set to YES, this keyword does not need to be set. Keyword AF includes the functionality of keyword TSCH.	NO
Convert Nulls to Blanks	CVNL	This keyword controls whether or not nulls (zero bytes) in character string data should be returned from the server to the client, or should be converted to blanks. Both fixed length and variable length character data are affected by this keyword. If this keyword is set to YES, then nulls will be converted to blanks. If this keyword is set to NO, then nulls will not be converted.	NO
Convert Strings to Params	LGPA	This keyword controls whether or not long strings should be converted to LONG VARCHAR parameters. This conversion is needed because some applications produce literals that are longer than can be handled by the host database. If this keyword is set to NO, each SQL string will not be scanned for long strings. If this keyword is set to YES, then all long strings will be converted to parameter markers.	NO
Convert Timestamps	CVTS	This keyword controls whether or not timestamp values should be converted to other types. If this keyword is set to YES, timestamps that appear to be dates will be converted to dates and timestamps that appear to be times will be converted to times. These conversions are required to circumvent bugs in several products including MS Access and Crystal Reports. If this keyword is set to NO, timestamps will not be modified.	NO
Correlation ID Support	COID	This keyword controls whether or not SQLGetInfo should return information about Correlation IDs. If this keyword is set to YES, SQLGetInfo will return information about Correlation IDs. If this keyword is set to NO, no information about Correlation IDs will be returned. This keyword must be set to NO to enable some ODBC tools to work.	YES

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
Count() Fix Type	COFX	This keyword controls how COUNT (column name) SQL functions are fixed. Microsoft Access uses the COUNT(column name) function two ways, both of which are wrong. In some cases, COUNT(column name) really means COUNT (DISTINCT column name). In other cases it means COUNT(*). If this keyword is set to DISTINCT, the DISTINCT keyword will be inserted. If this keyword is set to ASTERISK, the column name will be replaced with a "*". If this keyword is set to NONE, no changes will be made. Note: COUNT(column name) functions are never changed unless Access compatibility mode is active.	DISTINCT
Create Table Index Automatically	CRIN	This keyword controls whether or not an index is automatically created when you create a table with a primary key or unique constraint. This index enforces the uniqueness. To disable this option, set the value to NO.	YES
Current Degree	SEDG	This keyword is used to set the initial Current Degree value on the host. The only possible values for this keyword are ANY or 1. No other values are supported at this time. If this keyword is set to a non-blank, non-null value, a SET CURRENT DEGREE statement is issued after DSNALI OPEN processing is completed.	
Current Packageset	SEPK	This keyword is used to set the initial Current Packageset value on the host. This value must be eighteen or fewer characters long. If this keyword is set to a non-blank, non-null value, a SET CURRENT PACKAGESET statement is issued after DSNALI OPEN processing is completed. The SET statement assigns the specified value to the CURRENT PACKAGESET special register.	
Current Rules	SERL	This keyword is used to set the initial Current Rules value on the host. The only possible values for this keyword are DB2 or STD. No other values are supported at this time. If this keyword is set to a non-blank, non-null value, a SET CURRENT RULES statement is issued after DSNALI OPEN processing is completed.	
Cursor Commit/Rollback Behavior	CRBH	This keyword controls what value to return for the cursor commit and rollback behavior SQLGetInfo request. DELETE will close cursors and delete prepared statements. To use the cursor again, the application must re-prepare and re-execute the statement. CLOSE will only close cursors. For prepared statements, the application can call SQLExecute on the statement without calling SQLPrepare again.	DELETE
Customer Secret Key Date Format	SKEY DTFM	No description is available. This keyword is used to specify how ODBC dates are converted to character strings. If this keyword is set to ODBC, the standard ODBC/ISO format yyyy-mm-dd will be used. If this keyword is set to UK, then the dd-mm-yyyy format will be used. If this keyword is set to EUR, then the dd.mm.yyyy format will be used.	ODBC

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
DB2 Version String	D2VR	This keyword can overwrite the host returned DB2 version string data. This string should be in the following format: x.y.z -- where x, y and z are all numeric digits. x is the DB2 version byte. y is the DB2 modification level. z is the DB2 release level. Both the version byte and the modification level must be set. If the release level is not set, then 0 is assumed. For example, 2.3 is for DB2 version 2, mod level 3, and 4.1.1 is for DB2 version 4, mod level 1, and rel 1.	
DBCS Mode	DBMD	This keyword is used to specify how DBCS data is handled. If BINARY is specified, all DBCS data will be treated as binary data. Note: DBCS data will still be converted from host formats to client formats, however in BINARY mode, data will be described as binary even though it is not stored in binary columns on the host. If CHARACTER is specified, all DBCS data will be treated as character data even if it is stored in GRAPHIC columns. DEFAULT and GRAPHIC are not supported at this time.	DEFAULT
DBCS Remove Blanks	DBQO	This keyword controls whether or not blanks are removed from within quoted strings for double byte languages (such as Chinese, Japanese, and Korean). If this keyword is set to YES, blanks, that are within quoted strings in SQL being sent to the host, will be removed. If this keyword is NO, blanks, that are inside quoted strings in SQL being sent to the host, will not be changed.	YES
Default Column Names	DFCL	This keyword controls whether or not the driver will assign default column names (of the form COLnnn, where nnn is the column number) when the DBMS does not return names for the columns.	NO
Decimal Point is Comma	DCCM	This keyword controls if the decimal point is a comma or not. If this keyword is set to YES, a comma is considered to be the decimal point. If this keyword is set to NO, a period is considered to be the decimal point. This keyword should only be set to YES, if the host uses a comma as the decimal point. This keyword will cause commas to be converted to periods before data is passed to the Auto-Static SQL conversion facility.	NO
Defer Prepare for CALLS	DFPR	This keyword controls whether or not to defer prepare of CALL statements. If this keyword is set to YES, prepare is always deferred for CALL statements. If this keyword is set to NO, prepare is only deferred for CALLs that have one or more parameter markers. In other words, prepare is always deferred for CALLs that have parameter markers. If this keyword is set to YES, prepare is deferred for CALLs that do not have parameter markers.	NO
Disable Async Option	NOAS	This keyword is used to indicate whether or not asynchronous execution will be allowed. Setting this option to YES will prevent all asynchronous processing.	YES
Disable CTL3D	NO3D	This keyword is used to indicate whether or not CTL3D processing will occur. Setting this option to YES will prevent the use of the CTL3D DLL with the connection dialog. This option must be set to YES for some applications if an old copy of CTL3D is installed.	NO

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
Disable Prepare/Open Optimization	DIPO	This keyword disables the prepare/open optimization and consequently adds a network round-trip (for the SQLExecute). The RDO layer on top of VB issues two SQLPrepares and only one SQLExecute for each select statement. This causes two result-sets to be returned for every query.	NO
Disable SQLCancel Tracing	DIPO	This keyword disables the prepare/open optimization and consequently adds a network round-trip (for the SQLExecute). The RDO layer on top of VB issues two SQLPrepares and only one SQLExecute for each select statement. This causes two result-sets to be returned for every query.	NO
Disable All Prompts	NOPM	This keyword when set to YES will disable all interactive prompts or informational message boxes. This feature is required when the driver is being called from a NT service, a UNIX daemon process, or any server type applications which can not be interrupted.	NO
Display Wait Cursor	WACU	This keyword indicates whether or not to update the cursor (with an hourglass). You may observe a dramatic improvement in performance by setting this option to NO, which results in no updates to the cursor and no display of the hourglass.	NO
Don't Get Passwords from ODBC.INI	NOPW	This keyword controls whether or not the password should be obtained from the ODBC.INI file. If this keyword is set to YES, then the password must be provided by the connection string or the connection dialogs. Note: If this keyword is set to YES, any password value in the ODBC.INI will be ignored. If this keyword is set to NO, the password will be obtained from the ODBC.INI file in some cases.	NO
Don't Get Userids from ODBC.INI	NOUS	This keyword controls whether or not the userid should be obtained from the ODBC.INI file. If this keyword is set to YES, the userid must be provided by the connection string or the connection dialogs. Note: If this keyword is set to YES, any userid value in the ODBC.INI will be ignored. If this keyword is set to NO, the userid will be obtained from the ODBC.INI file in some cases.	NO
DTS Plan File	DSFL	This keyword is used to specify the name of the local plan file used to convert dynamic SQL to static SQL. The plan file name should end with .pln and can be a full path name, if need be. This keyword should only be used if the local plan file does not have the default name or is not located in the working directory of the application. The default name of the local plan file is always planname.pln, where planname is the name of the DB2 plan used by the application.	

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
Duplicate Socket Descriptor	DPSO	This keyword is a non-negative integer that is used to return a new and unused file descriptor for the connecting socket. The new socket number will be a descriptor number that is either equal to or greater than this number. The old descriptor will be closed after the replacement. This keyword is only enabled for the UNIX platforms, and is designed to allow C run-time applications to get around the 255 maximum open file handle limitation in studio library. For example, a good value to set this keyword to is 256.	0
Enable Multitasking	MT	This keyword is used to control the multitasking capabilities of Shadow Direct. This feature will keep your screen from locking up while ODBC calls are pending. However, this feature should be used with extreme caution! Many applications will not operate correctly with this feature activated.	NO
Extended Cursor Pool	EXCU	This keyword controls whether or not the Extended Cursor Pool should be used. If this keyword is set to YES, the Extended Cursor Pool will be used. The Extended Cursor Pool is much larger than the standard cursor pool. Note: The Extended Cursor Pool can not be used with Auto-Static SQL. If this keyword is set to NO, the normal cursor pool will be used.	NO
Fast Logon	FALG	This keyword controls whether or not Fast Logon should be used. If this keyword is set to YES, Logon processing will be done with a minimum of network I/O. This is only possible if the host server supports Fast Logon. Failures will result if Fast Logon is requested and the host server does not support it. The ability of the server to handle Fast Logon must be verified before this option is used. If this keyword is set to NO, then normal Logon processing will be used	NO
Fix INSERT Statements	FXIS	This keyword controls whether or not the VALUES clause of each INSERT statement should be scanned for dates, times, and timestamps, and whether or not quotes should be added. If this keyword is set to YES, quotes will be added to dates, times, and timestamps as need be. If this keyword is set to NO, INSERT statements will be not modified.	NO
Fix DB2 Outer Joins	OJFX	If this keyword is set, parentheses around ON clauses of DB2 outer joins will be removed. While this syntax is valid ANSI SQL, DB2 cannot handle it in its current release (version 5). This kind of SQL is generated by ODBC tools such as Brio.	NO
Fix Oracle 8.0.5 Bug	OJFX	This field is used to fix the oracle8.0.5 bug.	NO
Fix String Length	STFX	This keyword fixes a specific bug that can occur in some products, like Microsoft Access, where the string length is incorrectly set to zero when the SQL_CHAR keyword only contains blanks. If the keyword AF (MS Access Compatibility) is set to YES, this keyword does not need to be set. Keyword AF includes the functionality of keyword STFX.	NO
Floating to Character Digits	MXDG	This keyword controls the maximum number of digits that should be generated for a floating point to character conversion. This keyword should only be specified if the default value is not acceptable.	6

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
GAIJI Extension Support	GAJI	This keyword enables GAIJI extension support in a codepage	YES
GAIJI Extension Table Name	GATB	This keyword sets the GAIJI extension table name. There is no default name. If this keyword is left blank, a server-side default table name will be used. This keyword has a max keyword size of 4 characters.	
German NLS support	NLGR	This keyword is set to YES to resolve certain problems handling SQL in the German language environment. This keyword should not be set for any other reason.	NO
Host Debugging Values	HODB	This keyword is used to control the debugging of host programs (generally Stored Procedures). This keyword is either set to the type of language the host program was written in, or NONE. The host program language must be specified correctly so that the program can be invoked with the correct debugging options. The supported values are COBOL, PLI, C, and C++.	NONE
Host User Parm	USERPARM	This keyword is sent to the host as part of the logon information. The host uses this keyword to complete logon to the host security system or host databases.	
Identifier Quote Option	IDQO	This keyword is used to show what identifier quote character should be returned to the application using SQLGetInfo with SQL_IDENTIFIER_QUOTE_CHAR. This keyword is needed to fix certain application bugs.	BLANK
Ignore High Bound Column Errors	IGHI	This keyword controls whether or not SQLFetch should ignore high bound column errors. If this keyword is set to YES, SQLFetch will ignore column binding errors. If this keyword is set to NO, column binding errors will be handled normally. This keyword must be set to YES to enable some ODBC tools (such as Microsoft Excel) to work.	NO
Ignore Underscore Characters	IGUN	This keyword controls whether or not the underscore character should be considered as a wild card or a regular character. If this keyword is set to YES, underscore (“_”) will be handled as a normal byte. If this keyword is set to NO, underscore will be treated as a wildcard character that matches any other single byte. This keyword is used to solve problems in which a table or procedure name actually has an underscore in the name and the underscore is misinterpreted as a wildcard.	NO
Keep Literal Quotes	KEQU	This keyword is used to keep quotes around and embedded in string literals. If this keyword is set to YES, quotes around and embedded in string literals will be retained. Note: This only applies to string literals passed to Stored Procedures, including MDI Stored Procedures. If this keyword is set to NO, then quotes will be removed from string literals.	NO

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
Language ID	LGID	This keyword is used to specify the language to be used. The possible values are Arabic (ARB), Simplified Chinese (CHS), Traditional Chinese (CHT), Danish (DAN), Default (DFT), German (DEU), U.S. English Compatibility (ENC), U.K. English (ENG), U.S. English (ENU), Modern Spanish (ESN), Castilian Spanish (ESP), Finish (FIN), French (FRA), Canadian French (FRC), Icelandic (ISL), Italian (ITA), Japanese (JPL), Japanese Latin Extended (JPX), Korean (KOR), Micro Decisionware (MDI), Dutch (NLD), Norwegian (NOR), PeopleSoft English (PPS), Portuguese (PTG), Swedish (SVE) and Turkish (TUR).	ENU
Long Data Fix	LGFX	This keyword controls whether or not a large number should be returned for the length and precision of LONG VARCHAR keywords, rather than the actual length and precision. This fix is needed to resolve certain problems in ODBCdirect GetChunk processing. If this keyword is set to YES, a large number will be returned for the length and precision of LONG VARCHAR keywords. If this keyword is set to NO, the actual values will be returned for the length and precision of LONG VARCHAR keywords.	NO
Lotus Approach Compatibility	LAFX	This keyword is set to YES to resolve certain problems with Lotus Approach. This keyword should not be set for any other reason.	NO
LZ Compression	LZCM	If this keyword is set to YES and the server supports this feature, a variant of Lempel-Ziv compression will be used to compress all data flow between the client and the server.	YES
Maximum Buffer Size	MXBU	This keyword sets the driver side maximum communication buffer size for all data exchanges. The default for this keyword is 0 which currently sets the maximum buffer size to 100,000 bytes (32K bytes in 16-bit windows). Note: This value will be used to negotiated with the server-side limit. The actual runtime buffer size could be smaller, but not larger.	0
Message Type	MGTY	This keyword controls the type of messages used for message oriented connections. It is only used if a message connection mode is employed. The default value is NETWORK, which creates a TCP/IP or LU 6.2 session or conversation for each message. When the link type is set to TCP/IP, users can also choose the UDP message type as the messaging protocol. Using UDP datagrams for messaging usually incurs lower network overhead, however UDP communication is not reliable and is discouraged when connecting over a WAN. HTTP and HTTPS methods will tunnel the communication session inside the HTTP protocol stream to allow the session to be established over firewalls and HTTP proxies. HTTPS is the Netscape secured HTTP protocol which implements SSL and can be used when a secured channel between the two endpoints is required.	NETWORK

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
MS Access Compatibility	AF	If this keyword is set to YES, all host decimal data values are returned to the ODBC client application without any trailing zeroes to the right of the decimal point. This keyword must be set to YES in order for Microsoft Access to work with DB2 decimal data. This keyword may be needed with any other product that uses the Microsoft Jet Database Engine, such as Visual Basic.	NO
MTS Security SID Type	MSMT	This keyword is the security SID option when using the MS Transaction Server based network authentication/single sig-on model.	DIRECT CALLER SID
MultiNet TCP/IP Compatibility	MUNT	This keyword is needed to support applications using the Multi-Net TCP/IP product for host access.	NO
Network Authentication	SECU	When the DOMAIN-base authentication is selected, the driver will verify if the userid associated with the current process has been authenticated by a domain-based system. For the Win32 platforms, this requires the user to first log on to a NT domain. For the UNIX platforms, the local machine must be a member of a NIS domain, and the password database used to authenticate the user must be NIS-mapped. The MTS-based authentication model allows a MTS server COM component to pass on the account name of the client applications invoking the object. MyNet setting is used with the CKS MyNet single signon product support. The SNA setting works with the MS SNA server single signon facility. The ODBC client must be a MS/SNA CPI-C client.	NONE
No Nulls	NONL	This keyword controls whether or not zero length strings should be returned. If this keyword is set to YES, zero length strings will be replaced by one blank. If this keyword is set to NO, zero length strings will not be modified.	NO
Non-blocking Network I/O	NBIO	This keyword controls whether or not TCP/IP network I/O operations should be blocking. If this keyword is set to YES, non-blocking network I/O operations will be used and the Windows message queue will be used to wait for I/O operation completion. If this keyword is set to NO, blocking network I/O operations will be used.	NO
ODBC Client Code Page	ODPG	This keyword is used to specify the ODBC client code page. The default is a set of UNIX code pages for UNIX and Windows Latin 1 (ANSI) for Windows 95/NT. Windows Latin 1 is also known as ISO 8859. Specifying LATIN1 will force the use of the Windows Latin 1 code page in any environment. Specifying UNIX will force the use of the UNIX code pages in any environment.	DEFAULT
One-Way Messages	ONWY	This keyword controls whether or not responses to messages should be sent to the ODBC client. If this keyword is set to YES, one-way message processing will be used. This means that no responses will be returned to ODBC client applications. If this keyword is set to NO, normal responses will be returned for messages. Note: This keyword can only be used with messages.	NO

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
Operation Timeout Value	OPTM	This keyword controls the timeout value (in seconds) for all client network operations (i.e. query preparation, execution, retrieving result set) AFTER the connection has been established. This value should never be negative, but can be zero. This value has no effect in UDP messaging. Note: The use of this keyword may cause some unpredictable results in many environments. This option is only available in the TCP/IP and the MQ link types, but is applicable to both the permanent and the messaging modes.	0
Optimal Row Count	OPRW	This keyword limits the number of rows that will be returned from the host each time a request for rows is made. Since any number of requests for rows can be made for one query, this value will have no effect on the total number of rows returned by a query. This value can be used to control the size of the buffers used to return rows from the host.	0
Parameter Handling	PAHN	This keyword is used to control how parameters are handled when they are passed by ODBC applications. In some cases, special processing is needed to fix values passed by some applications. For example, Crystal Reports passes parameters using obsolete values that must be modified before they can be used. This keyword should be set to INPUT for Crystal Reports. This keyword should be set to LENGTH to fix some application parameter length errors. This value should be set to ADO to fix some ADO related errors. If this keyword is set to NONE, the parameter values passed by applications are not changed.	NONE
PowerBuilder Compatibility	PBFU	This keyword is set to YES to resolve certain problems with PowerBuilder. This keyword should not be set for any other reason.	NO
Procedure Name Filter	PRNF	This keyword is used to filter the procedure names returned by the product. This keyword is used only if it is set to a non-blank value and if the host application does not provide a procedure name filter string. The procedure name filter string from either the application or from this keyword restricts the information returned by procedure catalog inquiries (SQLProcedures and SQLProcedureColumns). Only rows that match the procedure name pattern provided will be returned. If this keyword is set to a single percent sign, then all procedures will be returned. There is no default value for this keyword.	
Procedure Owner Filter	PROF	This keyword is used to filter the procedure owners returned by the product. This keyword is used only if it is set to a non-blank value and if the host application does not provide a procedure owner filter string. The procedure owner filter string from either the application or from this keyword restricts the information returned by procedure catalog inquiries (SQLProcedures and SQLProcedureColumns). Only rows with owner ID's that match the procedure owner pattern provided will be returned. If this keyword is set to a single percent sign, then all procedures will be returned. This keyword defaults to the current userid or alternate userid.	

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
Procedure Owner Handling	PROW	This keyword is used to control how procedure owner values are returned to ODBC applications. If this keyword is set to NULL, the actual owner value is used as a prefix to the procedure name and the owner keyword is returned as a null value. If this keyword is set to EMPTY, the actual owner value is used as a prefix to the procedure name and the owner keyword is returned as an empty string. If this keyword is set to NONE, the procedure owner value is not changed.	NONE
Process Escapes	PRES	This keyword is used to turn off escape clause processing. PRES=NO will set escape clause processing off and may result in improved performance by reducing CPU overhead. This keyword is optional and, if not specified, will default to PRES=YES.	YES
Proxy Server Information	PXSR	When HTTP tunneling is used for messaging the server, and the HTTP stream must first connect to a proxy server, this keyword should be set to identify the proxy server's hostname (or IP address) and the port number. The accepted formats for this keyword are MY_PROXY:80 or 10.22.33.44:1200. The first part of the value is the hostname or the IP address of the proxy server and the second part is the port number on which the proxy server is listening. Note: For some proxy server setups, user authentication might be required.	
Proxy Server Userid	PXUS	When HTTP tunneling messaging mode is enabled, and the HTTP proxy server requires user authentication, this keyword allow the user to specify the proxy user ID (or name) string. At this moment, Basic HTTP Access authentication is the only method supported	
Proxy Server Password	PXPW	When HTTP tunneling messaging mode is enabled, and the HTTP proxy server requires user authentication, this keyword allows the user to specify the proxy user password string. At this moment, Basic HTTP Access authentication is the only method supported.	
Qualifier Name Separator	QUNA	This keyword is set to YES if SQLGetInfo to return a period when it is called to obtain the SQL_QUALIFIER_NAME_SEPARATOR. This keyword is provided to fix certain application bugs.	NO
RDO Compatibility	RDFX	This keyword is set to YES to resolve certain problems with Remote Data Objects. This keyword should not be set for any other reason.	NO
Read Only	RDON	This keyword is set to YES to make the data source read-only. Setting the data source to read-only will not actually prevent update operations. However, setting this keyword to read-only will prevent any index information from being returned to the application. This will generally prevent any updates from being attempted. Note: Setting this keyword to YES will greatly improve the performance of some applications using the standard DB2 catalogs.	NO
Result-set Cache Size	RSBK	This keyword indicates the number of SQL statements contained by the result-set cache. Note: A higher value translates into more disk space.	256

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
Result-set Caching	RSCH	If this keyword is set to YES, the ODBC driver will cache result-sets of SQL statements. When the same SQL statement is encountered, results are retrieved from a local cache if the underlying tables have not changed since the last cache update.	NO
Result-set Caching Interval	RSIN	This keyword specifies an interval in seconds, within which stale results are tolerated. If result-set caching is enabled, a previously cached SQL statement is encountered, and the number of seconds since the result-set was stored is within this interval, the result-set will be returned without host processing. A value of zero will guarantee no stale results: all statements will be submitted to the host, where the underlying tables will be checked for changes. A value of -1 will always retrieve from a local cache if one is available. Note: Tolerances greater than zero save network round-trips, but a value of zero is still useful since, if the underlying tables haven't changed, the result-set is still retrieved from a local cache.	0
Return code if no rows affected	NODA	This keyword controls what the return code should be set to, when an INSERT/UPDATE/DELETE does not affect any rows. If this keyword is set to NODATA, the return code from an SQL operation that does affect any rows will be SQL_NO_DATA_FOUND. If this keyword is set to INFO, the return code will be set to SQL_SUCCESS_WITH_INFO (in accordance with the ODBC specification). This keyword can also be set to SUCCESS or ERROR. In these cases, the return code will be SQL_SUCCESS or SQL_ERROR.	INFO
Return Global Tables	RTGL	This keyword controls whether or not global temporary tables will be returned as normal tables. If this keyword is set to YES, global temporary tables will be treated and described as normal tables. If this keyword is set to NO, global temporary tables will be handled as system tables. This keyword is provided to allow standard ODBC tools to be used with global temporary tables. Many of these tools bypass global temporary tables that are described as global temporary tables. However, the same tools will work correctly with global temporary tables if they are described as normal tables.	YES
Return Logon Messages	LGMG	This keyword controls whether or not non-error Logon messages should be returned to ODBC client applications. If this keyword is set to YES, and a message is returned to the application, the return code will be set to SQL_SUCCESS_WITH_INFO and the error message text will be available using the Solderer function (or some higher level API). If this keyword is set to NO, the message will be displayed using a dialog box.	NO

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
Return System Tables	RTSY	This keyword controls whether or not system tables (SYSIBM) will be returned as normal tables. If this keyword is set to YES, system tables will be treated and described as normal tables. If this keyword is set to NO, system tables will be handled as system tables. This keyword is provided to allow standard ODBC tools to be used with system tables. Many of these tools bypass system tables that are described as system tables. However, the same tools will work correctly with system tables if they are described as normal tables.	NO
Remove Equals	RMEQ	This keyword controls whether or not the equals byte should be removed from MDI keyword names as part of MDI RSP (Remote Stored Procedure) invocation. If this keyword is set to YES, the equals byte will be removed from each keyword name. If this keyword is set to NO, the equals will not be removed from the keyword name. Note: This keyword only applies to MDI RSPs invoked using TSQL 0/1/2 syntax. It does not apply to MDI RSPs invoked using the extended ODBC CALL syntax.	NO
SBCS Mode	SBMD	This keyword is used to specify how SBCS data is handled. SBCS strings are assumed to contain a mixture of single byte characters and double byte characters. Each set of DBCS characters starts with a SO and end with a SI byte. If BLANK is specified, all SO/SI bytes will be converted to blanks. If DELETE is specified, all SO/SI bytes will be deleted from each string. Note: DEFAULT is not supported at this time.	DEFAULT
Secured Channel Protocol	SSLX	This keyword allows the user to select a secured channel protocol to encrypt and authenticate the client/server session. The default is the standard NEON secured data stream which provides optimal performance but does not support strong encryption and authentication. SSL2 (SSL version 2) and SSL3 (SSL version 3) are based on the SSL standard.	STANDARD
SQLGetData out of order	GTDA	This keyword is used to control whether or not SQLGetData calls should be allowed out of order. Normally, SQLGetData calls can only be made in ascending column order starting after the last bound column. However, if this keyword is set to YES, SQLGetData calls can be made in any order. If this keyword is set to NO, the normal restrictions on SQLGetData calls are enforced.	NO
SQL ParamOptions support	PAOP	This keyword controls whether or not we should support SQLParamOptions in a limited form. If this keyword is set to YES, SQLParamOptions is supported to a limited extent. If this keyword is set to NO, SQLParamOptions is not supported at all. If this keyword is set to SERVER, the server codes will decide whether this support is enabled.	YES
SQL Server Support	SQVA	This keyword is used to control whether or not SQLTypeInfo calls should show that variable data is 255 bytes long. DB2 variable data is actually only 254 bytes long. However, some applications do not support this. If this keyword is set to YES, SQLTypeInfo will return 255 as the length of variable data. If this keyword is set to NO, SQLTypeInfo will return 254. NOTE: If this keyword is set to YES, the behavior of SQLCancel is modified to circumvent vendor implementation errors.	NO

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
Suppress Decimal Trailing Zeros	STZO	This keyword controls whether or not trailing zeros should be removed from decimal keywords. The default is NO. If the keyword AF (MS Access Compatibility) is set to YES, this keyword does not need to be set. Keyword AF includes the functionality of keyword STZO.	NO
Sybase Compatibility	SYFU	If this keyword is set to YES, a set of Sybase fixes will be used. These fixes bypass bugs in the Sybase ODBC support. If this keyword is set to NO, the Sybase fixes will not be used.	NO
System Engineering Value	SEVL	The System Engineering Value is set to various values to obtain diagnostic and debugging data. This keyword should only be used at the specific request of the Technical Support staff.	0
Table Name Filter	TBFL	This keyword is used to filter the table names returned by the product. This keyword is only used if it is set to a non-blank value and if the host application does not provide a table name filter string. The table name filter string from either the application or from this keyword restricts the information returned by catalog inquiries (SQLTables, SQLColumns, SQLTablePrivileges). Only rows that match the table name pattern provided will be returned. If this keyword is set to a single percent sign, all tables will be returned. There is no default value for this keyword. Example: TBFL=STAFF will cause only information from the table labeled STAFF to be displayed; TBFL=STA% will cause information from all tables that begin with STA to be displayed; TBFL=STAFF BOOK% will cause information from the table labeled STAFF and all tables that begin with BOOK to be displayed.	
Table Owner/Synonym Option	SYOP	This keyword is used to show how table owners should be handled for tables that are actually synonyms. Some of these choices are required to make specific desktop productivity tools work with the NEON ODBC driver.	ZERO
Table Qualifier Option	TQOP	This keyword is used to specify how Table Qualifiers should be returned to ODBC application programs. Some databases use Table Qualifiers as part of the overall name of each table. In other words, two tables can have exactly the same name in all other respects, if their Table Qualifiers are different. The TQOP option is used to control how Table Qualifier information is returned to ODBC applications. There are three possible values: NORMAL, NULL, and ZERO.	NORMAL
Transaction Name	TRNA	This keyword is used to specify the transaction name. The use of the transaction name is application specific. It can be up to eight (8) bytes long. The transaction name is padded on the right with blanks.	
Truncate Literal String	TRLT	This keyword is used to turn off literal string truncation. TRLT=NO will not truncate the literal string even if the string length is greater than 20.	YES

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
Unload Communication DLLs	FEDL	This keyword is used to determine whether the communication DLL (i.e. WINSOCKET.DLL or CPIC.DLL) should be unloaded after an application disconnects. By default, Shadow Direct always unloads the communication DLL after the application disconnects. Setting this keyword to NO, if unloading the communication DLL while the communication service is still active, causes errors. Setting this keyword to NO will prevent the DLL from being unloaded by the ODBC driver.	YES
Uppcase all character data	UPCH	This keyword controls whether or not all character data sent to the host should be converted to upper case. If this keyword is set to YES, all character data will be converted to upper case. If this keyword is set to NO, character data will not be converted to upper case.	NO
Uppcase double quote string	UPDB	This keyword controls whether or not strings in double quotes should be converted to upper case. If this keyword is set to YES, strings in double quotes will be converted to upper case. If this keyword is set to NO, strings in double quotes will not be modified. Strings must be converted to upper case in some cases because DB2 treats table names in double quotes as literals.	NO
UDP Maximum Retry Attempts	MXRT	This keyword controls the maximum number of UDP retries for an unacknowledged request. This value must be a non-negative number. If this keyword is set to ZERO, the UDP request retry is disabled.	10
UDP Datagram Retry Timeout	RTTM	This keyword controls the initial timeout period before an UDP request is retransmitted. This timeout value is doubled on successive retries; however, this value will not increase without bound. The maximum timeout allowable is 30 seconds. If the maximum retry attempts is set to ZERO, this keyword will have no effect.	30
Visual Age Compatibility	VAFX	This keyword is set to YES to resolve certain problems with Visual Age. This keyword should not be set for any other reason.	NO
Wollongong Pathworks 3.1 Support	WGPH	This keyword is needed when running in the Wollongong Pathworks 3.1 16-bit TCP/IP environment	NO
X/OPEN XA Support	XAEN	This keyword should be set if the target server will participate in an X/OPEN XA transaction. The value, TWO-PHASE, should be used in a situation where the transaction involves more than one resource manager, and the target server supports the two-phase commit protocol. The value TWO-PHASE also allows for the one-phase commit scenario whenever the coordinator thinks it is appropriate. The default is NONE, which implies that no part of the transaction will use the XA protocol.	NONE
X/OPEN XA Transaction Manager	XAOP	This keyword is used to set the XA Transaction Manager type. This option must be set correctly when participating in a distributed transaction coordinated by a monitor. MSDTC is the Microsoft Distributed Transaction Manager used mostly by the MTS applications. Two possible values are available for a BEA Tuxedo coordinated transaction. TUXEDO-TMS is to identify that the ODBC connection will be owned by the Tuxedo TMS server. TUXEDO-SQL is to identify that the ODBC connection will be owned by the Tuxedo SQL application server. The default is NONE which means that this is not a distributed transaction.	NONE

Table 31. Driver keywords (continued)

Description	Keyword	(Format) Explanation	Default
Zero Active Statements	ZEAS	This keyword controls whether or not zero should be returned for the number of active statements. If this keyword is set to YES, zero will be returned for the number of active statements supported by the ODBC driver. If this keyword is set to NO, the actual number of supported active statements will be returned.	NO
Zero Column Names	ZECL	This keyword is set to YES if column names should be set to binary zeros. This is a performance optimization for production applications.	NO

Driver keyword settings

Since all of keywords are not able to be set using all the methods described in the previous section, the following table references how they can be set, and whether the keyword is optional (if it is checked) or not (in other words, whether it needs to be included in either the ODBC.INI file or a connection string).

Table 32. Driver keywords and settings

Keyword	ODBC.INI	Connect String
ALCD	X	X
ALUS	X	X
APPC	X	X
APPL	X	X
AT	X	X
AF	X	X
BMPA ¹	X	X
BYDB	X	X
CD	X	X
CGFX	X	X
CLOR	X	X
CNDG		X
CPFX		X
CTLV		X
DBD	X	X
Description	X	X
DP	X	X
DRIVER		X
DSFL	X	X
DSN		X
FEDL	X	X
HD	X	X
HOST	X	X
IDQO	X	X
LAFX	X	X
LGID	X	X
LGPA	X	X
LINK	X	X
LKES	X	X
LKTH	X	X
LLUA ²	X	X
LNDN		X
LNID		X

Table 32. Driver keywords and settings (continued)

Keyword	ODBC.INI	Connect String
MODE ²	X	X
MR	X	X
MT	X	X
MXBU	X	X
NB10		
NEONTRACE	X	X
NOAS	X	X
NO3D	X	X
NOUS	X	X
NOPW	X	X
NWPW		X
PBFU	X	X
PGNA ³	X	X
PGPA	X	X
PLAN	X	X
PLUA ⁴	X	X
PORT ⁵	X	X
PRES	X	X
PSB ⁶	X	X
PWD	X	X
QUNA	X	X
RDON	X	X
RO	X	X
SECU	X	X
SIDEINFO	X	X
SRN	X	X
SUBSYS	X	X
SYOP	X	X
TQOP	X	X
UID	X	X
USERPARM	X	X
VAFX	X	X
WACU	X	X
ZECL	X	X

1. The BMP name keyword is required for single-threaded IMS connections.
2. This field is required for OS/2 and LU 6.2 host connections.
3. The PGNA keyword is required for RPC connections to the host.
4. This field is required for OS/2 and LU 6.2 host connections.
5. The port number is required for TCP/IP host connections.
6. The PSB name is required for multi-threaded IMS host connections.

Appendix D. Business object samples

This appendix details the sample business objects that are included with the connector for MAS. Business object samples are given for the following connectors:

- DB2 Connector
- VSAM Connector
- NATURAL Connector
- CICS Connector
- IDMS(DB) Connector
- ADABAS Connector
- IMS_TM Connector

DB2 Connector

Business object used for testing:

```
[ReposCopy]
Version = 3.0.0
[End]
[BusinessObjectDefinition]
Name = DB2XWORLDS_EVENTS
Version = 3.0.0
AppSpecificInfo = TN=SCJAB.XWORLDS_EVENTS
```

```
[Attribute]
Name = EVENT_ID
Type = Integer
MaxLength = 255
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=EVENT_ID
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = OBJECT_KEY
Type = Integer
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=OBJECT_KEY
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = OBJECT_NAME
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=OBJECT_NAME
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = EVENT_PRIORITY
```

```

Type = Integer
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=EVENT_PRIORITY
IsRequiredServerBound = false
[End]

[Attribute]
Name = EVENT_STATUS
Type = Integer
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=EVENT_STATUS
IsRequiredServerBound = false
[End]

[Attribute]
Name = OBJECT_VERB
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=OBJECT_VERB
IsRequiredServerBound = false
[End]

[Attribute]
Name = ObjectEventId
Type = String
MaxLength =
IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]

[Verb]
Name = Create
[End]

[Verb]
Name = Delete
[End]

[Verb]
Name = Retrieve
[End]

[Verb]
Name = Update
[End]
[End]

```

VSAM Connector

Business object used for testing:

```

Version = 3.0.0
[End]
[BusinessObjectDefinition]
Name = VSAM_DEMOVSAM
Version = 3.0.0

```



```

AppSpecificInfo = TN=DEMOVSAM

[Attribute]
Name = FILEREC
Type = String
MaxLength = 80
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=FILEREC
IsRequiredServerBound = false
[End]

[Attribute]
Name = STAT
Type = String
MaxLength = 80
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=STAT
IsRequiredServerBound = false
[End]

[Attribute]
Name = NUMB
Type = String
MaxLength = 80
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=NUMB
IsRequiredServerBound = false
[End]

[Attribute]
Name = NAME
Type = String
MaxLength = 80
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=NAME
IsRequiredServerBound = false
[End]

[Attribute]
Name = ADDR
Type = String
MaxLength = 80
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=ADDR
IsRequiredServerBound = false
[End]

[Attribute]
Name = PHONE
Type = String
MaxLength = 80
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=PHONE
IsRequiredServerBound = false
[End]

```

```

[Attribute]
Name = AMOUNT
Type = String
MaxLength = 80
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=AMOUNT
IsRequiredServerBound = false
[End]

[Attribute]
Name = ObjectEventId
Type = String
MaxLength =
IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]

[Verb]
Name = Create
[End]

[Verb]
Name = Delete
[End]

[Verb]
Name = Retrieve
[End]

[Verb]
Name = Update
[End]
[End]

```

NATURAL Connector

Business object used for testing:

```

[ReposCopy]
Version = 3.0.0
[End]
[BusinessObjectDefinition]
Name = NaturalWrap
Version = 3.0.0
AppSpecificInfo = WRAPPER=true

[Attribute]
Name = TestNatural
Type = TestNatural
ContainedObjectVersion = 3.0.0
Relationship = Containment
Cardinality = N
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]

[Attribute]
Name = Func
Type = String

```

```
MaxLength = 255
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=Func
DefaultValue = SEND
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = SerInOut
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=SerInOut
DefaultValue= ACISP03,SP03MIN1,SP03MOT1
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = Param1
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=Param1
DefaultValue = FIND
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = Param2
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=Param2
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = ObjectEventId
Type = String
MaxLength =
IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]
```

```
[Verb]
Name = Create
[End]
```

```
[Verb]
Name = Delete
[End]
```

```
[Verb]
Name = Retrieve
[End]
```

```
[Verb]
```

```
Name = Update
[End]
[End]
```

```
[BusinessObjectDefinition]
Name = TestNatural
Version = 3.0.0
AppSpecificInfo = TN=Dummy
```

```
[Attribute]
Name = FuncC
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = true
IsRequired = false
AppSpecificInfo = CN=FuncC:FK=Func
DefaultValue = SEND
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = SerInpOutC
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = true
IsRequired = false
AppSpecificInfo = CN=SerInpOutC:FK=SerInpOut
DefaultValue = ACISP05B,SP03MIN1,SP03MOT1
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = Param1C
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = true
IsRequired = false
AppSpecificInfo = CN=Param1C:FK=Param1
DefaultValue = FIND
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = Param2C
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = true
IsRequired = false
AppSpecificInfo = CN=Param2C:FK=Param2
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = RetrieveSP
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = SPN=shadow_aci;RS=true;
IP=FuncC:SerInpOutC:Param1C:Param2C
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = REcInfo
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=RECORD_INFO
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = PersonnelId
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=PERSONNEL_ID
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = LastName
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=LAST_NAME
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = FirstName
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=FIRST_NAME
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = MiddleI
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=MIDDLE_I
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = Addr01
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=ADDRESS_LINE001
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = Addr02
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=ADDRESS_LINE002
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = Addr03
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=ADDRESS_LINE003
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = City
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=CITY
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = ZIP
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=ZIP
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = Dept
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=DEPT
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = JobTitle
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=JOB_TITLE
IsRequiredServerBound = false
[End]
```

```
[Attribute]
```

```

Name = ObjectEventId
Type = String
MaxLength =
IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]

[Verb]
Name = Create
[End]

[Verb]
Name = Delete
[End]

[Verb]
Name = Retrieve
[End]

[Verb]
Name = Update
[End]
[End]

```

CICS Connector

Business object used for testing:

```

Version = 3.0.0
[End]
[BusinessObjectDefinition]
Name = CICSUPD_WRAP
Version = 3.0.0
AppSpecificInfo = WRAPPER=true

[Attribute]
Name = child
Type = CICSUPD_CHILD
ContainedObjectVersion = 1.0.0
Relationship = Containment
Cardinality = N
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_IN_UPD
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=CA_IN_UPD
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_IN_LID
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = false

```

```

IsRequired = false
AppSpecificInfo = CN=CA_IN_LID
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_IN_NAME
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=CA_IN_NAME
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_IN_PHONE
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=CA_IN_PHONE
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_IN_EMAIL
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=CA_IN_EMAIL
IsRequiredServerBound = false
[End]

[Attribute]
Name = ObjectEventId
Type = String
MaxLength =
IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]

[Verb]
Name = Create
[End]

[Verb]
Name = Delete
[End]

[Verb]
Name = Retrieve
[End]

[Verb]
Name = Update
[End]
[End]
[BusinessObjectDefinition]
Name = CICSUPD_CHILD
Version = 3.0.0

```


AppSpecificInfo = TN=Dummy

```
[Attribute]
Name = DeleteSP
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = SPN=CICSEX.NESPUPD;RS=true;
  IP=CA_IN_UPD:CA_IN_LID:CA_IN_NAME:CA_IN_PHONE:CA_IN_EMAIL
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = UpdateSP
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = SPN=CICSEX.NESPUPD;RS=true;
  IP=CA_IN_UPD:CA_IN_LID:CA_IN_NAME:CA_IN_PHONE:CA_IN_EMAIL
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = CreateSP
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = SPN=CICSEX.NESPUPD;RS=true;
  IP=CA_IN_UPD:CA_IN_LID:CA_IN_NAME:CA_IN_PHONE:CA_IN_EMAIL
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = BeforeRetrieveSP
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = SPN=CICSEX.WBIRETRV;RS=false
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = CA_IN_UPD
Type = String
MaxLength = 2
IsKey = true
IsForeignKey = true
IsRequired = false
AppSpecificInfo = CN=CA_IN_UPD:FK=CA_IN_UPD
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = CA_IN_LID
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = true
IsRequired = false
```

```

AppSpecificInfo = CN=CA_IN_LID:CA_IN_LID
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_IN_NAME
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = true
IsRequired = false
AppSpecificInfo = CN=CA_IN_NAME:CA_IN_NAME
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_IN_PHONE
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = true
IsRequired = false
AppSpecificInfo = CN=CA_IN_PHONE:CA_IN_PHONE
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_IN_EMAIL
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = true
IsRequired = false
AppSpecificInfo = CN=CA_IN_EMAIL:CA_IN_EMAIL
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_MESSAGE
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = true
AppSpecificInfo = CN=CA_MESSAGE
DefaultValue = test
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_STATUS
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=CA_STATUS
IsRequiredServerBound = false
[End]

[Attribute]
Name = ObjectEventId
Type = String
MaxLength =
IsKey = false
IsForeignKey = false
IsRequired = false

```

```

IsRequiredServerBound = false
[End]

[Verb]
Name = Create
[End]

[Verb]
Name = Delete
[End]

[Verb]
Name = Retrieve
[End]

[Verb]
Name = Update
[End]
[End]

[BusinessObjectDefinition]
Name = CICSEX_RET_WRAP
Version = 3.0.0
AppSpecificInfo = WRAPPER=true

[Attribute]
Name = Child
Type = CICSEX_RET
ContainedObjectVersion = 1.0.0
Relationship = Containment
Cardinality = N
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_IN_LID
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=CA_IN_LID
IsRequiredServerBound = false
[End]

[Attribute]
Name = ObjectEventId
Type = String
MaxLength =
IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]

[Verb]
Name = Create
[End]

[Verb]
Name = Delete
[End]

```

```

[Verb]
Name = Retrieve
[End]

[Verb]
Name = Update
[End]
[End]
[BusinessObjectDefinition]
Name = CICSEX_RET
Version = 3.0.0
AppSpecificInfo = TN=Dummy

[Attribute]
Name = RetrieveSP
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = SPN=CICSEX.NESPEL;RS=true;IP=CA_IN_LID
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_IN_LID
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = true
IsRequired = false
AppSpecificInfo = CN=CA_IN_LID:FK=CA_IN_LID
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_OUT_NAME
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=CA_OUT_NAME
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_OUT_PHONE
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=CA_OUT_PHONE
IsRequiredServerBound = false
[End]

[Attribute]
Name = CA_OUT_EMAIL
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=CA_OUT_EMAIL
IsRequiredServerBound = false
[End]

```

```

[Attribute]
Name = ObjectEventId
Type = String
MaxLength =
IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]

[Verb]
Name = Create
[End]

[Verb]
Name = Delete
[End]

[Verb]
Name = Retrieve
[End]

[Verb]
Name = Update
[End]
[End]

```

IDMS(DB) Connector

Business object used for testing:

```

Version = 3.0.0
[End]
[BusinessObjectDefinition]
Name = IDMSBO
Version = 3.0.0
AppSpecificInfo = WRAPPER=true

[Attribute]
Name = Dummy
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=Dummy
IsRequiredServerBound = false
[End]

[Attribute]
Name = IDMSChildBO
Type = IDMSChildBO
ContainedObjectVersion = 1.0.0
Relationship = Containment
Cardinality = N
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]

[Attribute]
Name = ObjectEventId
Type = String
MaxLength =

```

```

IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]

[Verb]
Name = Create
[End]

[Verb]
Name = Delete
[End]

[Verb]
Name = Retrieve
[End]

[Verb]
Name = Update
[End]
[End]
[BusinessObjectDefinition]
Name = IDMSChildBO
Version = 3.0.0
AppSpecificInfo = TN=Dummy

[Attribute]
Name = RetrieveSP
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = SPN=SDCOIM;RS=true
IsRequiredServerBound = false
[End]

[Attribute]
Name = INVCODE
Type = Integer
MaxLength = 255
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=INVCODE
IsRequiredServerBound = false
[End]

[Attribute]
Name = PARTNUMBER
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=PARTNUMBER
IsRequiredServerBound = false
[End]

[Attribute]
Name = PARTDESC
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false

```

```

AppSpecificInfo = CN=PARTDESC
IsRequiredServerBound = false
[End]

[Attribute]
Name = ObjectEventId
Type = String
MaxLength =
IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]

[Verb]
Name = Create
[End]

[Verb]
Name = Delete
[End]

[Verb]
Name = Retrieve
[End]

[Verb]
Name = Update
[End]
[End]

```

ADABAS Connector

Business object used for testing:

```

Version = 3.0.0
[End]
[BusinessObjectDefinition]
Name = ADABAS_VEHICLES
Version = 3.0.0
AppSpecificInfo = TN=Vehicles

[Attribute]
Name = REG_NUM
Type = String
Cardinality = 1
MaxLength = 255
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=REG_NUM
IsRequiredServerBound = false
[End]

[Attribute]
Name = CHASSIS_NUM
Type = String
Cardinality = 1
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=CHASSIS_NUM
IsRequiredServerBound = false
[End]

[Attribute]

```

```
Name = PERSONNEL_ID
Type = Integer
Cardinality = 1
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=PERSONNEL_ID
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = MAKE
Type = String
Cardinality = 1
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=MAKE
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = MODEL
Type = String
Cardinality = 1
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=MODEL
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = COLOR
Type = String
Cardinality = 1
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=COLOR
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = YEAR
Type = Integer
Cardinality = 1
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=YEAR
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = CLASS
Type = String
Cardinality = 1
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
```



```
AppSpecificInfo = CN=CLASS
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = LEASE_PUR
Type = String
Cardinality = 1
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=LEASE_PUR
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = DATE_ACQ
Type = String
Cardinality = 1
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=DATE_ACQ
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = CURR_CODE
Type = String
Cardinality = 1
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=CURR_CODE
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = MAINT_COST_1
Type = Integer
Cardinality = 1
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=MAINT_COST_1
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = MAINT_COST_2
Type = Integer
Cardinality = 1
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=MAINT_COST_2
DefaultValue = 0
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = MAINT_COST_3
```

```

Type = Integer
Cardinality = 1
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=MAINT_COST_3
DefaultValue = 0
IsRequiredServerBound = false
[End]

[Attribute]
Name = ObjectEventId
Type = String
Cardinality = 1
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]

[Verb]
Name = Create
[End]

[Verb]
Name = Delete
[End]

[Verb]
Name = Retrieve
[End]

[Verb]
Name = Update
[End]
[End]

```

IMS_TM Connector

Business object used for testing:

```

Version = 3.0.0
[End]
[BusinessObjectDefinition]
Name = IMSTMBO
Version = 3.0.0
AppSpecificInfo = TN=Dummy

[Attribute]
Name = ADDFLD
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=ADDFLD
IsRequiredServerBound = false
[End]

[Attribute]
Name = CHGFLD
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false

```

```
IsRequired = false
AppSpecificInfo = CN=CHGFLD
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = REMFLD
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=REMFLD
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = REVFLD
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=REVFLD
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = D05751
Type = String
MaxLength = 255
IsKey = true
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=D05751
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = D05757
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=D05757
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = D05752
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=D05752
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = D05754
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
```

```
AppSpecificInfo = CN=D05754
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = D05755
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=D05755
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = D05756
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=D05756
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = D15333
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=D15333
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = D15334
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=D15334
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = D07786
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=D07786
IsRequiredServerBound = false
[End]
```

```
[Attribute]
Name = F00028
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = CN=F00028
```

```

IsRequiredServerBound = false
[End]

[Attribute]
Name = RetrieveSP
Type = String
MaxLength = 255
IsKey = false
IsForeignKey = false
IsRequired = false
AppSpecificInfo = SPN=ims.imssp;IO=ADDFLD:CHGFLD:REMFLD:REVFLD:
D05751:D05757:D05752:D05754:D05755:D05756:D15333:D15334:D07786:F00028
IsRequiredServerBound = false
[End]

[Attribute]
Name = ObjectEventId
Type = String
MaxLength =
IsKey = false
IsForeignKey = false
IsRequired = false
IsRequiredServerBound = false
[End]

[Verb]
Name = Create
[End]

[Verb]
Name = Delete
[End]

[Verb]
Name = Retrieve
[End]

[Verb]
Name = Update
[End]
[End]

```

Appendix E. Support for null and blank values

This appendix details different pass and fail scenarios where the key value in a business object is blank or null. This appendix also contains the functional changes required for blank or null business object values.

Pass and fail scenarios

If a key value in a business object is blank or has a null value in the database, then build the where clause with the "is null" type instead of the "=" operator type.

IBM recommends that business objects have at least one key attribute that does not have a blank value.

The following scenario is a parent object with one key that has a null value. This scenario fails under these conditions.

Table 33. Customer

Attribute	Type
cid	Integer (Key)
name	String
comments	String

The following scenario is a parent object with two keys and one key has a null value. This scenario passes under these conditions.

Table 34. Customer

Attribute	Type
cid	Integer (Key)
name	String
comments	String

In scenario two, build the retrieve query by selecting cid, name, and comments from customer, where cid=1000 and name is set to null.

The following scenario is a parent object with one child object in a container object with a foreign key reference. This scenario fails under these conditions.

Table 35. Customer

Attribute	Type
cid	Integer (Key)
name	String (Key)
comments	String
Address	Address
Aid	Integer (Key) ASI:FK=cid
Acity	String

Table 35. Customer (continued)

Attribute	Type
Azip	String

If cid contains a null value, then build the retrieve query by selecting Aid, Acity, and Azip from address. Set the value of Aid to null.

The following scenario is a parent object with one child object in a container object with two key references. This scenario passes under these conditions.

Table 36. Customer

Attribute	Type
cid	Integer (Key)
name	String
comments	String
Address	Address
Aid	Integer (Key) ASI:FK=cid
Acity	String (Key) ASI:FK=name
Azip	String

If name has a null value, then build the Retrieve query by selecting Aid, Acity, and Azip from address, where Aid=Cid and Acity has a null value.

Functionality

If the connector encounters a blank value on a key, it then compares that value with the UseNull value in the attribute. If the value is true, then the connector adds null value to the query. This affects the following verb operations:

- Retrieve
- RetrieveBy Content
- Update
- Delete

Notices

IBM may not offer the products, services, or features discussed in this document in all 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 give 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.

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 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 Burlingame Laboratory Director
IBM Burlingame Laboratory
577 Airport Blvd., Suite 800

Burlingame, CA 94010
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 measurement 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 necessarily 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.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples may 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.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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

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

IBM
the IBM logo
AIX
CrossWorlds
DB2
DB2 Universal Database
Domino
Lotus
Lotus Notes
MQIntegrator
MQSeries
Tivoli
WebSphere

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

MMX, Pentium, and ProShare are trademarks or registered trademarks of Intel Corporation in the United States, other countries, or both.

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

Other company, product or service names may be trademarks or service marks of others.

Mainframe Adapter Suite includes software developed by the Eclipse Project (<http://www.eclipse.org>).



IBM WebSphere Business Integration Adapter Framework V2.4.



Printed in USA